Approximating Extent Measures of Points*

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

February 9, 2003

Abstract

We present a general technique for approximating various descriptors of the extent of a set P of n points in $\mathbb{R}^d$. For a given extent measure μ and a parameter $\varepsilon > 0$, it computes in time $O(n + 1/\varepsilon^{O(1)})$ a subset $Q \subseteq P$ of size $1/\varepsilon^{O(1)}$, with the property that $(1-\varepsilon)\mu(P) \leq \mu(Q) \leq \mu(P)$. The specific applications of our technique include ε-approximation algorithms for (i) computing diameter, width, and smallest bounding box, ball, and cylinder of P, (ii) maintaining all the previous measures for a set of moving points, and (iii) fitting spheres and cylinders through a point set P. Our algorithms are considerably simpler, and faster in many cases, than the known algorithms.

1 Introduction

Motivated by a variety of applications, considerable work has been done on measuring various descriptors of the extent of a set P of n points in $\mathbb{R}^d$. We refer to such measures as *extent measures* of P. Roughly speaking, an extent measure of P either computes certain statistics of P itself or it computes certain statistics of a (possibly nonconvex) geometric shape (e.g. sphere, box, cylinder, etc.) enclosing P. Examples of the former include computing the kth largest distance between pairs of points in P, and the examples of the latter include computing the smallest radius of a sphere (or cylinder), the minimum volume (or surface area) of a box, and the smallest width of a slab (or a spherical or cylindrical shell) that contain P. Although P is assumed to be stationary in most of the work done so far, there has been some recent work on maintaining extent measures of a set of moving points [AGHV01].

*Research by the first author is supported by NSF under grants CCR-00-86013 EIA-98-70724, EIA-01-31905, and CCR-02-04118, and by a grant from the U.S.–Israel Binational Science Foundation. Research by the second author is supported by NSF CAREER award CCR-0132901.

A preliminary version of the paper appeared as: (i) P. K. Agarwal and S. Har-Peled, Maintaining the approximate extent measures of moving points, *Proc. 12th ACM-SIAM Sympos. Discrete Algorithms*, pages 148–157, 2001, and (ii) S. Har-Peled and K. R. Varadarajan, Approximate shape fitting via linearization, *Proc. 42nd Annu. IEEE Sympos. Found. Comput. Sci.*, pages 66–73, 2001.

†Department of Computer Science, Box 90129, Duke University, Durham NC 27708-0129; pankaj@cs.duke.edu; http://www.cs.duke.edu/~pankaj/

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

§Department of Computer Science, The University of Iowa Iowa City, IA 52242-1419; kvaradar@cs.uiowa.edu; http://www.cs.uiowa.edu/~kvaradar/

Shape fitting, a fundamental problem in computational geometry, computer vision, machine learning, data mining, and many other areas, is closely related to computing extent measures. A widely used shape-fitting problem asks for finding a shape that best fits P under some "fitting" criterion. A typical criterion for measuring how well a surface γ fits P, denoted as $\mu(P, \gamma)$, is the maximum distance between a point of P and its nearest point on γ, i.e., $\mu(P, \gamma) = \max_{p \in P} \min_{q \in \gamma} d(p, q)$. Then one can define the extent measure of P to be $\mu(P) = \min_\gamma \mu(P, \gamma)$, where the minimum is taken over a family of surfaces (such as points, lines, hyperplanes, spheres, etc.). For example, the problem of finding the minimum radius sphere (resp. cylinder) enclosing P is the same as finding the point (resp. line) that fits P best, and the problem of finding the smallest width slab (resp. spherical shell, cylindrical shell)[1] is the same as finding the hyperplane (resp. sphere, cylinder) that fits P best.

The exact algorithm for computing extent measures are generally expensive, e.g., the best known algorithms for computing the smallest volume bounding box or tetrahedron containing P in $\mathbb{R}^3$ require $O(n^3)$ time. Consequently, attention has shifted to developing approximation algorithms [BH01, ZS02]. Despite considerable work, no unified theory has evolved for computing extent measures approximately. Ideally, one would like to argue that for any extent measure μ and for any given parameter ε, there exists a subset $Q \subset P$ of size $1/\varepsilon^{O(1)}$ so that $\mu(Q) \geq (1 - \varepsilon)\mu(P)$. No such result is known except in a few special cases. It is known that an arbitrary convex body C can be approximated by a convex polytope Q so that the Hausdorff distance between C and Q is at most $\varepsilon \cdot \mathrm{diam}(\Gamma)$ and so that Q is either defined as the convex hull of a set of $1/\varepsilon^{O(1)}$ points or the intersection of a set of $1/\varepsilon^{O(1)}$ halfspaces. If the given extent measure μ of P is the same as that of $\mathcal{CH}(P)$, (e.g., diameter and width), then one can approximate $\mathcal{CH}(P)$ by Q, compute $\mu(Q)$, and argue that $\mu(Q)$ approximates $\mu(P)$. Although this approach has been used for computing a few extent measures of P [BH01, Cha02], it does not always work, especially if the extent measure μ is defined in terms of a nonconvex shape (such as spherical shell) containing P.

This paper is a step toward the aforementioned goal of developing a unified theory for approximating extent measures. We introduce the notion of an *ε-approximation* of a point set P. Roughly speaking, a subset $Q \subseteq P$ is called an ε-approximation of P if for every slab W containing Q, the expanded slab $(1 + \varepsilon)W$ contains P. One of the main results of the paper is an $O(n + 1/\varepsilon^{d-1})$-time algorithm for computing an ε-approximation of P of size $O(1/\varepsilon^{d-1})$ or an $O(n + 1/\varepsilon^{3(d-1)/2})$-time algorithm for computing an ε-approximation of size $O(1/\varepsilon^{(d-1)/2})$. We call an extent measure μ *faithful* if there exists a constant $\alpha > 0$ such that for any ε-approximation Q of P, $\mu(Q) \geq (1 - \alpha\varepsilon)\mu(P)$. Our algorithm for computing an ε-approximation immediately gives an $O(n + 1/\varepsilon^{O(1)})$ time algorithm for computing faithful measures approximately. In order to handle unfaithful measures, we introduce the notion of ε-approximations for a family of functions. Let $\mathcal{F}$ be a family of $(d-1)$-variate functions. We define the *extent* of $\mathcal{F}$ at a point $x \in \mathbb{R}^{d-1}$ to be $\mathfrak{I}_\mathcal{F}(x) = \max_{f \in \mathcal{F}} f(x) - \min_{f \in \mathcal{F}} f(x)$. We call a subset $\mathcal{G} \subseteq \mathcal{F}$ an ε-approximation of $\mathcal{F}$ if $\mathfrak{I}_\mathcal{G}(x) \geq (1 - \varepsilon)\mathfrak{I}_\mathcal{F}(x)$ for all $x \in \mathbb{R}^{d-1}$. Using our result on ε-approximation of points and the linearization technique, we show that we can compute in $O(n + 1/\varepsilon^{O(1)})$ time an ε-approximation of $\mathcal{F}$ of size $O(1/\varepsilon^{r\sigma})$ if each f_i is of the form $g_i^{1/r}$, where g_i is a polynomial, r is a positive integer, $\sigma = \min\{d - 1, k/2\}$,

[1] A *slab* is a region lying between two parallel hyperplanes; a *spherical shell* is the region lying between two concentric spheres; a *cylindrical shell* is the region lying between two coaxial cylinders.

and k is the dimension of linearization for g_i's (see Section 4 for the definition of k). Our algorithms for computing ε-approximations can be adapted to handle insertions and deletions of points (or functions) efficiently, see Section 5. If we only insert points, we can maintain an ε-approximation using only $(\log(n)/\varepsilon)^{O(1)}$ space.

We show that many extent-measure problems can be formulated as computing $\min_x \mathfrak{I}_{\mathcal{F}}(x)$, where $\mathcal{F}$ is obtained by transforming each input point to a function. Specific applications of our technique include the following:

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

$$\sigma(\ell, r, R) = \left\{p \in \mathbb{R}^d \mid r \leq d(x, p) \leq R\right\},$$

where $d(x, p)$ is the Euclidean distance between the points p and line x. The *width* of $\sigma(x, r, R)$ is $R - r$. In the *ε-approximate spherical-shell* problem, we are given a set P of n points and a parameter $\varepsilon > 0$, and we want to compute an spherical shell containing P whose width is at most $(1+\varepsilon)$ times the width of the minimum-width spherical shell containing P.

This problem, motivated by applications in computational metrology, has been widely studied; see [AAHS00, AS98, Cha02] and the references therein. The best known exact algorithm runs in $O(n^{3/2+\delta})$ time in $\mathbb{R}^2$, and in $O(n^{3-1/19+\delta})$ time in $\mathbb{R}^3$, for any $\delta > 0$. The best known ε-approximation algorithm, proposed by Chan [Cha02], takes $O(n+1/\varepsilon^{d^2/4})$ time. Our technique leads to an $O(n + 1/\varepsilon^{3d})$-time algorithm for the d-dimensional ε-approximate spherical-shell problem, thereby improving Chan's algorithm.

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 two co-axial cylinders of radii r and R with ℓ as their axis, i.e.,

$$\Sigma(\ell, r, R) = \left\{p \in \mathbb{R}^d \;\middle|\; r \leq d(\ell, p) \leq R\right\},$$

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.

Agarwal *et al.* [AAS01] 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, our technique gives an $(1+\varepsilon)$-approximation algorithm that runs in $O(n + 1/\varepsilon^{O(d^2)})$ time in $\mathbb{R}^d$, a significant improvement over their algorithm.

Maintaining faithful measures of moving points. Let P be a set of n points in $\mathbb{R}^d$, each point moving independently. Many applications call for maintaining extent measures of P as the points move with time. For example, various indexing structures, which

answer range-searching queries or nearest-neighbor queries on P, need an algorithm for maintaining the smallest orthogonal box containing P [AAE00, SJLL00, PAH02]. Agarwal *et al.* [AGHV01] have described kinetic data structures for maintaining a number of extent measures of points moving in the plane. They also show that most of these extent measures are expensive to maintain — the diametral pair of a set of points, each moving with a fixed velocity in the plane, can change $\Omega(n^2)$ times, and no subcubic bound is known on the number of triples defining the smallest enclosing ball of a set of points moving in the plane. This has raised the question whether faster approximation algorithms exist for maintaining an extent-measure of a set of moving points.

For any $\varepsilon > 0$, we say that a subset $Q \subseteq P$ ε-approximates P with respect to measure μ if $(1-\varepsilon)\mu(P(t)) \leq \mu(Q(t))$ for every t. We show that our techniques can compute an ε-approximation of size $1/\varepsilon^{O(1)}$ for numerous extent measures. For any set P of points in $\mathbb{R}^d$ with linear motion, our technique 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, directional width. If we want to maintain an ε-approximation of the smallest orthogonal box enclosing P, the size of Q can be reduced to $O(1/\sqrt{\varepsilon})$, for any fixed dimension. These results generalize to algebraic motion and to "non-convex" measures such as minimum-width spherical/cylindrical shell. Our scheme can also allow efficient insertions into and deletions from the set P. Note that the ε-approximation does not change with time unless the trajectory of a point changes. These results must be contrasted with the schemes for maintaining the exact extent measures, which require at least quadratic updates.

Maintaining faithful measures in a streaming model. Motivated by various applications, the need for analyzing and processing massive data in real time has led to a flurry of activity related to performing computations on a *data stream*. The goal is to maintain a summary of the input data using little space and processing time, as the data objects arrives. The efficiency of an algorithm in this model is measured in terms of the size of the working space and the time spent on performing the computation on a new data object. See [MP80, GKS01, GMMO00, KMS02, CDH+02] and references therein for recent algorithms developed in the data-stream model. Our technique can be adapted to maintain various extent measures approximately in the streaming model. Specifically, an ε-approximation of a stream of points in $\mathbb{R}^d$ can be maintained using a data structure of size $O(\log^d(n)/\varepsilon^{(d-1)/2})$. If we allow $O((\log^d n + 1/\varepsilon^{d-1})/\varepsilon^{(d-1)/2})$ (resp. $O(1/\varepsilon^{3(d-1)/2})$) amortized time to process each new point, the data structure can maintain an ε-approximation of size $O(1/\varepsilon^{(d-1)/2})$ (resp. $O(\log^d(n)/\varepsilon^{(d-1)/2})$). The same result holds for ε-approximations of linear functions. Consequently, we can maintain an ε-approximation of the diameter of a stream of points in $\mathbb{R}^d$ in amortized $O((\log^d n + 1/\varepsilon^{d-1})/\varepsilon^{(d-1)/2})$ time using $O(\log^d(n)/\varepsilon^{(d-1)/2})$ space.

The paper is organized as follows. In Section 2, we formally define ε-approximations for points and functions and make a few simple observations about them. In Section 3, we show that if $\mathcal{F}$ is a set of linear functions, there is a small subset $\mathcal{G} \subseteq \mathcal{F}$ whose extent approximates the extent of $\mathcal{F}$. Section 4 shows that this property is also true for polynomials and related functions, using linearization. Section 5 shows that our technique can be dynamized. In Section 6, we apply these ideas to the problems mentioned above.

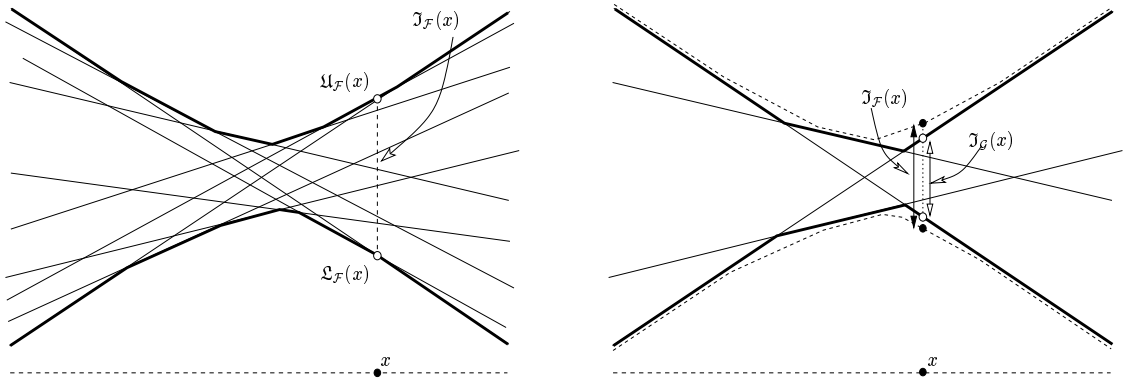

Figure 1. (i) Lower and upper envelopes and the extent of a family of linear functions; the extent at any point is the length of the vertical segment connecting lower and upper envelopes. (ii) An ε-approximation $\mathcal{G}$ of $\mathcal{F}$; dashed edges denote the envelopes of $\mathcal{F}$, and the thick lines denote the envelopes of $\mathcal{G}$.

2 Preliminaries

Envelopes and extent. Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a set of n $(d-1)$-variate functions defined over $x = (x_1, \ldots, x_{d-1}) \in \mathbb{R}^{d-1}$. The *lower envelope* of $\mathcal{F}$ is the graph of the function $\mathfrak{L}_{\mathcal{F}} : \mathbb{R}^{d-1} \to \mathbb{R}$ defined as $\mathfrak{L}_{\mathcal{F}}(x) = \min_{f \in \mathcal{F}} f(x)$. Similarly, the *upper envelope* of $\mathcal{F}$ is the graph of the function $\mathfrak{U}_{\mathcal{F}} : \mathbb{R}^{d-1} \to \mathbb{R}$ defined as $\mathfrak{U}_{\mathcal{F}}(x) = \max_{f \in \mathcal{F}} f(x)$. The *extent* $\mathfrak{I}_{\mathcal{F}} : \mathbb{R}^{d-1} \to \mathbb{R}$ of $\mathcal{F}$ is defined as

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

Let $\varepsilon > 0$ be a parameter, and let Δ be a subset of $\mathbb{R}^{d-1}$. We say that a subset $\mathcal{G} \subseteq \mathcal{F}$ is an *ε-approximation* of the extent of $\mathcal{F}$ within Δ if

$$(1-\varepsilon)\mathfrak{I}_{\mathcal{F}}(x) \leq \mathfrak{I}_{\mathcal{G}}(x)$$

for each $x \in \Delta$. Obviously, $\mathfrak{I}_{\mathcal{G}}(x) \leq \mathfrak{I}_{\mathcal{F}}(x)$, as $\mathcal{G} \subseteq \mathcal{F}$. If $\Delta = \mathbb{R}^{d-1}$, we say that $\mathcal{G}$ is an ε-approximation of the extent of $\mathcal{F}$.

Lemma 2.1 *Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a family of $(d-1)$-variate functions, $\varphi(x), \psi(x)$ two other $(d-1)$-variate functions, and $\varepsilon > 0$ a parameter. Let $\widehat{f}_i(x) = \varphi(x) + \psi(x) f_i(x)$, and set $\widehat{\mathcal{F}} = \left\{\widehat{f}_i \mid 1 \leq i \leq n\right\}$. If $\mathcal{K}$ is an ε-approximation of $\mathcal{F}$ within a region $\Delta \subseteq \mathbb{R}^{d-1}$, then $\widehat{\mathcal{K}} = \left\{\widehat{f}_i \mid f_i \in \mathcal{K}\right\}$ is an ε-approximation of $\widehat{\mathcal{F}}$ within Δ.*

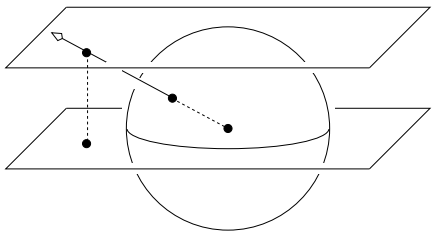

Figure 2. Representing a direction in $\mathbb{R}^d$.

Proof: For any $x \in \Delta$,

$$\begin{aligned}
(1-\varepsilon)\mathfrak{I}_{\widehat{\mathcal{F}}}(x) &= (1-\varepsilon)\left[\max_{\widehat{f}_i\in\widehat{\mathcal{F}}} \widehat{f}_i(x) - \min_{\widehat{f}_i\in\widehat{\mathcal{F}}} \widehat{f}_i(x)\right] \\
&= (1-\varepsilon)\left[\max_{f_i\in\mathcal{F}}(\varphi(x)+\psi(x)f_i(x)) - \min_{f_i\in\mathcal{F}}(\varphi(x)+\psi(x)f_i(x))\right] \\
&= (1-\varepsilon)\psi(x)\left[\max_{f_i\in\mathcal{F}} f_i(x) - \min_{f_i\in\mathcal{F}} f_i(x)\right] \\
&\leq \psi(x)\left[\max_{f_i\in\mathcal{K}} f_i(x) - \min_{f_i\in\mathcal{K}} f_i(x)\right] \\
&= \max_{f_i\in\mathcal{K}}(\varphi(x)+\psi(x)f_i(x)) - \min_{f_i\in\mathcal{K}}(\varphi(x)+\psi(x)f_i(x)) \\
&= \mathfrak{I}_{\widehat{\mathcal{K}}}(x).
\end{aligned}$$

Hence $\widehat{\mathcal{K}}$ is an ε-approximation of $\widehat{\mathcal{F}}$. ∎

Directions. With a slight abuse of notation, we will not distinguish between a vector in $\mathbb{R}^d$ and the corresponding point in $\mathbb{R}^d$. Let $\mathbb{P}$ denote the (projective) hyperplane $x_d = 1$ in $\mathbb{R}^d$, and let $\mathbb{S}^{d-1}$ represent the sphere of directions in $\mathbb{R}^d$. Normally, a direction in $\mathbb{R}^d$ is represented as a point in $\mathbb{S}^{d-1}$. However, we will not distinguish between directions $x \in \mathbb{S}^{d-1}$ and $-x \in \mathbb{S}^{d-1}$, therefore we can represent a direction $u^* \in \mathbb{S}^{d-1}$ as a point $u \in \mathbb{R}^{d-1}$, with the interpretation that $\widetilde{u} = (u, 1) \in \mathbb{P}$ is the central projection of the unit vector u^*; see Figure 2. Although this representation has the drawback that a direction lying in the plane $x_d = 0$ maps to a point at infinity and thus needs a special treatment, we will use this representation as it will be convenient for our applications and the drawback is a minor technicality that we can ignore. However, at some places we will also represent a direction as a point $x \in \mathbb{S}^{d-1}$. For an arbitrary nonzero vector $v \in \mathbb{R}^d$, we will use $\phi(v) \in \mathbb{S}^{d-1}$ to denote the direction corresponding to the unit vector $v/\|v\|$. Namely, for $u \in \mathbb{R}^{d-1}$, $u^* = \phi(\widetilde{u})$.

Directional width. We can define the concept of extent for a set of points. For any non-zero vector $x \in \mathbb{R}^d$ and a point set $P \subseteq \mathbb{R}^d$, we define

$$\overline{\omega}(x, P) = \max_{p\in P} \langle x, p\rangle - \min_{p\in P} \langle x, p\rangle ,$$

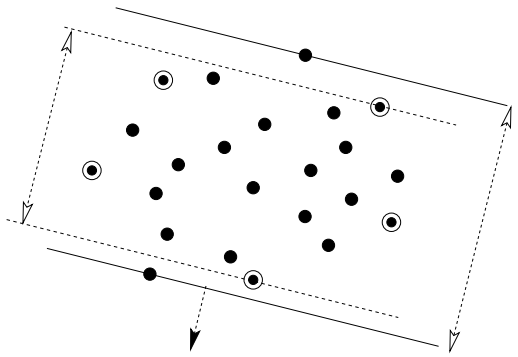

Figure 3. A point set P, its ε-approximation Q (points with double circles), and their directional widths.

where $\langle \cdot, \cdot \rangle$ is the inner product. For any set P of points in $\mathbb{R}^d$ and any $u \in \mathbb{R}^{d-1}$, we define the *directional width* of P in direction u, denoted by $\omega(u, P)$, to be

$$\omega(u, P) = \overline{\omega}(\widetilde{u}, P).$$

It is also called the *u-breadth* of P, see [GK92]. Let $\varepsilon > 0$ be a parameter, and let $\Delta \subseteq \mathbb{R}^d$. A subset $Q \subseteq P$ is called an *ε-approximation* of P within $\Delta \subseteq \mathbb{R}^{d-1}$ if for each $u \in \Delta$,

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

Clearly, $\omega(u, Q) \leq \omega(u, P)$. If $\Delta = \mathbb{R}^{d-1}$, we call Q an ε-approximation of P. Note that $\omega(u, Q) \geq (1 - \varepsilon)\omega(u, P)$ if and only if for every $0 \neq \lambda \in \mathbb{R}$, $\overline{\omega}(\lambda\widetilde{u}, Q) \geq (1 - \varepsilon)\overline{\omega}(\lambda\widetilde{u}, P)$.

Table 1 summarizes the notation used in this paper.

Arrangement. The *arrangement* of a collection $\mathcal{J}$ of m hyperplanes in $\mathbb{R}^d$, denoted as $\mathcal{A}(\mathcal{J})$, is the decomposition of the space into relatively open connected cells of dimensions $0, \ldots, d$ induced by $\mathcal{J}$, where each cell is a maximal connected set of points lying in the intersection of a fixed subset of $\mathcal{J}$. The complexity of $\mathcal{A}(\mathcal{J})$ is defined to be the number of cells of all dimensions in the arrangement. It is well known that the complexity of $\mathcal{A}(\mathcal{J})$ is $O(m^d)$ [AS00]. A set $\mathcal{J}$ of hyperplanes is *k-uniform* if $\mathcal{J}$ consists of k families of either parallel hyperplanes or hyperplanes that share a $(d-2)$-flat. In this case, each cell of $\mathcal{A}(\mathcal{J})$ has at most $2k$ facets. The notion of arrangement can be extended to a family of (hyper-)surfaces in $\mathbb{R}^d$. If $\mathcal{G}$ is a family of m algebraic surfaces of bounded maximum degree, then the complexity of the arrangement is $O(m^d)$.

Lemma 2.2 *For any $\varepsilon > 0$, there is a set $\mathcal{J}$ of $O(1/\varepsilon)$ $d(d-1)$-uniform hyperplanes in $\mathbb{R}^{d-1}$ so that for any two points u, v lying in the (closure of the) same cell of $\mathcal{A}(\mathcal{J})$,*

$$\|u^* - v^*\| \leq \varepsilon.$$

Proof: Partition the boundary of the hypercube $\mathbb{C} = [-1, +1]^d$ in $\mathbb{R}^d$ into small $(d-1)$-dimensional "hypercubes" of diameter ε, by laying a uniform $(d-1)$-dimensional axis-parallel grid on each facet of $\mathbb{C}$; see Figure 4 (i). Each such grid is formed by $(d-1)$ families of

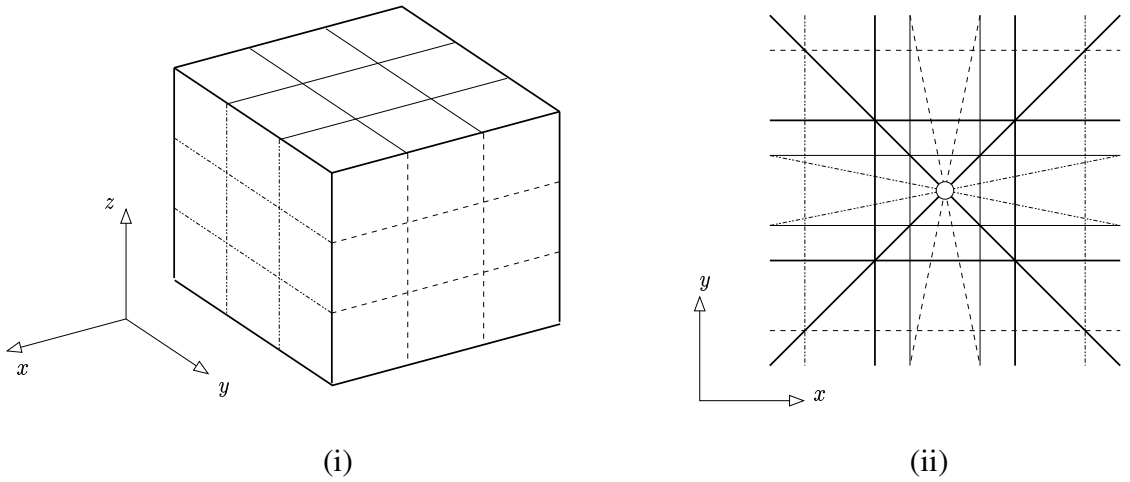

Figure 4. (i) Grid drawn on each facet of the unit cube $\mathbb{C}$ ($d = 3$). (ii) Lines in $\mathcal{J}$; thick lines correspond to the grid lines of the edges of $\mathbb{C}$, and solid (resp. dashed, dashed-dotted) lines correspond to the grid on the face normal to the z-axis (resp. y-axis, x-axis). The grid lines parallel to the z-axis on $\mathbb{C}$ map to the lines passing through the origin, and the grid lines parallel to the x-axis (resp. y-axis) map to lines parallel to the x-axis (resp. y-axis).

parallel $(d-2)$-flats. We extend each such $(d-2)$-flat f into a $(d-1)$-hyperplane $\hat{f}$, by considering the unique hyperplane that passes through it and the origin, and then intersect it with $\mathbb{P}$ (namely, the resulting $(d-2)$-flat lies on the projective hyperplanes $\mathbb{P} : x_d = 1$ in $\mathbb{R}^d$ and as such can be regarded as a hyperplane in $\mathbb{R}^{d-1}$). Because of symmetry, the $(d-2)$-flats $x_i = -1, x_j = \delta$ (i.e., the intersection of hyperplanes $x_i = -1$ and $x_j = \delta$) and $x_i = 1, x_j = -\delta$ map to the same hyperplane, so it suffices to extend the $(d-2)$-flats of the grid on the "front" facets of $\mathbb{C}$, i.e., the facets with $x_i = 1$ for $1 \leq i \leq d$. We claim that the resulting set composed of $d(d-1)$ families of uniform hyperplanes is the desired set of hyperplanes.

Formally, let $F(i, j, \beta)$ denote the $(d-2)$-flat $x_i = 1, x_j = \beta$ in $\mathbb{R}^d$. Set $\gamma = \left\lceil 4\sqrt{d}/\varepsilon \right\rceil$, and for integers i, j, l, let

$$\mathcal{F} = \{F(i, j, l/\gamma) \mid 1 \leq i \neq j \leq d, -\gamma \leq l \leq \gamma\}.$$

For a $(d-2)$-flat $F \in \mathcal{F}$ not passing through origin, let $\eta(F)$ be the $(d-2)$-hyperplane in $\mathbb{R}^{d-1}$ defined as

$$\eta(F) = \{x \in \mathbb{R}^{d-1} \mid (x, 1) \in \text{aff}(F \cup \{0\}) \cap \mathbb{P}\}.$$

In other words, $\eta(F)$ is the $(d-2)$-hyperplane in $\mathbb{R}^{d-1}$ corresponding to the intersection of $\mathbb{P}$ with the $(d-1)$-hyperplane $\text{aff}(F \cup \{0\})$. We set $\mathcal{J} = \{\eta(F) \mid F \in \mathcal{F}\}$. See Figure 4(ii). Clearly $\mathcal{J}$ is a $d(d-1)$-uniform family of hyperplanes because for any fixed pair i, j, either all hyperplanes $\eta(F(i, j, l/\gamma))$ in $\mathcal{F}$ are parallel or all of them pass through a $(d-3)$-flat.

Let f be a face of $\partial\mathbb{C} \cap \mathcal{A}(\mathcal{J})$. For any $u, v \in f$, we have $\|u\|, \|v\| \geq 1$ and $\|uv\| \leq \varepsilon/4$, which implies that $\|u^* - v^*\| \leq \angle uov \leq \varepsilon$. Hence, $\mathcal{A}(\mathcal{J}_{d-1})$ is the required partition. ∎

Remark 2.3 An interesting open question is to obtain a tight bound on the minimum number of uniform families of hyperplanes needed to achieve the partition of Lemma 2.2.

Agarwal and Matoušek [AM] has shown that the number of families is at least $2d-3$, and they conjecture this bound to be tight.

Duality. Let $\mathcal{H} = \{h_1, \ldots, h_n\}$ be a family of $(d-1)$-variate linear functions and $\varepsilon > 0$ a parameter. We define a *duality* transformation that maps the $(d-1)$-variate function (or a hyperplane in $\mathbb{R}^d$) $h : x_d = a_1x_1 + a_2x_2 + \cdots + a_{d-1}x_{d-1} + a_d$ to the point $h^\star = (a_1, a_2, \ldots, a_{d-1}, a_d)$ in $\mathbb{R}^d$. Let $\mathcal{H}^\star = \{h^\star \mid h \in \mathcal{H}\}$. The following lemma is immediate from the definition of duality.

Lemma 2.4 *Let $\mathcal{H} = \{h_1, \ldots, h_n\}$ be a family of $(d-1)$-variate linear functions and $\varepsilon > 0$ a parameter. A subset $\mathcal{K}^* \subseteq \mathcal{H}^*$ is an ε-approximation of $\mathcal{H}^*$ within a region $\Delta \subseteq \mathbb{R}^{d-1}$ if and only if $\mathcal{K}$ is an ε-approximation of $\mathcal{H}$ within Δ.*

3 Approximating the Extent of Linear Functions

In this section we describe algorithms for computing ε-approximations of the extent of a set of linear functions whose size depends only on ε and d. We first show that if we can compute an ε-approximation of (the directional width) of a "fat" point set contained in in the unit hypercube $\mathbb{C}$, then we can also compute an ε-approximation of an arbitrary point set. We then describe fast algorithms for computing ε-approximations of fat point sets. Finally, we use Lemma 2.4 to construct ε-approximations of the extent of linear functions.

Reduction to fat point set. We begin by proving a simple lemma, which will be crucial for reducing the problem of computing an ε-approximation to a fat point set.

Lemma 3.1 *Let $T(x) = Mx + b$ be an affine transformation from $\mathbb{R}^d$ to $\mathbb{R}^d$, where $M \in \mathbb{R}^{d\times d}$ is non-singular and $b \in \mathbb{R}^d$, let P be a point-set in $\mathbb{R}^d$, and let $\Delta \subseteq \mathbb{R}^{d-1}$. Define $\widehat{M}(\Delta) = \{u \in \mathbb{R}^{d-1} \mid \phi(M^T\widetilde{u}) \in \Delta^*\}$; where $\widetilde{u} = (u, 1)$ and $\phi(x) = x/\|x\|$, as defined above, and $\Delta^* = \{x^* \mid x \in \Delta\}$. Then $Q \subseteq P$ ε-approximates P if and only if $T(Q)$ ε-approximates $T(P)$ within $\widehat{M}(\Delta)$.*

Proof: For any vectors $t \in \mathbb{R}^d$ and $u \in \mathbb{R}^{d-1}$, it is easily seen that $\omega(u, P) = \omega(u, P + t)$. Therefore we can consider $T(x)$ to be an affine transformation with $b = 0$. Obviously, for any vector $x \in \mathbb{R}^d$,

$$\langle x, Mp\rangle = x^T Mp = \left\langle M^T x, p\right\rangle.$$

Therefore for any $z \in \widehat{M}(\Delta)$,

$$\begin{aligned}
\omega(z, M(Q)) &= \overline{\omega}(\widetilde{z}, M(Q)) \\
&= \max_{q\in Q} \langle \widetilde{z}, Mq\rangle - \min_{q\in Q} \langle \widetilde{z}, Mq\rangle \\
&= \max_{q\in Q} \left\langle M^T\widetilde{z}, q\right\rangle - \min_{q\in Q} \left\langle M^T\widetilde{z}, q\right\rangle \\
&= \overline{\omega}(M^T\widetilde{z}, Q).
\end{aligned}$$

Since $z \in \widehat{M}(\Delta)$, we have $\phi(M^T\widetilde{z}) \in \Delta^*$, which implies that

$$\overline{\omega}(M^T\widetilde{z}, Q) \geq (1-\varepsilon)\overline{\omega}(M^T\widetilde{z}, P).$$

Hence,

$$\omega(z, M(Q)) \geq (1-\varepsilon)\overline{\omega}(M^T\widetilde{z}, P) = (1-\varepsilon)\overline{\omega}(\widetilde{z}, M(P)) = (1-\varepsilon)\omega(z, T(P)).$$

■

We call P *α-fat*, for $\alpha \leq 1$, if there exist a point $x \in \mathbb{R}^d$ and a hypercube $\overline{\mathbb{C}}$ centered at origin so that $p + \overline{\mathbb{C}} \supset \mathcal{CH}(P) \supset p + \alpha\overline{\mathbb{C}}$.

Lemma 3.2 *Let P be a set of n points in $\mathbb{R}^d$, and let ε be a parameter. There exists a linear non-singular transform T such that $T(P)$ is α_d-fat, where α_d is a constant depending only on d.*

Proof: Using the algorithm of Barequet and Har-Peled [BH01], we compute in $O(n)$ time two concentric, homotetic 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) $B' \subseteq \mathcal{CH}(P) \subseteq B$.

Let $R \in \mathbb{R}^{d\times d}$ be the rotation transform and let $t \in \mathbb{R}^d$ the translation vector such that $R(B+t)$ is the orthogonal box centered at the origin. Finally, let S be the scaling transform that maps $R(B+t)$ to $\mathbb{C}$. Set $T(x) = (S \cdot R)x + (S \cdot R)t$. By construction, the point set $P' = T(P)$ is a_d-fat. This completes the proof of the first part of the lemma. It is easy to verify that $M = S \cdot R$ is non-singular. ■

Lemmas 3.1 and 3.2 imply that it suffices to describe an algorithm for computing an ε-approximation of an α-fat point set for some $\alpha < 1$. Without loss of generality, we assume that $\mathbb{C} \supset P \supset [-\alpha, \alpha]^d$. The following simple lemma, which follows immediately from the observation that for any point $q \in \partial\mathcal{CH}(P)$ and for any $u \in \mathbb{R}^d$, $\langle u, q\rangle \geq \alpha \|u\|$, will be useful for our analysis.

Lemma 3.3 *Let $P \subset \mathbb{C}$ be a set of n points in $\mathbb{R}^d$, which is α-fat. For any $x \in \mathbb{R}^d$, $\overline{\omega}(x, P) \geq 2\alpha \|x\|$.*

A weaker bound on ε-approximation. Next, we prove a weaker bound on the size of an ε-approximation for a fat point set.

Lemma 3.4 *Let P be a α-fat point set contained in $\mathbb{C} = [-1, +1]^d$, and let $\varepsilon > 0$ be a parameter. Suppose P' is a point set with the following property: for any $p \in P$, there is a $p' \in P'$ such that $d(p, p') \leq \varepsilon\alpha$. Then $(1-\varepsilon)\overline{\omega}(x, P) \leq \overline{\omega}(x, P')$ for any $x \in \mathbb{R}^d$.*

Proof: By Lemma 3.3, $\overline{\omega}(x, P) \geq 2\alpha \|x\|$. Let $p, q \in P$ be two points such that

$$\overline{\omega}(x, \{p, q\}) = \overline{\omega}(x, P) \geq 2\alpha \|x\|,$$

and let $p', q' \in P'$ be two points such that $d(p, p'), d(q, q') \leq \varepsilon\alpha$.

Let $w = p - q$ and $w' = p' - q'$. Then

$$\|w - w'\| \leq \|p - p'\| + \|q - q'\| \leq 2\varepsilon\alpha.$$

Moreover,

$$\begin{aligned}\overline{\omega}(x, \{p, q\}) &= \max\{\langle p, x\rangle, \langle q, x\rangle\} - \min\{\langle p, x\rangle, \langle q, x\rangle\} \\ &= |\langle p, x\rangle - \langle q, x\rangle| = |\langle w, x\rangle|.\end{aligned}$$

Similarly, $\overline{\omega}(x, \{p', q'\}) = |\langle w', x\rangle|$.

$$\begin{aligned}\overline{\omega}(x, P) - \overline{\omega}(x, P') &\leq \overline{\omega}(x, \{p, q\}) - \overline{\omega}(x, \{p', q'\}) \\ &= |\langle w, x\rangle| - |\langle w', x\rangle| \\ &\leq |\langle w - w', x\rangle| \leq \|w - w'\| \cdot \|x\| \\ &\leq 2\varepsilon\alpha\|x\| \\ &\leq \varepsilon\overline{\omega}(x, P).\end{aligned}$$

■

Using the above lemma, we can construct an ε-approximation of a fat point set as follows.

Lemma 3.5 *Let P be a α-fat point set contained in $\mathbb{C}$. For any $\varepsilon > 0$, we can compute, in $O(n + 1/(\alpha\varepsilon)^{d-1})$ time, a subset $Q \subseteq P$ of $O(1/(\alpha\varepsilon)^{d-1})$ points that ε-approximates P.*

Proof: We consider the d-dimensional grid $\mathbb{Z}$ of size $\delta = \frac{\varepsilon}{6\sqrt{d}}\alpha$. That is,

$$\mathbb{Z} = \{(\delta i_1, \ldots, \delta i_d) \mid i_1, \ldots, i_d \in \mathbb{Z}\}.$$

For each d-tuple $I = (i_1, \ldots, i_d)$, let C_I the cell (in x_d-direction) of $\mathbb{Z}$ of the form $[\delta i_1, \delta(i_1 + 1)] \times \cdots \times [\delta i_{d-1}, \delta(i_{d-1} + 1)] \times [\delta r, \delta(r + 1)]$, $r \in \mathbb{Z}$, that contains a point of P; if none of the cells in this column contains a point of P, we can define C_I^-, C_I^+ to be any cell in the column. Let $\mathcal{P} = \bigcup_I (P \cap (C_I^- \cup C_I^+))$, i.e., the subset of points that lie in the cells C_I^- and C_I^+. Clearly, by construction, the Hausdorff distance between $\mathcal{CH}(P)$ and $\mathcal{CH}(\mathcal{P})$ is smaller than $\alpha\varepsilon/6$. Thus, arguing as in the proof of Lemma 3.4, we have that for any point $u \in \mathbb{R}^{d-1}$, $\omega(u, P)$ is realized by a pair of vertices of $\mathcal{CH}(P)$, and therefore $\omega(u, P) - (\varepsilon\alpha/3) \leq \omega(u, \mathcal{P})$; namely, $(1 - \varepsilon/3)\omega(u, P) \leq \omega(u, \mathcal{P})$.

For each $(d-1)$-tuple I, we choose one point from $P \cap C_I^-$ and another point from $C_I^+ \cap P$ and add both of them to Q. Since $P \subseteq \mathbb{C} = [-1, +1]^d$, $|Q| = O(1/(\alpha\varepsilon)^{d-1})$; Q can be constructed in $O(n + 1/(\alpha\varepsilon)^{d-1})$ time, assuming that the ceiling operation (i.e., $\lceil\ \rceil$) can be performed in constant time. We have chosen in Q one point of P from grid cells C_I^-, C_I^+, for every $(d-1)$-tuple I, which contained a point of P. Therefore for every point $p \in \mathcal{P}$, there is a point $q \in Q$ with the property that $d(p, q) \leq \varepsilon/6$. Hence, by Lemma 3.4,

$$(1 - \varepsilon)\omega(u, P) \leq (1 - \varepsilon/3)^2\omega(u, P) \leq (1 - \varepsilon/3)\omega(u, \mathcal{P}) \leq \omega(u, Q),$$

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

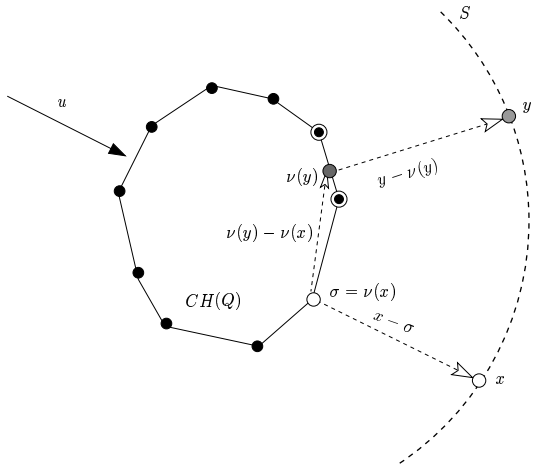

Figure 5. Illustration of the proof of Lemma 3.6; σ is the farthest vertex of $\mathcal{CH}(Q)$ in direction u; the two double-circles denote $b(y)$ (for $d = 2$).

A stronger bound on ε-approximation. Dudley [Dud74] and Bronshteyn and Ivanov [BI76] have shown that given a convex body C, which is contained in a unit ball in $\mathbb{R}^d$, and a parameter $\varepsilon > 0$, one can compute a convex polytope C' so that the Hausdorff distance between C and C' is at most ε. Dudley represents C' as the intersection of $O(1/\varepsilon^{(d-1)/2})$ halfspaces and Bronshteyn and Ivanov represent C' as the convex hull of a set of $O(1/\varepsilon^{(d-1)/2})$ points. In the next lemma we use a variant of the construction in [BI76] to generate a set $O(1/\varepsilon^{(d-1)/2})$ points that ε-approximates P.

Lemma 3.6 *Let P be a α-fat point set in $\mathbb{C}$. For any $\varepsilon > 0$, we can compute, in $O(n + 1/(\alpha\varepsilon)^{3(d-1)/2})$ time, a subset $Q \subseteq P$ of $O(1/(\alpha\varepsilon)^{(d-1)/2})$ points that ε-approximates P.*

Proof: Let S be the sphere of radius $\sqrt{d}+1$ centered at the center of the unit hypercube $\mathbb{C}$ 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 3.5, we compute a set $Q' \subseteq P$ of $O(1/(\alpha\varepsilon)^{d-1})$ points that ε-approximates P. Let $\delta = \sqrt{\varepsilon\alpha/2}$. We compute a set $\mathcal{I}$ of $O(1/\delta^{d-1}) = O(1/(\alpha\varepsilon)^{(d-1)/2})$ points on the sphere S such that for any point x on sphere S (e.g., using the construction in the proof of Lemma 2.2), there is a point $y \in \mathcal{I}$ such that $||x - y|| \leq \delta$. For each point $y \in \mathcal{I}$, we then compute the point $\nu(y)$ on $\mathcal{CH}(Q')$ that is closest to y. Using the algorithm of Gärtner [Gär95], this can be done for each y in time $O(Q') = O(1/(\alpha\varepsilon)^{d-1})$. Gärtner'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 $b(y)$. Set $Q = \bigcup_{y\in\mathcal{I}} b(y)$. It takes $O(1/(\alpha\varepsilon)^{3(d-1)/2})$ time to compute Q. We now argue that Q ε-approximates P.

Fix a direction $u \in \mathbb{R}^{d-1}$, and let $u^* \in \mathbb{S}^{d-1}$ be the unit vector $\phi(\widetilde{u})$. Let $\sigma \in Q'$ be the point that maximizes $\langle u^*, q'\rangle$ over all $q' \in Q'$. Suppose the ray emanating from σ in direction u^* hits S at a point x. Then σ is the unique point on $\mathcal{CH}(Q')$ nearest to x, i.e., $\sigma = \nu(x)$, because the hyperplane normal to the vector $\sigma - x$ supports $\mathcal{CH}(Q')$ at σ and separates x from Q'. Moreover,

$$\frac{x - \nu(x)}{\|x - \nu(x)\|} = \phi(x - \nu(x)) = \phi(x - \sigma) = u^* \quad \text{and} \quad ||x - \nu(x)|| \geq 1 = ||u^*||. \tag{1}$$

Let $y \in \mathcal{I}$ be such that $||x - y|| \leq \delta$. Since $\nu(y)$ is the closest point to y in $\mathcal{CH}(Q')$, the hyperplane normal to $y - \nu(y)$ and passing through $\nu(y)$ separates y and $\nu(x)$, therefore

$$0 \leq \langle y - \nu(y), \nu(y) - \nu(x) \rangle . \tag{2}$$

See Figure 5. Note that for any $a, b \in \mathbb{R}^d$, $2\langle a, b \rangle \leq \|a\|^2 + \|b\|^2$, therefore

$$\langle a, b \rangle - \|b\|^2 \leq \|a\|^2 . \tag{3}$$

Namely,

$$\begin{aligned}
0 \leq \max_{q' \in Q'} \langle u^*, q' \rangle - \max_{q \in Q} \langle u^*, q \rangle &\leq \langle u^*, \sigma \rangle - \langle u^*, \nu(y) \rangle = \langle u^*, \nu(x) - \nu(y) \rangle \\
&\leq \langle x - \nu(x), \nu(x) - \nu(y) \rangle \qquad \text{(using (1))} \\
&\leq \langle x - \nu(x), \nu(x) - \nu(y) \rangle + \langle y - \nu(y), \nu(y) - \nu(x) \rangle \\
&\qquad \text{(using (2))} \\
&\leq \langle x - \nu(x) - (y - \nu(y)), \nu(x) - \nu(y) \rangle \\
&= \langle x - y, \nu(x) - \nu(y) \rangle - \|\nu(x) - \nu(y)\|^2 \\
&\leq \|x - y\|^2 \qquad \text{(using (3))} \\
&\leq \delta^2 = \alpha\varepsilon/2.
\end{aligned}$$

Hence,

$$\max_{q \in Q} \langle \widetilde{u}, q \rangle \geq \max_{q' \in Q'} \langle \widetilde{u}, q' \rangle - \frac{\alpha\varepsilon}{2} \|\widetilde{u}\| .$$

Similarly, we have

$$\min_{q \in Q} \langle \widetilde{u}, q \rangle \leq \min_{q \in Q'} \langle \widetilde{u}, q \rangle + \frac{\alpha\varepsilon}{2} \|\widetilde{u}\| .$$

Using Lemma 3.3, we obtain

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

■

Combining Lemmas 3.5 and 3.6 with Lemma 3.2, we obtain the following result.

Theorem 3.7 *Let P be a point set in $\mathbb{R}^d$, and let $\varepsilon > 0$ be a parameter. We can compute in $O(n + 1/\varepsilon^{d-1})$ time an ε-approximation of P of size $O(1/\varepsilon^{d-1})$, or in $O(n + 1/\varepsilon^{3(d-1)/2})$ time an ε-approximation of P of size $O(1/\varepsilon^{(d-1)/2})$.*

Combining this theorem with Lemma 2.4, we obtain the following.

Theorem 3.8 *Let $\mathcal{H}$ be a set of n $(d-1)$-variate linear functions, and let $\varepsilon > 0$ be a parameter. We can compute in $O(n + 1/\varepsilon^{d-1})$ time an ε-approximation of $\mathcal{H}$ of size $O(1/\varepsilon^{d-1})$, or in $O(n + 1/\varepsilon^{3(d-1)/2})$ time an ε-approximation of size $O(1/\varepsilon^{(d-1)/2})$.*

A decomposition based bound. Next, we show that we can decompose $\mathbb{R}^{d-1}$ into cells so that a pair of points ε-approximates a point set within each cell of the decomposition.

Lemma 3.9 *Let P be an α-fat point set contained in $\mathbb{C}$, and let $\varepsilon > 0$ be a parameter. We can compute, in $O(n + 1/(\alpha\varepsilon)^{3(d-1)/2})$ time, a set $\mathcal{J}$ of $O(1/(\alpha\varepsilon))$ $d(d-1)$-uniform hyperplanes in $\mathbb{R}^{d-1}$ with the following property: for any cell $\Delta \in \mathcal{A}(\mathcal{J})$, there are two points p_Δ, p'_Δ such that $\{p_\Delta, p'_\Delta\}$ ε-approximates P inside Δ.*

Proof: We first use Lemma 3.6 to compute a subset Q of $O(1/(\alpha\varepsilon)^{(d-1)/2})$ points, which is an $(\varepsilon/2)$-approximation of P. We compute a set $\mathcal{J}$ of $O(1/(\alpha\varepsilon))$ hyperplanes, using Lemma 2.2, such that for any two points u, v in the same cell of $\mathcal{A}(\mathcal{J})$,

$$\|u^* - v^*\| \leq \frac{\varepsilon}{4\sqrt{d}}\alpha.$$

We choose any point u_Δ from each cell $\Delta \in \mathcal{A}(\mathcal{J})$ and compute the points p_Δ and p'_Δ, by examining each point in Q, that achieve $\max_{q\in Q} \langle {u^*}_\Delta, q\rangle$ and $\min_{q\in Q} \langle {u^*}_\Delta, q\rangle$, respectively. We associate the points p_Δ and p'_Δ with Δ.

By Lemma 2.2, $\mathcal{A}(\mathcal{J})$ can be computed in $O(n+1/(\alpha\varepsilon)^{d-1})$ time. We spend $O(1/(\alpha\varepsilon)^{(d-1)/2})$ time at each cell $\Delta \in \mathcal{A}(\mathcal{J})$ to compute p_Δ, p'_Δ. So the total running time of the algorithm is $O(n + 1/(\alpha\varepsilon)^{3(d-1)/2})$.

We now argue that $\{p_\Delta, p'_\Delta\}$ is an ε-approximation of P within Δ. Let $u = u_\Delta$, $p = p_\Delta$, $p' = p'_\Delta$. Let v be another point in Δ, 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 $Q \subset \mathbb{C}$, $\|p - q\| \leq 2\sqrt{d}$.

$$\begin{aligned}
\langle v^*, p\rangle &= \langle u^*, p\rangle + \langle v^* - u^*, p\rangle \\
&\geq \langle u^*, q\rangle + \langle v^* - u^*, p\rangle \\
&= \langle v^*, q\rangle - \langle v^* - u^*, q\rangle + \langle v^* - u^*, p\rangle \\
&\geq \langle v^*, q\rangle - \|v^* - u^*\| \cdot \|p - q\| \\
&\geq \langle v^*, q\rangle - \frac{\varepsilon\alpha}{4\sqrt{d}} \cdot 2\sqrt{d} \\
&\geq \langle v^*, q\rangle - \frac{\varepsilon\alpha}{2} = \langle v^*, q\rangle - \frac{\varepsilon\alpha}{2}\|v^*\|.
\end{aligned}$$

Therefore $\langle \widetilde{v}, p\rangle \geq \langle \widetilde{v}, q\rangle - (\varepsilon\alpha/2)\|\widetilde{v}\|$, which implies that

$$\begin{aligned}
\omega(v, \{p_\Delta, p'_\Delta\}) &\geq \omega(v, Q) - \varepsilon\alpha\|\widetilde{v}\| \geq (1 - \varepsilon/2)\omega(v, P) - (\varepsilon/2)\omega(v, P) \\
&\geq (1 - \varepsilon)\omega(v, P)
\end{aligned}$$

This completes the proof of the lemma. ∎

Theorem 3.10 *Let P be a set of n points in $\mathbb{R}^d$, and let $\varepsilon > 0$ be a parameter. We can compute, in $O(n+1/\varepsilon^{3(d-1)/2})$ time, a set $\mathcal{J}$ of $O(1/\varepsilon)$ $d(d-1)$-uniform hyperplanes in $\mathbb{R}^{d-1}$ with the following property: for any cell $\Delta \in \mathcal{A}(\mathcal{J})$, there are two points p_Δ, p'_Δ such that $\{p_\Delta, p'_\Delta\}$ ε-approximates P inside Δ.*

Proof: By Lemma 3.2, let $T(x) = M(x) + b$ be the nonsingular transform so that $T(P)$ is α_d-fat. Using Lemma 3.9, we compute a set H of $O(1/\varepsilon)$ $(d-2)$-hyperplanes in $\mathbb{R}^{d-1}$ so that for any cell $\Delta \in \mathcal{A}(H)$, $\{T(q_\Delta), T(q'_\Delta)\}$ ε-approximates $T(P)$. For a hyperplane $h \in H$, let h' be the $(d-1)$-hyperplane containing $\widetilde{h} = \{\widetilde{x} \mid x \in h\}$ and passing through origin, and let $\hat{h} = M^T\widetilde{h} \cap \mathbb{P}$. We set $\mathcal{J} = \{\hat{h} \mid h \in H\}$. Since M is nonsingular (composition of rotation and scaling), $\mathcal{A}(\mathcal{J})$ is isomorphic to $\mathcal{A}(H)$ and each cell $\Delta \in \mathcal{A}(H)$ maps to $\widehat{M}(\Delta)$ in $\mathcal{A}(\mathcal{J})$. Hence, we associate q_Δ, q'_Δ with $\widehat{M}(\Delta)$. The proof now follows immediately from Lemma 3.1. ∎

Finally, using Lemma 2.4 we conclude the following.

Theorem 3.11 *Given a family $\mathcal{H}$ of n $(d-1)$-variate linear functions and a parameter $\varepsilon > 0$, we can compute in $O(n + 1/\varepsilon^{3(d-1)/2})$ time a family $\mathcal{J}$ of $O(1/\varepsilon)$ $d(d-1)$-uniform hyperplanes in $\mathbb{R}^{d-1}$ with the following property: for each cell $\Delta \in \mathcal{A}(\mathcal{J})$, there are two associated linear functions, $h'_\Delta, h''_\Delta \in \mathcal{H}$ that ε-approximate $\mathcal{H}$ inside Δ.*

4 ε-Approximations for Polynomials and Their Variants

Extent of polynomials. Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a family of $(d-1)$-variate polynomials and $\varepsilon > 0$ a parameter. We use the linearization technique [AM94, YY85] to compute ε-approximations for $\mathcal{F}$.

Let $f(x, a)$ be a $(d+p-1)$-variate polynomial, $x \in \mathbb{R}^{d-1}$ and $a \in \mathbb{R}^p$, such that $f_i(x) \equiv f(x, a^i)$ for some $a^i \in \mathbb{R}^p$. There always exists such a polynomial for $\mathcal{F}$. Suppose we can express $f(x, a)$ in the form

$$f(x, a) = \psi_0(a) + \psi_1(a)\varphi_1(x) + \cdots + \psi_k(a)\varphi_k(x), \tag{4}$$

where $\psi_0, \ldots, \psi_k$ are p-variate polynomials and $\varphi_1, \ldots, \varphi_k$ are $(d-1)$-variate polynomials. We define the map $\varphi : \mathbb{R}^{d-1} \to \mathbb{R}^k$

$$\varphi(x) = (\varphi_1(x), \ldots, \varphi_k(x)).$$

Then the image $\Gamma = \left\{\varphi(x) \mid x \in \mathbb{R}^{d-1}\right\}$ of $\mathbb{R}^{d-1}$ is a $(d-1)$-dimensional surface in $\mathbb{R}^k$, and for any $a \in \mathbb{R}^p$, $f(x, a)$ maps to a k-variate linear function

$$h_a(y_1, \ldots, y_k) = \psi_0(a) + \psi_1(a)y_1 + \cdots + \psi_k(a)y_k$$

in the sense that for any $x \in \mathbb{R}^{d-1}$, $f(x, a) = h_a(\varphi(x))$. We refer to k as the *dimension* of linearization. The simplest way to express the polynomial $f(x, a)$ in the form (4) is to write f as a sum of monomials in $x_1, \ldots, x_{d-1}$ with its coefficients being polynomials in $a_1 \ldots, a_p$. Then each monomial in the $x_1, \ldots, x_{d-1}$ corresponds to one function φ_i, and its coefficient is the corresponding function ψ_i. However, this method does not necessarily give a linearization of the smallest dimension. For example, let $f(x_1, x_2, a_1, a_2, a_3)$ be the square of the distance between a point $(x_1, x_2) \in \mathbb{R}^2$ and a circle with center (a_1, a_2) and radius a_3, which is the 5-variate polynomial

$$f(x_1, x_2, a_1, a_2, a_3) = a_3^2 - (x_1 - a_1)^2 - (x_2 - a_2)^2 .$$

A straightforward application of the above method yields a linearization of dimension 4. However, f can be written in the form

$$f(x_1, x_2, a_1, a_2, a_3) = [a_3^2 - a_1^2 - a_2^2] + [2a_1x_1] + [2a_2x_2] - [x_1^2 + x_2^2], \tag{5}$$

thus, setting

$$\begin{aligned} \psi_0(a) = a_3^2 - a_1^2 - a_2^2, \quad \psi_1(a) = 2a_1, \quad \psi_2(a) = 2a_2, \quad \psi_3(a) = -1, \\ \varphi_1(x) = x_1, \quad \varphi_2(x) = x_2, \quad \varphi_3(x) = x_1^2 + x_2^2, \end{aligned}$$

we get a linearization of dimension 3. It corresponds to the well-known "lifting" transform to the unit paraboloid. Agarwal and Matoušek [AM94] describe an algorithm that computes a linearization of the smallest dimension.

Returning to the problem of computing an ε-approximation of $\mathcal{F}$, let $\mathcal{H} = \{h_{a^i} \mid 1 \le i \le n\}$. Let $\mathcal{K}$ be an ε-approximation of $\mathcal{H}$ within a region $\Delta \in \mathbb{R}^k$. Since $f_i(x) = h_{a^i}(\varphi(x))$ for any $x \in \mathbb{R}^{d-1}$, $\mathcal{G} = \{f_i \mid h_{a^i} \in \mathcal{K}\}$ is an ε-approximation of $\mathcal{F}$ within the region $\varphi^{-1}(\Delta \cap \Gamma)$, where $\varphi^{-1}(\gamma) = \left\{x \in \mathbb{R}^{d-1} \mid \varphi(x) \in \gamma\right\}$, for $\gamma \in \mathbb{R}^k$, is the pre-image of γ in $\mathbb{R}^{d-1}$. Hence, by Theorem 3.8, we obtain the following.

Theorem 4.1 *Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a family of $(d-1)$-variate polynomials that admits a linearization of dimension k, and let $\varepsilon > 0$ be a parameter. We can compute an ε-approximation of $\mathcal{F}$ of size $O(1/\varepsilon^k)$ in time $O(n + 1/\varepsilon^k)$, or an ε-approximation of size $O(1/\varepsilon^{k/2})$ in time $O(n + 1/\varepsilon^{3k/2})$.*

For a $(k-1)$-dimensional hyperplane h in $\mathbb{R}^k$, let h^{-1} denote the pre-image $\varphi^{-1}(h \cap \Gamma)$ in $\mathbb{R}^{d-1}$; h^{-1} is an $(d-2)$-dimensional algebraic variety, whose degree depends on the maximum degree of a polynomial in $\mathcal{F}$ and on d. Using Theorem 3.11, we can prove the following.

Theorem 4.2 *Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a family of $(d-1)$-variate polynomials of bounded maximum degree that admits a linearization of dimension k, and let $\varepsilon > 0$ be a parameter. We can compute in time $O(n + 1/\varepsilon^{3k/2})$ a family $\mathcal{G}$ of $O(1/\varepsilon)$ algebraic varieties, whose degrees depend on d and the maximum degree of a polynomial in $\mathcal{F}$, so that for any cell τ of $\mathcal{A}(\mathcal{G})$, there are two polynomials $f_\tau, f'_\tau \in \mathcal{F}$ that ε-approximate $\mathcal{F}$ within τ.*

Proof: Let $\mathcal{H}$ be the linearization of $\mathcal{F}$ of dimension k. By Theorem 3.11, we can compute in $O(n+1/\varepsilon^{3k/2})$ time a set $\mathcal{K}$ of $O(1/\varepsilon)$ $(k-1)$-dimensional hyperplanes in $\mathbb{R}^k$ such that for any cell Δ of $\mathcal{A}(\mathcal{K})$, there exist two hyperplanes h_Δ, h'_Δ that ε-approximate $\mathcal{H}$ within Δ. Set $\mathcal{G} = \{h^{-1} \mid h \in \mathcal{K}\}$. Each cell τ in $\mathcal{A}(\mathcal{G})$ is the pre-image $\varphi^{-1}(\Delta \cap \Gamma)$ of some cell $\Delta \in \mathcal{A}(\mathcal{H})$. For each cell $\tau \in \mathcal{A}(\mathcal{G})$, which is the pre-image of $\Delta \cap \Gamma$, we set $f_\tau = h_\Delta^{-1}$ and $f'_\tau = h'^{-1}_\Delta$. It is easily seen that $\{f_\tau, f'_\tau\}$ ε-approximates $\mathcal{F}$ within τ. ∎

Since $\bigcup_{\tau \in \mathcal{A}(\mathcal{G})} \{f_\tau, f'_\tau\}$ is an ε-approximation of $\mathcal{F}$ and $\mathcal{A}(\mathcal{G})$ has $O(1/\varepsilon^{d-1})$ cells [AS00], combining this observation with Theorem 4.1 we can conclude the following.

Theorem 4.3 *Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a family of $(d-1)$-variate polynomials that admits a linearization of dimension k, and let $\varepsilon > 0$ be a parameter. We can compute in time $O(n + 1/\varepsilon^{3k/2})$ an ε-approximation of $\mathcal{F}$ of size $O(1/\varepsilon^\sigma)$, where $\sigma = \min\{d-1, k/2\}$.*

Unlike an arrangement of hyperplanes, it is not known whether an arrangement of m algebraic surfaces in $\mathbb{R}^d$, each of constant degree, can be decomposed into $O(m^d)$ Tarski cells.[2] However, such a decomposition is feasible for the surfaces in Theorem 4.1. Indeed, by construction in the proof of Lemma 2.2, each cell Δ in $\mathcal{A}(\mathcal{K})$ has $O(1)$ faces, so its pre-image $\varphi^{-1}(\Delta \cap \Gamma)$ also has $O(1)$ complexity. We can further refine it into $O(1)$ Tarski cells. Hence, we can decompose $\mathcal{A}(\mathcal{G})$ into $O(1/\varepsilon^{d-1})$ Tarski cells.

Theorem 4.4 *Let $\mathcal{F} = \{f_1, \ldots, f_n\}$ be a family of $(d-1)$-variate polynomials that admits a linearization of dimension k, and let $\varepsilon > 0$ be a parameter. We can compute in time $O(n + 1/\varepsilon^{3k/2})$ a decomposition Ξ of $\mathbb{R}^{d-1}$ into $O(1/\varepsilon^{d-1})$ Tarski cells with the following property: for each cell τ in Ξ, there are two polynomials $f_\tau, f'_\tau \in \mathcal{F}$ that ε-approximate $\mathcal{F}$ within τ.*

Remark 4.5 Note that the result of Theorem 4.4 is somewhat surprising. In particular, it implies that if $\mathcal{F}$ is a family of polynomial defined over a single variable (i.e., $d = 2$), then the extent of $\mathcal{F}$ has an approximation of size $O(1/\varepsilon)$. We use this observation in Theorem 6.6.

Fractional powers of polynomials. We now consider the problem of ε-approximating a family of functions $\mathcal{F} = \left\{(f_1)^{1/r}, \ldots, (f_n)^{1/r}\right\}$, where $r \geq 1$ is an integer and each f_i is a polynomial of some bounded degree. This case is considerably harder than handling polynomials because they can not be linearized directly. In certain special cases this can be overcome by special considerations of the functions at hand [AAHS00, Cha02]. We, however, prove here that it is enough to compute an $O(\varepsilon^r)$-approximation of the polynomials inside the roots. We need the following lemma.

Lemma 4.6 *Let $0 < \varepsilon < 1$ be a parameter, and let $\delta = (\varepsilon/2(r-1))^r$. If we have $0 \leq a \leq A \leq B \leq b$ and $B - A \geq (1-\delta)(b-a)$, then*

$$B^{1/r} - A^{1/r} \geq (1-\varepsilon)(b^{1/r} - a^{1/r}).$$

Proof: First, observe that for any x, y and for any integer $r \geq 0$,

$$x^r - y^r = (x-y)(x^{r-1} + x^{r-2}y + \cdots + xy^{r-2} + y^{r-1}), \tag{6}$$

and for any $0 \leq p \leq 1$,

$$x^p + y^p \geq (x+y)^p. \tag{7}$$

[2] A k-dimensional semialgebraic set is called a *Tarski cell* if it is homeomorphic to a k-dimensional ball and it is defined by constant number of polynomial inequalities, each of which has bounded degree.

Using (6),

$$\begin{aligned} B^{1/r} - A^{1/r} &= (B-A) \Big/ \left(\sum_{i=0}^{r-1} A^{i/r} B^{1-(i+1)/r} \right) \\ &\geq (1-\delta)(b-a) \Big/ \left(\sum_{i=0}^{r-1} A^{i/r} B^{1-(i+1)/r} \right) \\ &\geq (1-\delta)(b^{1/r} - a^{1/r}) \left(\sum_{i=0}^{r-1} a^{i/r} b^{1-(i+1)/r} \right) \Big/ \left(\sum_{i=0}^{r-1} A^{i/r} B^{1-(i+1)/r} \right). \end{aligned}$$

Therefore, for $0 \leq i < r$,

$$\begin{aligned} a^{i/r} b^{1-(i+1)/r} &\geq a^{i/r} B^{1-(i+1)/r} \\ &\geq a^{i/r} B^{1-(i+1)/r} + \delta^{i/r} B^{1-1/r} - \delta^{i/r} B^{1-1/r} \\ &\geq (a^{i/r} + (\delta B)^{i/r}) B^{1-(i+1)/r} - \delta^{i/r} B^{1-1/r} \\ &\geq (a + \delta B)^{i/r} B^{1-(i+1)/r} - \delta^{i/r} B^{1-1/r} \\ &\qquad \text{(Using (7) since } i < r) \\ &\geq A^{i/r} B^{1-(i+1)/r} - \delta^{i/r} B^{1-1/r}. \end{aligned}$$

The last inequality holds because, by our assumption,

$$\begin{aligned} B - A \geq (1-\delta)(b-a) &\Rightarrow (1-\delta)a + \delta B \geq (1-\delta)B + A \\ &\Rightarrow a + \delta B \geq A. \end{aligned}$$

Hence,

$$\begin{aligned} \sum_{i=0}^{r-1} a^{i/r} b^{1-(i+1)/r} &\geq \sum_{i=1}^{r-1} (A^{i/r} B^{1-(i+1)/r} - \delta^{i/r} B^{1-1/r}) + B^{1-1/r} \\ &\geq \sum_{i=1}^{r-1} A^{i/r} B^{1-(i+1)/r} + (1 - (r-1)\delta^{1/r}) B^{1-1/r}) \\ &\geq (1 - (r-1)\delta^{1/r}) \sum_{i=0}^{r-1} A^{i/r} B^{1-(i+1)/r}. \end{aligned}$$

Putting everything together,

$$\begin{aligned} B^{1/r} - A^{1/r} &\geq (1-\delta)(b^{1/r} - a^{1/r})(1 - (r-1)\delta^{1/r}) \\ &\geq (1 - (\varepsilon/2(r-1))^r)(1 - \varepsilon/2)(b^{1/r} - a^{1/r}) \\ &\geq (1-\varepsilon)(b^{1/r} - a^{1/r}). \end{aligned}$$

■

Theorem 4.7 *Let $\mathcal{F} = \left\{ (f_1)^{1/r}, \ldots, (f_n)^{1/r} \right\}$ be a family of $(d-1)$-variate functions (over $x = (x_1, \ldots, x_{d-1}) \in \mathbb{R}^{d-1}$), where $r \geq 1$ is an integer and each f_i is a polynomial that*

is non-negative for every $x \in \mathbb{R}^k$, and let $\varepsilon > 0$ be a parameter. Suppose f_i's admit a linearization of dimension k. We can compute an ε-approximation of $\mathcal{F}$ of size $O(1/\varepsilon^{rk})$ in time $O(n + 1/\varepsilon^{rk})$, or an ε-approximation of size $O(1/\varepsilon^{r\sigma})$, where $\sigma = \min\{d-1, k/2\}$, in $O(n + 1/\varepsilon^{3rk/2})$ time.

Proof: Let $\mathcal{F}^r = \{f_1, \ldots, f_n\}$. Let $\mathcal{G}^r \subseteq \mathcal{F}^r$ be a δ-approximation of $\mathcal{F}^r$ for $\delta = (\varepsilon/2(r-1))^r$. By Lemma 4.6, $\mathcal{G} = \left\{(f_i)^{1/r} \mid f_i \in \mathcal{G}^r\right\}$ is an ε-approximation of $\mathcal{F}$. The claim now follows from Theorem 4.1. ∎

Similarly, by Theorem 4.2, we can prove the following.

Theorem 4.8 *Let $\mathcal{F} = \left\{(f_1)^{1/r}, \ldots, (f_n)^{1/r}\right\}$ be a family of $(d-1)$-variate functions (over $x = (x_1, \ldots, x_k) \in \mathbb{R}^{d-1}$), where each f_i is a polynomial that is non-negative for every $x \in \mathbb{R}^{d-1}$, and let $\varepsilon > 0$ be a parameter. Suppose f_i's admit a linearization of dimension k. We compute in $O(n + 1/\varepsilon^{3rk/2})$ time, a family of $O(1/\varepsilon^r)$ $(d-2)$-dimensional surfaces in $\mathbb{R}^{d-1}$, so that for each cell $\Delta \in \mathcal{A}(\mathcal{G})$ there are two associated functions $f'_\Delta, f''_\Delta \in \mathcal{F}$ that ε-approximate $\mathcal{F}$ within Δ.*

5 Dynamization

In this section we show that we can adapt our algorithm for maintaining an ε-approximation of a set of points or a set of linear functions under insertions and deletions. We describe the algorithm for a set P of points in $\mathbb{R}^d$. We assume the existence of an algorithm $\mathbb{A}$ that can compute a δ-approximation of a subset $S \subseteq P$ of size $O(1/\delta^k)$ in time $O(|P| + f(\delta))$. We will use $\mathbb{A}$ to maintain an ε-approximation dynamically. We first describe a dynamic data structure of linear size that handles both insertions and deletion. Next, we describe another data structure that uses $O((\log(n)/\varepsilon)^{O(1)})$ space and handles each insertion in $O((\log(n)/\varepsilon)^{O(1)})$ amortized time.

A fully dynamic data structure. We assume that each point in P has a unique id. Using this id as the key, we store P in a 2-4-tree T of height at most $2\log_2 n$; each point of P is stored at a leaf of T. For a node $v \in T$, let $P_v \subseteq P$ be the subset of points stored at the leaves in the subtree rooted at v. We also associate a subset $Q_v \subseteq P_v$ with v, which is defined recursively, as follows. Set $\delta = \varepsilon/3h$, where h is the height of T. If v is a leaf, then $Q_v = P_v$. For an internal node v with w and z as its children, Q_v is a δ-approximation of $Q_w \cup Q_z$ of size $O(1/\delta^k)$, computed using algorithm $\mathbb{A}$. Our construction ensures that for a node at height i (leaves have height 0), Q_v is an $(\varepsilon i/(2h))$-approximation of P_v since $(1 + \varepsilon/3h)^i \leq (1 + \varepsilon i/(2h))$. Therefore the subset Q_u associated with the root u of T is an $(\varepsilon/2)$-approximation of P. Finally, we maintain an $(\varepsilon/2)$-approximation Q of Q_u of size $O(1/\varepsilon^k)$ using algorithm $\mathbb{A}$; Q is an ε-approximation of P.

Suppose we want to delete a point p_i from P. We find the leaf z that stores p_i, delete that leaf, and recompute Q_v at all ancestors v of z in a bottom-up manner. At each ancestor v, with x and w as its children, we compute, in $O(((\log n)/\varepsilon)^k + f((\log n)/\varepsilon))$ time, a δ-approximation of $Q_w \cup Q_x$ using algorithm $\mathbb{A}$. Finally, we recompute, in time $O((1/\varepsilon)^k + f(1/\varepsilon^k))$, an $(\varepsilon/2)$-approximation Q of Q_u. The total time spent is thus $O((\log^{k+1} n)/\varepsilon^k + f((\log n)/\varepsilon)\log n)$. We can insert a point in the same way. Finally, if the height h of T

changes, we recompute Q_v's at all nodes of T with the new value of h. Since the height of T changes after at least $n_0/2$ updates, where n_0 is the number of points in T when its height changed the last time, updating the tree costs $O(((\log n)/\varepsilon)^k + f((\log n)/\varepsilon))$ time per update operation. Hence, we obtain the following.

Theorem 5.1 *Let P be a set of points in $\mathbb{R}^d$, and let $\varepsilon > 0$ be a parameter. Suppose we can compute an ε-approximation of a subset $S \subseteq P$ of size $O(1/\varepsilon^k)$ in time $O(|S| + f(\varepsilon))$ time, then we can maintain an ε-approximation of P of size $O(1/\varepsilon^k)$ under insertion/deletion in time $O((\log^{k+1} n)/\varepsilon^k + f((\log n)/\varepsilon)\log n)$ per update operation.*

Remark 5.2 A weakness of our approach is that insertion or deletion of a point can change the ε-approximation completely. It would be desirable to develop a dynamic data structure that causes $O(1)$ change in the ε-approximation after insertion or deletion of a point.

Corollary 5.3 *Let $\mathcal{F}$ be a set of functions, and let $\varepsilon > 0$ be a parameter. Suppose we can compute an ε-approximation of a subset $\mathcal{G} \subseteq \mathcal{F}$ of size $O(1/\varepsilon^k)$ in time $O(|\mathcal{G}| + f(\varepsilon))$ time, then we can maintain an ε-approximation of $\mathcal{F}$ of size $O(1/\varepsilon^k)$ under insertion/deletion in time $O(((\log n)/\varepsilon)^k + f((\log n)/\varepsilon)\log n)$ per update operation.*

An insertion-only data structure. Suppose we are receiving a stream of points $p_1, p_2, \ldots$ in $\mathbb{R}^d$. Given a parameter $\varepsilon > 0$, we wish to maintain an ε-approximation of the n points received so far. (Note, that our analysis is in term of n, the number of points inserted into the data structure. However, n does not need to be specified in advance. In particular, if n is specified in advance, a slightly simpler solution arises using the techniques described above.) We use the dynamization technique of Bentley-Saxe [BS80], as follows: Let $P = \langle p_1, \ldots, p_n \rangle$ be the sequence of points that we have received so far. For an integer $j \geq 0$, let $\rho_j = \varepsilon/cj^2$, where $c > 0$ is a constant, and set $\delta_j = \prod_{l=0}^{j}(1 + \rho_l)$. We partition P into $u \leq \lceil \log_2 n \rceil$ subsets $P_1, \ldots P_u$. For each i, we ensure that $|P_i| = 2^j$ for some $j < \lceil \log_2 n \rceil$. We refer to j as the *rank* of P_i. We maintain the invariant that the ranks of all P_i's are distinct. Formally, a subset of rank j exists in the partition if and only if the jth rightmost bit in the binary representation of n is 1.

Unlike the standard Bentley-Saxe technique, we do not maintain each P_i explicitly. Instead, for a subset P_i of rank j, we maintain a δ_j-approximation Q_i of P_i. Since

$$\delta_j = \prod_{l=1}^{j}\Big(1 + \frac{\varepsilon}{cl^2}\Big) \leq \exp\Big(\sum_{l}\frac{\varepsilon}{cl^2}\Big) \leq \exp\Big(\frac{\pi^2\varepsilon}{6c}\Big) \leq 1 + \frac{\varepsilon}{2},$$

provided c is chosen sufficiently large, Q_i is an $(\varepsilon/2)$-approximation of P_i. Therefore $Q = \bigcup_{i=1}^{u} Q_j$ is an $(\varepsilon/2)$-approximation of P. We can either return Q as an ε-approximation or compute an $(\varepsilon/3)$-approximation of Q of size $O(1/\varepsilon^k)$ using algorithm $\mathbb{A}$.

At the arrival of the next point p_{n+1}, the data structure is updated as follows. We set $P_0 = \{p_{n+1}\}$. The rank of P_0 is 0, so we set $Q_0 = P_0$ as a $(1/\delta_0)$-approximation of P_0. Next, if there are two δ_j-approximations Q_x, Q_y of rank j, for some $j \leq \lceil \log_2(n+1) \rceil$, we compute a $(1 + \rho_j)$-approximation Q_z of $Q_x \cup Q_y$ using algorithm $\mathbb{A}$, set the rank of Q_z to $j + 1$, and discard the sets Q_x and Q_y. By construction, Q_z is a δ_{j+1}-approximation of $P_z = P_x \cup P_y$ of

size $O(1/\rho_{j+1}^k) = O(j^{2k}/\varepsilon^k)$ and $|P_z| = 2^{j+1}$. We repeat this step until the ranks of all Q_i's are distinct. Hence,

$$|Q| = \sum_{i=1}^{u} |Q_i| \leq \sum_{j=0}^{\lceil \log_2 n \rceil} O(1/\rho_j^k) = \sum_{j=0}^{\lceil \log_2 n \rceil} O\left(\frac{j^{2k}}{\varepsilon^k}\right) = O\left(\frac{\log^{2k+1} n}{\varepsilon^k}\right).$$

For any fixed $j \geq 0$, a δ_j-approximation of a subset P_i of rank j is constructed after every 2^j insertions, therefore the amortized time spent in updating Q after inserting a point is

$$\sum_{j=0}^{\lceil \log_2 n \rceil} \left(\frac{1}{2}\right)^j O\left(\frac{j^{2k}}{\varepsilon^k} + f\left(\frac{\varepsilon}{cj^2}\right)\right) = O\left(\frac{1}{\varepsilon^k} + \sum_{j=1}^{\lceil \log_2 n \rceil} \left(\frac{1}{2}\right)^j f\left(\frac{\varepsilon}{cj^2}\right)\right).$$

If $f(x)$ is bounded by a polynomial in $1/x$, then the above expression is bounded by $O(1/\varepsilon^k + f(\varepsilon))$. If we also construct an $\varepsilon/3$-approximation of Q, we spend an additional $O(\log^{2k+1}(n)/\varepsilon^k + f(\varepsilon/3))$ time. Hence, we obtain the following.

Theorem 5.4 *Let P be a stream of points in $\mathbb{R}^d$, and let $\varepsilon > 0$ be a parameter. Suppose we can compute an ε-approximation of a subset $S \subseteq P$ of size $O(1/\varepsilon^k)$ in time $O(|S| + f(\varepsilon))$ time, where $f(\cdot)$ is bounded by a polynomial. Then we can maintain an ε-approximation of P using a data structure of size $O(\log^{2k+1}(n)/\varepsilon^k)$. The size of the ε-approximation is $O(\log^{2k+1}(n)/\varepsilon^k)$ if we allow $O(1/\varepsilon^k + f(\varepsilon))$ amortized time to insert a point, and the size is $O(1/\varepsilon^k)$ is we allow $O(\log^{2k+1}(n)/\varepsilon^k + f(\varepsilon))$ amortized time.*

The following is an immediate corollary of Theorem 3.7.

Corollary 5.5 *Let P be a stream of points in $\mathbb{R}^d$, and let $\varepsilon > 0$ be a parameter. We can maintain an ε-approximation of P using a data structure of size $O(\log^d(n)/\varepsilon^{(d-1)/2})$. The size of the ε-approximation is $O(\log^d(n)/\varepsilon^{(d-1)/2})$ if we allow $O(1/\varepsilon^{3(d-1)/2})$ amortized time to insert a point, and the size is $O(1/\varepsilon^{(d-1)/2})$ is we allow $O((\log^d(n) + 1/\varepsilon^{d-1})/\varepsilon^{(d-1)/2})$ amortized time.*

6 Applications

In this section we present a few specific applications of the results on ε-approximations obtained in Sections 3 and 4. We begin by describing approximation algorithms for computing faithful extent measures, and then showing that our technique can be extended to maintaining faithful measures of moving points. Next, we describe approximation algorithms for computing two nonfaithful measures, namely computing the minimum width of spherical and cylindrical shells that contain a set of points.

6.1 Approximating faithful extent measures

A function $\mu(\cdot)$ defined over a finite set P of points is called a *faithful measure* if (i) for any $P \subseteq \mathbb{R}^d$, $\mu(P) \geq 0$, and (ii) there exists a constant (depending on μ) $c \geq 0$, so that for any

ε-approximation Q of P, $(1-c\varepsilon)\mu(P) \leq \mu(Q) \leq \mu(P)$. Examples of faithful measures are common and include diameter, width, radius of the smallest enclosing ball, volume of the minimum bounding box, volume of $\mathcal{CH}(P)$, and surface area of $\mathcal{CH}(P)$. A common property of all these measures is that $\mu(P) = \mu(\mathcal{CH}(P))$. As mentioned in Section 3, if $P \subset \mathbb{C}$, then the Hausdorff distance between $\partial(\mathcal{CH}(P))$ and $\partial(\mathcal{CH}(Q))$, for an ε-approximation Q of P, is at most $O(\varepsilon)$, which implies that $\mu(P) = \mu(\mathcal{CH}(P)) \geq \mu(Q) \geq (1-\varepsilon)\mu(P)$. For a given point set P, a faithful measure μ, and a parameter $\varepsilon > 0$, we can compute a value $\bar{\mu}$, $(1-\varepsilon)\mu(P) \leq \bar{\mu} \leq \mu(P)$ by first computing an ε-approximation Q of P and then using an exact algorithm for computing $\mu(Q)$. Using Theorems 3.7 and 3.10 we obtain the following.

Theorem 6.1 *Given a set P of n points in $\mathbb{R}^d$, a faithful measure μ that can be computed in n^α time, and a parameter $\varepsilon > 0$, we can compute, in time $O(n + f(\varepsilon))$, a value $\bar{\mu}$ so that $(1-\varepsilon)\mu(P) \leq \bar{\mu} \leq \mu(P)$, where $f(\varepsilon) = \min\left\{1/\varepsilon^{\alpha(d-1)}, 1/\varepsilon^{3(d-1)/2} + 1/\varepsilon^{\alpha(d-1)/2}\right\}$. Moreover, P can be stored in a dynamic data structure that can update μ in amortized time*

$$\min\left\{\frac{\log^d n}{\varepsilon^{d-1}} + \frac{1}{\varepsilon^{\alpha(d-1)}}, \frac{\log^{3d/2-1/2} n}{\varepsilon^{3(d-1)/2}} + \frac{1}{\varepsilon^{\alpha(d-1)/2}}\right\}$$

if a point is inserted into or deleted from P.

For example, since the diameter of a point set in $\mathbb{R}^d$ can be trivially computed in $O(n^2)$ time, we can compute an ε-approximation of the diameter in $O(n + 1/\varepsilon^{3(d-1)/2})$ time. Similarly, we can compute in $O(n + 1/\varepsilon^3)$ time an ε-approximation of the volume of the smallest box or simplex enclosing a set of n points in $\mathbb{R}^3$, as the exact algorithms for these problems take $O(n^3)$ time [BH01, O'R85]. For all of the measures mentioned in the beginning of this section, algorithms with similar running time (even slightly better in some cases) are already known [BH01, Cha02]. However, our technique is general and does not require us to carefully inspect the problem at hand to develop an approximation algorithm.

We can use Corollary 5.5 for maintaining faithful extent measures of a stream of points in $\mathbb{R}^d$ using $O(\log^d(n)/\varepsilon^{(d-1)/2})$ space. For instance, an ε-approximation of the diameter of the stream can be maintained by spending $O((1/\varepsilon^{d-1} + \log^d n)/\varepsilon^{(d-1)/2})$ amortized time at each incoming point.

6.2 Maintaining faithful measures of moving points

Next we show that our technique can be extended to maintain various extent measures of a set of moving points. Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{R}^d$, each moving independently. Let $p_i(t) = (p_{i1}(t), \ldots, p_{id}(t))$ denote the position of point p_i at time t. Set $P(t) = \{p_i(t) \mid 1 \leq i \leq n\}$. If each p_{ij} is a polynomial of degree at most r, we say that the motion of P has *degree* r. We call the motion of P *linear* if $r = 1$ and *algebraic* if r is bounded by a constant.

Given a parameter $\varepsilon > 0$, we call a subset $Q \subseteq P$ an *ε-approximation* of P if for any direction $u \in \mathbb{R}^{d-1}$,

$$(1-\varepsilon)\omega(u, P(t)) \leq \omega(u, Q(t)) \qquad \text{for all } t \in \mathbb{R}.$$

We first show that a small ε-approximation of P can be computed efficiently and then discuss how to use it to maintain a faithful measure of P approximately as the points move, assuming that the trajectories of points are algebraic and do not change over time. Finally, we show how to update the ε-approximation if we allow the trajectories of points to be inserted or deleted.

Computing an ε-approximation. First let us assume that the motion of P is linear, i.e., $p_i(t) = a_i + b_i t$, for $1 \le i \le n$, where $a_i, b_i \in \mathbb{R}^d$. For a direction $u = (u_1, \ldots, u_{d-1}) \in \mathbb{R}^{d-1}$, we define a d-variate polynomial

$$\begin{aligned} f_i(u,t) &= \langle p_i(t), \widetilde{u} \rangle = \langle a_i + b_i t, \widetilde{u} \rangle \\ &= \sum_{i=1}^{d-1} a_{ij} u_j + \sum_{i=1}^{d-1} b_{ij} \cdot (t u_j) + a_{id} + b_{id} t. \end{aligned}$$

Set $\mathcal{F} = \{f_1, \ldots, f_n\}$. Then

$$\omega(u, P(t)) = \max_i \langle p_i(t), \widetilde{u} \rangle - \min_i \langle p_i(t), \widetilde{u} \rangle = \max_i f_i(u,t) - \min_i f_i(u,t) = \mathfrak{I}_{\mathcal{F}}(u,t).$$

Since $\mathcal{F}$ is a family of d-variate polynomials, which admits a linearization of dimension $2d-1$ (there are $2d-1$ monomials), using Theorem 4.1, we conclude the following.

Theorem 6.2 *Given a set P of n points in $\mathbb{R}^d$, each moving linearly, and a parameter $\varepsilon > 0$, we can compute an ε-approximation of P of size $O(1/\varepsilon^{2d-1})$ in $O(n + 1/\varepsilon^{2d-1})$ time, or an ε-approximation of size $O(1/\varepsilon^{d-1/2})$ in $O(n + 1/\varepsilon^{3(d-1/2)})$ time.*

If the degree of motion of P is $r > 1$, we can write the d-variate polynomial $f_i(u,t)$ as:

$$f_i(u,t) = \langle p_i(t), \widetilde{u} \rangle = \left\langle \sum_{j=0}^{r} a_{ij} t^j, \widetilde{u} \right\rangle = \sum_{j=0}^{r} \langle a_{ij} t^j, \widetilde{u} \rangle$$

where $a_{ij} \in \mathbb{R}^d$. A straightforward extension of the above argument shows that f_i's admit a linearization of dimension $(r+1)d-1$. Using Theorems 4.1 and 4.3, we obtain the following.

Theorem 6.3 *Given a set P of n moving points in $\mathbb{R}^d$ whose motion has degree r and a parameter $\varepsilon > 0$, we can compute an ε-approximation of P of size $O(1/\varepsilon^{(r+1)d-1})$ in $O(n + 1/\varepsilon^{(r+1)d-1})$ time, or of size $O(1/\varepsilon^d)$ in $O(n + 1/\varepsilon^{3((r+1)d-1)/2})$ time.*

Remark 6.4 By Corollary 5.3, if we can compute in time $O(n + f(\varepsilon))$ an ε-approximation of size $O(1/\varepsilon^k)$ of a set P of n moving points in $\mathbb{R}^d$, then we can update it in time $O(((\log n)/\varepsilon)^k + f((\log n)/\varepsilon) \log n)$ per insertion/deletion of a point.

Kinetic data structures. As in Section 6.1, we can use an ε-approximation of P to maintain various faithful extent measure of P approximately, as the points in P move. Namely, we first compute an ε-approximation Q of P and then maintain the desired measure for Q. Note that Q does not depend on the underlying measure. Agarwal *et al.* [AGHV01] have described kinetic data structures for maintaining various extent measures, including diameter, width, area (or perimeter) of the smallest enclosing rectangle, of a set of points moving algebraically in the plane. Plugging their technique on Q, we can, for example, construct a kinetic data structure of size $O(|Q|)$ that maintains a pair (q, q') with the property that

$$d(q(t), q'(t)) = \operatorname{diam}(Q(t)) \geq (1 - \varepsilon) \operatorname{diam}(P(t)).$$

The pair (q, q') is updated $O(|Q|^{2+\delta})$ times, for any $\delta > 0$, and the data structure can be updated in $O(\log |Q|)$ time at each such event. Similar bounds hold for width, area of the smallest enclosing rectangle, etc.

Theorem 6.5 *Let P be a set of n points moving in the plane, and let $\varepsilon > 0$ be a parameter. If P is moving linearly, then after $O(n+1/\varepsilon^3)$ preprocessing, we can construct a kinetic data structure of size $O(1/\varepsilon^{3/2})$ so that an ε-approximation of diameter, width, or the area (or perimeter) of the smallest enclosing rectangle of P can be maintained. The data structure processes $O(1/\varepsilon^{3+\delta})$ events, for an arbitrarily small constant $\delta > 0$, and each such event requires $O(\log(1/\varepsilon))$ time. If the motion of P has degree r, then the preprocessing time is $O(1/\varepsilon^{3r-2})$, the size of the data structure is $O(1/\varepsilon^2)$, and the number of events is $O(1/\varepsilon^{4+\delta})$.*

In some cases, the size of the ε-approximation that we use to maintain a faithful measure can be improved by reducing the problem to a lower dimensional problem. For example, let $\mathcal{B}(t) = \mathcal{B}(P(t))$ denote the smallest orthogonal box containing $P(t)$, and let $\mathcal{B}^\varepsilon(t) = (1+\varepsilon)B(t)$, scaled with respect to the center of $\mathcal{B}(t)$. We call a box $\hat{B}(t)$ an ε-approximation of $\mathcal{B}(t)$ if $\mathcal{B}(t) \subseteq \hat{\mathcal{B}}(t) \subseteq \mathcal{B}^\varepsilon(t)$. Let Q be an ε-approximation of P, then $\mathcal{B}(Q(t)) \supseteq (1 - \varepsilon)\mathcal{B}(P(t))$, therefore we can compute an ε-approximation of size $O(1/\varepsilon^{d-1/2})$ (if points are moving linearly) and maintain its bounding box. However, one can do better using the following observation.

For $1 \leq i \leq d$, let $P^j(t) = \{p_{ij}(t) \mid 1 \leq i \leq n\}$. Then $\mathcal{B}(t) = \beta_1(t) \times \cdots \times \beta_d(t)$, where $\beta_j(t)$ is the smallest interval containing $P^j(t)$. Hence, the problem of maintaining $\mathcal{B}(t)$ reduces to maintaining the smallest interval containing $P^j(t)$, for each $j \leq d$, using Theorem 4.4 (see also Remark 4.5). We thus compute an ε-approximation Q^j of each P^j and maintain the smallest interval containing Q^j; the latter can be accomplished by maintaining the maximum and minimum of Q^j, using a kinetic tournament tree described in [BGH99]. The data structure processes $O(|Q^j| \log |Q^j|)$ events, and each event requires $O(\log^2 |Q_j|)$ time. Since $P^j(t)$ is a set of n points moving in $\mathbb{R}$, using Theorem 6.3 and putting everything together, we obtain the following.

Theorem 6.6 *Let P be a set of n points moving in $\mathbb{R}^d$, and let $\varepsilon > 0$ be a parameter. If P is moving linearly, then after $O(n+1/\varepsilon)$ preprocessing, we can construct a kinetic data structure of size $O(1/\sqrt{\varepsilon})$ that maintains an ε-approximation of the smallest orthogonal box containing P; the data structure processes $O((1/\sqrt{\varepsilon}) \log(1/\varepsilon))$ events, and each event takes $O(\log^2(1/\varepsilon))$*

time. If the motion of P has degree r, then the preprocessing time is $O(n+1/\varepsilon^{3r/2})$, the size of the data structure is $O(1/\varepsilon)$, the number of events is $O((1/\varepsilon)\log(1/\varepsilon))$, and each event takes $O(\log^2(1/\varepsilon))$ time.

The data structures described above assume that the trajectories of each point is specified in the beginning and it remains fixed. However in most of the applications, we know only a part of the trajectory, and it changes with time. We can handle trajectory updates using the dynamization technique described in Section 5. Since the ε-approximation Q of P being maintained by our algorithm can change significantly after an update operation, we simply reconstruct the kinetic data structure on Q. If we can prove a bound on how much Q changes after an update operation, a kinetic data structure that supports efficient updates can improve the efficiency of our algorithm.

6.3 Minimum-width spherical shell

Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{R}^d$. As defined in Section 1, a spherical shell is (the closure of) the region bounded by two concentric spheres: the width of the shell is the difference of their radii. Let $d(x,p)$ be the Euclidean distance between x and p, and let $f_p(x) = d(x,p)$. Set $\mathcal{F} = \{f_{p_i} \mid p_i \in P\}$. Let $w(S,x)$ denote the width of the thinnest spherical shell centered at x that contains S, and let $w^* = w^*(S) = \min_{x\in\mathbb{R}^d} w(S,x)$ be the width of the thinnest spherical shell containing S. Then

$$w(S,x) = \max_{p\in P} d(x,p) - \min_{p\in P} d(x,p) = \max_{f_p\in\mathcal{F}} f_p(x) - \min_{f_p\in\mathcal{F}} f_p(x) = \mathfrak{I}_{\mathcal{F}}(x).$$

Therefore, $w^* = \min_{x\in\mathbb{R}^d} \mathfrak{I}_{\mathcal{F}}(x)$. It thus suffices to compute an ε-approximation of $\mathcal{F}$. Set

$$g_p(x) = f_p(x)^2 = \|x\|^2 - 2\langle x, p\rangle + \|p_i\|^2.$$

As shown in Section 4 (for $d=2$), $\mathcal{G} = \{g_{p_i} \mid p_i \in P\}$ admits a linearization of dimension $d+1$. However, let $g_i'(x) = g_{p_i}(x) - \|x\|^2$. By Lemma 2.1, an ε-approximation of $\mathcal{G}' = \{g_1' \ldots g_n'\}$ is also an ε-approximation of $\mathcal{G}$. Since $\mathcal{G}'$ admits a linearization of dimension d, by Theorem 4.1, we can compute an ε-approximation of $\mathcal{G}'$, and thus of $\mathcal{G}$, of size $O(1/\varepsilon^d)$ in time $O(1/\varepsilon^d)$. Hence, using Theorem 4.7 (with $r=2$), we can compute an ε-approximation Q of $\mathcal{F}$ of size $O(1/\varepsilon^{2d})$ and compute $w^*(Q)$ in time $1/\varepsilon^{O(d^2)}$ [AAHS00]. However, we can do better using Theorem 4.8. We construct in time $O(n+1/\varepsilon^{3d})$ time a decomposition Ξ of $\mathbb{R}^d$ into $O(1/\varepsilon^{2d})$ Tarski cells along with two functions f_Δ, f_Δ' for each $\Delta \in \Xi$ that $(\varepsilon/2)$-approximate $\mathcal{F}$ within Δ. For each cell $\Delta \in \Xi$, we compute $w_\Delta^* = \min_{x\in\Delta} |f_\Delta(x) - f_\Delta'(x)|$, and then compute $\bar{w} = \min_\Delta w_\Delta^*$ as well as a point $x^* \in \mathbb{R}^d$ that realizes $\bar{w}$. We return the smallest spherical shell centered at x^* that contains P. Note that $w^* \geq \bar{w} \geq (1-\varepsilon/2)\mathfrak{I}_F(x^*)$. Therefore

$$\mathfrak{I}_{\mathcal{F}}(x^*) \leq \frac{1}{1-\varepsilon/2}\bar{w} \leq (1+\varepsilon)w^*.$$

Hence, we obtain the following.

Theorem 6.7 *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 $(1+\varepsilon)w^*(S)$. We can also compute within the same time bound a subset $Q \subseteq S$ so that $w^*(Q) \geq (1-\varepsilon)w^*(S)$.*

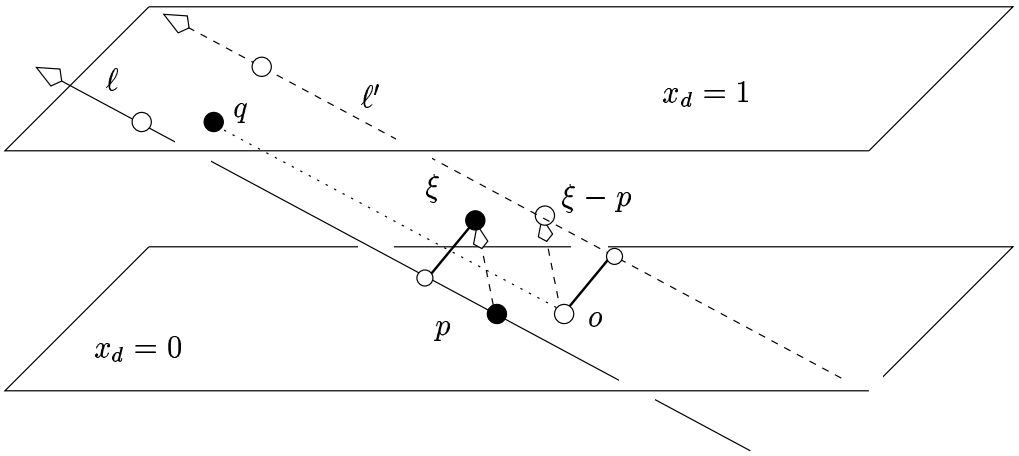

Figure 6. Parametrization of a line ℓ in $\mathbb{R}^3$ and its distance from a point ξ; the small hollow circle on ℓ is the point closest to ξ.

6.4 Minimum-width cylindrical shell

Let $P = \{p_1, \ldots, p_n\}$ be a set of n points in $\mathbb{R}^d$, and a parameter $\varepsilon > 0$. Let $w^* = w^*(P)$ denote the width of the thinnest cylindrical shell, the region lying between two co-axial cylinders, containing P. Let $d(\ell, p)$ denote the distance between a point $p \in \mathbb{R}^d$ and a line $\ell \subset \mathbb{R}^d$. If we fix a line ℓ, then the width of the thinnest cylindrical shell with axis ℓ and containing P is $\mu(\ell) = \max_{p \in P} d(\ell, p) - \min_{p \in P} d(\ell, p)$. A line $\ell \in \mathbb{R}^d$ can be represented by a $(2d-2)$-tuple $(x_1, \ldots, x_{2d-2}) \in \mathbb{R}^{2d-2}$:

$$\ell = \{p + tq \mid t \in \mathbb{R}\},$$

where $p = (x_1, \ldots, x_{d-1}, 0)$ is the intersection point of ℓ with the hyperplane $x_d = 0$ and $q = (x_d, \ldots, x_{2d-2}, 1)$ is the orientation of ℓ (i.e., q is the intersection point of the hyperplane $x_d = 1$ with the line parallel to ℓ and passing through the origin). The distance between ℓ and a point ξ is the same as the distance of the line $\ell' = \{(p - \xi) + tq \mid t \in \mathbb{R}\}$ from the origin; see Figure 6. The point y on ℓ closest to the origin satisfies $y = (p - \xi) + tq$ for some t, and at the same time $\langle y, q \rangle = 0$, which implies that

$$d(\xi, \ell) = \|y\| = \left\| (p - \xi) - \frac{[\langle p - \xi, q \rangle] q}{\|q\|^2} \right\|,$$

Define $f_i(\ell) = d(p_i, \ell)$, and set $\mathcal{F} = \{f_i \mid p_i \in P\}$. Then $w^* = \min_{x \in \mathbb{R}^{2d-2}} \mathfrak{I}_{\mathcal{F}}(x)$. Let $f_i'(x) = \|q\|^2 \cdot f_i(x)$, and set $\mathcal{F}' = \{f_1', \ldots, f_n'\}$. By Lemma 2.1, it suffices to compute an ε-approximation of $\mathcal{F}'$. Define $g_i = f_i'(x)^2$, and let $\mathcal{G} = \{g_1 \ldots, g_n\}$. Then g_i is a $(2d-2)$-variate polynomial and has $O(d^2)$ monomials. Therefore $\mathcal{G}$ admits a linearization of dimension $O(d^2)$. Now proceeding as for spherical shells and using Theorems 4.7 and 4.8, we can compute in $O(n + 1/\varepsilon^{O(d^2)})$ time a set $Q \subseteq P$ of $1/\varepsilon^{O(d^2)}$ points so that $w^*(P) \geq w^*(Q) \geq (1 - \varepsilon) w^*(P)$ as well as a cylindrical shell of width at most $(1 + \varepsilon) w^*(P)$ that contains P. Hence, we conclude the following.

Theorem 6.8 *Given a set P of n points in $\mathbb{R}^d$ and a parameter $\varepsilon > 0$, we can compute in*

$O(n + 1/\varepsilon^{O(d^2)})$ time a subset Q of $1/\varepsilon^{O(d^2)}$ points so that $w^*(Q) \geq (1-\varepsilon)w^*(Q)$ as well as a cylindrical shell containing P whose width is at most $(1+\varepsilon)w^*(P)$.

7 Conclusions

In this paper, we have presented a general technique for computing extent measures approximately. The new technique shows that for many extent measures μ, one can compute in time $O(n+1/\varepsilon^{O(1)})$ a subset Q of size $1/\varepsilon^{O(1)}$ so that $\mu(Q) \geq (1-\varepsilon)\mu(P)$. We then simply compute $\mu(Q)$. Such a subset Q is computed by combining convex-approximation techniques with duality and linearization techniques. Specific applications of our technique include near-linear approximation algorithms for computing minimum width spherical and cylindrical shells, a general technique for approximating faithful measures of stationary as well as moving points. We believe that there are numerous 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.

Interestingly enough, the dynamization and streaming technique presented in Section 5 seems to be generic and applicable to any optimization problem for which a small core-set exists. As such, we expect it to be widely applicable for other problems.

A possible direction for future research is to investigate how practical is this technique, and to improve it further. In particular, it seems that faster algorithms should exist for the problems of approximating the diameter and width of a point set. Another interesting direction for further research is to extent it to handle outliers. Some progress in this direction has recently been made in [HPW].

Acknowledgments

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

References

[AAE00] P. K. Agarwal, L. Arge, and J. Erickson. Indexing moving points. In *Proc. 19th ACM Sympos. Principles Database Syst.*, pages 175–186, 2000.

[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.

[AAS01] P. K. Agarwal, B. Aronov, and M. Sharir. Exact and approximation algorithms for minimum-width cylindrical shells. *Discrete Comput. Geom.*, 26(3):307–320, 2001.

[AGHV01] P. K. Agarwal, L. J. Guibas, J. Hershberger, and E. Veach. Maintaining the extent of a moving point set. *Discrete Comput. Geom.*, 26(3):353–374, 2001.

[AM] P. K. Agarwal and J. Matoušek. Discretizing the space of directions. in preparation.

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

[AS98] P. K. Agarwal and M. Sharir. Efficient algorithms for geometric optimization. *ACM Comput. Surv.*, 30:412–458, 1998.

[AS00] P. K. Agarwal and M. Sharir. Arrangements and their applications. In J.-R. Sack and J. Urrutia, editors, *Handbook of Computational Geometry*, pages 49–119. Elsevier Science Publishers B.V. North-Holland, Amsterdam, 2000.

[BGH99] J. Basch, L. J. Guibas, and J. Hershberger. Data structures for mobile data. *J. Algorithms*, 31(1):1–28, 1999.

[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.

[BI76] E. M. Bronshteyn and L. D. Ivanov. The approximation of convex sets by polyhedra. *Siberian Math. J.*, 16:852–853, 1976.

[BS80] J. L. Bentley and J. B. Saxe. Decomposable searching problems I: Static-to-dynamic transformation. *J. Algorithms*, 1:301–358, 1980.

[CDH+02] Y. Chen, G. Dong, J. Han, B. W. Wah, and J. Wang. Multi-dimensional regression analysis of time-series data streams. In *Proc. International Conference on Very Large Databases*, 2002.

[Cha02] T. M. Chan. Approximating the diameter, width, smallest enclosing cylinder and minimum-width annulus. *Internat. J. Comput. Geom. Appl.*, 12(2):67–85, 2002.

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

[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.

[GKS01] Sudipto Guha, Nick Koudas, and Kyuseok Shim. Data-streams and histograms. In *ACM Symposium on Theory of Computing*, pages 471–475, 2001.

[GMMO00] S. Guha, N. Mishra, R. Motwani, and L. O'Callaghan. Clustering data streams. In *Proc. 41th Annu. IEEE Sympos. Found. Comput. Sci.*, pages 359–366, 2000.

[HPW] S. Har-Peled and Y. Wang. Shape fitting with outliers. submitted for publication.

[KMS02] Flip Korn, S. Muthukrishnan, and D. Srivastava. Reverse nearest neighbour aggregates over data streams. In *Proc. 28th Conference on Very Large Databases*, 2002.

[MP80] J. I. Munro and M. S. Paterson. Selection and sorting with limited storage. *Theoretical Computer Science*, 12:315–323, 1980.

[O'R85] J. O'Rourke. Finding minimal enclosing boxes. *Internat. J. Comput. Inform. Sci.*, 14:183–199, 1985.

[PAH02] C. M. Procopiuc, P. K. Agarwal, and S. Har-Peled. STAR-Tree: An efficient self-adjusting index for moving objects. In D. M. Mount and C. Stein, editors, *Proc. 4th Workshop Algorithm Eng. Exper.*, volume 2409 of *Lect. Notes in Comp. Sci.*, pages 178–193, 2002.

[SJLL00] S. Saltenis, C. S. Jensen, S. Leutenegger, and M. Lopez. Indexing the positions of continuously moving objects. In *ACM SIGMOD Int. Conf. on Manag. of Data*, pages 331–342, 2000.

[YY85] A. C. Yao and F. F. Yao. A general approach to D-dimensional geometric queries. In *Proc. 17th Annu. ACM Sympos. Theory Comput.*, pages 163–168, 1985.

[ZS02] Y. Zhou and S. Suri. Algorithms for a minimum volume enclosing simplex in three dimensions. *SIAM J. Comput.*, 31(5):1339–1357, 2002.

Appendix: Summary of Notations

$\mathbb{R}^d$	d-dimensioanl Euclidean space
$\mathbb{S}^{d-1}$	Unit sphere in $\mathbb{R}^d$
$\mathbb{P}$	$(d-1)$-dimensional projective plane $x_d = 1$
$\mathbb{C}$	d-dimensional unit hypercube $[-1,+1]^d$
$\mathfrak{U}_{\mathcal{F}}$	Upper envelope of $\mathcal{F}$
$\mathfrak{L}_{\mathcal{F}}$	Lower envelope of $\mathcal{F}$
$\mathfrak{I}_{\mathcal{F}}$	Extent of $\mathcal{F}$
$\mathcal{A}(\mathcal{J})$	Arrangement of $\mathcal{J}$
$\mathcal{CH}(S)$	Convex hull of S
$\phi(v)$	$v/\|v\|$, $v \in \mathbb{R}^d$
$\widetilde{u}$	$(u,1) \in \mathbb{P}$, $u \in \mathbb{R}^{d-1}$
u^*	$\phi(\widetilde{u}) \in \mathbb{S}^{d-1}$, $u \in \mathbb{R}^{d-1}$.
$\overline{\omega}(x,P)$	$\max_{p\in P}\langle x,p\rangle - \min_{p\in P}\langle x,p\rangle$, $x \in \mathbb{R}^d, P \subseteq \mathbb{R}^d$
$\omega(u,P)$	$\overline{\omega}(\widetilde{u},P)$, $u \in \mathbb{R}^{d-1}$, $P \subseteq \mathbb{R}^d$

Table 1. Summary of notations used in the paper.