On the Combinatorial Complexity of Approximating Polytopes

Sunil Arya
Department of Computer Science and Engineering
Hong Kong University of Science and Technology
Clear Water Bay, Kowloon, Hong Kong
arya@cse.ust.hk

Guilherme D. da Fonseca
Université d'Auvergne and LIMOS
Clermont-Ferrand, France
guilherme.dias_da_fonseca@udamail.fr

David M. Mount
Department of Computer Science and
Institute for Advanced Computer Studies
University of Maryland
College Park, Maryland 20742
mount@cs.umd.edu

December 21, 2016

Abstract

Approximating convex bodies succinctly by convex polytopes is a fundamental problem in discrete geometry. A convex body K of diameter $\text{diam}(K)$ is given in Euclidean d-dimensional space, where d is a constant. Given an error parameter $\varepsilon > 0$, the objective is to determine a polytope of minimum combinatorial complexity whose Hausdorff distance from K is at most $\varepsilon \cdot \text{diam}(K)$. By combinatorial complexity we mean the total number of faces of all dimensions of the polytope. A well-known result by Dudley implies that $O(1/\varepsilon^{(d-1)/2})$ facets suffice, and a dual result by Bronshteyn and Ivanov similarly bounds the number of vertices, but neither result bounds the total combinatorial complexity. We show that there exists an approximating polytope whose total combinatorial complexity is $\widetilde{O}(1/\varepsilon^{(d-1)/2})$, where $\widetilde{O}$ conceals a polylogarithmic factor in $1/\varepsilon$. This is a significant improvement upon the best known bound, which is roughly $O(1/\varepsilon^{d-2})$.

Our result is based on a novel combination of both old and new ideas. First, we employ Macbeath regions, a classical structure from the theory of convexity. The construction of our approximating polytope employs a new stratified placement of these regions. Second, in order to analyze the combinatorial complexity of the approximating polytope, we present a tight analysis of a width-based variant of Bárány and Larman's *economical cap covering*. Finally, we use a deterministic adaptation of the witness-collector technique (developed recently by Devillers *et al.*) in the context of our stratified construction.

Keywords: Convex polytopes, polytope approximation, combinatorial complexity, Macbeath regions

1 Introduction

Approximating general convex bodies by convex polytopes is a fundamental geometric problem. It has been extensively studied in the literature under various formulations. (See Bronstein [14] for a survey.) Consider a convex body K, that is, a closed, convex set of bounded diameter, in Euclidean d-dimensional space. At issue is the structure of the simplest polytope P that approximates K.

There are various ways to define the notions of "simplest" and "approximates." Our notion of approximation will be based on the *Hausdorff metric*, that is, the maximum distance between a point in the boundary of P or K and the boundary of the other body. Normally, approximation error is defined relative to K's diameter. It will simplify matters to assume that K has been uniformly scaled to unit diameter. For a given error $\varepsilon > 0$, we say that a polytope P is an *ε-approximating polytope* to K if the Hausdorff distance between K and P is at most ε. The simplicity of an approximating polytope P will be measured in terms of its *combinatorial complexity*, that is, the total number of k-faces, for $0 \le k \le d - 1$. For the purposes of stating asymptotic bounds, we assume that the dimension d is a constant.

The bounds given in the literature for convex approximation are of two common types [14]. In both cases, the bounds hold for all $\varepsilon \le \varepsilon_0$, for some $\varepsilon_0 > 0$. In *nonuniform bounds*, the value of ε_0 depends on K (for example, on K's maximum curvature). Such bounds are often stated as holding "in the limit" as ε approaches zero, or equivalently as the combinatorial complexity of the approximating polytope approaches infinity. Examples include bounds by Gruber [20], Clarkson [16], and others [11, 26, 27]. Our interest is in *uniform bounds*, where the value of ε_0 is independent of K. Examples include the results of Dudley [18] and Bronshteyn and Ivanov [13]. Such bounds hold without any assumptions on K.

Dudley showed that, for $\varepsilon \le 1$, any convex body K of unit diameter can be ε-approximated by a convex polytope P with $O(1/\varepsilon^{(d-1)/2})$ facets. This bound is known to be optimal in the worst case and is achieved when K is a Euclidean ball (see, e.g., [14]). Alternatively, Bronshteyn and Ivanov showed the same bound holds for the number of vertices, which is also the best possible. No convex polytope approximation is known that attains both bounds simultaneously.[1]

Establishing good uniform bounds on the combinatorial complexity of convex polytope approximations is a major open problem. The Upper-Bound Theorem [24] implies that a polytope with n vertices (resp., facets) has total combinatorial complexity $O(n^{\lfloor d/2 \rfloor})$. Applying this to the results of either Dudley or Bronshteyn and Ivanov directly yields a bound of $O(1/\varepsilon^{(d^2-d)/4})$ on the combinatorial complexity of an ε-approximating polytope. Better uniform bounds without d^2 in the exponent are known, however. Consider a uniform grid Ψ of points with spacing $\Theta(\varepsilon)$, and let P denote the convex hull of $\Psi \cap K$. It is easy to see that P is an ε-approximating polytope for K. The combinatorial complexity of any lattice polytope[2] is known to be $O(V^{(d-1)/(d+1)})$, where V is the volume of the polytope [2, 9]. This implies that P has combinatorial complexity $O(1/\varepsilon^{d(d-1)/(d+1)}) \approx O(1/\varepsilon^{d-2})$. While this is significantly better than the bound provided by the Upper-Bound Theorem, it is still much larger than the lower bound of $\Omega(1/\varepsilon^{(d-1)/2})$.

We show that this gap can be dramatically reduced. In particular, we establish an upper bound on the combinatorial complexity of convex approximation that is optimal up to a polylogarithmic factor in $1/\varepsilon$.

---

[1]Jeff Erickson noted that both bounds can be attained simultaneously but at the cost of sacrificing convexity [16].

[2]A *lattice polytope* is the convex hull of any set of points with integer coordinates.

Theorem 1.1. *Let $K \subset \mathbb{R}^d$ be a convex body of unit diameter, where d is a fixed constant. For all sufficiently small positive ε (independent of K) there exists an ε-approximating convex polytope P to K of combinatorial complexity $O(1/\widehat{\varepsilon}^{(d-1)/2})$, where $\widehat{\varepsilon} = \varepsilon / \log(1/\varepsilon)$.*

This is within a factor of $O(\log^{(d-1)/2}(1/\varepsilon))$ of the aforementioned lower bound. Our approach employs a classical structure from the theory of convexity, called *Macbeath regions* [23]. Macbeath regions have found numerous uses in the theory of convex sets and the geometry of numbers (see Bárány [8] for an excellent survey). They have also been applied to a small but growing number of results in the field of computational geometry (see, e.g., [3, 5, 6, 12]). Our construction of the approximating polytope uses a new stratified placement of these regions. In order to analyze the combinatorial complexity of the approximating polytope, in Section 3 we present a tight analysis of a width-based variant of Bárány and Larman's economical cap covering. This result plays a central role in our recent work on approximate polytope membership queries [4] and may find use in other applications. Finally, we employ a deterministic version of the witness-collector technique, developed recently by Devillers *et al.* [17], in the context of our stratified construction.

The paper is organized as follows. In Section 2, we define concepts related to Macbeath regions and present some of their key properties. In Section 3, we prove the width-based economical cap covering lemma. The stratified placement of the Macbeath regions and the bound on the combinatorial complexity of approximating polytopes follow in Section 4. We conclude with several open problems in Section 5.

2   Geometric Preliminaries

Recall that K is a convex body of unit diameter in $\mathbb{R}^d$. Let ∂K denote its boundary. Let O denote the origin of $\mathbb{R}^d$, and for $x \in \mathbb{R}^d$ and $r \geq 0$, let $B^r(x)$ denote the Euclidean ball of radius r centered at x. It will be convenient to first map K to a convenient form. We say that a convex body K is in *canonical form* if $B^{1/2d}(O) \subseteq K \subseteq B^{1/2}(O)$. Given a parameter $0 < \gamma \leq 1$, we say that a convex body K is γ-*fat* if there exist concentric Euclidean balls B and B', such that $B \subseteq K \subseteq B'$, and $\mathrm{radius}(B)/\mathrm{radius}(B') \geq \gamma$. Thus, a body in canonical form is $(1/d)$-fat and has diameter $\Theta(1)$. We will refer to point O as the *center* of K.

The following lemma shows that, up to constant factors, the problem of approximating an arbitrary convex body can be reduced to approximating a convex body in canonical form. The proof follows from a combination of John's Theorem [22] and Lemma 3.1 of Agarwal *et al.* [1] and is included for completeness.

Lemma 2.1. *Let K be a convex body of unit diameter in $\mathbb{R}^d$. There exists a non-singular affine transformation T such that $T(K)$ is in canonical form and if P is any (ε/d)-approximating polytope to $T(K)$, then $T^{-1}(P)$ is an ε-approximating polytope to K.*

Proof. Let E denote a maximum volume ellipsoid enclosed within K (that is, the John ellipsoid). Since K is of unit diameter, E's semi-principal axes are all of length at most $1/2$. Consider a frame centered at E's center and whose axes coincide with E's semi-principal axes. Let T be an affine transformation that maps this frame's origin to the origin of the space, and scales all of the frame's basis vectors to length $1/2d$. This affine transformation maps E to $B^{1/2d}(O)$. Since each of the frame's basis vectors is scaled from a length of at most $1/2$ to a length of $1/2d$, it follows that T maps any vector v to a vector of length at least $\|v\|/d$. Thus, T^{-1} maps any vector v to a vector

of length at most $d\|v\|$. Therefore, if P is any (ε/d)-approximating polytope to $T(K)$, $T^{-1}(P)$ is an ε-approximating polytope to $T^{-1}(T(K)) = K$, as desired. □

We assume henceforth that K is given in canonical form and that ε has been appropriately scaled. This scaling only affects the constant factors in our asymptotic bounds.

A *cap* C is defined to be the nonempty intersection of the convex body K with a halfspace H (see Figure 1(a)). Let h denote the hyperplane bounding H. We define the *base* of C to be $h \cap K$. The *apex* of C is any point in the cap such that the supporting hyperplane of K at this point is parallel to h. The *width* of C is the distance between h and this supporting hyperplane. Given any cap C of width w and a real $\lambda \geq 0$, we define its λ-*expansion*, denoted C^λ, to be the cap of K cut by a hyperplane parallel to and at distance λw from this supporting hyperplane. (Note that $C^\lambda = K$, if λw exceeds the width of K along the defining direction.) An easy consequence of convexity is that, for $\lambda \geq 1$, C^λ is a subset of the region obtained by scaling C by a factor of λ about its apex. It follows that, for $\lambda \geq 1$, $\mathrm{vol}(C^\lambda) \leq \lambda^d \cdot \mathrm{vol}(C)$. For a given $\varepsilon > 0$, let $K(\varepsilon) \subset K$ denote the points of K within distance at most ε from ∂K (equivalently, the union of all ε-width caps).
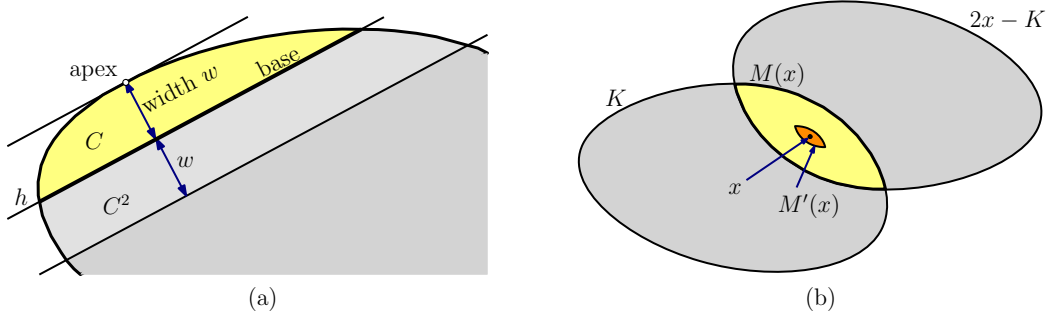

Figure 1: (a) Cap concepts and (b) Macbeath regions.

Given a point $x \in K$ and real parameter $\lambda \geq 0$, the *Macbeath region* $M^\lambda(x)$ (also called an *M-region*) is defined as:
$$M^\lambda(x) \;=\; x + \lambda((K - x) \cap (x - K)).$$

It is easy to see that $M^1(x)$ is the intersection of K and the reflection of K around x (see Figure 1(b)), and so $M^1(x)$ is centrally symmetric about x. $M^\lambda(x)$ is a scaled copy of $M^1(x)$ by the factor λ about x. We refer to x as the *center* of $M^\lambda(x)$ and to λ as its *scaling factor*. As a convenience, we define $M(x) = M^1(x)$ and $M'(x) = M^{1/5}(x)$.

We begin with two lemmas that encapsulate relevant properties of Macbeath regions. Both were proved originally by Ewald, Larman, and Rogers [19], but our statements follow the forms given by Brönnimann, Chazelle, and Pach [12]. (Lemmas 2.2 and 2.3 below are restatements of Lemmas 2.5 and 2.6 from [12], respectively.)

Lemma 2.2. *Let K be a convex body. If $x, y \in K$ such that $M'(x) \cap M'(y) \neq \emptyset$, then $M'(y) \subseteq M(x)$.*

Lemma 2.3. *Let $K \subset \mathbb{R}^d$ be a convex body in canonical form, and let $\Delta_0 = 1/(6d)$ be a constant. Let C be a cap of K of width at most Δ_0. Let x denote the centroid of the base of this cap. Then $C \subseteq M^{3d}(x)$.*

The following lemma is an immediate consequence of the definition of Macbeath region.

4

Lemma 2.4. *Let K be a convex body and $\lambda > 0$. If x is a point in a cap C of K, then $M^\lambda(x) \cap K \subseteq C^{1+\lambda}$. Furthermore, if $\lambda \leq 1$, then $M^\lambda(x) \subseteq C^{1+\lambda}$.*

The next lemma is useful in situations when we know that a Macbeath region partially overlaps a cap of K. It allows us to conclude that a constant factor expansion of the cap will fully contain the Macbeath region.

Lemma 2.5. *Let K be a convex body. Let C be a cap of K and x be a point in K such that $C \cap M'(x) \neq \emptyset$. Then $M'(x) \subseteq C^2$.*

Proof. Let y be any point in $C \cap M'(x)$. Since $M'(x) \cap M'(y) \neq \emptyset$ obviously holds, we can apply Lemma 2.2 to conclude that $M'(x) \subseteq M(y)$. By Lemma 2.4 (with $\lambda = 1$), $M(y) \subseteq C^2$. It follows that $M'(x) \subseteq C^2$. $\qquad\square$

Next, we give two straightforward lemmas dealing with scaling of centrally symmetric convex bodies. As Macbeath regions are centrally symmetric, these lemmas will be useful to us in conjunction with their standard properties. A proof of Lemma 2.6 appears in Bárány [7]. For any centrally symmetric convex body A, define A^λ to be the body obtained by scaling A by a factor of λ about its center.

Lemma 2.6. *Let $\lambda \geq 1$. Let A and B be centrally symmetric convex bodies such that $A \subseteq B$. Then $A^\lambda \subseteq B^\lambda$.*

Lemma 2.7. *Let $\lambda \geq 1$. Let A be a centrally symmetric convex body. Let A' be the body obtained by scaling A by a factor of λ about any point in A. Then $A' \subseteq A^{2\lambda-1}$.*

Proof. We take the origin to be at the center of A. Let A' be the body obtained by scaling A by a factor of λ about a point $a \in A$. Any point u in A' is of the form $a + \lambda(x - a)$, where $x \in A$. This can be expressed as

$$(2\lambda - 1)\left[\frac{\lambda}{2\lambda - 1}x + \frac{\lambda - 1}{2\lambda - 1}(-a)\right].$$

Since $\lambda \geq 1$, the point $(\lambda/(2\lambda - 1))x + ((\lambda - 1)/(2\lambda - 1))(-a)$ lies on the segment joining x and $-a$. Since both x and $-a$ lie within A, it follows that $u \in A^{2\lambda-1}$, as desired. $\qquad\square$

The following lemma is an easy consequence of Lemmas 2.3 and 2.7.

Lemma 2.8. *Let $\lambda \geq 1$ and let K, C, and x be as defined in Lemma 2.3. Then $C^\lambda \subseteq M^{3d(2\lambda-1)}(x)$.*

Proof. By Lemma 2.3, $C \subseteq M^{3d}(x)$. Recall that C^λ is contained within the region obtained by scaling C by a factor of λ about its apex. Applying Lemma 2.7 (applied to $M^{3d}(x)$ and the apex point), it follows that $C^\lambda \subseteq M^{3d(2\lambda-1)}(x)$. $\qquad\square$

The well known Lemma 2.2 states that if two $(1/5)$-shrunken Macbeath regions have a nonempty intersection, then a constant factor expansion of one contains the other [12,19]. We show next that this holds for the associated caps as well. (Note that this does not hold in general for overlapping caps. If two caps C_1 and C_2 have a nonempty intersection, there is no constant β that guarantees that $C_1 \subseteq C_2^\beta$.)

Lemma 2.9. *Let Δ_0 be the constant of Lemma 2.3 and let $\lambda \geq 1$ be any real. There exists a constant $\beta \geq 1$ such that the following holds. Let $K \subset \mathbb{R}^d$ be a convex body in canonical form. Let C_1 and C_2 be any two caps of K of width at most Δ_0. Let x_1 and x_2 denote the centroids of the bases of the caps C_1 and C_2, respectively. If $M'(x_1) \cap M'(x_2) \neq \emptyset$, then $C_1^\lambda \subseteq C_2^{\beta\lambda}$.*

Proof. By Lemma 2.8, $C_1^\lambda \subseteq M^\alpha(x_1)$, where $\alpha = 3d(2\lambda - 1)$. Since $M'(x_1)$ and $M'(x_2)$ overlap, by Lemma 2.2, $M'(x_1) \subseteq M(x_2)$. By definition, $M'(x_1) = M^{1/5}(x_1)$ and so $M^\alpha(x_1) = (M'(x_1))^{5\alpha}$. Since $M'(x_1)$ and $M(x_2)$ are centrally symmetric bodies and $M'(x_1) \subseteq M(x_2)$, by Lemma 2.6, it follows that $(M'(x_1))^{5\alpha} \subseteq M^{5\alpha}(x_2)$. Putting it together, we obtain

$$C_1^\lambda \subseteq M^\alpha(x_1) = (M'(x_1))^{5\alpha} \subseteq M^{5\alpha}(x_2).$$

By Lemma 2.4, $M^{5\alpha}(x_2) \cap K \subseteq C_2^{1+5\alpha}$. Since $C_1^\lambda \subseteq M^{5\alpha}(x_2)$ and $C_1^\lambda \subseteq K$, we have $C_1^\lambda \subseteq M^{5\alpha}(x_2) \cap K \subseteq C_2^{1+5\alpha}$. Recalling that $\alpha = 3d(2\lambda - 1)$, we have $C_1^\lambda \subseteq C_2^{30d\lambda}$. This proves the lemma for constant $\beta = 30d$. $\square$

3 Economical Cap Covering

In this section we present a tight analysis of a width-based variant of Bárány and Larman's economical cap covering [10]. The lemma applies generally to any convex body K that has constant diameter and is γ-fat for some constant γ (where the constants may depend on d). The proof of this lemma follows from the ideas in [7, 10, 19]. Our principal contribution is an optimal bound of $O(1/\varepsilon^{(d-1)/2})$ on the number of bodies needed.

Lemma 3.1 (Width-based economical cap covering lemma). *Let $\varepsilon > 0$ be a sufficiently small parameter. Let $K \subset \mathbb{R}^d$ be a convex body in canonical form. There exists a collection $\mathcal{R}$ of $k = O(1/\varepsilon^{(d-1)/2})$ disjoint centrally symmetric convex bodies $R_1, \ldots, R_k$ (see Figure 2(a)) and associated caps $C_1, \ldots, C_k$ such that the following hold (for some constants β and λ, which depend only on d):*

1. *For each i, C_i is a cap of width $\beta\varepsilon$, and $R_i \subseteq C_i \subseteq R_i^\lambda$.*

2. *Let C be any cap of width ε. Then there is an i such that $R_i \subseteq C$ and $C_i^{1/\beta^2} \subseteq C \subseteq C_i$ (see Figure 2(b)).*
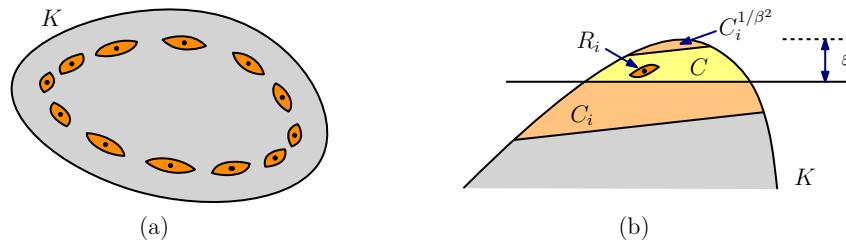

Figure 2: Illustrating Lemma 3.1.

The R_i's in this lemma are Macbeath regions with scaling factor $1/5$. Since any cap of width ε is contained in some cap C_i, it follows that the C_i's together cover $K(\varepsilon)$. Further, from Property 1,

6

we can see that the sum of the volume of the C_i's is no more than a constant times the volume of $K(\varepsilon)$. It is in this sense that the C_i's constitute an *economical* cap covering.

It is worth mentioning that Property 2 is stronger than similar properties given previously in the literature in the following sense. For any cap of width ε, we show not merely that it is contained within some cap C_i of the cover, but it is effectively "sandwiched" between two caps with parallel bases, each of width $\Theta(\varepsilon)$.

A key technical contribution of our paper is the following lemma. It will help us bound the number of bodies needed in the width-based cap covering lemma. Because of its broader utility, this lemma is given in a slightly more general form than is needed here.

Lemma 3.2. *Let $K \subset \mathbb{R}^d$ be a convex body in canonical form. Let $0 < \delta \le \Delta_0/2$, where Δ_0 is the constant of Lemma 2.3. Let $\mathcal{C}$ be a set of caps, whose widths lie between δ and 2δ, such that the Macbeath regions $M'(x)$ centered at the centroids x of the bases of these caps are disjoint. Then $|\mathcal{C}| = O(1/\delta^{(d-1)/2})$.*

Our proof of Lemma 3.2 will require the following geometric observation, which is a straightforward extension of Dudley's convex approximation construction (see Lemma 4.4 of [18]). It is similar to other results based on Dudley's construction (including Lemma 3.6 of [1] and Lemma 23.12 of [21]). We will present the proof for the sake of completeness. Let S denote the sphere of radius 2 centered at the origin O, which we call the *Dudley sphere*. Given vectors u and v, let $\langle u, v \rangle$ denote their dot product and let $\|u\| = \langle u, u \rangle^{1/2}$ denote u's Euclidean length.

Lemma 3.3. *Let K be a convex body that lies within a unit sphere centered at the origin, and let $0 < \delta \le 1$. Let x' and y' be two points of S. Let x and y be the points of ∂K that are closest to x' and y', respectively. Let h denote the supporting hyperplane at x orthogonal to the segment xx'. Let C denote the cap cut from K by a hyperplane parallel to and at distance δ from h. If $y \notin C$, then $\|x' - y'\| \ge \sqrt{\delta}$.*

Proof. Before starting the proof, we recall a technical result (Lemma 4.3) from Dudley [18], which states that given vectors x, y, u, v in $\mathbb{R}^d$ such that $\langle x - y, u \rangle \ge 0$ and $\langle x - y, v \rangle \le 0$, $\|(x + u) - (y + v)\| \ge \max(\|x - y\|, \|u - v\|)$. This follows from the observation that

$$\|(x + u) - (y + v)\|^2 \;=\; \|x - y\|^2 + \|u - v\|^2 + 2 \langle x - y, u - v \rangle \;\ge\; \|x - y\|^2 + \|u - v\|^2.$$

Returning to the proof, suppose towards a contradiction that $y \notin C$ but $\|x' - y'\| < \sqrt{\delta}$. Let $u = x' - x$ and $v = y' - y$, and let $\widehat{u} = u/\|u\|$ and $\widehat{v} = v/\|v\|$ (see Figure 3). Clearly, $\|(x + u) - (y + v)\| = \|x' - y'\| < \sqrt{\delta}$. A direct consequence of convexity is that $\langle x - y, u \rangle \ge 0$ and $\langle x - y, v \rangle \le 0$, and so by the above result it follows that $\|x - y\|$ and $\|u - v\|$ are both less than $\sqrt{\delta}$. Clearly, u and v are of at least unit length, and thus $\|\widehat{u} - \widehat{v}\| \le \|u - v\| < \sqrt{\delta}$. Let θ denote the angle between $\widehat{u}$ and $\widehat{v}$. Since $\|x' - y'\| < \sqrt{\delta} \le 1$ and the radius of S is 2, it follows that $\theta < \pi/2$.

Consider the right triangle whose hypotenuse is xy and whose third vertex is the orthogonal projection of y onto the supporting hyperplane h, which we denote by z. Letting $\phi = \angle zxy$, it follows from convexity that $\phi \le \theta$. (This is because any supporting hyperplane through y cannot pass below x.) Because $\theta < \pi/2$, $\sin \theta \ge \sin \phi$. Also, since $y \notin C$, we have $\|z - y\| > \delta$, and therefore

$$\sin \theta \;\ge\; \sin \phi \;=\; \frac{\|z - y\|}{\|x - y\|} \;>\; \frac{\delta}{\sqrt{\delta}} \;=\; \sqrt{\delta}.$$
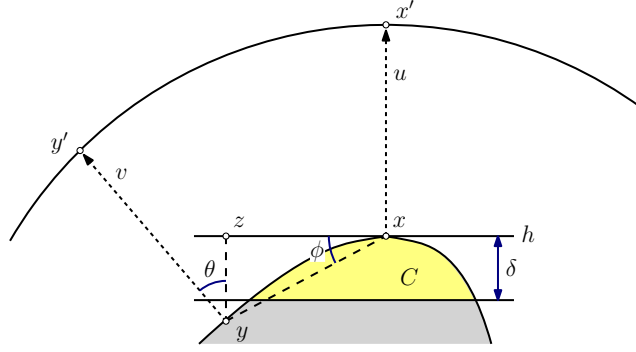
Figure 3: Illustrating Lemma 3.3.

Observe that $\|\widehat{u} - \widehat{v}\|$ is the length of a chord of a unit circle that subtends an arc of angle θ, and therefore $\|\widehat{u} - \widehat{v}\| = 2\sin\frac{\theta}{2}$. Given our earlier bound on this distance, we obtain the following contradiction:

$$\sqrt{\delta} \; < \; \sin\theta \; = \; 2\sin\frac{\theta}{2}\cos\frac{\theta}{2} \; \leq \; 2\sin\frac{\theta}{2} \; = \; \|\widehat{u} - \widehat{v}\| \; < \; \sqrt{\delta}.$$

$\square$

We are now ready to present the proof of Lemma 3.2.

Proof. (of Lemma 3.2) Let A be the set of disjoint Macbeath regions $M'(x)$ described in the lemma. For each region $M'(x)$, let $C(x)$ denote the cap whose base centroid point generates $M'(x)$. We begin by pruning A to obtain a subset B, which to within constant factors has the same cardinality as A. We construct B incrementally as follows. Initially B is the empty set. In each step, from among the Macbeath regions that still remain in A, we choose a Macbeath region $M'(x)$ that has the smallest volume, and insert it into B. We then prune all the Macbeath regions from A that intersect the cap $C^4(x)$. We continue in this manner until A is exhausted.

We claim that in each step, we prune a constant number of Macbeath regions from A. Let $M'(x)$ denote the Macbeath region inserted into B in this step. If $M'(y)$ is a Macbeath region that is pruned in this step, then $M'(y)$ intersects the cap $C^4(x)$. It then follows from Lemma 2.5 that $M'(y) \subseteq C^8(x)$. Note that

$$\mathrm{vol}(C^8(x)) \; \leq \; 8^d \mathrm{vol}(C(x)) \; = \; O(\mathrm{vol}(C(x))).$$

Since $C(x)$ is of width at most $2\delta \leq \Delta_0$, we may apply Lemma 2.3, which yields $C(x) \subseteq M^{3d}(x)$. It follows that

$$\mathrm{vol}(M(x)) \; \geq \; \mathrm{vol}(C(x))/(3d)^d \; = \; \Omega(\mathrm{vol}(C(x))).$$

Recall that each Macbeath region pruned has volume greater than or equal to the volume of $M'(x)$. It follows that the volume of each Macbeath region pruned is $\Omega(\mathrm{vol}(M(x))) = \Omega(\mathrm{vol}(C(x)))$. Since the pruned Macbeath regions are disjoint and contained in a region of volume $O(\mathrm{vol}(C(x)))$, a straightforward packing argument implies that the number of Macbeath regions pruned is $O(1)$.

The claim immediately implies that $|A| = O(|B|)$. In the remainder of the proof, we will show that $|B| = O(1/\delta^{(d-1)/2})$, which will complete the proof.

Let X denote the set of centers of the Macbeath regions of B, that is, $X = \{x : M'(x) \in B\}$. We map each point $x \in X$ to a point x' on the Dudley sphere such that xx' is normal to the base of the cap $C(x)$. We claim that the distance between any pair of the projected points x' on the Dudley sphere is at least $\sqrt{\delta}$. Note that this claim would imply the desired bound on $|B|$ and complete the proof.

To see this claim, consider any two Macbeath regions $M'(x)$ and $M'(y)$ in the set B. Without loss of generality, suppose that $M'(y)$ is inserted into B after $M'(x)$. By our construction, it follows that y is not contained in $C^4(x)$ (because otherwise $M'(y)$ would intersect $C^4(x)$ and would have been pruned after inserting $M'(x)$ into B). We now consider two cases, depending on whether or not x is contained in $C(y)$.

Case 1: $(x \notin C(y))$ Consider the convex body K' that is the closure of $K \setminus (C(x) \cup C(y))$ (outlined in red in Figure 4(a)). Note that x and y are on the boundary of the convex body K' and these are the points of $\partial K'$ that are closest to x' and y', respectively. Next, consider the cap of K' whose apex is x and width is δ. Call this cap $C'(x)$. Since the width of $C(x)$ is at least δ, and $y \notin C^4(x)$, it is easy to see that $y \notin C'(x)$. Applying Lemma 3.3 to the convex body K' and the points x', y', x, and y, it follows that $\|x'y'\| \geq \sqrt{\delta}$.
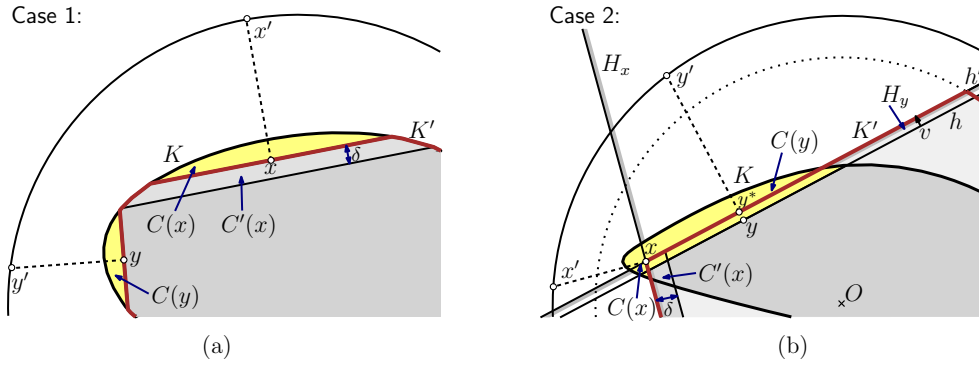

Figure 4: Cases arising in the proof of Lemma 3.2. (Figure not to scale.)

Case 2: $(x \in C(y))$ Let h denote the hyperplane that forms the base of $C(y)$ (see Figure 4(b)). Let h' denote the hyperplane parallel to h that passes through x. Let v denote the vector normal to h, whose magnitude is the distance between h and h'. Note that $h' = h + v$. Since $C(y)$ is a cap of width at most 2δ, the magnitude of the translation vector v is at most 2δ. Let $y^* = y + v$. Let H_y denote the halfspace bounded by h' that contains the origin. Let H_x denote the halfspace that contains the origin and whose boundary is the hyperplane forming the base of $C(x)$. Define the convex body K' as the intersection of H_x and H_y and a ball of unit radius centered at the origin. Note that x and y^* lie on the boundary of K' (since $\|Ox\| < 1$ and $\|Oy^*\| < 1$; $\|Ox\| < 1$ holds trivially since $x \in K$ and $K \subseteq B^{1/2}(O)$, and $\|Oy^*\| \leq \|Oy\| + \|yy^*\| \leq 1/2 + 2\delta \leq 1/2 + 2\Delta_0 < 1$).

Further, the points x and y^* are the points of $\partial K'$ that are closest to x' and y', respectively. Next, consider the cap of K' whose apex is x and width is δ and whose base is parallel to the base of $C(x)$. Call this cap $C'(x)$. Recall that $y \notin C^4(x)$, the width of $C(x)$ is at least δ, and the distance between y and y^* is at most 2δ. It follows that y^* is at distance bigger than $3\delta - 2\delta = \delta$ from the hyperplane passing through the base of $C(x)$. Since the distance between the hyperplanes passing through the bases of $C(x)$ and $C'(x)$, respectively, is δ, it follows that $y^* \notin C'(x)$. Applying

9

Lemma 3.3 to the convex body K' and the points x', y', x, and y^*, it follows that the distance between x' and y' is at least $\sqrt{\delta}$. This establishes the above claim and completes the proof. $\quad\square$

The remainder of this section is devoted to proving Lemma 3.1.

Proof. Assume that $\varepsilon \le \Delta_0$, where Δ_0 is the constant of Lemma 2.3. Let $\beta = 30d$ be the constant of Lemma 2.9. Let $\mathcal{C}$ be a maximal set of caps, each of width ε/β, such that the $(1/5)$-scaled Macbeath regions centered at the centroids of the bases of these caps are disjoint. Let $A_1, \ldots, A_k$ denote the caps of $\mathcal{C}$. Let x_i denote the centroid of the base of cap A_i. With each cap A_i, we associate a convex body $R_i = M'(x_i)$ and a cap $C_i = A_i^{\beta^2}$. We will show that the convex bodies R_i and caps C_i satisfy the properties given in the lemma.

By Lemma 3.2, $|\mathcal{C}| = O(1/\varepsilon^{(d-1)/2})$, which implies the desired upper bound on k. Since C_i is a β^2-expansion of A_i, its width is $\beta\varepsilon$. To prove Property 1, it remains to show that $M'(x_i) \subseteq C_i \subseteq (M'(x_i))^\lambda$. By Lemma 2.4, $M'(x_i) \subseteq A_i^{6/5}$. Since $A_i^{6/5} \subseteq A_i^{\beta^2} = C_i$, we obtain $M'(x_i) \subseteq C_i$. Also, applying Lemma 2.8, we obtain

$$C_i \;=\; A_i^{\beta^2} \;\subseteq\; M^{3d(2\beta^2-1)}(x_i) \;=\; (M'(x_i))^{15d(2\beta^2-1)} \;\subseteq\; (M'(x_i))^\lambda,$$

where $\lambda = 30d\beta^2$. Thus, $M'(x_i) \subseteq C_i \subseteq (M'(x_i))^\lambda$.

To show Property 2, let C be any cap of width ε. Let x denote the centroid of the base of $C^{1/\beta}$. By maximality of $\mathcal{C}$, there must be a Macbeath region $M'(x_i)$ that has a nonempty intersection with $M'(x)$ (note x_i may be the same as point x). Applying Lemma 2.2, it follows that $M'(x_i) \subseteq M(x)$. By Lemma 2.4, $M(x) \subseteq C^{2/\beta}$. Putting it together, we obtain $M'(x_i) \subseteq M(x) \subseteq C^{2/\beta} \subseteq C$, which establishes the first part of Property 2.

It remains to show that $C_i^{1/\beta^2} \subseteq C \subseteq C_i$. Since $M'(x_i) \cap M'(x) \ne \emptyset$, we can apply Lemma 2.9 to caps A_i and $C^{1/\beta}$ (for $\lambda = 1$) to obtain $A_i \subseteq (C^{1/\beta})^\beta$. Applying Lemma 2.9 again to caps $C^{1/\beta}$ and A_i (for $\lambda = \beta$), we obtain $(C^{1/\beta})^\beta \subseteq A_i^{\beta^2}$. Thus $A_i \subseteq C \subseteq A_i^{\beta^2}$. Recalling that $C_i = A_i^{\beta^2}$, we obtain $C_i^{1/\beta^2} \subseteq C \subseteq C_i$, as desired. $\quad\square$

4 Polytope Approximation

In this section, we will show how to obtain an ε-approximating convex polytope P of low combinatorial complexity. Let K be a convex body in canonical form. Our strategy is as follows. First, we build a set $\mathcal{R}$ of disjoint centrally symmetric convex bodies lying within K and close to its boundary. These bodies will possess certain key properties to be specified later. For each $R \in \mathcal{R}$, we select a point arbitrarily from this body, and let S denote this set of points. The approximation P is defined as the convex hull of S. In Lemma 4.10, we will prove that P is an ε-approximation of K and, in Lemma 4.11, we will apply a deterministic variant of the witness-collector approach [17] to show that P has low combinatorial complexity.

Before delving into the details, we provide a high-level overview of the witness-collector method, adapted to our context. Let $\mathcal{H}$ denote the set of all halfspaces in $\mathbb{R}^d$. We define a set $\mathcal{W}$ of regions called *witnesses* and a set $\mathcal{C}$ of regions called *collectors*, which satisfy the following properties:

(1) Each witness of $\mathcal{W}$ contains a point of S in its interior.

(2) Any halfspace $H \in \mathcal{H}$ either contains a witness $W \in \mathcal{W}$ or $H \cap S$ is contained in a collector $C \in \mathcal{C}$.

(3) Each collector $C \in \mathcal{C}$ contains a constant number of points of S.

The key idea of the witness-collector method is encapsulated in the following lemma.

Lemma 4.1. *Given a set of witnesses and collectors satisfying the above properties, the combinatorial complexity of the convex hull P of S is $O(|\mathcal{C}|)$.*

Proof. We map each face f of P to any maximal subset $S_f \subseteq S$ of affinely independent points on f. Note that this is a one-to-one mapping and $|S_f| \leq d$. In order to bound the combinatorial complexity of P it suffices to bound the number of such subsets S_f.

For a given face f, let H be any halfspace such that $H \cap P = f$. Clearly H does not contain any witness since otherwise, by Property 1, it would contain a point of S in its interior. By Property 2, $H \cap S$ is contained in some collector $C \in \mathcal{C}$. Thus $S_f \subseteq C$. Since $|S_f| \leq d$, it follows that the number of such subsets S_f that are contained in any collector C is at most

$$\sum_{1 \leq j \leq d} \binom{|C \cap S|}{j} \; = \; O(|C \cap S|^d) \; = \; O(1),$$

where in the last step we have used the fact that $|C \cap S| = O(1)$ (Property 3). Summing over all the collectors, it follows that the total number of sets S_f, and hence the combinatorial complexity of P, is $O(|\mathcal{C}|)$. $\qquad\square$

A natural choice for the witnesses and collectors would be the convex bodies R_i and the caps C_i, respectively, from Lemma 3.1. Unfortunately, these bodies do not work for our purposes. The main difficulty is that Property 3 could fail, since a cap C_i could intersect a non-constant number of bodies of $\mathcal{R}$, and hence contain a non-constant number of points of S. (To see this, suppose that K is a cylinder in 3-dimensional space. A cap of width $\Theta(\varepsilon)$ that is parallel to the circular flat face of K intersects $\Omega(1/\sqrt{\varepsilon})$ bodies, which will be distributed around the circular boundary of this face.) In this section, we show that it is possible to construct a set of witnesses and collectors that satisfy all the requirements by scaling and translating the convex bodies from Lemma 3.1 into a stratified placement according to their volumes. The properties we obtain are specified below in Lemma 4.5.

We begin with some easy geometric facts about a convex body K in canonical form. For any point $x \in K$, define $\delta(x)$ to be the minimum distance from x to any point on ∂K. Further, define the *ray-distance* of a point x to the boundary as follows. Consider the ray emanating from O and passing through x. Let p denote the intersection of this ray with ∂K. We define $\mathrm{ray}(x) = \|xp\|$. Clearly $\mathrm{ray}(x) \geq \delta(x)$. Lemma 4.2 shows that these two quantities are the same to within a constant factor.

Lemma 4.2. *Let $K \subset \mathbb{R}^d$ be a convex body in canonical form. For any point $x \in K$, $\mathrm{ray}(x) \leq d \cdot \delta(x)$.*

Proof. Let p denote the intersection with ∂K of the ray emanating from O and passing through x (see Figure 5(a)). Let K' denote the convex hull of the point p and the ball $B^{1/2d}(O)$. By convexity,
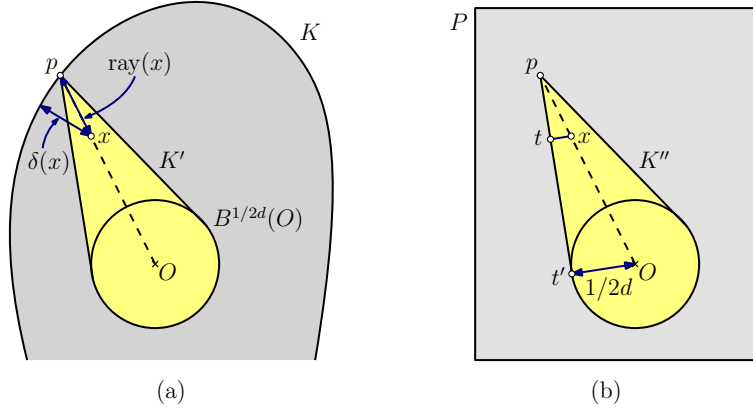
11

Figure 5: Illustrating Lemma 4.2.

K' contains the segment Op and $K' \subseteq K$. It follows that the distance between x and $\partial K'$ is a lower bound on $\delta(x)$.

To compute the distance between x and $\partial K'$, consider any 2-flat P containing the line Op and let $K'' = K' \cap P$ (see Figure 5(b)). By symmetry, the distance between x and $\partial K'$ is the same as the distance between x and $\partial K''$. Note that $\partial K''$ consists of a portion of a circle of radius $1/(2d)$ centered at O, and the two tangents to this circle from point p. It is straightforward to see that the points of $\partial K''$ that are closest to x lie on the two tangent lines (one on each tangent). Let t' denote the point where one of these tangents touches the circle, and let t denote the point on segment pt' that is closest to x. Since triangles $\triangle Ot'p$ and $\triangle xtp$ are similar, we have $\|xp\|/\|xt\| = \|Op\|/\|Ot'\|$. Since $\|Op\| \leq 1/2$ and $\|Ot'\| \geq 1/(2d)$, we have $\|xp\|/\|xt\| \leq d$. That is, $\mathrm{ray}(x) = \|xp\| \leq d\|xt\| \leq d \cdot \delta(x)$, as desired. $\qquad\square$

The following technical lemma gives upper and lower bounds on the volume of a cap of width α.

Lemma 4.3. *Let $K \subset \mathbb{R}^d$ be a convex body in canonical form and let $\alpha < 1$ be a positive real. Then the volume of any cap C of width α is $O(\alpha)$ and $\Omega(\alpha^d)$.*

Proof. Let h_1 be the hyperplane passing through the base of C and let h_2 be the parallel hyperplane passing through the apex x of C. Since C is contained in the intersection of ball $B^{1/2}(O)$ with the slab bounded by h_1 and h_2, it follows that $\mathrm{vol}(C) = O(\alpha)$.

To prove the lower bound, let y denote the point where the ray Ox intersects the base of the cap. We have $\mathrm{ray}(y) = \|xy\| \geq \alpha$. By Lemma 4.2, we have $\delta(y) \geq \mathrm{ray}(y)/d$. It follows that $\delta(y) \geq \alpha/d$. Note that the ball of radius $\delta(y)$ centered at y is contained within K and half this ball lies within the cap C. Therefore, $\mathrm{vol}(C) = \Omega(\alpha^d)$. $\qquad\square$

The following lemma states that containment of caps is preserved if the halfspaces defining both caps are consistently scaled about a point that is common to both caps.

Lemma 4.4. *Let K be a convex body and let $\lambda \geq 1$. Let C_1 and C_2 be two caps of K such that $C_1 \subseteq C_2$. Let H_1 and H_2 be the defining halfspaces of C_1 and C_2, respectively. Let H_1' and H_2' be the halfspaces obtained by scaling H_1 and H_2, respectively, by a factor of λ about p, where p is any point in $K \cap C_1$. Let C_1' and C_2' be the caps $K \cap H_1'$ and $K \cap H_2'$, respectively. Then $C_1' \subseteq C_2'$.*

12

Proof. Given λ and p, consider the affine transformation $f(q) = \lambda(q - p) + p$, which scales space by a factor of λ about p. Thus, $H'_1 = f(H_1)$ and $H'_2 = f(H_2)$, Since $p \in K$ and $\lambda \geq 1$, it follows directly from convexity that $K \subseteq f(K)$. Given any halfspace H such that $p \in K \cap H$, it follows that $K \cap f(H) = K \cap f(K \cap H)$. Since, $C_1 \subseteq C_2$, we have $f(K \cap H_1) \subseteq f(K \cap H_2)$, and thus,

$$C'_1 \;=\; K \cap f(H_1) \;=\; K \cap f(K \cap H_1) \;\subseteq\; K \cap f(K \cap H_2) \;=\; K \cap f(H_2) \;=\; C'_2,$$

as desired. $\qquad\square$

Our choice of witnesses and collectors will be based on the following lemma. Specifically, the convex bodies $R_1, \ldots, R_k$, will play the role of the witnesses and the regions $C_1, \ldots, C_k$, will play the role of the collectors. The lemma strengthens Lemma 3.1, achieving the critical property that any collector C_i intersects only a constant number of convex bodies of $\mathcal{R}$. As each witness set R_i will contain one point, this ensures that a collector contains only a constant number of input points (Property 3 of the witness-collector system). This strengthening is achieved at the expense of only an extra polylogarithmic factor in the number of collectors needed, compared with Lemma 3.1. Also, the collectors are no longer simple caps, but have a more complex shape as described in the proof (this, however, has no adverse effect in our application).

Lemma 4.5. *Let $\varepsilon > 0$ be a sufficiently small parameter, and $\widehat{\varepsilon} = \varepsilon/\log(1/\varepsilon)$. Let $K \subset \mathbb{R}^d$ be a convex body in canonical form. There exists a collection $\mathcal{R}$ of $k = O(1/\widehat{\varepsilon}^{(d-1)/2})$ disjoint centrally symmetric convex bodies $R_1, \ldots, R_k$ and associated regions $C_1, \ldots, C_k$ such that the following hold:*

1. *Let C be any cap of width ε. Then there is an i such that $R_i \subseteq C$.*

2. *Let C be any cap. Then there is an i such that either (i) $R_i \subseteq C$ or (ii) $C \subseteq C_i$.*

3. *For each i, the region C_i intersects at most a constant number of bodies of $\mathcal{R}$.*

As mentioned earlier, our proof of this lemma is based on a stratified placement of the convex bodies from Lemma 3.1, which are distributed among $O(\log(1/\varepsilon))$ layers that lie close to the boundary of K. Let $\alpha = c_1 \varepsilon/\log(1/\varepsilon)$, where c_1 is a suitable constant to be specified later. We begin by applying Lemma 3.1 to K using $\varepsilon = \alpha$. This yields a collection $\mathcal{R}'$ of $k = O(1/\alpha^{(d-1)/2})$ disjoint centrally symmetric convex bodies $\{R'_1, \ldots, R'_k\}$ and associated caps $\mathcal{C}' = \{C'_1, \ldots, C'_k\}$. Our definition of the convex bodies R_i and regions C_i required in Lemma 4.5 will be based on R'_i and C'_i, respectively. In particular, the convex body R_i will be obtained by translating a scaled copy of R'_i into an appropriate layer, based on the volume of R'_i.

Before describing the construction of the layers, it will be convenient to group the bodies in $\mathcal{R}'$ based on their volumes. We claim that the volume of any convex body R'_i lies between $c_2 \alpha^d$ and $c_3 \alpha$ for suitable constants c_2 and c_3. By Property 1 of Lemma 3.1, $R'_i \subseteq C'_i \subseteq (R'_i)^\lambda$ and C'_i has width $\beta\alpha$, for constants β and λ depending only on d. By Lemma 4.3, the volume of C'_i is $O(\alpha)$ and $\Omega(\alpha^d)$. Since $\mathrm{vol}(R'_i) = \Theta(\mathrm{vol}(C'_i))$, the desired claim follows.

We partition the set $\mathcal{R}'$ of convex bodies into t groups, where each group contains bodies whose volumes differ by a factor of at most 2. More precisely, for $0 \leq j \leq t-1$, group j consists of bodies in $\mathcal{R}'$ whose volume lies between $c_3\alpha/2^j$ and $c_3\alpha/2^{j+1}$. The lower and upper bound on the volume of bodies in $\mathcal{R}'$ implies that the number of groups t can be expressed as $\lfloor c_4 \log(1/\alpha) \rfloor$ for a suitable constant c_4 (depending on c_2 and c_3).
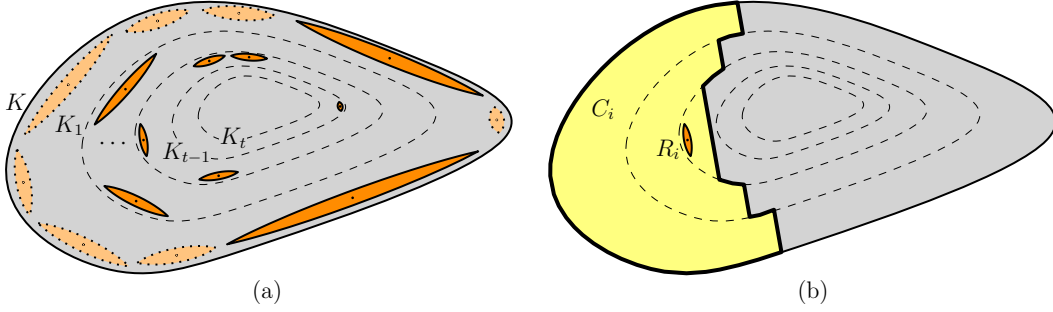
13

Figure 6: (a) Stratified placement of the bodies R_i and (b) the region C_i corresponding to a body R_i. (Figure not to scale.)

Next we describe how the layers are constructed. We will construct t layers corresponding to the t groups of $\mathcal{R}'$. Let $\gamma = 1 - 4d\beta\alpha$. For $0 \leq j \leq t$, let T_j denote the linear transformation that represents a uniform scaling by a factor of γ^j about the origin, and let $K_j = T_j(K)$ (see Figure 6(a)). Note that $K_0 = T_0(K) = K$. For $0 \leq j \leq t - 1$, define layer j, denoted L_j, to be the difference $K_j \setminus K_{j+1}$. Whenever we refer parallel supporting hyperplanes for two bodies K_i and K_j, we assume that both hyperplanes lie on the same side of the origin.

The following lemma describes some straightforward properties of these layers and the scaling transformations. In particular, the lemma shows that the t layers lie close to the boundary of K (within distance ε) and each layer has a "thickness" of $\Theta(\alpha)$.

Lemma 4.6. *Let $\varepsilon > 0$ be a sufficiently small parameter. For sufficiently small constant c_1 in the definition of α (depending on c_4, β, and d), the layered decomposition and the scaling transformations described above satisfy the following properties:*

(a) For $0 \leq j \leq t - 1$, the distance between parallel supporting hyperplanes of K_j and K_{j+1} is at most $2d\beta\alpha$.

(b) For $0 \leq j \leq t - 1$, the distance between parallel supporting hyperplanes of K_j and K_{j+1} is at least $\beta\alpha$.

(c) The distance between parallel supporting hyperplanes of K and K_t is at most ε.

(d) For $0 \leq j \leq t$, the scaling factor for T_j is at least $1/2$ and at most 1.

(e) For $0 \leq j \leq t$, T_j preserves volumes up to a constant factor.

(f) For $0 \leq j \leq t$, and any point $p \in K$, the distance between p and $T_j(p)$ is at most $2jd\beta\alpha$.

Proof. To prove (a), let h_1, h_2 denote parallel supporting hyperplanes of K_j, K_{j+1}, respectively. Since K is in canonical form, and the scaling factor of the transformation T_j is at most 1, it follows that h_1 is at distance at most $1/2$ from the origin. Since h_2 is the hyperplane obtained by scaling h_1 by a factor of $1 - 4d\beta\alpha$ about the origin, it follows that the distance between h_1 and h_2 is at most $2d\beta\alpha$.

To prove (c), let h_1, h_2 denote parallel supporting hyperplanes of K, K_t, respectively. The upper bound of (a) implies that the distance between h_1 and h_2 is at most $2td\beta\alpha$. Recall that

14

$t \leq c_4 \log(1/\alpha)$ and $\alpha = c_1\, \varepsilon / \log(1/\varepsilon)$. By choosing a sufficiently small constant c_1 in the definition of α (depending on d, c_4 and β), we can ensure that the distance between h_1 and h_2 is at most $2td\beta\alpha \leq \varepsilon$.

In the rest of this proof, we will assume that c_1 in the definition of α is sufficiently small, so (c) holds. To prove (d), note that we only need to show the lower bound on the scaling factor of T_j, since the upper bound is obvious. Again, let h_1, h_2 denote parallel supporting hyperplanes of K, K_t, respectively. Since K is in canonical position, h_1 is at distance at least $1/(2d)$ from the origin. Recall that T_t maps h_1 to h_2 and, as shown above, the distance between h_1 and h_2 is at most ε. It follows that the scaling factor of T_t is at least $1 - \varepsilon/(1/2d) = 1 - 2d\varepsilon$. By choosing ε sufficiently small, we can ensure that the scaling factor of T_t is at least $1/2$. Clearly, this lower bound on the scaling factor also applies to any transformation T_j, $0 \leq j \leq t$. This proves (d). Note that (e) is an immediate consequence.

To prove (b), let h_1, h_2 denote parallel supporting hyperplanes of K_j, K_{j+1}, respectively. Let h_1', h_2', denote the corresponding supporting hyperplanes of K, K_1, respectively. That is, $h_1 = T_j(h_1')$ and $h_2 = T_j(h_2')$. Since K is in canonical form, h_1' is at distance at least $1/(2d)$ from the origin. As h_2' is obtained by scaling h_1' by a factor of $1 - 4d\beta\alpha$ about the origin, it follows that the distance between h_1' and h_2' is at least $2\beta\alpha$. Since $h_1 = T_j(h_1')$ and $h_2 = T_j(h_2')$ and, by (d), the scaling factor of T_j is at least $1/2$, (b) follows.

Finally, to prove (f), note that the distance of p from the origin is at most $1/2$. It follows that applying T_1 to p moves it closer to the origin by a distance of at most $2d\beta\alpha$. Since $T_j = (T_1)^j$, (f) follows. $\square$

We are now ready to define the regions R_i and C_i required in Lemma 4.5. Suppose that R_i' is in group j and let $C_i' = K \cap H_i'$, where H_i' is a halfspace. We define $R_i = T_j(R_i')$. In order to define C_i, we first define caps $C_{i,r}$ of K_r as $C_{i,r} = K_r \cap T_j(H_i')$ for $0 \leq r \leq j$. We then define

$$ C_i = \bigcup_{r=0}^{j} C_{i,r}^{\sigma} \cap L_r, $$

where $\sigma = 4d\beta^2$. (See Figure 6(b).)

In Lemma 4.7, we show that the regions R_i are contained in layer j if R_i' is in group j. In Lemma 4.8, we establish Properties 1 and 2 of Lemma 4.5. Finally, in Lemma 4.9, we establish Property 3 of Lemma 4.5.

Lemma 4.7. *Let $R_i \in \mathcal{R}$. If R_i' is in group j, then $C_{i,j} = T_j(C_i')$ and $R_i \subseteq C_{i,j} \subseteq L_j$.*

Proof. Let H_i' denote the halfspace as defined above, that is, $C_i' = K \cap H_i'$. By definition, $C_{i,j} = K_j \cap T_j(H_i') = T_j(K \cap H_i') = T_j(C_i')$. By Property 1 of Lemma 3.1, $R_i' \subseteq C_i'$ and C_i' is a cap of K of width $\beta\alpha$. By Lemma 4.6(b), the distance between any parallel supporting hyperplanes of K and K_1, respectively, is at least $\beta\alpha$. It follows that $R_i' \subseteq C_i' \subseteq L_0 = K \setminus K_1$. Applying the transformation T_j to all these sets yields $R_i \subseteq C_{i,j} \subseteq L_j = K_j \setminus K_{j+1}$. This completes the proof. $\square$

Lemma 4.8. *Let C be any cap of K. Then there is an i such that either (i) $R_i \subseteq C$ or (ii) $C \subseteq C_i$. Furthermore, if the width of C is ε, then (i) holds.*
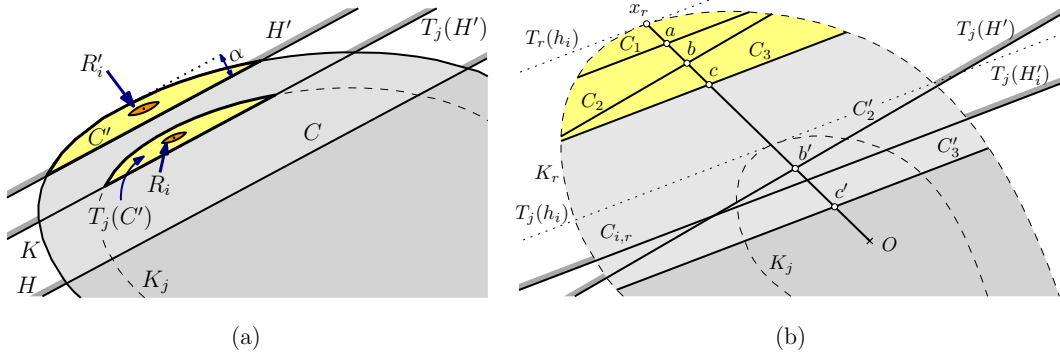
15

Figure 7: Proof of Lemma 4.8 (a) Case 1 and (b) Case 2. (Figure not to scale.)

Proof. Let $C' \subseteq C$ be the cap of width α, whose base is parallel to the base of C. Let H and H' denote the defining halfspaces of C and C', respectively. By Property 2 of Lemma 3.1, there is an i such that $R'_i \subseteq C'$. Suppose that R'_i is in group j. We consider two cases, depending on whether $T_j(H') \subseteq H$ or $H \subset T_j(H')$. To complete the proof of the lemma, we will show that in the former case, $R_i \subseteq C$ and, in the latter case, $C \subseteq C_i$. Additionally, we will show that if C has width ε, then the former case holds (implying that $R_i \subseteq C$).

Case 1: $T_j(H') \subseteq H$**.** Arguing as in the proof of Lemma 4.7 (but with C' in place of C'_i), we have $R_i \subseteq T_j(C') = K_j \cap T_j(H') \subseteq L_j$ (see Figure 7(a)). Observe that $K_j \cap T_j(H') \subseteq K \cap H = C$. Therefore $R_i \subseteq C$.

Also, by Lemma 4.6(c), the distance between any parallel supporting hyperplanes of K and K_t is at most ε. Since $K_j \cap T_j(H') \subseteq L_j$, it follows that the width of cap $K \cap T_j(H')$ is at most ε. Therefore, if C has width ε, then $T_j(H') \subseteq H$ and Case 1 holds.

Case 2: $H \subset T_j(H')$**.** Recall that we need to show that $C \subseteq C_i$. Clearly, it suffices to show that $K \cap T_j(H') \subseteq C_i$ since $C = K \cap H \subset K \cap T_j(H')$. In turn, the definition of C_i implies that it suffices to show that for $0 \le r \le j$, $T_j(H') \cap K_r \subseteq C^\sigma_{i,r}$.

By Property 2 of Lemma 3.1, there is an i such that $(C'_i)^\phi \subseteq C' \subseteq C'_i$, where $\phi = 1/\beta^2$. By Property 1 of Lemma 3.1, the widths of the caps $(C'_i)^\phi$ and C'_i are α/β and $\beta\alpha$, respectively. Recall that H'_i denotes the defining halfspace for the cap C'_i. Also, let x denote the apex of C'_i, and let h_i denote the supporting hyperplane to K passing through x and parallel to C'_i's base.

Let C_1, C_2, and C_3 denote the caps of K_r obtained by applying the transformation T_r to the caps $(C'_i)^\phi, C'$, and C'_i, respectively (see Figure 7(b)). We have $C_1 \subseteq C_2 \subseteq C_3$. Let a, b and c denote the point of intersection of the bases of the caps C_1, C_2 and C_3, respectively, with the line segment Ox. Let b' denote the point of intersection of the base of the cap $K \cap T_j(H')$ with the segment Ox. Let x_r denote the point $T_r(x)$. Consider scaling caps C_2 and C_3 as described in Lemma 4.4, about the point x_r with scaling factor $\rho = \|b'x_r\|/\|bx_r\|$. Let C'_2 and C'_3 denote the caps of K_r obtained from C_2 and C_3, respectively, through this transformation. By Lemma 4.4, $C'_2 \subseteq C'_3$. Our choice of the scaling factor implies that C'_2 is the cap $T_j(H') \cap K_r$. We claim that $C'_3 \subseteq C^\sigma_{i,r}$. Note that this claim would imply that $T_j(H') \cap K_r \subseteq C^\sigma_{i,r}$, and complete the proof.

To prove the above claim, we first show that $\rho = O(j - r + 1)$. Observe that $\rho = (\|b'b\| +$
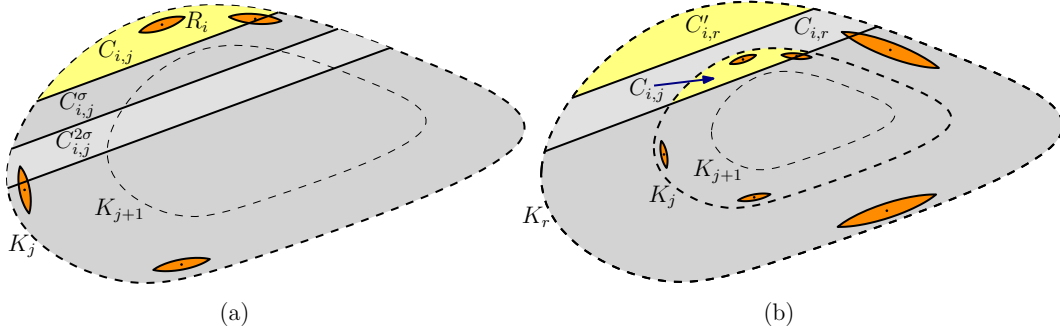
16

Figure 8: Proof of Lemma 4.9. (Figure not to scale.)

$\|bx_r\|)/\|bx_r\| = \|b'b\|/\|bx_r\| + 1$. We have

$$\|bx_r\| \;\geq\; \|ax_r\| \;\geq\; \mathrm{width}(C_1) \;\geq\; \frac{\mathrm{width}((C_i')^\phi)}{2} \;\geq\; \frac{\alpha}{2\beta},$$

where in the third inequality, we have used Lemma 4.6(d) and the fact that $C_1 = T_r((C_i')^\phi)$. Also, since $T_{j-r}(b) = b'$, it follows from Lemma 4.6(f) that $\|b'b\|$ is at most $2(j - r)d\beta\alpha$. Substituting the derived bounds on $\|b'b\|$ and $\|bx_r\|$, we obtain $\rho \leq 4d\beta^2(j - r) + 1$.

Recall that C_3' and $C_{i,r}$ are caps of K_r defined by parallel halfspaces. To prove that $C_3' \subseteq C_{i,r}^\sigma$, it therefore suffices to show that $\mathrm{width}(C_3')/\mathrm{width}(C_{i,r}) \leq \sigma$. We have

$$\mathrm{width}(C_3') \;=\; \rho \cdot \mathrm{width}(C_3) \;\leq\; \rho \cdot \mathrm{width}(C_i') \;=\; \rho\beta\alpha,$$

where in the second step, we have used Lemma 4.6(d) and the fact that $C_3 = T_r(C_i')$. Also, it is easy to see that the width of $C_{i,r}$ is the sum of the width of the cap $T_j(C_i')$ and the distance between the hyperplanes $T_r(h_i)$ and $T_j(h_i)$. Since $\mathrm{width}(C_i') = \beta\alpha$, by Lemma 4.6(d), the width of the cap $T_j(C_i')$ is at least $\beta\alpha/2$. Also, by Lemma 4.6(b), the distance between the hyperplanes $T_r(h_i)$ and $T_j(h_i)$ is at least $(j - r)\beta\alpha$. It follows that the width of $C_{i,r}$ is at least $\beta\alpha/2 + (j - r)\beta\alpha = (j - r + 1/2)\beta\alpha$. Thus,

$$\frac{\mathrm{width}(C_3')}{\mathrm{width}(C_{i,r})} \;\leq\; \frac{\rho\beta\alpha}{(j - r + 1/2)\beta\alpha} \;=\; \frac{\rho}{j - r + 1/2} \;\leq\; \frac{4d\beta^2(j - r) + 1}{j - r + 1/2} \;\leq\; 4d\beta^2 \;=\; \sigma,$$

as desired. $\qquad\square$

Lemma 4.9. *For each i, the region C_i intersects $O(1)$ bodies of $\mathcal{R}$.*

Proof. Suppose that R_i' is in group j. Recall that $R_i = T_j(R_i')$, $C_i' = K \cap H_i'$ and $C_i = \bigcup_{r=0}^{j}(C_{i,r}^\sigma \cap L_r)$. We begin by bounding the number of bodies of $\mathcal{R}$ that overlap $C_{i,j}^\sigma \cap L_j$. (See Figure 8(a).) By Lemma 4.7, $C_{i,j} = T_j(C_i')$ and $R_i \subseteq C_{i,j} \subseteq L_j$. By Property 1 of Lemma 3.1, we have $C_i' \subseteq (R_i')^\lambda$, which implies that $\mathrm{vol}(R_i') = \Omega(\mathrm{vol}(C_i'))$. Recall that all the bodies of $\mathcal{R}'$ in group j have the same volumes to within a factor of 2, and so they all have volumes $\Omega(\mathrm{vol}(C_i'))$. By Lemma 4.6(e), the scaling transformations used in our construction preserve volumes to within a constant factor. Also, recall that the bodies of $\mathcal{R}$ in layer j are scaled copies of the bodies of $\mathcal{R}'$ in group j. It follows that the bodies of $\mathcal{R}$ in layer j all have volumes $\Omega(\mathrm{vol}(C_{i,j}))$.

17

Next, we assert that any body of $\mathcal{R}$ that overlaps $C_{i,j}^\sigma \cap L_j$ is contained within the cap $C_{i,j}^{2\sigma}$. To prove this, recall from the proof of Lemma 3.1 that the bodies of $\mathcal{R}'$ are $(1/5)$-scaled disjoint Macbeath regions with respect to K. It follows that the bodies of $\mathcal{R}$ in layer j are $(1/5)$-scaled disjoint Macbeath regions with respect to K_j. By Lemma 2.5, it now follows that any body of $\mathcal{R}$ that overlaps $C_{i,j}^\sigma \cap L_j$ is contained within the cap $C_{i,j}^{2\sigma}$. Since $\mathrm{vol}(C_{i,j}^{2\sigma}) = O(\mathrm{vol}(C_{i,j}))$, and all bodies of $\mathcal{R}$ in layer j have volumes $\Omega(\mathrm{vol}(C_{i,j}))$, it follows by a simple packing argument that the number of bodies of $\mathcal{R}$ that overlap $C_{i,j}^\sigma \cap L_j$ is $O(1)$.

Next we bound the number of bodies of $\mathcal{R}$ that overlap $C_{i,r}^\sigma \cap L_r$, where $0 \le r < j$. (See Figure 8(b).) Recall that $C_{i,r} = K_r \cap T_j(H_i')$. Roughly speaking, we will show that the volume of $C_{i,r}$ exceeds the volume of $C_{i,j}$ by a factor that is at most polynomial in $j - r$, while the volume of the bodies in layer r exceeds the volume of the bodies in layer j by a factor that is exponential in $j - r$. This will allow us to show that the number of bodies of $\mathcal{R}$ that overlap C_i is bounded by a constant. We now present the details.

Define $C_{i,r}' = T_r(C_i')$. Recall that $C_{i,j} = T_j(C_i')$. By Lemma 4.6(e), T_j and T_r preserve volumes up to constant factors, and so $\mathrm{vol}(C_{i,r}') = \Theta(\mathrm{vol}(C_{i,j}))$. Since the width of C_i' is $\beta\alpha$, by Lemma 4.6(d), it follows that the width of $C_{i,r}'$ is at least $\beta\alpha/2$. Also, the width of $C_{i,r}$ is upper bounded by the distance between parallel supporting hyperplanes of K_r and K_{j+1} which by Lemma 4.6(a) is at most $2d\beta\alpha(j-r+1)$. It follows that the width of $C_{i,r}$ is $O(j-r+1)$ times the width of $C_{i,r}'$. Recalling that, for $\lambda \ge 1$, the volume of a λ-expansion of a cap is at most λ^d times the volume of the cap, it follows that $\mathrm{vol}(C_{i,r}) = O((j-r+1)^d) \cdot \mathrm{vol}(C_{i,r}') = O((j-r+1)^d) \cdot \mathrm{vol}(C_{i,j})$.

Next, recall that the volume of the bodies of $\mathcal{R}'$ in group r exceeds the volume of the bodies of $\mathcal{R}'$ in group j by a factor of $\Omega(2^{j-r+1})$. It follows from Lemma 4.6(e) and our construction that the volume of the bodies of $\mathcal{R}$ in layer r exceeds the volume of the bodies of $\mathcal{R}$ in layer j by a factor of $\Omega(2^{j-r+1})$. For the same reasons as discussed above, any body of $\mathcal{R}$ that overlaps $C_{i,r}^\sigma \cap L_r$ is contained within $C_{i,r}^{2\sigma}$, and $\mathrm{vol}(C_{i,r}^{2\sigma}) = O(\mathrm{vol}(C_{i,r}))$. Putting this together with the upper bound on $\mathrm{vol}(C_{i,r})$ shown above, we have $\mathrm{vol}(C_{i,r}^{2\sigma}) = O((j-r+1)^d) \cdot \mathrm{vol}(C_{i,j})$. By a simple packing argument, it follows that the ratio of the number of bodies of $\mathcal{R}$ that overlap $C_{i,r}^\sigma \cap L_r$ to the number of bodies of $\mathcal{R}$ that overlap $C_{i,j}^\sigma \cap L_j$ is $O((j-r+1)^d/2^{j-r+1})$. Recall that the number of bodies of $\mathcal{R}$ that overlap $C_{i,j}^\sigma \cap L_j$ is $O(1)$. It follows that the number of bodies of $\mathcal{R}$ that overlap $C_i = \bigcup_{r=0}^j (C_{i,r}^\sigma \cap L_r)$ is on the order of $\sum_{0 \le r \le j} (j-r+1)^d/2^{j-r+1} = O(1)$, as desired. $\qquad \square$

Let S be a set of points containing one point inside each body of $\mathcal{R}$ defined in Lemma 4.5 and no other points.

Lemma 4.10. *The polytope $P = \mathrm{conv}(S)$ is an ε-approximation of K.*

Proof. A set of points S *stabs* every cap of width ε if every such cap contains at least one point of S. It is well known that if a set of points $S \subset K$ stabs all caps of width ε of K, then $\mathrm{conv}(S)$ is an ε-approximation of K [13]. Let C be a cap of width ε. By Lemma 4.5, Property 1, there is a convex body $R_i \subseteq C$. Since S contains a point that is in R_i, we have that the cap C is stabbed. $\quad \square$

To bound the combinatorial complexity of $\mathrm{conv}(S)$, and hence conclude the proof of Theorem 1.1, we use the witness-collector approach [17].

Lemma 4.11. *The number of faces of $P = \mathrm{conv}(S)$ is $O(1/\widehat{\varepsilon}^{(d-1)/2})$.*

Proof. Define the witness set $\mathcal{W} = R_1, \ldots, R_k$ and the collector set $\mathcal{C} = C_1, \ldots, C_k$, where the R_i's and C_i's are as defined in Lemma 4.5. As there is a point of S in each body R_i, Property 1 of the witness-collector method is satisfied. To prove Property 2, let H be any halfspace. If H does not intersect K, then Property 2 of the witness-collector method holds trivially. Otherwise let $C = K \cap H$. By Property 2 of Lemma 4.5, there is an i such that either $R_i \subseteq C$ or $C \subseteq C_i$. It follows that H contains witness R_i or $H \cap S$ is contained in collector C_i. Thus Property 2 of the witness-collector method is satisfied. Finally, Property 3 of Lemma 4.5 implies Property 3 of the witness-collector method. Thus, we can apply Lemma 4.1 to conclude that the number of faces of P is $O(|\mathcal{C}|) = O(k)$, which proves the lemma. □

5   Conclusions and Open Problems

We considered the problem of ε-approximating a convex body $K \subset \mathbb{R}^d$ by a polytope P of small combinatorial complexity. We proved an upper bound of $\widetilde{O}(1/\varepsilon^{(d-1)/2})$ to the combinatorial complexity, almost a square-root improvement over the previous bound of $O(1/\varepsilon^{d(d-1)/(d+1)}) \approx O(1/\varepsilon^{d-2})$. Our bound is optimal up to logarithmic factors. Two natural questions arise. First, can the logarithmic factors be removed or is there a fundamental reason why they appear? Second, our construction is much more complex than the ones of Dudley or Bronshteyn and Ivanov. Can we show that those simpler constructions also attain a low combinatorial complexity or find a counterexample? Furthermore, our bounds are purely existential. While our construction can be turned into an algorithm, there are a number of nontrivial technical issues that would need to be handled in order to obtain an efficient solution.

Our bounds are presented as a function of ε, but a natural question is whether it is possible to obtain bounds that are sensitive to the polytope being approximated. One may consider finding the polytope of minimum combinatorial complexity that approximates a given polytope K as an optimization problem. Approximation algorithms for minimizing the number of vertices of an ε-approximating polytope are well known [15, 25], but we know of no similar results for minimizing the combinatorial complexity.

Acknowledgments

References

[1] P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Approximating extent measures of points. *J. Assoc. Comput. Mach.*, 51:606–635, 2004.

[2] G. E. Andrews. A lower bound for the volumes of strictly convex bodies with many boundary points. *Trans. Amer. Math. Soc.*, 106:270–279, 1963.

[3] S. Arya, G. D. da Fonseca, and D. M. Mount. Optimal area-sensitive bounds for polytope approximation. In *Proc. 28th Annu. Sympos. Comput. Geom.*, pages 363–372, 2012.

[4] S. Arya, G. D. da Fonseca, and D. M. Mount. Optimal approximate polytope membership. In *Proc. 28th Annu. ACM-SIAM Sympos. Discrete Algorithms*, 2017. (To appear).

[5] S. Arya, T. Malamatos, and D. M. Mount. The effect of corners on the complexity of approximate range searching. *Discrete Comput. Geom.*, 41:398–443, 2009.

[6] S. Arya, D. M. Mount, and J. Xia. Tight lower bounds for halfspace range searching. *Discrete Comput. Geom.*, 47:711–730, 2012.

[7] I. Bárány. Intrinsic volumes and f-vectors of random polytopes. *Math. Ann.*, 285:671–699, 1989.

[8] I. Bárány. The technique of M-regions and cap-coverings: A survey. *Rend. Circ. Mat. Palermo*, 65:21–38, 2000.

[9] I. Bárány. Extremal problems for convex lattice polytopes: A survey. *Contemp. Math.*, 453:87–103, 2008.

[10] I. Bárány and D. G. Larman. Convex bodies, economic cap coverings, random polytopes. *Mathematika*, 35:274–291, 1988.

[11] K. Böröczky, Jr. Approximation of general smooth convex bodies. *Adv. Math.*, 153:325–341, 2000.

[12] H. Brönnimann, B. Chazelle, and J. Pach. How hard is halfspace range searching. *Discrete Comput. Geom.*, 10:143–155, 1993.

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

[14] E. M. Bronstein. Approximation of convex sets by polytopes. *J. Math. Sci.*, 153(6):727–762, 2008.

[15] K. L. Clarkson. Algorithms for polytope covering and approximation. In *Proc. Third Internat. Workshop Algorithms Data Struct.*, pages 246–252, 1993.

[16] K. L. Clarkson. Building triangulations using ε-nets. In *Proc. 38th Annu. ACM Sympos. Theory Comput.*, pages 326–335, 2006.

[17] O. Devillers, M. Glisse, and X. Goaoc. Complexity analysis of random geometric structures made simpler. In *Proc. 29th Annu. Sympos. Comput. Geom.*, pages 167–176, 2013.

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

[19] G. Ewald, D. G. Larman, and C. A. Rogers. The directions of the line segments and of the r-dimensional balls on the boundary of a convex body in Euclidean space. *Mathematika*, 17:1–20, 1970.

[20] P. M. Gruber. Asymptotic estimates for best and stepwise approximation of convex bodies I. *Forum Math.*, 5:521–537, 1993.

[21] S. Har-Peled. *Geometric approximation algorithms*. Number 173 in Mathematical surveys and monographs. American Mathematical Society, 2011.

[22] F. John. Extremum problems with inequalities as subsidiary conditions. In *Studies and Essays Presented to R. Courant on his 60th Birthday*, pages 187–204. Interscience Publishers, Inc., New York, 1948.

[23] A. M. Macbeath. A theorem on non-homogeneous lattices. *Ann. of Math.*, 56:269–293, 1952.

[24] P. McMullen. The maximum numbers of faces of a convex polytope. *Mathematika*, 17:179–184, 1970.

[25] J. S. B. Mitchell and S. Suri. Separation and approximation of polyhedral objects. *Comput. Geom. Theory Appl.*, 5:95–114, 1995.

[26] R. Schneider. Polyhedral approximation of smooth convex bodies. *J. Math. Anal. Appl.*, 128:470–474, 1987.

[27] L. F. Toth. Approximation by polygons and polyhedra. *Bull. Amer. Math. Soc.*, 54:431–438, 1948.