

Near-Optimal ε -Kernel Construction and Related Problems*

Sunil Arya^{†1}, Guilherme D. da Fonseca², and David M. Mount^{‡3}

- 1 Department of Computer Science and Engineering, The Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong
arya@cse.ust.hk
- 2 Université Clermont Auvergne and LIMOS, Clermont-Ferrand, France
fonseca@isima.fr
- 3 Department of Computer Science and Institute for Advanced Computer Studies, University of Maryland, College Park, MD, USA
mount@cs.umd.edu

Abstract

The computation of (i) ε -kernels, (ii) approximate diameter, and (iii) approximate bichromatic closest pair are fundamental problems in geometric approximation. In this paper, we describe new algorithms that offer significant improvements to their running times. In each case the input is a set of n points in \mathbb{R}^d for a constant dimension $d \geq 3$ and an approximation parameter $\varepsilon > 0$. We reduce the respective running times

- (i) from $O((n + 1/\varepsilon^{d-2}) \log \frac{1}{\varepsilon})$ to $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(d-1)/2+\alpha})$,
- (ii) from $O((n + 1/\varepsilon^{d-2}) \log \frac{1}{\varepsilon})$ to $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(d-1)/2+\alpha})$, and
- (iii) from $O(n/\varepsilon^{d/3})$ to $O(n/\varepsilon^{d/4+\alpha})$,

for an arbitrarily small constant $\alpha > 0$. Result (i) is nearly optimal since the size of the output ε -kernel is $\Theta(1/\varepsilon^{(d-1)/2})$ in the worst case.

These results are all based on an efficient decomposition of a convex body using a hierarchy of Macbeath regions and contrast with previous solutions, which decompose space using quadtrees and grids. By further application of these techniques, we also show that it is possible to obtain near-optimal preprocessing times for the most efficient data structures to approximately answer queries for (iv) nearest-neighbor searching, (v) directional width, and (vi) polytope membership.

1998 ACM Subject Classification F.2.2 Geometrical problems and computations

Keywords and phrases Approximation, diameter, kernel, coresets, nearest neighbor, polytope membership, bichromatic closest pair, Macbeath regions

Digital Object Identifier 10.4230/LIPIcs.SoCG.2017.10

1 Introduction

In this paper we present new faster algorithms to several fundamental geometric approximation problems involving point sets in d -dimensional space. In particular, we present approximation algorithms for ε -kernels, diameter, bichromatic closest pair, and the minimum bottleneck spanning tree. Our results arise from a recently developed shape-sensitive approach to approximating convex bodies, which is based on the classical concept of Macbeath

* A full version of the paper is available at <http://arxiv.org/abs/1703.10868>.

[†] Research supported by the Research Grants Council of Hong Kong, China under project number 610012.

[‡] Research supported by NSF grant CCF-1618866.



regions. This approach has been applied to computing area-sensitive bounds for polytope approximation [5], polytope approximations with low combinatorial complexity [6], answering approximate polytope-membership queries [7], and approximate nearest-neighbor searching [7]. The results of [7] demonstrate the existence of data structures for these query problems but did not discuss preprocessing in detail. We complete the story by presenting efficient algorithms for building data structures for three related queries: approximate polytope membership, approximate directional width, and approximate nearest-neighbors.

Throughout, we assume that the dimension d is a constant. Our running times will often involve expressions of the form $1/\varepsilon^\alpha$. In such cases, $\alpha > 0$ is constant that can be made arbitrarily small. The approximation parameter ε is treated as an asymptotic variable that approaches 0. We assume throughout that $\varepsilon < 1$, which guarantees that $\log \frac{1}{\varepsilon} > 0$.

In Section 1.1, we present our results for ε -kernels, diameter, bichromatic closest pair, and minimum bottleneck tree. In Section 1.2, we present our results for the data structure problems. In Section 1.3, we give an overview of the techniques used.

Concurrently and independently, Timothy Chan has reported complexity bounds that are very similar to our results [18]. Remarkably, the computational techniques seem to be very different, based on Chebyshev polynomials.

1.1 Static Results

Kernel. Given a set S of n points in \mathbb{R}^d and an approximation parameter $\varepsilon > 0$, an ε -coreset is an (ideally small) subset of S that approximates some measure over S (see [2] for a survey). Given a nonzero vector $v \in \mathbb{R}^d$, the *directional width* of S in direction v , $\text{width}_v(S)$ is the minimum distance between two hyperplanes that enclose S and are orthogonal to v . A *coreset for the directional width* (also known as an ε -kernel and as a *coreset for the extent measure*) is a subset $Q \subseteq S$ such that $\text{width}_v(Q) \geq (1 - \varepsilon) \text{width}_v(S)$, for all $v \in \mathbb{R}^d$. Kernels are among the most fundamental constructions in geometric approximation, playing a role similar to that of convex hulls in exact computations. Kernels have been used to obtain approximation algorithms to several problems such as diameter, minimum width, convex hull volume, minimum enclosing cylinder, minimum enclosing annulus, and minimum-width cylindrical shell [1, 2].

The concept of ε -kernels was introduced by Agarwal et al. [1]. The existence of ε -kernels with $O(1/\varepsilon^{(d-1)/2})$ points is implied in the works of Dudley [19] and Bronshteyn and Ivanov [16], and this is known to be optimal in the worst case. Agarwal et al. [1] demonstrated how to compute such a kernel in $O(n + 1/\varepsilon^{3(d-1)/2})$ time, which reduces to $O(n)$ when $n = \Omega(1/\varepsilon^{3(d-1)/2})$. While less succinct ε -kernels with $O(1/\varepsilon^{d-1})$ points can be constructed in time $O(n)$ for all n [1, 14], no linear-time algorithm is known to build an ε -kernel of optimal size. Hereafter, we use the term ε -kernel to refer exclusively to an ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$.

Chan [17] showed that an ε -kernel can be constructed in $O((n + 1/\varepsilon^{d-2}) \log \frac{1}{\varepsilon})$ time, which is nearly linear when $n = \Omega(1/\varepsilon^{d-2})$. He posed the open problem of obtaining a faster algorithm. A decade later, Arya and Chan [11] showed how to build an ε -kernel in roughly $O(n + \sqrt{n}/\varepsilon^{d/2})$ time using discrete Voronoi diagrams. In this paper, we attain the following near-optimal construction time.

► **Theorem 1.1.** *Given a set S of n points in \mathbb{R}^d and an approximation parameter $\varepsilon > 0$, it is possible to construct an ε -kernel of S with $O(1/\varepsilon^{(d-1)/2})$ points in $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(d-1)/2+\alpha})$ time, where α is an arbitrarily small positive constant.*

Because the worst-case output size is $O(1/\varepsilon^{(d-1)/2})$, we may assume that n is at least this large, for otherwise we can simply take S itself to be the kernel. Since $1/\varepsilon^\alpha$ dominates

$\log \frac{1}{\varepsilon}$, the above running time can be expressed as $O(n/\varepsilon^\alpha)$, which is nearly linear given that α can be made arbitrarily small.

Diameter. An important application of ε -kernels is to approximate the diameter of a point set. Given n data points, the *diameter* is defined to be the maximum distance between any two data points. An ε -approximation of the diameter is a pair of points whose distance is at least $(1 - \varepsilon)$ times the exact diameter. There are multiple algorithms to approximate the diameter [1, 3, 11, 13, 17]. The fastest running times are $O((n + 1/\varepsilon^{d-2}) \log \frac{1}{\varepsilon})$ [17] and roughly $O(n + \sqrt{n}/\varepsilon^{d/2})$ [11]. The algorithm from [17] essentially computes an ε -kernel Q and then determines the maximum value of $\text{width}_v(Q)$ among a set of $k = O(1/\varepsilon^{(d-1)/2})$ directions v by brute force [1]. Discrete Voronoi diagrams [11] permit this computation in roughly $O(n + \sqrt{n}/\varepsilon^{d/2})$ time. Therefore, combining the kernel construction of Theorem 1.1 with discrete Voronoi diagrams [11], we reduce n to $O(1/\varepsilon^{(d-1)/2})$ and obtain an algorithm to ε -approximate the diameter in roughly $O(n + 1/\varepsilon^{3d/4})$ time. However, we show that it is possible to obtain a much faster algorithm, as presented in the following theorem.

► **Theorem 1.2.** *Given a set S of n points in \mathbb{R}^d and an approximation parameter $\varepsilon > 0$, it is possible to compute an ε -approximation to the diameter of S in $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(d-1)/2+\alpha})$ time.*

Bichromatic Closest Pair. In the *bichromatic closest pair* (BCP) problem, we are given n points from two sets, designated red and blue, and we want to find the closest red-blue pair. In the ε -approximate version, the goal is to find a red-blue pair of points whose distance is at most $(1 + \varepsilon)$ times the exact BCP distance. Approximations to the BCP problem were introduced in [23], and the most efficient randomized approximation algorithm runs in roughly $O(n/\varepsilon^{d/3})$ expected time [11]. We present the following result.

► **Theorem 1.3.** *Given n red and blue points in \mathbb{R}^d and an approximation parameter $\varepsilon > 0$, there is a randomized algorithm that computes an ε -approximation to the bichromatic closest pair in $O(n/\varepsilon^{d/4+\alpha})$ expected time.*

Euclidean Trees. Given a set S of n points in \mathbb{R}^d , a *Euclidean minimum spanning tree* is the spanning tree with vertex set S that minimizes the sum of the edge lengths, while a *Euclidean minimum bottleneck tree* minimizes the maximum edge length. In the approximate version we respectively approximate the sum and the maximum of the edge lengths. A minimum spanning tree is a minimum bottleneck tree (although the converse does not hold). However, an approximation to the minimum spanning tree is not necessarily an approximation to the minimum bottleneck tree. A recent approximation algorithm to the Euclidean minimum spanning tree takes roughly $O(n \log n + n/\varepsilon^2)$ time, regardless of the (constant) dimension [9]. On the other hand, the fastest algorithm to approximate the minimum bottleneck tree takes roughly $O((n \log n)/\varepsilon^{d/3})$ expected time [11]. The algorithm uses BCP to simultaneously attain an approximation to the minimum bottleneck and the minimum spanning trees. We prove the following theorem.

► **Theorem 1.4.** *Given n points in \mathbb{R}^d and an approximation parameter $\varepsilon > 0$, there is a randomized algorithm that computes a tree T that is an ε -approximation to both the Euclidean minimum bottleneck and the Euclidean minimum spanning trees in $O((n \log n)/\varepsilon^{d/4+\alpha})$ expected time.*

1.2 Data Structure Results

Polytope membership. Let P denote a convex polytope in \mathbb{R}^d , represented as the bounded intersection of n halfspaces. The *polytope membership problem* consists of preprocessing P so that it is possible to determine efficiently whether a given query point $q \in \mathbb{R}^d$ lies within P . In the ε -approximate version, we consider an expanded convex body $K \supset P$. A natural way to define this expansion would be to consider the set of points that lie within distance $\varepsilon \cdot \text{diam}(P)$ of P , thus defining a body whose Hausdorff distance from P is $\varepsilon \cdot \text{diam}(P)$. However, this definition has the shortcoming that it is not sensitive to the directional width of P . Instead, we define K as follows. For any nonzero vector $v \in \mathbb{R}^d$, consider the two supporting hyperplanes for P that are normal to v . Translate each of these hyperplanes outward by a distance of $\varepsilon \cdot \text{width}_v(P)$, and consider the closed slab-like region lying between them. Define K to be the intersection of this (infinite) set of slabs. This is clearly a stronger approximation than the Hausdorff-based definition. An ε -approximate polytope membership query (ε -APM query) returns a positive result if the query point q is inside P , a negative result if q is outside K , and may return either result otherwise.¹

We recently proposed an optimal data structure to answer approximate polytope membership queries, but efficient preprocessing remained an open problem [7]. In this paper, we present a similar data structure that not only attains optimal storage and query time, but can also be preprocessed in near-optimal time.


► **Theorem 1.5.** *Given a convex polytope P in \mathbb{R}^d represented as the intersection of n halfspaces and an approximation parameter $\varepsilon > 0$, there is a data structure that can answer ε -approximate polytope membership queries with query time $O(\log \frac{1}{\varepsilon})$, space $O(1/\varepsilon^{(d-1)/2})$, and preprocessing time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(d-1)/2+\alpha})$.*

Directional width. Applying the previous data structure in the dual space, we obtain a data structure for the following ε -approximate directional width problem, which is closely related to ε -kernels. Given a set S of n points in a constant dimension d and an approximation parameter $\varepsilon > 0$, the goal is to preprocess S to efficiently ε -approximate $\text{width}_v(S)$, for a nonzero query vector v . We present the following result.

► **Theorem 1.6.** *Given a set S of n points in \mathbb{R}^d and an approximation parameter $\varepsilon > 0$, there is a data structure that can answer ε -approximate directional width queries with query time $O(\log^2 \frac{1}{\varepsilon})$, space $O(1/\varepsilon^{(d-1)/2})$, and preprocessing time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(d-1)/2+\alpha})$.*

Nearest Neighbor. Let S be a set of n points in \mathbb{R}^d . Given any $q \in \mathbb{R}^d$, an ε -approximate nearest neighbor (ANN) of q is any point of S whose distance from q is at most $(1 + \varepsilon)$ times the distance to q 's closest point in S . The objective is to preprocess S in order to answer such queries efficiently. Data structures for approximate nearest neighbor searching (in fixed dimensions) have been proposed by several authors, offering space-time tradeoffs (see [7] for an overview of the tradeoffs). Applying the reduction from approximate nearest neighbor to approximate polytope membership established in [4] together with Theorem 1.5, we obtain the following result, which matches the best bound [7] up to an $O(\log \frac{1}{\varepsilon})$ factor in the query time, but offers faster preprocessing time.

¹ Our earlier works on ε -APM queries [4, 7] use the weaker Hausdorff form to define the problem, but the solutions presented there actually achieve the stronger direction-sensitive form.



■ **Figure 1** Two levels of the ellipsoid hierarchy.

► **Theorem 1.7.** *Given a set S of n points in \mathbb{R}^d , an approximation parameter $\varepsilon > 0$, and m such that $\log \frac{1}{\varepsilon} \leq m \leq 1/(\varepsilon^{d/2} \log \frac{1}{\varepsilon})$, there is a data structure that can answer Euclidean ε -approximate nearest neighbor queries with query time $O(\log n + (\log \frac{1}{\varepsilon})/(m \cdot \varepsilon^{d/2}))$ space $O(nm)$, and preprocessing time $O(n \log n \log \frac{1}{\varepsilon} + nm/\varepsilon^\alpha)$.*

1.3 Techniques

In contrast to previous kernel constructions, which are based on grids and the execution of Bronshteyn and Ivanov's algorithm, our construction employs a classical structure from the theory of convexity, called *Macbeath regions* [24]. Macbeath regions, which will be defined in Section 2.1, have found numerous uses in the theory of convex sets and the geometry of numbers (see Bárány [12] for an excellent survey). They have also been applied to several problems in the field of computational geometry. However, most previous results were either in the form of lower bounds [8, 10, 15] or focused on existential results [5, 6, 20, 25].

In [7] the authors introduced a data structure employing a hierarchy of ellipsoids based on Macbeath regions to answer approximate polytope membership queries, but the efficient computation of the hierarchy was not considered. In this paper, we show how to efficiently construct the Macbeath regions that form the basis of this hierarchy.

Let P denote a convex polytope in \mathbb{R}^d . Each level i in the hierarchy corresponds to a δ_i -approximation of the boundary of P by a set of $O(1/\delta_i^{(d-1)/2})$ ellipsoids, where $\delta_i = \Theta(1/2^i)$. Each ellipsoid is sandwiched between two Macbeath regions and has $O(1)$ children, which correspond to the ellipsoids of the following level that approximate the same portion of the boundary (see Figure 1). The hierarchy starts with $\delta_0 = \Theta(1)$ and stops after $O(\log \frac{1}{\delta})$ levels when $\delta_i = \delta$, for a desired approximation δ . We present a simple algorithm to construct the hierarchy in $O(n + 1/\delta^{3(d-1)/2})$ time. The polytope P can be presented as either the intersection of n halfspaces or the convex hull of n points. We present the relevant background in Section 3.

Our algorithm to compute an ε -kernel in time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(d-1)/2+\alpha})$ (Theorem 1.1) is based on a bootstrapping process. Since the time to build the ε -approximation hierarchy for the convex hull is prohibitively high, we use an approximation parameter $\delta = \varepsilon^{1/3}$ to build a δ -approximation hierarchy in $O(n + 1/\varepsilon^{(d-1)/2})$ time. By navigating through this hierarchy, we partition the n points among the leaf Macbeath ellipsoids in $O(n \log \frac{1}{\varepsilon})$ time, discarding points that are too far from the boundary. We then compute an (ε/δ) -kernel for the set of points in each leaf ellipsoid and return the union of the kernels computed.

Given an algorithm to compute an ε -kernel in $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{t(d-1)})$ time, the previous procedure produces an ε -kernel in $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{t'(d-1)})$ time, where $t' = (4t + 1)/6$. By bootstrapping the construction a constant number of times, the value of t decreases from 1 to a value that is arbitrarily close to $\frac{1}{2}$. (This accounts for the $O(1/\varepsilon^\alpha)$ factors in our running times.) The construction and its analysis are presented in Section 4.

In Section 5, we use our kernel construction in the dual space to efficiently build a polytope membership data structure, proving Theorem 1.5. The key idea is to compute multiple kernels in order to avoid examining the whole polytope when building each Macbeath region. Again, we use bootstrapping to obtain a near-optimal preprocessing time. The remaining theorems follow from Theorems 1.1 and 1.5, together with several known reductions.

2 Geometric Preliminaries

Consider a convex body K in d -dimensional space \mathbb{R}^d . Let ∂K denote the boundary of K . Let O denote the origin of \mathbb{R}^d . Given a parameter $0 < \gamma \leq 1$, we say that K is γ -fat if there exist concentric Euclidean balls B and B' , such that $B \subseteq K \subseteq B'$, and $\text{radius}(B)/\text{radius}(B') \geq \gamma$. We say that K is fat if it is γ -fat for a constant γ (possibly depending on d , but not on ε).

Unless otherwise specified, the notion of ε -approximation between convex bodies will be based on the direction-sensitive definition given in Section 1.2. We say that a convex body K' is an *absolute* ε -approximation to another convex body K if they are within Hausdorff error ε of each other. Further, we say that K' is an *inner* (resp., *outer*) approximation if $K' \subseteq K$ (resp., $K' \supseteq K$).

Let B_0 denote a ball of radius $r_0 = \frac{1}{2}$ centered at the origin. For $0 < \gamma \leq 1$, let γB_0 denote the concentric ball of radius $\gamma r_0 = \frac{\gamma}{2}$. We say that a convex body K is in γ -canonical form if it is nested between γB_0 and B_0 . A body in γ -canonical form is γ -fat and has diameter $\Theta(1)$. We will refer to point O as the *center* of P .

For any point $x \in K$, define $\delta(x)$ to be minimum distance from x to any point on ∂K . For the sake of ray-shooting queries, it is useful to define a ray-based notion of distance as well. Given $x \in K$, define the *ray-distance* of x to the boundary, denoted $\text{ray}(x)$, as follows. Consider the intersection point p of ∂K and the ray emanating from O that passes through x . We define $\text{ray}(x) = \|xp\|$. The following utility lemma will be helpful in relating distances to the boundary.


► **Lemma 2.1.** *Given a convex body K in γ -canonical form:*

- (a) *For any point $x \in P$, $\delta(x) \leq \text{ray}(x) \leq \delta(x)/\gamma$.*
- (b) *Let h be a supporting hyperplane of K . Let p be any point inside K at distance at most some distance w from h , where $w \leq \gamma/4$. Let p' denote the intersection of the ray Op and h . Then $\|pp'\| \leq 2w/\gamma$.*
- (c) *Let p be any point on the boundary of K , and let h be a supporting hyperplane at p . Let h' denote the hyperplane obtained by translating h in the direction of the outward normal by some distance w . Let p' denote the intersection of the ray Op with h' . Then $\|pp'\| \leq w/\gamma$.*

We omit the straightforward proof. The lower bound on $\text{ray}(x)$ for part (a) is trivial, and the upper bound follows by a straightforward adaption of Lemma 4.2 of [6]. Part (b) is an adaptation of Lemma 2.11 of [7], and part (c) is similar.

2.1 Caps and Macbeath Regions

Much of the material in this section has been presented in [6, 7]. We include it here for the sake of completeness. Given a convex body K , a *cap* C is defined to be the nonempty intersection of K with a halfspace (see Figure 2(a)). Let h denote the hyperplane bounding this halfspace. 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 , denoted $\text{width}(C)$, is the distance between h and this supporting hyperplane. Given any cap



■ **Figure 2** (a) Cap concepts and (b) Macbeath regions.

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.)

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 2(b)). Clearly, $M^1(x)$ is centrally symmetric about x , and $M^\lambda(x)$ 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 refer to the latter as the *shrunk* Macbeath region.

We now present a few lemmas that encapsulate key properties of Macbeath regions. The first lemma shows that if two shrunk Macbeath regions have a nonempty intersection, then a constant factor expansion of one contains the other [7, 15, 21].

► **Lemma 2.2.** *Let K be a convex body, and let $\lambda \leq 1/5$ be any real. If $x, y \in K$ such that $M^\lambda(x) \cap M^\lambda(y) \neq \emptyset$, then $M^\lambda(y) \subseteq M^{4\lambda}(x)$.*

The following lemma shows that all points in a shrunk Macbeath region have similar distances from the boundary of K . The proof appears in [7].

► **Lemma 2.3.** *Let K be a convex body. If $x \in K$ and $x' \in M'(x)$, then $4\delta(x)/5 \leq \delta(x') \leq 4\delta(x)/3$.*

For any $\delta > 0$, define the δ -erosion of a convex body K , denoted $K(\delta)$, to be the closed convex body formed by removing from K all points lying within distance δ of ∂K . The next lemma bounds the number of disjoint Macbeath regions that can be centered on the boundary of $K(\delta)$. The proof appears in [7].

► **Lemma 2.4.** *Consider a convex body $K \subset \mathbb{R}^d$ in γ -canonical form for some constant γ . Define $\Delta_0 = \frac{1}{2}(\gamma^2/(4d))^d$. For any fixed constant $0 < \lambda \leq 1/5$ and real parameter $\delta \leq \Delta_0$, let \mathcal{M} be a set of disjoint λ -scaled Macbeath regions whose centers lie on the boundary of $K(\delta)$. Then $|\mathcal{M}| = O(1/\delta^{(d-1)/2})$.*

2.2 Shadows of Macbeath regions

Shrunk Macbeath regions reside within the interior of the convex body, but it is useful to identify the portion of the body’s boundary that this Macbeath region will be responsible for approximating. For this purpose, we introduce the shadow of a Macbeath region. Given a

convex body K that contains the origin O and a region $R \subseteq K$, we define the *shadow* of R (with respect to K), denoted $\text{shadow}(R)$, to be the set of points $x \in K$ such that the line segment Ox intersects R .

We also define a set of *normal directions* for R , denoted $\text{normals}(R)$. Consider the set of all hyperplanes that support K at some point in the shadow of R . Define $\text{normals}(R)$ to be the set of outward unit normals to these supporting hyperplanes. Typically, the region R in our constructions will be a (scaled) Macbeath region or an associated John ellipsoid (as defined in Section 3), close to the boundary of K . The following lemma captures a salient feature of these shadows, namely, that the shadow of a Macbeath region $M'(x)$ can be enclosed in an ellipsoid whose width in all normal directions is $O(\delta(x))$. The proof is presented in the full version.)

► **Lemma 2.5.** *Let $K \subset \mathbb{R}^d$ be a convex body in γ -canonical form for some constant γ . Let $x \in K$ be a point at distance δ from the boundary of K , where $\delta \leq \Delta_0$. Let $M = M'(x)$, $S = \text{shadow}(M)$, $N = \text{normals}(M)$, and $\widehat{M} = M^{4/\gamma}(x)$. Then:*

- (a) $S \subseteq \widehat{M}$.
- (b) $\text{width}_v(S) \leq c_1 \delta$ for all $v \in N$. Here c_1 is the constant $8/(3\gamma)$.
- (c) $\text{width}_v(\widehat{M}) \leq c_2 \delta$ for all $v \in N$. Here c_2 is the constant $160/(3\gamma^2)$.

2.3 Representation Conversions

Convex sets are naturally described in two ways, as the convex hull of a discrete set of points and as the intersection of a discrete set of halfspaces. Some computational tasks are more easily performed using one representation or the other, and hence it will be useful to convert between them. Also, when approximate representations suffice, it will be useful to prune a large set down to a smaller size. In this section we will present a few technical utilities to perform these conversions. We refer the reader to the full version for the missing proofs.

Given an n -element point set in \mathbb{R}^d , Chan showed that it is possible to construct an ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$ in time $O(n + 1/\varepsilon^{d-1})$ [17]. The following lemma shows that, by applying Chan's construction, it is possible to concisely approximate the convex hull of n points as the intersection of halfspaces.

► **Lemma 2.6.** *Let $\gamma < 1$ be a positive constant, and $\varepsilon > 0$ be a real parameter. Let P be a polytope in γ -canonical form represented as the convex hull of n points. In $O(n + 1/\varepsilon^{d-1})$ time it is possible to compute a polytope P' represented as the intersection of $O(1/\varepsilon^{(d-1)/2})$ halfspaces such that P' is an inner absolute ε -approximation of P .*

The following lemma is useful when representing polytopes by the intersection of halfspaces.

► **Lemma 2.7.** *Let $\gamma < 1$ be a positive constant, and $\varepsilon > 0$ be a real parameter. Let P be a polytope in γ -canonical form represented as the intersection of n halfspaces. In $O(n + 1/\varepsilon^{d-1})$ time it is possible to compute a polytope P' represented as the intersection of $O(1/\varepsilon^{(d-1)/2})$ halfspaces such that P' is an outer absolute ε -approximation of P .*

► **Remark.** Theorem 1.1 shows that an ε -kernel of size $O(1/\varepsilon^{(d-1)/2})$ can be computed in time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(d-1)/2+\alpha})$. The construction time in Lemma 2.7 (which is derived in the full version) is asymptotically dominated by the time needed to construct an ε -kernel. Therefore, the construction time can be reduced to this quantity.

3 Hierarchy of Macbeath Ellipsoids

The data structure presented in [7] for the approximate polytope membership problem is based on constructing a hierarchy of ellipsoids. In this section, we present a variant of this structure, which will play an important role in our constructions.

For a Macbeath region $M^\lambda(x)$, we denote its circumscribing John ellipsoid by $E^\lambda(x)$, which we call a *Macbeath ellipsoid*. Since Macbeath regions are centrally symmetric and the constant in John's Theorem [22] is \sqrt{d} for centrally symmetric bodies, we have $E^\lambda(x) \subseteq M^{\lambda\sqrt{d}}(x)$. Recall the constant $\Delta_0 = \frac{1}{2}(\gamma^2/4d)^d$ defined in the statement of Lemma 2.4, and define $\lambda_0 = 1/(20d)$. We omit the proof of the following lemma due to space limitations. (We caution the reader that in the lemmas of this section, the value of n used in the application of the lemma may differ from the original input size.)

► **Lemma 3.1.** *Let $\gamma < 1$ be a positive constant, and let $0 < \delta \leq \Delta_0$ be a real parameter. Let P be a polytope in γ -canonical form, represented as the intersection of n halfspaces. In $O(n/\delta^{d-1} + 1/\delta^{3(d-1)/2})$ time, we can construct a DAG structure satisfying the following properties:*

- (a) *The total number of nodes (including leaves), and the total space used by the DAG are each $O(1/\delta^{(d-1)/2})$.*
- (b) *Each leaf is associated with an ellipsoid $E^{4\lambda_0\sqrt{d}}(x)$, where $x \in \partial P(\delta)$. The union of the ellipsoids associated with all the leaves covers $\partial P(\delta)$.*
- (c) *Given a query ray Oq , in $O(\log \frac{1}{\delta})$ time, we can find a leaf node such that the associated ellipsoid intersects this ray.*

Given a convex body K and query point q , an *absolute ε -APM* query returns a positive result if q lies within K , a negative result if q is at distance at least ε from K , and otherwise it may return either result. After a small enhancement, this DAG can be used for answering absolute ε -APM queries for a polytope P in γ -canonical form. We assume that P is represented as the intersection of a set H of n halfspaces. We invoke the above lemma for $\delta = \varepsilon\gamma/(2c_1)$, where c_1 is the constant of Lemma 2.5(b). We then associate each leaf of the DAG with a halfspace as follows. Let x denote the center of the leaf ellipsoid and let p denote the intersection of the ray Ox with ∂P . Let $h \in H$ denote any supporting halfspace of P (containing P) at p . We store h with this leaf. By exhaustive search, we can determine h in $O(n)$ time, so the total time for this step is $O(n/\varepsilon^{(d-1)/2})$. Asymptotically, this does not affect the time it takes to construct the data structure. Given a query point q , we answer queries by first determining a leaf whose ellipsoid intersects the ray Oq . By Lemma 3.1(c), this takes $O(\log \frac{1}{\varepsilon})$ time. We return a positive answer if and only if q is contained in the associated halfspace.

The following lemma summarizes the result, whose proof is presented in the full version.

► **Lemma 3.2.** *Let $\gamma < 1$ be a positive constant, and let $\varepsilon > 0$ be a real parameter. Let P be a polytope in γ -canonical form, represented as the intersection of n halfspaces. In $O(n/\varepsilon^{d-1} + 1/\varepsilon^{3(d-1)/2})$ time, we can construct a data structure that uses $O(1/\varepsilon^{(d-1)/2})$ space and answers absolute ε -APM queries in $O(\log \frac{1}{\varepsilon})$ time.*

4 Kernel Construction

In this section we establish Theorem 1.1 by showing how to build an ε -kernel efficiently. The input to an ε -kernel construction consists of the approximation parameter ε and a set S of n points. Our algorithm is based on a bootstrapping strategy. We assume that we

have access to an algorithm that can construct an ε -kernel of $O(1/\varepsilon^{(d-1)/2})$ size in time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(1/2+\beta)(d-1)})$, where $\beta > 0$ is a parameter. Recall that the size of the kernel is asymptotically optimal in the worst case. We will present a method for improving the running time of this algorithm. Recall that Chan [17] gave an algorithm for constructing kernels of optimal size which runs in time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{d-1})$. By setting $\beta = \frac{1}{2}$, this will form the basis of our bootstrapping, which is described below. Throughout, let $\delta = \varepsilon^{1/3}$.

1. Fatten the input point set S by computing an affine transformation that maps S to S' , such that $\text{conv}(S')$ is in γ -canonical form for some constant γ . By standard results (see, e.g., the journal version of [4]), this can be done in $O(n)$ time.
2. Using Lemma 2.6, build a polytope P , represented as the intersection of $O(1/\delta^{(d-1)/2})$ halfspaces, such that P is an inner absolute δ -approximation of $\text{conv}(S')$. This step takes $O(n + 1/\delta^{d-1}) = O(n + 1/\varepsilon^{(d-1)/3})$ time.
3. Apply Lemma 3.1 to construct a DAG structure for P using the parameter δ . Replacing n in the statement of the lemma by $O(1/\delta^{(d-1)/2})$, it follows that this step takes $O(1/\delta^{3(d-1)/2}) = O(1/\varepsilon^{(d-1)/2})$ time.
4. By Lemma 3.1(c), for each point $p \in S'$, find a leaf of the DAG such that the associated ellipsoid $E^{4\lambda_0\sqrt{d}}(x)$ intersects the ray Op . Recall that $x \in \partial P(\delta)$. This takes $O(\log \frac{1}{\delta})$ per point. In $O(1)$ additional time, determine whether p lies in the shadow of this ellipsoid (with respect to $\text{conv}(S')$). If so, associate p with this ellipsoid, and otherwise discard it. All the points of S' can be processed in time $O(n \log \frac{1}{\delta}) = O(n \log \frac{1}{\varepsilon})$.
5. For each leaf ellipsoid of the DAG, build a $(c_3\varepsilon/\delta)$ -kernel for the points of S' that lie in its shadow, where c_3 is a suitably small constant that will be selected later. This kernel is computed using the aforementioned algorithm that computes the ε -kernel of a point set of size n in time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(1/2+\beta)(d-1)})$. The size of the $O(\varepsilon/\delta)$ -kernel computed for each shadow is $O((\delta/\varepsilon)^{(d-1)/2})$ and the time required is $O(n_i \log \frac{\delta}{\varepsilon} + (\delta/\varepsilon)^{(1/2+\beta)(d-1)})$, where n_i denotes the number of points of S' in the shadow. Summed over all the shadows, it follows that the total time required is

$$O\left(n \log \frac{\delta}{\varepsilon} + \left(\frac{1}{\delta}\right)^{\frac{d-1}{2}} \left(\frac{\delta}{\varepsilon}\right)^{\left(\frac{1}{2}+\beta\right)(d-1)}\right) = O\left(n \log \frac{1}{\varepsilon} + \left(\frac{1}{\varepsilon}\right)^{\left(\frac{1}{2}+\frac{2\beta}{3}\right)(d-1)}\right).$$

Here we have used the facts that each point of S' is assigned to at most one shadow and the total number of shadows, which is bounded by the number of leaves in the DAG, is $O(1/\delta^{(d-1)/2})$.

6. Let $S'' \subseteq S'$ be the union of the kernels computed in the previous step. Since the number of shadows is $O(1/\delta^{(d-1)/2})$ and the size of the kernel for each shadow is $O((\delta/\varepsilon)^{(d-1)/2})$, it follows that $|S''| = O(1/\varepsilon^{(d-1)/2})$. Apply the inverse of the affine transformation computed in Step 1 to the points of S'' , and output the resulting set of points as the desired ε -kernel for S .

We have shown that the size of the output kernel is $O(1/\varepsilon^{(d-1)/2})$, as desired. The running time of Step 5 dominates the time complexity. Our next lemma establishes the correctness of this construction.

► **Lemma 4.1.** *The construction yields an ε -kernel.*

Proof. Throughout this proof, for a given convex body K , we use $M_K(x)$, $E_K(x)$, and $\delta_K(x)$ to denote the quantities $M(x)$, $E(x)$, and $\delta(x)$ with respect to K . Let $P' = \text{conv}(S')$. By standard results on fattening, it suffices to show that $\text{conv}(S'')$ is an absolute $O(\varepsilon)$ -approximation of P' . Let v be an arbitrary direction. Consider the extreme point p of S'

in direction v . Clearly $p \in \partial P'$. Recall that P is an inner δ -approximation of P' , and the ellipsoids associated with the leaves of the DAG cover the boundary of $P(\delta)$. Thus, there must be an ellipsoid $E = E_P^{4\lambda_0\sqrt{d}}(x)$, $x \in \partial P(\delta)$, such that p is assigned to the shadow of E in Step 4. Note that this shadow and all shadows throughout this proof are assumed to be with respect to the polytope P' (and not P). We claim that $\text{width}_v(\text{shadow}(E)) \leq 2c_1\delta$, where c_1 is the constant of Lemma 2.5(b). Assuming this claim for now, let us complete the proof of the lemma. Recall that in Step 5, we built a $(c_3\varepsilon/\delta)$ -kernel for all the points of S' that are assigned to the shadow of E , and S'' includes all the points of this kernel. It follows that the distance between the supporting hyperplanes of $\text{conv}(S')$ and $\text{conv}(S'')$ in direction v is at most $(c_3\varepsilon/\delta) \cdot \text{width}_v(\text{shadow}(E)) \leq (c_3\varepsilon/\delta) \cdot (2c_1\delta) = 2c_1c_3\varepsilon$. By choosing c_3 sufficiently small, we can ensure that this quantity is smaller than any desired constant times ε , which proves the lemma.

It remains to show that $\text{width}_v(\text{shadow}(E)) \leq 2c_1\delta$. Recall that

$$E = E_P^{4\lambda_0\sqrt{d}}(x) \subseteq M_P^{4\lambda_0d}(x) = M'_P(x).$$

Furthermore, since $P \subseteq P'$, a straightforward consequence of the definition of Macbeath regions is that $M'_P(x) \subseteq M'_{P'}(x)$. To simplify the notation, let M denote $M'_{P'}(x)$. Putting it together, we obtain $E \subseteq M$. Thus $\text{shadow}(E) \subseteq \text{shadow}(M)$, which implies that $\text{width}_v(\text{shadow}(E)) \leq \text{width}_v(\text{shadow}(M))$. By Lemma 2.5(b),

$$\text{width}_v(\text{shadow}(M)) \leq c_1\delta_{P'}(x).$$

Using the triangle inequality and the fact that P is an inner δ -approximation of P' , we obtain $\delta_{P'}(x) \leq \delta_P(x) + \delta = 2\delta$. Thus $\text{width}_v(\text{shadow}(E)) \leq \text{width}_v(\text{shadow}(M)) \leq 2c_1\delta$, as desired. \blacktriangleleft

We are now ready to establish the main result of this section.

Proof. (of Theorem 1.1) Our proof is based on a constant number of applications of the algorithm from this section. It suffices to show that there is an algorithm that can construct an ε -kernel of $O(1/\varepsilon^{(d-1)/2})$ size in time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(1/2+\beta')(d-1)})$, where $\beta' = \alpha/(d-1)$.

We initialize the bootstrapping process by Chan's algorithm [17], which has $\beta = \frac{1}{2}$. Observe that the value of β is initially $\frac{1}{2}$ and falls by a factor of $\frac{2}{3}$ with each application of the algorithm. It follows that after $O(\log \frac{1}{\alpha})$ applications, we will obtain an algorithm with the desired running time. This completes the proof. \blacktriangleleft

5 Approximate Polytope Membership

In this section we show how to obtain a data structure for approximate polytope membership, proving Theorem 1.5. Our best data structure for APM achieves query time $O(\log \frac{1}{\varepsilon})$ with storage $O(1/\varepsilon^{(d-1)/2})$ and preprocessing time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(d-1)/2+\alpha})$. As with kernels, our construction here is again based on a bootstrapping strategy. To initialize the process, we will use a data structure that achieves the aforementioned query time with the same storage but with preprocessing time $O(n + 1/\varepsilon^{3(d-1)/2})$. The data structure is based on Lemma 3.2. Recall that the input is a polytope represented as the intersection of n halfspaces.

We begin by “fattening” the input polytope. As before, we use an affine transformation to map the input polytope to a polytope P' that is in γ -canonical form. This step takes $O(n)$ time [4]. By standard results, it suffices to build a data structure for answering absolute $O(\varepsilon)$ -APM queries with respect to P' (see, e.g., Lemma 7.1 of the journal version of [4]).

Next, we apply Lemma 2.7 to construct an outer absolute $O(\varepsilon)$ -approximation P of P' , where P is represented as the intersection of $O(1/\varepsilon^{(d-1)/2})$ halfspaces. This step takes $O(n + 1/\varepsilon^{d-1})$ time. Finally, we use Lemma 3.2 to construct a data structure for answering absolute $O(\varepsilon)$ -APM queries with respect to P . Replacing n in the statement of the lemma by $O(1/\varepsilon^{(d-1)/2})$, it follows that this step takes $O(1/\varepsilon^{3(d-1)/2})$ time.

The total construction time is $O(n + 1/\varepsilon^{3(d-1)/2})$. To answer a query, we map the query point using the same transformation used to fatten the polytope, and then use the data structure constructed above to determine whether the resulting point lies in polytope P . Subject to an appropriate choice of constant factors, the correctness of this method follows from the fact that P is an outer absolute $O(\varepsilon)$ -approximation of P' .

We summarize this result in the following lemma.

► **Lemma 5.1.** *Let $\varepsilon > 0$ be a real parameter and let P be a polytope, represented as the intersection of n halfspaces. In $O(n + 1/\varepsilon^{3(d-1)/2})$ time, we can construct a data structure that uses $O(1/\varepsilon^{(d-1)/2})$ space and answers ε -APM queries in $O(\log \frac{1}{\varepsilon})$ time.*

We now present the details of our bootstrapping approach. We assume that for a parameter $\beta > 0$, in time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(1/2+\beta)(d-1)})$ we can construct a data structure that can answer ε -APM queries in $O(\log \frac{1}{\varepsilon})$ time with $O(1/\varepsilon^{(d-1)/2})$ storage. We present a method for constructing a new data structure that matches the same storage and query time but has a lower preprocessing time. Throughout, let $\delta = \varepsilon^{\beta/(1+\beta)}$.

1. As in the kernel construction, first fatten the input polytope by applying an affine transformation that maps the input polytope to a polytope P' that is in γ -canonical form. By standard results (see, e.g., [4]), this step takes $O(n)$ time, and it suffices to build a data structure for answering absolute $O(\varepsilon)$ -APM queries with respect to P' .
2. Using Lemma 2.7, build an outer absolute $O(\varepsilon)$ -approximation of P' , denoted P , which is represented as the intersection of $O(1/\varepsilon^{(d-1)/2})$ halfspaces. By the remark following Lemma 2.7, this step takes $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(d-1)/2+\alpha})$ time.
3. Apply Lemma 3.1 to construct a DAG structure for P using the parameter δ . Replacing n in the statement of the lemma by $O(1/\varepsilon^{(d-1)/2})$, it follows that this step takes $O((1/\delta)^{d-1} \cdot (1/\varepsilon)^{(d-1)/2})$ time.
4. For each leaf of the DAG, construct an APM data structure as follows. Let $E = E^{4\lambda_0\sqrt{d}}(x)$ denote the ellipsoid associated with the leaf. Let R denote the minimum enclosing hyperrectangle of the ellipsoid $E^{4/\gamma}(x)$. We will see later that R contains the shadow of E (with respect to P), and its width in any direction in $\text{normals}(E)$ is at most $c_2 d \delta = O(\delta)$, where c_2 is the constant in Lemma 2.5(c).

Using the aforementioned algorithm, construct an APM data structure for this region with approximation parameter $c_3 \varepsilon / \delta$, where c_3 is a sufficiently small constant that we will select later. Note that each such region can be expressed as the intersection of $n_i = O(1/\varepsilon^{(d-1)/2})$ halfspaces, namely, all the halfspaces defining P together with the $2d$ halfspaces defined by the facets of R . The construction time of the APM data structure for each leaf is

$$O\left(n_i \log \frac{\delta}{\varepsilon} + \left(\frac{\delta}{\varepsilon}\right)^{\left(\frac{1}{2}+\beta\right)(d-1)}\right) = O\left(\left(\frac{1}{\varepsilon}\right)^{\frac{d-1}{2}} \log \frac{\delta}{\varepsilon} + \left(\frac{\delta}{\varepsilon}\right)^{\left(\frac{1}{2}+\beta\right)(d-1)}\right),$$

and the space used is $O((\delta/\varepsilon)^{(d-1)/2})$. Since there are $O(1/\delta^{(d-1)/2})$ leaves, it follows that the total space is $O(1/\varepsilon^{(d-1)/2})$, and the total construction time is the product of $O(1/\delta^{(d-1)/2})$ and the above construction time for each leaf.

Summing up the time over all the four steps, we obtain a total construction time on the order of

$$\left(n \log \frac{1}{\varepsilon} + \left(\frac{1}{\varepsilon} \right)^{\frac{d-1}{2} + \alpha} \right) + \left(\frac{1}{\delta} \right)^{d-1} \left(\frac{1}{\varepsilon} \right)^{\frac{d-1}{2}} + \left(\frac{1}{\delta} \right)^{\frac{d-1}{2}} \cdot \left(\left(\frac{1}{\varepsilon} \right)^{\frac{d-1}{2}} \log \frac{\delta}{\varepsilon} + \left(\frac{\delta}{\varepsilon} \right)^{\left(\frac{1}{2} + \beta \right)(d-1)} \right).$$

Recalling that $\delta = \varepsilon^{\beta/(1+\beta)}$ and assuming that the constant α is much smaller than β , it follows that the construction time is

$$O \left(n \log \frac{1}{\varepsilon} + \left(\frac{1}{\varepsilon} \right)^{\left(\frac{1}{2} + \frac{\beta}{1+\beta} \right)(d-1)} \right).$$

We answer queries as follows. We apply the affine transformation of Step 1 to the input query point to obtain a point q . Recall that it suffices to answer absolute $O(\varepsilon)$ -APM queries for q with respect to P' . As P is an outer absolute $O(\varepsilon)$ -approximation of P' , it suffices to answer absolute $O(\varepsilon)$ -APM queries for q with respect to P . To answer this query, we identify a leaf of the DAG such that the associated ellipsoid E intersects the ray Oq . This takes time $O(\log \frac{1}{\delta})$. Let y denote an intersection point of this ray with the ellipsoid E . If q lies on the segment Oy , then q is declared as lying inside P . Otherwise we return the answer we get for query q using the APM data structure we built for this leaf. It takes time $O(\log \frac{\delta}{\varepsilon})$ to answer this query. Including the time to locate the leaf, the total query time is $O(\log \frac{1}{\varepsilon})$. Our next lemma shows that queries are answered correctly.

► **Lemma 5.2.** *The query procedure returns a valid answer to the ε -APM query.*

Proof. We borrow the terminology from the query procedure given above. As mentioned, it suffices to show that our algorithm correctly answers absolute $O(\varepsilon)$ -APM queries for q with respect to the polytope P . Recall that we identify a leaf of the DAG whose associated ellipsoid $E = E^{4\lambda_0\sqrt{d}}(x)$ intersects the ray Oq . Recall that y is a point on the intersection of the ray Oq with E . Clearly, if q lies on segment Oy , then $q \in P$ and q is correctly declared as lying inside P .

It remains to show that queries are answered correctly when $\|Oq\| > \|Oy\|$. In this case, we handle the query using the APM data structure we built for the leaf. Recall that this structure is built for the polytope formed by intersecting P with the smallest enclosing hyperrectangle R of the ellipsoid $E^{4/\gamma}(x)$. It suffices to show: (i) $\text{shadow}(E) \subseteq R$ and (ii) $\text{width}_v(R) \leq c_2d\delta$ for all $v \in \text{normals}(E)$, where c_2 is the constant in Lemma 2.5(c).

To establish (i), recall that $M^\lambda(x) \subseteq E^\lambda(x) \subseteq M^{\lambda\sqrt{d}}(x)$ for any $\lambda > 0$. Using this fact, it follows that $M^{4/\gamma}(x) \subseteq E^{4/\gamma}(x) \subseteq M^{4\sqrt{d}/\gamma}(x)$. By Lemma 2.5(a), $\text{shadow}(E) \subseteq M^{4/\gamma}(x)$. Thus $\text{shadow}(E) \subseteq E^{4/\gamma}(x) \subseteq R$, which proves (i). To prove (ii), note that $R \subseteq E^{4\sqrt{d}/\gamma}(x)$, since R is the smallest enclosing hyperrectangle of $E^{4/\gamma}(x)$. Also $E^{4\sqrt{d}/\gamma}(x) \subseteq M^{4d/\gamma}(x)$. Thus $R \subseteq M^{4d/\gamma}(x)$. By Lemma 2.5(c), $\text{width}_v(M^{4/\gamma}(x)) \leq c_2\delta$ for all $v \in \text{normals}(M'(x))$. Since $R \subseteq M^{4d/\gamma}(x)$ and $E \subseteq M'(x)$, it follows that $\text{width}_v(R) \leq c_2d\delta$ for all $v \in \text{normals}(E)$.

We return to showing that queries are correctly answered when $\|Oq\| > \|Oy\|$. We consider two possibilities depending on whether q is inside or outside P . If $q \in P$ then $q \in \text{shadow}(E)$. By part (i) of the above claim, $\text{shadow}(E) \subseteq R$, and thus $q \in P \cap R$. It follows that the APM structure built for the leaf will declare this point as lying inside $P \cap R$, and hence the overall algorithm will correctly declare that q lies in P .

Finally, we consider the case when $q \notin P$. To complete the proof, we need to show that if the distance of q from the boundary of P is greater than ε , then q is declared as lying outside P . Let p denote the point of intersection of the ray Oq with ∂P , let h denote a hyperplane

supporting P at p , and let v denote the outward normal to h . Recall by part (i) of the claim that $\text{shadow}(E) \subseteq R$. It follows that h is a supporting hyperplane of $P \cap R$ at p . By part (ii) of the claim, $\text{width}_v(R) \leq c_2 d \delta$, and hence $\text{width}_v(P \cap R) \leq c_2 d \delta$. Recall that the APM data structure for the leaf is built using the approximation parameter $c_3 \varepsilon / \delta$ for some constant c_3 . By definition of APM query (in the standard, direction-sensitive sense), the absolute error allowed in direction v is at most $(c_3 \varepsilon / \delta) \cdot \text{width}_v(P \cap R) \leq (c_3 \varepsilon / \delta)(c_2 d \delta)$. By choosing c_3 sufficiently small we can ensure that this error is at most $\varepsilon \gamma$. To make this more precise, let h' denote the hyperplane parallel to h (outside P), and at distance $\varepsilon \gamma$ from it. Consider the halfspace bounded by h' and containing P . By the definition of APM query, if q is not contained in this halfspace, then q would be declared as lying outside $P \cap R$, and the overall algorithm would declare q as lying outside P . Let p' denote the point of intersection of the ray Oq with h' . By Lemma 2.1(c), $\|pp'\| \leq (\varepsilon \gamma) / \gamma = \varepsilon$. Thus, if the distance of q from ∂P is greater than ε , then q cannot lie on segment pp' and q is correctly declared as lying outside P . This completes the proof of correctness. \blacktriangleleft

We now establish the main result of this section.

Proof. (of Theorem 1.5) Our proof is based on a constant number of applications of the method presented in this section. It suffices to show that there is a data structure with space and query time as in the theorem and preprocessing time $O(n \log \frac{1}{\varepsilon} + 1/\varepsilon^{(1/2+\beta')(d-1)})$, where $\beta' = \alpha/(d-1)$.

We initialize the bootstrapping process by the data structure described in the beginning of this section, which has $\beta = 1$. Recall that applying the method once changes the value of β to $\beta/(1+\beta)$. It is easy to show that after i applications, the value of β will fall to $1/(i+1)$. Thus, after $O(1/\alpha)$ applications, we will obtain a data structure with the desired preprocessing time. \blacktriangleleft

The remaining theorems follow from previous reductions. Theorem 1.2 follows from performing $O(1/\varepsilon^{(d-1)/2})$ width queries [3, 17] using Theorem 1.6. Theorem 1.3 is a consequence of Theorem 1.5 together with [4, Lemma 9.2 of the journal version] and the construction from [11, Theorem 3.2]. Theorem 1.4 follows from 1.3 using [11, Theorem 4.1]. Theorem 1.6 follows from Theorem 1.5 by using duality and binary search. Theorem 1.7 is a consequence of Theorem 1.5 and the reduction presented in [4, Lemma 9.3 of the journal version].

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 P. K. Agarwal, S. Har-Peled, and K. R. Varadarajan. Geometric approximation via coresets. In J. E. Goodman, J. Pach, and E. Welzl, editors, *Combinatorial and Computational Geometry*. MSRI Publications, 2005.
- 3 P. K. Agarwal, J. Matoušek, and S. Suri. Farthest neighbors, maximum spanning trees and related problems in higher dimensions. *Comput. Geom. Theory Appl.*, 1(4):189–201, 1992.
- 4 S. Arya, G. D. da Fonseca, and D. M. Mount. Approximate polytope membership queries. In *Proc. 43rd Annu. ACM Sympos. Theory Comput.*, pages 579–586, 2011. doi:10.1145/1993636.1993713.
- 5 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.
- 6 S. Arya, G. D. da Fonseca, and D. M. Mount. On the combinatorial complexity of approximating polytopes. In *Proc. 32nd Internat. Sympos. Comput. Geom.*, pages 11:1–11:15, 2016. doi:10.4230/LIPIcs.SocG.2016.11.

- 7 S. Arya, G. D. da Fonseca, and D. M. Mount. Optimal approximate polytope membership. In *Proc. 28th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 270–288, 2017.
- 8 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.
- 9 S. Arya and D. M. Mount. A fast and simple algorithm for computing approximate Euclidean minimum spanning trees. In *Proc. 27th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 1220–1233, 2016.
- 10 S. Arya, D. M. Mount, and J. Xia. Tight lower bounds for halfspace range searching. *Discrete Comput. Geom.*, 47:711–730, 2012. doi:10.1007/s00454-012-9412-x.
- 11 Sunil Arya and Timothy M. Chan. Better ε -dependencies for offline approximate nearest neighbor search, Euclidean minimum spanning trees, and ε -kernels. In *Proc. 30th Annu. Sympos. Comput. Geom.*, pages 416–425, 2014.
- 12 I. Bárány. The technique of M-regions and cap-coverings: A survey. *Rend. Circ. Mat. Palermo*, 65:21–38, 2000.
- 13 G. Barequet and S. Har-Peled. Efficiently approximating the minimum-volume bounding box of a point set in three dimensions. *J. Algorithms*, 38(1):91–109, 2001.
- 14 J. L. Bentley, M. G. Faust, and F. P. Preparata. Approximation algorithms for convex hulls. *Commun. ACM*, 25(1):64–68, 1982. doi:10.1145/358315.358392.
- 15 H. Brönnimann, B. Chazelle, and J. Pach. How hard is halfspace range searching. *Discrete Comput. Geom.*, 10:143–155, 1993.
- 16 E. M. Bronshteyn and L. D. Ivanov. The approximation of convex sets by polyhedra. *Siberian Math. J.*, 16:852–853, 1976.
- 17 T. M. Chan. Faster core-set constructions and data-stream algorithms in fixed dimensions. *Comput. Geom. Theory Appl.*, 35(1):20–35, 2006. doi:10.1016/j.comgeo.2005.10.002.
- 18 T. M. Chan. Applications of Chebyshev polynomials to low-dimensional computational geometry. In *Proc. 33rd Internat. Sympos. Comput. Geom.*, pages 26:1–15, 2017.
- 19 R. M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *J. Approx. Theory*, 10(3):227–236, 1974.
- 20 K. Dutta, A. Ghosh, B. Jartoux, and N. H. Mustafa. Shallow packings, semialgebraic set systems, Macbeath regions and polynomial partitioning. In *Proc. 33rd Internat. Sympos. Comput. Geom.*, pages 38:1–15, 2017.
- 21 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.
- 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 S. Khuller and Y. Matias. A simple randomized sieve algorithm for the closest-pair problem. *Information and Computation*, 118(1):34–37, 1995.
- 24 A. M. Macbeath. A theorem on non-homogeneous lattices. *Ann. of Math.*, 56:269–293, 1952.
- 25 N. H. Mustafa and S. Ray. Near-optimal generalisations of a theorem of Macbeath. In *Proc. 31st Internat. Sympos. on Theoret. Aspects of Comp. Sci.*, pages 578–589, 2014.