Optimal Area-Sensitive Bounds for Polytope Approximation*

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‡
Aix-Marseille Université and LIS
Marseille, France
guilherme.fonseca@lis-lab.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

Abstract

Approximating convex bodies is a fundamental question in geometry and has a wide variety of applications. Given a convex body K in $\mathbb{R}^d$ for fixed d, the objective is to minimize the number of vertices (alternatively, the number of facets) of an approximating polytope for a given Hausdorff error ε. The best known uniform bound, due to Dudley (1974), shows that $O((\mathrm{diam}(K)/\varepsilon)^{(d-1)/2})$ facets suffice. While this bound is optimal in the case of a Euclidean ball, it is far from optimal for "skinny" convex bodies.

A natural way to characterize a convex object's skinniness is in terms of its relationship to the Euclidean ball. Given a convex body K, define its *surface diameter* $\Delta_{d-1}(K)$ to be the diameter of a Euclidean ball of the same surface area as K. It follows from generalizations of the isoperimetric inequality that $\mathrm{diam}(K) \geq \Delta_{d-1}(K)$.

We show that, under the assumption that the width of the body in any direction is at least ε, it is possible to approximate a convex body using $O((\Delta_{d-1}(K)/\varepsilon)^{(d-1)/2})$ facets. This bound is never worse than the previous bound and may be significantly better for skinny bodies. The bound is tight, in the sense that for any value of Δ_{d-1}, there exist convex bodies that, up to constant factors, require this many facets.

The improvement arises from a novel approach to sampling points on the boundary of a convex body. We employ a classical concept from convexity, called Macbeath regions. We demonstrate that Macbeath regions in K and K's polar behave much like polar pairs. We then apply known results on the Mahler volume to bound their number.

---

*A preliminary version of this paper appeared in the 28th Symp. on Computational Geometry, 2012.

†Research supported by the Research Grants Council of Hong Kong, China under project numbers 610012 and 16213219.

‡Supported by the French ANR PRC grant ADDS (ANR-19-CE48-0005).

§Research supported by NSF grant CCF-1618866.

1 Introduction

Approximating convex bodies by polytopes is a fundamental problem in computational and combinatorial geometry. (See Bronstein [24] for a survey.) Given a convex body K in Euclidean d-dimensional space and a scalar $\varepsilon > 0$, the problem is to construct a polytope P of low combinatorial complexity that is ε-close to K according to some distance measure. There are various ways to measure combinatorial complexity and various distance measures. In this paper, distances are based on the Hausdorff metric and complexity is measured by the number of facets in the approximating polytope. A polytope P is said to be an *ε-approximation* to K if the Hausdorff distance [24] between K and P is at most ε. The approximation is *outer* if $K \subseteq P$ and is *inner* if $P \subseteq K$. Throughout, we assume that the dimension d is a constant, implying that our asymptotic forms conceal constant factors that depend on d.

There are two common types of approximation bounds in the literature. In both cases a positive parameter ε_0 is given, and the bounds hold for all $\varepsilon \leq \varepsilon_0$. Bounds are said to be *nonuniform* if the value of ε_0 depends on properties of K, and otherwise they are said to be *uniform*. Nonuniform bounds often hold subject to smoothness conditions on K's boundary (e.g., K's boundary is C^2 continuous). Examples of nonuniform bounds appear in works by Gruber [34], Clarkson [28], and others [14, 42, 51, 52]. In contrast, in uniform bounds the value of ε_0 is independent of K (but may depend on d). Such bounds hold without any additional smoothness assumptions. Examples include the results of Dudley [31] and Bronshteyn and Ivanov [23]. Our results are of this latter type.

Dudley [31] showed that any convex body K can be ε-approximated by a polytope P with $O((\mathrm{diam}(K)/\varepsilon)^{(d-1)/2})$ facets, where $\mathrm{diam}(K)$ denotes K's diameter and $0 < \varepsilon \leq \mathrm{diam}(K)$. Bronshteyn and Ivanov showed the same asymptotic bound holds for the number of vertices. Constants hidden in the O-notation depend only on d. These results have many applications, for example, in the construction of coresets [2, 5, 7].

As evidenced by the case where K is a Euclidean ball, the bounds given by both Dudley and Bronshteyn-Ivanov are tight in the worst case up to constant factors [24]. However, these bounds may be significantly suboptimal if K is "skinny". A convex body's skinniness can be measured based on its similarity to a Euclidean ball. We define the *volume diameter* of K, denoted $\Delta_d(K)$, to be the diameter of a Euclidean ball of the same volume as K. Its *surface diameter*, denoted $\Delta_{d-1}(K)$, is defined analogously for surface area. Up to constant factors, these quantities are closely related to the classical concepts of *quermassintegrals* and of *intrinsic volumes* of the convex body [43, 44]. It follows from generalizations of the isoperimetric inequality that $\mathrm{diam}(K) \geq \Delta_{d-1}(K) \geq \Delta_d(K)$ [44].

In this paper, we strengthen Dudley's bound by showing that the complexity of approximation can be made sensitive to K's skinniness as expressed in terms of surface diameter. Here is our main result.

Theorem 1. *Consider any convex body K in $\mathbb{R}^d$ and any $\varepsilon > 0$ such that the width of K in any direction is at least ε. There exists an outer ε-approximating polytope P for K whose number of facets is at most*

$$\left(\frac{c_d \Delta_{d-1}(K)}{\varepsilon}\right)^{\frac{d-1}{2}},$$

where c_d is a constant (depending on d) and $\Delta_{d-1}(K)$ is K's surface diameter.

The bound can equivalently be stated in terms of K's surface area as $c_d \sqrt{\mathrm{area}(K)}/\varepsilon^{(d-1)/2}$. The isoperimetric inequality implies that for a given diameter, the surface area is maximized for a Euclidean ball, implying that (up to constant factors) $\mathrm{area}(K) \le \mathrm{diam}(K)^{d-1}$. Thus, by taking K to be a ball of radius $r \geq \varepsilon$, the tightness of Dudley's bound implies that, up to constant factors depending on the dimension, the number of facets is at least $(r/\varepsilon)^{(d-1)/2} \approx \sqrt{\mathrm{area}(K)}/\varepsilon^{(d-1)/2}$. Thus, as a function of surface area or surface diameter alone the bound of Theorem 1 is tight up to constant factors.

The problem of shape-sensitive approximations was considered by Bonnet [19], who considered the problem in the uniform setting (which he terms the non-asymptotic, non-smooth case). His results are presented in terms of the intrinsic volume $V_i(K)$, for $1 \le i \le d$. Up to constant factors, we can define $\Delta_i(K)$ to be $V_i(K)^{1/i}$. Expressed in our notation, his results imply that there exists an ε-approximation to K with $O((\mathrm{diam}(K)\delta^{\beta}/\varepsilon)^{(d-1)/2})$ facets, where $\delta = \Delta_{(d-1)/2}(K)/\mathrm{diam}(K)$ and β is roughly $1/2d$. (He conjectures that the results hold for $\beta = 1$ and with $\Delta_{d-1}(K)$ in place of $\Delta_{(d-1)/2}(K)$, which basically matches the bound of Theorem 1.) We also considered the problem of area-sensitive approximation in an earlier work [6]. The bound presented there was worse by a factor of $\log^{O(1)} \varepsilon$. In Section 5 we will present a simple derivation of Theorem 1 in the nonuniform setting.

The width requirement in Theorem 1 seems to be necessary. Consider, for example, a $(d-2)$-dimensional unit ball B embedded within $\mathbb{R}^d$, and let B_δ denote its Minkowski sum with a d-dimensional ball of radius $\delta \ll \varepsilon$. By the optimality of Dudley's bound for Euclidean balls, $\Omega(1/\varepsilon^{(d-3)/2})$ facets are needed to approximate B, and hence this bound applies to B_δ as well. However, the surface area of B_δ can be made arbitrarily small as a function of δ. Of course, the width condition can always be satisfied by first taking the Minkowski sum of K with a Euclidean ball of radius $\varepsilon/2$.

An additional contribution of this paper, which may be of independent interest, involves our approach. We show that the problem of generating a Hausdorff ε-approximation to a convex body in $\mathbb{R}^d$ can be reduced to the problem of generating an ε-approximation to a convex function. More precisely, let Ψ denote a convex subset of $\mathbb{R}^{d-1}$ and consider a convex function $f : \Psi \to \mathbb{R}$. We say that a function $\widehat{f}$ is a (lower) *ε-approximation* to f if $f(x) - \varepsilon \le \widehat{f}(x) \le f(x)$, for all $x \in \Psi$. We show (in Lemma 3) how to map any convex body K in $\mathbb{R}^d$ to a set of $2d$ functions and associated domains $\{(f_1, \Psi_1), \ldots, (f_{2d}, \Psi_{2d})\}$ such that any set of lower ε-approximations to these functions can be combined to obtain a single outer ε-approximation to K. To achieve area sensitivity, our reduction guarantees that, up to constant factors, $\mathrm{area}(\Psi_i) \le \mathrm{area}(K)$. In Theorem 2 in Section 2.2, we present an area-sensitive bound on the complexity of a piecewise-linear ε-approximation to such a function. (Technically, our approximation applies to a subdomain that is infinitesimally close to Ψ_i.)

In a recent paper, we presented a volume-sensitive bound by proving the existence of an approximation with $O((\Delta_d/\varepsilon)^{(d-1)/2})$ facets [12]. As observed above, $\Delta_{d-1}(K) \ge \Delta_d(K)$, and hence that bound subsumes the area-sensitive bound presented here. However, we believe that the area-sensitive bounds are of intrinsic interest. By its nature, approximation applies near a body's boundary, and hence surface area is a relevant parameter. There are applications involving unbounded objects, such as in the generation of space-efficient minimization diagrams for approximate nearest-neighbor searching [1,36], where area-sensitivity is meaningful, but volume sensitivity is not. Also, our function-based approach may have additional applications. Rote applied the functional perspective to convex approximation in $\mathbb{R}^2$, resulting in a bound that matches ours for the $d = 2$

case [50, Theorem 3]. He posed the open problem of higher dimensional generalizations.

1.1 Notation and Assumptions

Throughout, K will denote a convex body in $\mathbb{R}^d$, that is, a compact convex subset with nonempty interior. Let us assume that K is positioned so that its centroid coincides with the origin. Let $\mathrm{vol}(K)$ denote its d-dimensional Hausdorff measure, $\mathrm{area}(K)$ denote its surface area measure, and ∂K denote its boundary. We assume that the dimension d is a constant. For $\alpha \geq 0$, αK denotes a uniform scaling of K about the origin, and for $p \in \mathbb{R}^d$, $K + p$ denotes the translation of K by p. Let $K \oplus r$ denote the *dilation* of K by r, that is, the Minkowski sum of K and a Euclidean ball of radius r centered at the origin. Let $K \ominus r$ denote the *erosion* of K by r, that is, the subset of K that lies at distance at least r from ∂K. The dilation and erosion of a convex set are both convex. A convex body is *smooth*, if at each boundary point, there exists a unique supporting hyperplane, and it is *strictly convex* if its boundary contains no line segment. Throughout, we use $\langle \cdot, \cdot \rangle$ to denote the standard inner (dot) product, and use $\| \cdot \| = \sqrt{\langle \cdot, \cdot \rangle}$ to denote the Euclidean norm.

Letting $\Psi \subseteq \mathbb{R}^{d-1}$ be an open convex set, consider a function $f : \Psi \to \mathbb{R}$. This function's *graph* is the set $\{(u; v) \in \mathbb{R}^{d-1} \times \mathbb{R} \,:\, v = f(u)\}$, and its *epigraph* is similarly defined where $v \geq f(u)$. If f is a convex function, its epigraph is a convex set in $\mathbb{R}^d$. To exploit the relationship between convex sets and epigraphs, we will often express a point $q \in \mathbb{R}^d$ as a pair $(u; v) \in \mathbb{R}^{d-1} \times \mathbb{R}$. If f is smooth, let $\nabla f = \left(\frac{\partial f}{\partial x_1}, \ldots, \frac{\partial f}{\partial x_{d-1}}\right)$ denote its gradient. Given a hyperplane $y = \langle a, x \rangle + b$, define its *lower* (resp., *upper*) *halfspace* to be the set of points $q = (u; v)$ where $v \leq \langle a, u \rangle + b$ (resp., $v \geq \langle a, u \rangle + b$). If $q = (u; v)$ is a point on f's graph and $\nabla f(u)$ is defined, then the *tangent hyperplane* for f at q is

$$h_f(q) \;=\; \{(x; y) \in \mathbb{R}^{d-1} \times \mathbb{R} \,:\, y = \langle \nabla f(u), x - u \rangle + v\}. \tag{1}$$

Given $\varepsilon > 0$, we say that $\widehat{f}$ is an *upper ε-approximation to f over a subdomain* $\Psi^- \subseteq \Psi$ if for all $x \in \Psi^-$, $f(x) \leq \widehat{f}(x) \leq f(x) + \varepsilon$. A *lower ε-approximation* is defined analogously with $f(x) - \varepsilon \leq \widehat{f}(x) \leq f(x)$.

To avoid specifying many real-valued constants that arise in our analysis, we will often hide them using asymptotic notation. For positive real x, we use the notation $O(x)$ or "$\lesssim x$" to denote a quantity whose value is at most cx for an appropriately chosen constant c. We use $\Omega(x)$ or "$\gtrsim x$" for a quantity that is at least cx. We use $\Theta(x)$ or "$\approx x$" to denote a quantity that lies within the interval $[cx, c'x]$ for appropriate constants c and c'. These constants may depend on the dimension d, but not on K or ε. Therefore, all such bounds hold uniformly.

1.2 Overview of Methods

It is well known that computing a Hausdorff approximation of a convex body K by a polytope can be reduced to sampling an appropriate set system involving surface patches on K's boundary (see, e.g., [22, 23, 27]). For example, an outer approximation can be obtained by first sampling a sufficiently dense set of points from K's boundary and then intersecting the halfspaces defined by the supporting hyperplanes at each of these points. The definition of "sufficiently dense" can be based on the concept of a hitting set for dual caps. Consider a point q' that is at distance ε from K, and let q be its closest point on the boundary of K (see Figure 1(a)). The set of points on the boundary of K that are visible to q' is called an *ε-dual cap*. (Dual caps will be defined formally in Section 2.3.) We say that a set P of points on K's boundary is a *hitting set* for dual caps if every ε-dual cap contains at least one point $p \in P$. It is easy to see that a polytope formed from the supporting hyperplanes at these points yields an ε-approximation because every point that is at distance at least ε from K (such as q') is separated from its closest boundary point by at least one of these supporting hyperplanes (such as the one at p).

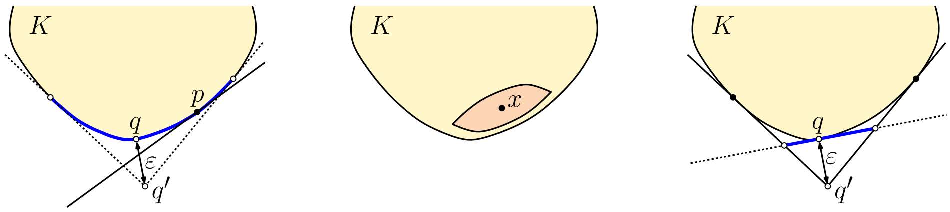

(a) (b) (c)

Figure 1: Overview of methods: (a) ε-dual caps and hitting, (b) shrunken Macbeath region, and (c) the base of a dual cap.

Prior attempts to obtain an area-sensitive approximation bound suffered from an additional polylogarithmic factor due to a relatively inefficient sampling method based on ϵ-nets for halfspace ranges [6]. In this paper, we develop a more efficient approach, which is based on *Macbeath regions*. This is a classical concept from the theory of convexity [40]. Given a convex body K and a point $x \in K$, a Macbeath region is a centrally symmetric body centered at a point of K that adheres locally to K's shape (see Figure 1(b)). One way to use Macbeath regions to build a hitting set is to first generate a maximal set of disjoint shrunken Macbeath regions that lie within distance of roughly ε of K's boundary. A constant number of points are then sampled from the neighborhood around each Macbeath region, and these points are projected onto K's boundary. It can be shown that this yields a valid hitting set and hence a valid approximating polytope.

Unfortunately, a direct application of this process does not yield a good bound on the complexity of the approximation. While a packing argument can be applied to bound the number of Macbeath regions of large volume, it is not easy to bound the number of Macbeath regions of small volume. However, the number of small Macbeath regions can be bounded by appealing to polarity. An interesting property of Macbeath regions is that Macbeath regions that have small volume in the original body, correspond to Macbeath regions in the polar body that have large volume.

In order to exploit this property, we introduce two intermediate structures. First, we show that K can be represented as the intersection of $2d$ unbounded convex sets, which we call *support sets* (see Figure 2 in Section 2.1). Each of these sets can be interpreted as the epigraph of a convex real-valued function over $\mathbb{R}^{d-1}$. To achieve area-sensitivity, we make further refinements to these functions. The result is a set of $2d$ functions and associated domains $\{(f_1, \Psi_1), \ldots, (f_{2d}, \Psi_{2d})\}$, which possess a number of nice properties. Each such function f_i is smooth and strictly convex. Each domain $\Psi_i \subset \mathbb{R}^{d-1}$ is bounded, open, and convex, such that $\mathrm{area}(\Psi_i) \lesssim \mathrm{area}(K)$. Each function has a dual representation in the form of its Legendre transform, denoted f_i^*, whose domain is $\mathbb{R}^{d-1}$. We show that each point on the graph of f_i has a unique *dual counterpart* on the graph of f_i^*. Finally, we show that by independently building approximations to each of these functions, these multiple approximations can be combined to obtain a single approximation to K. We abstract these properties in a concept, which we call a *U-shaped set*.

The notion of an ε-dual cap can be readily adapted to this functional context. We also define the notion of an ε-cap, and we show that the ε-dual cap induced by a point on the graph of f_i corresponds to the ε-cap of its dual counterpart in f_i^*. For each dual cap and each cap, we introduce a $(d-1)$-dimensional convex body called its *base*. A key result relating caps and dual caps is presented in Lemma 8, where we show that, subject to a projection and scaling, the base of a dual cap of f_i is the polar of the base of the associated cap of f_i^*.

This polar relationship is a central ingredient in our construction because it allows us to apply the classical concept of the *Mahler volume*, which implies that small dual caps in the epigraph of f_i correspond to large caps in the epigraph of the dual f_i^*. We show that this applies to the Macbeath regions in the vicinity of these caps and dual caps. We employ a two-pronged approach to sampling from both the primal and dual. Macbeath regions that are too small to be sampled from f_i will be sampled as large Macbeath regions in f_i^*. Observations in this spirit have been made in other contexts (see, e.g., [4, 10, 47]), but the functional formulation given here is both cleaner and more general.

The remainder of the paper is organized as follows. In Section 2 we introduce our functional approach to approximation, where in Lemma 3 we prove that independent approximations to these functions can be combined to obtain a single approximation for the original convex body. We also present Lemma 4, which establishes how a hitting set for dual caps yields an approximation. Next, in Section 3 we explore the dual relationships between caps and dual caps. There we present Lemma 8, which shows that the projected bases of caps and dual caps are polars of each other. In Section 4 we present our Macbeath-based sampling process. First in Section 4.2 we introduce Macbeath regions, and we present an adaptation of a well known result about Macbeath regions, called *cap covering* (which is proved later in Section 6). Then in Sections 4.3 and 4.4 we explain how our Macbeath-based machinery can be employed to build hitting sets for large and small dual caps, respectively. Finally, in Section 5 we present an additional result which shows that a bound similar to ours can be derived in the nonuniform setting.

2 Approximation Through a Functional Lens

2.1 Support Sets

In our analysis, we will use approximations of convex functions as intermediaries when working with approximations of convex bodies. We will employ a representation of our body K in $\mathbb{R}^d$ as the intersection of $2d$ functions and associated domains $\{(f_1, \Psi_1), \ldots, (f_{2d}, \Psi_{2d})\}$. In this section, we will describe this functional representation and its relevance to convex approximation.

Let $\Gamma^d = [-1,1]^d$ denote a hypercube in $\mathbb{R}^d$. This hypercube is bounded by $2d$ (closed) facets, which can be identified with the elements of the set $I^{\pm} = \{\pm 1, \ldots, \pm d\}$. In particular, for each $i \in I^{\pm}$, facet F_i consists of the vectors whose ith coordinate is either $+1$ or -1 depending on the sign of i, and all its other coordinates lie in the interval $[-1,1]$. Given any vector $u \in \mathbb{R}^d$, let $\varphi_K(u) = \sup_{x \in K} \langle x, u \rangle$ denote the support function of K. Given any $i \in I^{\pm}$, define K_i to be the largest convex set whose support function agrees with K's over the vectors in F_i, that is, $K_i = \{x \in \mathbb{R}^d : \langle x, u \rangle \leq \varphi_K(u), \forall u \in F_i\}$ (see Figure 2). Intuitively, K_i is the intersection of all the halfspaces that contain K and are orthogonal to an element of F_i. Because these facets cover all the possible directions, $K = \bigcap_{i \in I^{\pm}} K_i$. Also, subject to an appropriate rotation, each set K_i can be interpreted as the epigraph of a convex function over $\mathbb{R}^{d-1}$.

For the remainder of this section, let us fix any $i \in I^{\pm}$, and let X_i denote the $(d-1)$-dimensional

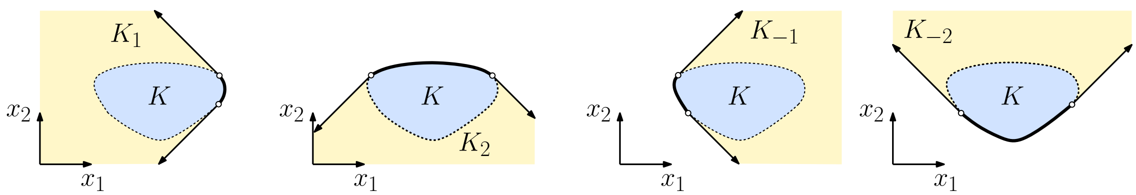

Figure 2: Support sets.

linear subspace spanned by the basis vectors other than the ith basis vector. We may identify X_i with $\mathbb{R}^{d-1}$, and we may express any $q = (q_1, \ldots, q_d) \in \mathbb{R}^d$ as a pair $(u; v) \in \mathbb{R}^{d-1} \times \mathbb{R}$, where $u = (q_j)_{j \neq |i|}$ and $v = -\mathrm{sign}(i) \cdot q_i$. Under this interpretation, K_i is the epigraph of a convex real-valued function over $\mathbb{R}^{d-1}$. Define $q^{\downarrow} = u$ to be the orthogonal projection of q onto $\mathbb{R}^{d-1}$, and define the projection of any set $Q \subseteq \mathbb{R}^d$, denoted $Q^{\downarrow}$, analogously. (The meaning of $Q^{\downarrow}$ is dependent on the choice of $i \in I^{\pm}$, but whenever we use it, the value of i will be clear from the context.)

From the functional perspective, the interesting portion of a convex set is the portion of the boundary that is visible from below. Define the *lower boundary* of a convex set Q, denoted $\overline{\partial} Q$, to be the set of boundary points of Q such that there exists a supporting hyperplane passing through this point whose outward normal vector points downwards, meaning that it has a strictly negative dot product with the vertical axis. Clearly, for every point $p' \in Q^{\downarrow}$, there is a point $p \in \overline{\partial} Q$ such that $p' = p^{\downarrow}$.

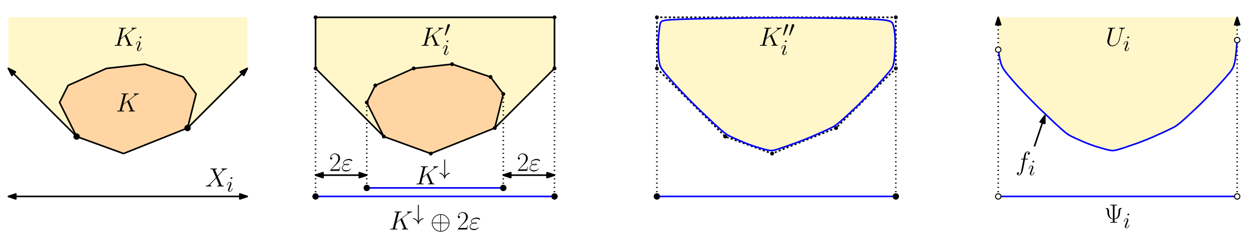

(a) (b) (c) (d)

Figure 3: Restricting and smoothing the support set.

The projection $K_i^{\downarrow}$ covers all of $\mathbb{R}^{d-1}$ (see Figure 3(a)), but for the sake of deriving area-sensitive bounds, it will be necessary to restrict its projection to one whose area more closely matches K's surface area. For any $\alpha \geq 0$, recall that $K^{\downarrow} \oplus \alpha$ denotes the subset of $\mathbb{R}^{d-1}$ that lies within distance α of $K^{\downarrow}$. To limit the horizontal extent, we intersect K_i with a vertical cylinder C_i whose horizontal cross section is $K^{\downarrow} \oplus 2\varepsilon$. To limit the vertical extent we take the portion of this set lying below a horizontal hyperplane that is sufficiently high to include the entire lower boundary of $K_i \cap C_i$. Call the resulting set K_i' (see Figure 3(b)). Clearly, this set is bounded, but we also want its lower boundary to be smooth. To do this, we apply a result due to Klee [38, Theorem 1.5], which shows that there is a convex body $K_i'' \subseteq K_i'$ that is arbitrarily close to K_i' and whose boundary is smooth and strictly convex (see Figure 3(c)). Klee defines closeness in terms of the Banach-Mazur distance where the the origin may be placed anywhere within K_i', but since K_i' is bounded, we may use Hausdorff distance instead.

Define $\Psi_i = (\overline{\partial} K_i'')^{\downarrow}$, that is, the vertical projection of the points on the lower boundary of K_i''. Clearly, Ψ_i is bounded and convex. We will show below that it is also open. Let $f_i : \mathbb{R}^{d-1} \to \mathbb{R}$

denote the function over Ψ_i whose graph is $\overline{\partial} K_i''$. It follows that f_i is smooth, strictly convex, and its range is bounded. From the definition of Ψ_i and the smoothness and strict convexity of f_i, the mapping $x \mapsto \nabla f_i(x)$ defines a homeomorphism between Ψ_i and $\mathbb{R}^{d-1}$. Define U_i to be the epigraph of f_i (see Figure 3(d)).

These sets U_i and the associated functions f_i will play a key role in our future constructions. We encapsulate their principal properties in the following concept.

Definition 1 (U-Shaped Function/Set)**.** *Given a bounded open, convex domain $\Psi \subset \mathbb{R}^{d-1}$, a function $f : \Psi \to \mathbb{R}$ is* U-shaped *if it is smooth, strictly convex, has a bounded range, and the function $x \mapsto \nabla f(x)$ is a homeomorphism between Ψ and $\mathbb{R}^{d-1}$. A set $U \subset \mathbb{R}^d$ is* U-shaped *if it is the epigraph of a U-shaped function.*

U-shaped functions are closely related to the concept of convex functions of Legendre type, as defined by Rockafellar [49].

For some of our results, we will require the additional property that the slope of our function is bounded by a constant over almost all of its domain. Given two positive scalars δ and γ, we say that a U-shaped function f over Ψ is (δ, γ)-*slope restricted* if for all $x \in \Psi \ominus \delta$, $\|\nabla f(x)\| \leq \gamma$. (Recall that $\Psi \ominus \delta$ is the erosion of Ψ by δ, and $\|\cdot\|$ is the Euclidean norm.) The following lemma summarizes the salient properties of the above construction.

Lemma 1. *Consider a convex body K in $\mathbb{R}^d$ and scalars $\delta > 0$ and $\gamma \geq 2\sqrt{d-1}$. For each $i \in I^{\pm}$, there exists a (δ, γ)-slope restricted U-shaped function f_i over the domain $\Psi_i = K^{\downarrow} \oplus 2\varepsilon$ such that f_i is an upper δ-approximation to the function whose epigraph is K_i over the subdomain $\Psi_i \ominus \delta$.*

Proof. The function f_i, domain Ψ_i, and set U_i are as defined in the above construction. Since i will be fixed throughout, to simplify notation, we will omit it throughout the proof. (Note that all instances of K in the proof refer to K_i.)

We first show that Ψ and f satisfy the conditions of being U-shaped. By construction, Ψ is bounded and convex, and f is smooth and strictly convex. By limiting the vertical extent in forming K_i', its range is bounded. Since the lower boundary consists of points whose normal has a strictly negative dot product with the vertical axis, the smoothness of K'' implies that Ψ is open. The smoothness of f implies that for each $x \in \Psi$, there is a well defined gradient $\nabla f(x)$. The normal vectors along the boundary of K'' vary continuously over all possible directions, and so it follows that for each vector $w \in \mathbb{R}^{d-1}$, there exists a point on $\overline{\partial} K''$ such that the gradient of the associated function is w. Continuity and uniqueness imply that this is a homeomorphism.

Next, to establish the slope restriction, let Ψ^- denote the erosion $\Psi \ominus \delta$. We assume that Klee's construction yields a convex body K'' that is within Hausdorff distance δ' of K', where $\delta' < \delta/\sqrt{5(d-1)}$.

The support-set construction implies that the vertical component of any outward unit normal vector for any point on the boundary of K dominates all the other coordinates in absolute value. If we view the boundary of K from a functional perspective and identify X with $\mathbb{R}^{d-1}$, this implies that if we move a distance δ'' along any basis vector of $\mathbb{R}^{d-1}$, the function value can increase by at most δ''. It follows that if we move a distance δ'' in any direction in $\mathbb{R}^{d-1}$, the function value can increase by at most $\delta''\sqrt{d-1}$.

We assert that for any point $q = (u; v) \in \overline{\partial} U$, where $u \in \Psi^-$, $\|\nabla f(u)\| \leq 2\sqrt{d-1}$ (see Figure 4(a)). Suppose to the contrary that $\|\nabla f(u)\| > 2\sqrt{d-1}$. To simplify notation, we will use ∇f in place of $\nabla f(u)$ throughout the proof. The equation of the supporting hyperplane at point q is given by $y = \langle \nabla f, x - u \rangle + v$. The upward-directed normal to this hyperplane in $\mathbb{R}^d$ is $(-\nabla f; 1)$, which we denote by n. Let $\widehat{n} = n/\|n\|$ be the associated unit normal. Note that $\|n\| = \sqrt{\|\nabla f\|^2 + 1}$.

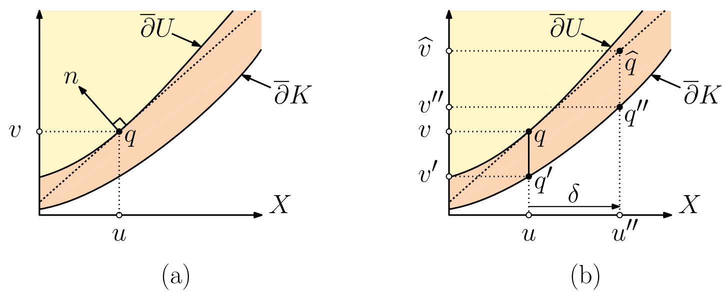

Figure 4: Proof of Lemma 1.

Let $q' = (u; v')$ be the point on the lower boundary of K that is vertically below q (see Figure 4(b)). Let u'' be the point obtained by translating u by distance δ in the direction of ∇f. Since $u \in \Psi^-$, it follows from the triangle inequality that $u'' \in \Psi$. Let $q'' = (u''; v'')$ and $\widehat{q} = (u''; \widehat{v})$ be the points on the lower boundary of K and the supporting hyperplane at point q, respectively. We assert that q'' is at distance greater than δ' from $\overline{\partial}U$, in violation of Klee's construction. This will complete the proof by contradiction.

To prove the above assertion, recall from our earlier remarks that v'' exceeds v' by at most $\delta\sqrt{d-1}$. Further, $\widehat{v}$ exceeds v by at least $\delta\|\nabla f\|$. It follows that

$$\widehat{v} - v'' \;\geq\; (\widehat{v} - v'') - (v - v') \;=\; (\widehat{v} - v) - (v'' - v') \;\geq\; \delta\left(\|\nabla f\| - \sqrt{d-1}\right).$$

By convexity, $\overline{\partial}U$ lies entirely above the supporting hyperplane at q, and thus the distance between q'' and $\overline{\partial}U$ is at least as large as the distance of q'' to the hyperplane, which is

$$\frac{\widehat{v} - v''}{\|n\|} \;\geq\; \frac{\delta\left(\|\nabla f\| - \sqrt{d-1}\right)}{\sqrt{\|\nabla f\|^2 + 1}}.$$

It is easy to verify that the quantity on the right hand side increases monotonically with $\|\nabla f\|$ in the range $\|\nabla f\| > 2\sqrt{d-1}$. Hence the distance between q'' and $\overline{\partial}U$ is at least

$$\begin{aligned} \frac{\delta\left(2\sqrt{d-1} - \sqrt{d-1}\right)}{\left(\left(2\sqrt{d-1}\right)^2 + 1\right)^{1/2}} \;&=\; \frac{\delta\sqrt{d-1}}{(4(d-1)+1)^{1/2}} \;\geq\; \frac{\delta\sqrt{d-1}}{(4(d-1)+(d-1))^{1/2}} \\ &=\; \frac{\delta}{\sqrt{5}} \;\geq\; \frac{\delta}{\sqrt{5(d-1)}} \;>\; \delta', \end{aligned}$$

as desired.

It remains to show that f is an upper δ-approximation to the function whose epigraph is K over the subdomain $\Psi \ominus \delta$. For this purpose, it suffices to show that the distance between q and q' is at most δ. As observed earlier, $\overline{\partial}U$ lies entirely above the supporting hyperplane at q, and thus the distance between q' and $\overline{\partial}U$ satisfies

$$\mathrm{dist}(q', \overline{\partial}U) \;\geq\; \frac{v - v'}{\|n\|} \;=\; \frac{v - v'}{\sqrt{\|\nabla f\|^2 + 1}}.$$

By Klee's construction, $\mathrm{dist}(q', \overline{\partial} U) \leq \delta'$. Since $\|\nabla f\| \leq 2\sqrt{d-1}$, we have

$$v - v' \;\leq\; \delta' \sqrt{\|\nabla f\|^2 + 1} \;<\; \frac{\delta}{\sqrt{5(d-1)}} \left((2\sqrt{d-1})^2 + 1 \right)^{1/2} \;\leq\; \delta,$$

which completes the proof. □

Since δ can be made arbitrarily small relative to ε, it will simplify our analysis to ignore its impact in our various inequalities involving quantities that are on the order of ε or more. We will abuse the term "infinitesimal" to refer to quantities that are functions of δ, independent of ε. Since K_i'' is infinitesimally close to K_i', it follows that $\Psi_i = (\overline{\partial} K_i'')^{\downarrow}$ is infinitesimally close to $(K_i')^{\downarrow} = K^{\downarrow} \oplus 2\varepsilon$.

The reason for restricting the extent of $K^{\downarrow}$ is to limit its projected area. The following lemma shows that, under our assumption that K's width is bounded below by ε, the projected area of each U_i is, up to constant factors, bounded by the area of K itself. Recall that "$\lesssim$" ignores constant factors.

Lemma 2. *If the width of K in any direction is at least ε, then for any $i \in I^{\pm}$, $\mathrm{area}(\Psi_i) \lesssim \mathrm{area}(K)$.*

Proof. By definition, $\mathrm{area}(\Psi_i) = \mathrm{area}(K^{\downarrow} \oplus 2\varepsilon)$. Thus, it suffices to show that $\mathrm{area}(K^{\downarrow} \oplus 2\varepsilon) \lesssim \mathrm{area}(K)$. It is well-known that for any convex body, the ratio of the distances of the body's centroid from any pair of supporting hyperplanes is at most d (see, e.g., [35]). It follows that a ball of radius $\varepsilon/(d+1)$ centered at the centroid of K is contained within K. Letting x denote this centroid, the $(d-1)$-dimensional ball of radius $\varepsilon/(d+1)$ centered at $x^{\downarrow}$ is contained within $K^{\downarrow}$. Thus, an expansion of $K^{\downarrow}$ by a factor of

$$\frac{2\varepsilon + \varepsilon/(d+1)}{\varepsilon/(d+1)} \;=\; 2d+3$$

about $x^{\downarrow}$ contains $K^{\downarrow} \oplus 2\varepsilon$. Since projection can only decrease surface areas, we have

$$\mathrm{area}(K^{\downarrow} \oplus 2\varepsilon) \;\leq\; (2d+3)^{d-1} \cdot \mathrm{area}(K^{\downarrow}) \;\lesssim\; \mathrm{area}(K^{\downarrow}) \;\leq\; \mathrm{area}(K),$$

as desired. □

2.2 Approximating Bodies and Functions

As mentioned earlier, our strategy for approximating K will be to represent its boundary as a set of $2d$ U-shaped functions as described in Section 2.1, construct approximations individually to these functions, and combine them into a single global approximation. In this section we will present further details.

Given $\varepsilon > 0$, recall the notion of a lower ε-approximation from Section 1.1. We say that a function is *polytopal*, if it is the pointwise maximum of a finite set of linear functions over Ψ. Given a U-shaped function f, we can generate such an approximation by sampling a sufficiently dense set of points on the graph of f and taking the supporting hyperplanes at each point (see Figure 5).

For any $i \in I^{\pm}$, the epigraph of a polytopal function $\widehat{f_i}$ can be interpreted as an (unbounded) convex polytope in the original space $\mathbb{R}^d$, subject to an appropriate transformation of the coordinate axes. Given a set of such functions for all $i \in I^{\pm}$, we define their *intersection* to be the intersection of these transformed epigraphs. This effectively "undoes" the support-set decomposition, combining

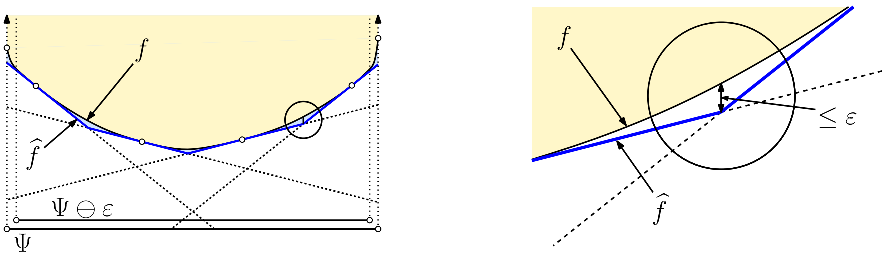

Figure 5: A polytopal lower ε-approximation to f over $\Psi \ominus \varepsilon$.

a set of $2d$ unbounded epigraphs into a single bounded convex body. The following lemma asserts that if we can approximate each function f_i to within a distance of ε of its domain's boundary, then their intersection yields an approximation to the original body.

Lemma 3. *Let K be a convex body in $\mathbb{R}^d$. For any $\varepsilon > 0$ and $i \in I^{\pm}$, let $\widehat{f_i}$ be a polytopal lower ε-approximation to f_i over $\Psi_i \ominus \varepsilon$. Then the intersection of these polytopal approximations results in a polytope that ε-approximates K.*

Proof. Let P denote the intersection of the epigraphs of the various approximations. Consider any $i \in I^{\pm}$. Since $\widehat{f_i}$ is a polytopal lower approximation, it is the maximum of a set of linear functions, each of which is pointwise not greater than f_i. This implies that the epigraph of $\widehat{f_i}$ contains the epigraph of f_i. Up to infinitesimal perturbations induced by the smoothing process, K'_i is contained within the epigraph of f_i, and hence so is K. Therefore, up to infinitesimals, K is contained within P.

To complete the proof, we show that any point s that lies at distance greater than ε from K does not lie within P. Let $q = (u; v)$ be the closest point of ∂K to s. Consider a ray emanating from the origin along the direction $s - q$, and recalling the notation from Section 2.1, let F_i denote the face of the unit hypercube $[-1, +1]^d$ that is hit by this ray, for some $i \in I^{\pm}$. (If the ray hits multiple faces, take any one of them.) Henceforth, projections will be taken orthogonal to X_i. Let $p = (a; b)$ denote the point on the segment $\overline{sq}$, such that $\|p - q\| = \varepsilon$ (see Figure 6).

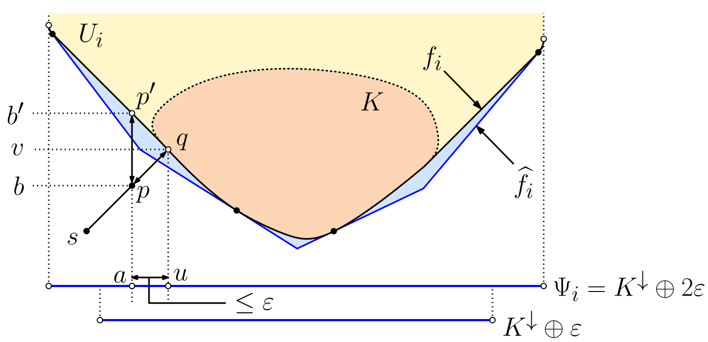

Figure 6: Proof of Lemma 3.

Since orthogonal projection cannot increase distances, $\|u - a\| \leq \varepsilon$. Since $u \in K^{\downarrow}$, we have $a \in K^{\downarrow} \oplus \varepsilon$. Recall that $\Psi_i = K^{\downarrow} \oplus 2\varepsilon$. Clearly, $a \in \Psi_i$ and (disregarding the infinitesimals of Klee's rounding) is at distance at least ε from $\partial\Psi_i$. Thus, $a \in \Psi_i \ominus \varepsilon$. Let $p' = (a; b')$ denote the point of $\overline{\partial} U_i$ that is vertically above p.

By local minimality considerations, q is the closest point to p on $\overline{\partial} K_i$. Since $\overline{\partial} K_i$ and $\overline{\partial} U_i$ differ infinitesimally, the distance from p to $\overline{\partial} U_i$ is infinitesimally close to ε. Thus, the distance from p to p' exceeds ε, that is, $b' - b > \varepsilon$. Since $\widehat{f_i}$ is a lower ε-approximation to f_i, we have

$$b \;<\; b' - \varepsilon \;=\; f_i(a) - \varepsilon \;\leq\; \widehat{f_i}(a),$$

which implies that p lies below $\widehat{f_i}$. Thus, q lies within the epigraph of $\widehat{f_i}$ and p lies outside it, and by the convexity of the epigraph, s lies outside as well. Therefore, s is external to P, which completes the proof. $\square$

Next, we show that such a polytopal approximation exists for the functions arising from our support-set construction. The proof will be presented in Section 4 below.

Theorem 2. *Given positive scalars γ, δ, and ε, where $\delta \ll \varepsilon$ and $\gamma \geq 1$, consider any (δ, γ)-slope restricted U-shaped set in $\mathbb{R}^d$ defined by a function f over a domain Ψ such that the width of Ψ in any direction is at least ε. There exists a polytopal lower ε-approximation to f over $\Psi \ominus \varepsilon$ of complexity at most*

$$\left(\frac{c_d \gamma^{d-1} \Delta_{d-1}(\Psi)}{\varepsilon} \right)^{\frac{d-1}{2}},$$

where c_d is a constant (depending only on d) and $\Delta_{d-1}(\Psi)$ is Ψ's surface diameter.

To see that this implies Theorem 1, first observe that by Lemma 2, for each $i \in I^{\pm}$, $\text{area}(\Psi_i) \lesssim \text{area}(K)$. This implies that $\Delta_{d-1}(\Psi_i) \lesssim \Delta_{d-1}(K)$. Therefore, the total number of facets in the approximating polytope for K is equal to the sum over $i \in I^{\pm}$ of the complexities of the functional approximations. By the above theorem, this yields an overall size of

$$\sum_{i \in I^{\pm}} \left(\frac{c_d \Delta_{d-1}(\Psi_i)}{\varepsilon} \right)^{\frac{d-1}{2}} \;\lesssim\; \sum_{i \in I^{\pm}} \left(\frac{c_d \Delta_{d-1}(K)}{\varepsilon} \right)^{\frac{d-1}{2}} \;=\; \left(\frac{c'_d \Delta_{d-1}(K)}{\varepsilon} \right)^{\frac{d-1}{2}},$$

where $c'_d = c_d \cdot (2d)^{2/(d-1)}$. By Lemma 3 the intersection of these local approximations yields an ε-approximation to K. This establishes Theorem 1. For the remainder of the paper, we focus on proving Theorem 2.

2.3 Dual Caps and Approximation

In this section, we will show how the problem of computing a lower ε-approximation to a U-shaped function can be reduced to computing a small set of points that hit a collection of cap-like objects. Given a U-shaped set U, recall that $\overline{\partial} U$ denotes the lower boundary of U. Since its associated function f is smooth, each $q = (u; v) \in \overline{\partial} U$ has a unique supporting hyperplane $h(q)$ as given in Eq. (1) of Section 1.1.

A key element of our construction is called a *dual cap*. For any point $q = (u; v) \in \mathbb{R}^d$ and $\gamma \in \mathbb{R}$, define $q + \gamma = (u; v + \gamma)$ to be the vertical translate of q by distance γ. For any set $H \subseteq \mathbb{R}^d$, define $H + \gamma$ analogously. Given a U-shaped set U and $q \in \overline{\partial} U$, define the *$\varepsilon$-dual cap* induced by q, denoted $D_q(U)$, to be the portion of U's lower boundary that is "visible" to $q - \varepsilon$ (see Figure 7(a)). Formally, a point $p \in \overline{\partial} U$ is on $D_q(U)$ if its supporting hyperplane $h(p)$ separates $q - \varepsilon$ from U. Letting $q = (u; v)$ and $p = (a; b)$, we define

$$D_q(U) \;=\; \{ p \in \overline{\partial} U \;:\; v - \varepsilon \leq \langle \nabla f(a), u - a \rangle + b \}$$

(see Figure 7(b)). Since we will be using the same value of ε for all dual caps henceforth, we will often omit reference to ε when discussing dual caps.

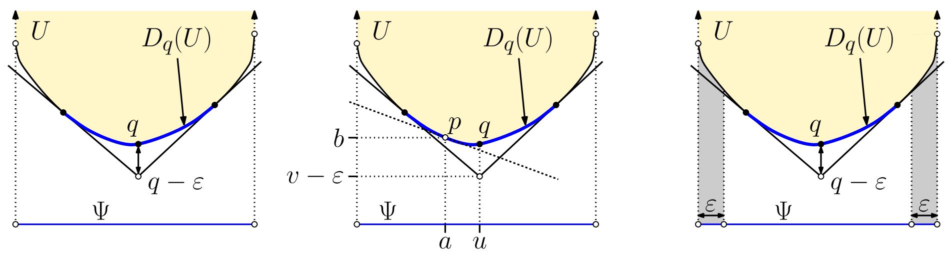

(a) (b) (c)

Figure 7: (a) ε-dual caps, (b) formal definition, and (c) useful dual caps.

It will be problematic to analyze dual caps for which the vertical projection of the defining point is very close to the boundary of Ψ. Recalling that $\Psi \ominus \varepsilon$ denotes the erosion of Ψ by ε, we restrict attention to dual caps that are induced by points whose projection is ε-far from Ψ's boundary (see Figure 7(c)).

Definition 2 (Useful Point/Dual Cap)**.** *Given a U-shaped set U and $\varepsilon > 0$, we say that a point $q = (u; v)$ is* useful *if $u \in \Psi \ominus \varepsilon$. We say that an ε-dual cap of U is* useful *if it is induced by a useful point.*

Given a U-shaped set U, $\varepsilon > 0$, and a set $\mathcal{D}$ of dual caps, we say that a discrete set $Q \subset \overline{\partial} U$ is a *hitting set* for $\mathcal{D}$ if each dual cap of $\mathcal{D}$ contains at least one point of Q. Consider any hitting set for the set of all useful dual caps. For each point $q = (u; v) \in \overline{\partial} U$ where $u \in \Psi \ominus \varepsilon$, q is useful and therefore there exists a point p of the hitting set that lies within $D_q(U)$. By the definition of ε-dual cap, the supporting hyperplane at p separates q from $q - \varepsilon$. Thus we have the following.

Lemma 4. *Let U be a U-shaped set in $\mathbb{R}^d$. For any $\varepsilon > 0$, let $Q \subset \overline{\partial} U$ be any hitting set for the set of all useful ε-dual caps of U. For each $p \in Q$, let $H^+(p)$ denote the closed upper halfspace bounded by p's supporting hyperplane. Then $P = \bigcap_{p \in Q} H^+(p)$ is a polytopal lower ε-approximation to f over $\Psi \ominus \varepsilon$.*

3 Caps and Dual Caps

With Lemmas 3 and 4, we have reduced the problem of computing an ε Hausdorff approximation of any U-shaped set to that of hitting all of its useful ε-dual caps. In our analysis we will use the fact that these sets are (δ, γ)-slope restricted, for arbitrarily small δ and constant γ. In the remainder of the paper, we will focus exclusively on this latter problem. Our approach involves the application of duality. In this section, we will explore the structure of U-shaped sets and dual caps under duality.

3.1 Dual Transforms

Our results are based on two commonly used dual transforms in geometry. These transforms map points to hyperplanes and vice versa, while preserving point-hyperplane incidences and convexity.

The first transform is the well known polar transform, and the second is an adaptation of the Legendre transform to our setting.

Given a convex body $K \subseteq \mathbb{R}^d$ that contains the origin in its interior, its *polar*, denoted here K°, is defined to be $\{u \in \mathbb{R}^d : \langle u, v\rangle \le 1, \forall v \in K\}$ (see Figure 8). Given $\alpha > 0$, let αK° denote a scaling of K° by a factor of α. Clearly, $\alpha K^\circ = \{u \in \mathbb{R}^d : \langle u, v\rangle \le \alpha, \forall v \in K\}$.

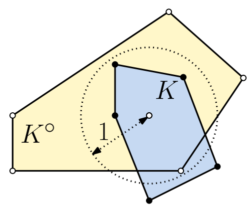

Figure 8: The polar body of K.

We will make use of an important result from the theory of convex sets, which states that, given a convex body K, the product $\mathrm{vol}(K) \cdot \mathrm{vol}(K^\circ)$, which is called K's *Mahler volume*, is bounded below by a constant depending only on the dimension (see, e.g., [20, 39, 48]). Henceforth, we denote this constant by μ_d.

Lemma 5 (Kuperberg [39])**.** *Given a convex body $K \subseteq \mathbb{R}^d$ whose interior contains the origin,*

$$\mathrm{vol}(K) \cdot \mathrm{vol}(K^\circ) \;\ge\; \left(\frac{\pi}{2e}\right)^d \frac{(d+1)^{d+1}}{(d!)^2}.$$

Next, let us define the *projective dual transform* [26]. By expressing any $q \in \mathbb{R}^d$ as $(u; v) \in \mathbb{R}^{d-1} \times \mathbb{R}$, its *projective dual* is the hyperplane $q^* : y = \langle u, x\rangle - v$ (see Figure 9(a)). The dual also acts on (nonvertical) hyperplanes, mapping the hyperplane $p^* : y = \langle a, x\rangle - b$ to the dual point $p = (a; b)$. It is easy to see that this transform preserves point-hyperplane incidences and it reverses vertical distances, in the sense that point q is at vertical distance γ above a hyperplane p^* if and only if p is at vertical distance γ above q^*.

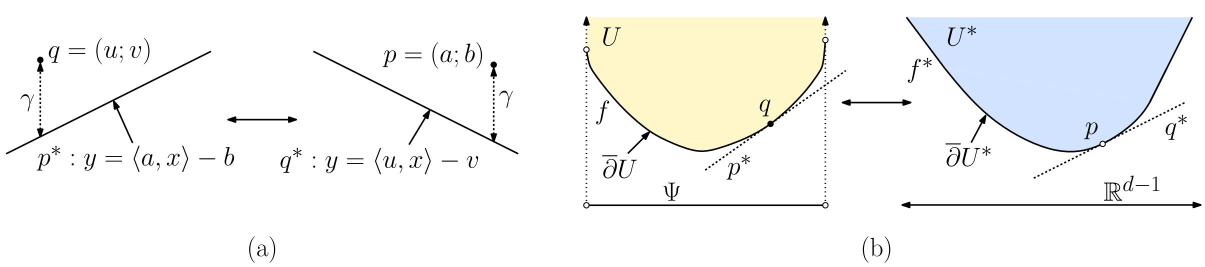

Figure 9: The projective dual.

This is closely related to the Legendre transform. Given a smooth real-valued function f over an open domain in $\mathbb{R}^{d-1}$ where ∇f is one-to-one, the *Legendre transform* of f is the function $f^* : \mathbb{R}^{d-1} \to \mathbb{R}$ defined

$$f^*(a) \;=\; \langle a, (\nabla f)^{-1}(a)\rangle - f((\nabla f)^{-1}(a)) \tag{2}$$

(see, e.g., [49]). To see this relationship, consider a U-shaped set U that is defined as the epigraph of a smooth, strictly convex function f over a domain Ψ. Given $q = (u; v)$ on U's lower boundary, we have $u \in \Psi$ and $v = f(u)$. By Eq. (1), the unique supporting hyperplane for U at q can be expressed as $y = \langle a, x \rangle - b$, where $a = \nabla f(u)$ and $b = \langle a, u \rangle - f(u)$. Recall from the definition of U-shaped sets that the gradient defines a homeomorphism between Ψ and $\mathbb{R}^{d-1}$. Let f^* be the real-valued function over $\mathbb{R}^{d-1}$, defined by $f^*(a) = b$, and let U^* denote the epigraph of f^* (see Figure 9(b)). Clearly, $(\nabla f)^{-1}$ is well defined over $\mathbb{R}^{d-1}$. Substituting $u = (\nabla f)^{-1}(a)$ into $f^*(a) = \langle a, u \rangle - f(u)$ yields Eq. (2), which establishes the equivalence. Letting $p = (a; b)$, we refer to p as q's *dual counterpart* and vice versa. The following lemma summarizes these observations.

Lemma 6. *Let U be a U-shaped set defined as the epigraph of function f over a domain Ψ. The above construction yields a smooth, strictly convex function $f^* : \mathbb{R}^{d-1} \to \mathbb{R}$ such that for each $q = (u; v) \in \overline{\partial} U$, there is a unique point $p = (a; b) \in \overline{\partial} U^*$ (its dual counterpart), such that*

(i) p^* (resp., q^*) is the supporting hyperplane for U (resp., U^*) at q (resp., p),
(ii) $\nabla f(u) = a$ and $\nabla f^*(a) = u$.

Proof. By convexity, for all points $q' = (u'; v') \in \overline{\partial} U$, q lies on or above the supporting hyperplane through q'. Letting $y = \langle a', x \rangle - b'$ denote this hyperplane, we have $v \geq \langle a', u \rangle - b'$, or equivalently $b' \geq \langle u, a' \rangle - v$. Therefore, every point $p' = (a'; b') \in \overline{\partial} U^*$ lies on or above q^*. Since q^* passes through p, it follows that q^* is the supporting hyperplane for U^* at p and $\nabla f^*(a) = u$. The reverse relationships hold symmetrically. □

Note that U^* is an unbounded convex set, and hence it is not U-shaped. We say that a point $p \in \overline{\partial} U^*$ and the cap it induces is *useful* if it is the dual counterpart of a useful point $q \in \overline{\partial} U$.

3.2 Dual Caps and Caps in the Dual

In Section 2.2, we demonstrated the relevance of dual caps of U-shaped sets to approximation. In this section, we explore how dual caps are manifest in the dual setting. Consider a U-shaped set U and let q be any point on $\overline{\partial} U$. Let U^* be the corresponding dual set, and let p denote q's dual counterpart on $\overline{\partial} U^*$. By Lemma 6(i), q^* is the supporting hyperplane for U^* at p. Let H denote the closed lower halfspace bounded by q^*, and let $H + \varepsilon$ denote the upwards vertical translate of H by distance ε. Define the *ε-cap* induced by p, denoted $C_p(U^*)$, to be

$$C_p(U^*) \;=\; \overline{\partial} U^* \cap (H + \varepsilon).$$

(See Figure 10. In contrast to standard usage, our caps consist only of boundary points.) The quantity ε is called the *width* of the cap. As with dual caps, unless otherwise specified, caps are of width ε.

It is easy to see that the hyperplane bounding $H + \varepsilon$ is the projective dual of $q - \varepsilon$, and so it is natural to expect that the dual cap $D_q(U)$ and the cap $C_p(U^*)$ are related. The next lemma establishes this relationship.

Lemma 7. *Given a positive real ε, a U-shaped set U, and a point $q \in \overline{\partial} U$, let $p \in \overline{\partial} U^*$ be q's dual counterpart. Then the points of the ε-cap $C_p(U^*)$ are the dual counterparts of the points of the ε-dual cap $D_q(U)$, and vice versa.*

Proof. Let $q = (u; v)$ and $p = (a; b)$. By Lemma 6(i), the supporting line for U through q is $y = \langle a, x \rangle - b$, and the supporting line for U^* through p is $y = \langle u, x \rangle - v$. The halfspace $H + \varepsilon$ in the cap definition is defined by the inequality $y \leq \langle u, x \rangle - v + \varepsilon$.

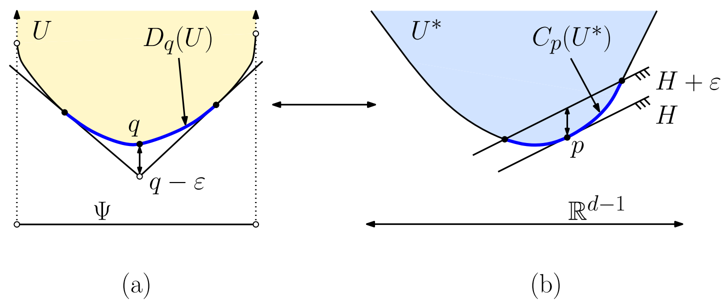

Figure 10: (a) An ε-dual cap in U and (b) its corresponding ε-cap in U^*.

Consider any pair of points $q' \in \overline{\partial} U$ and $p' \in \overline{\partial} U^*$, where p' is the dual counterpart of q'. Let $q' = (u'; v')$ and $p' = (a'; b')$. By the definition of dual cap, $q' \in D_q(U)$ if and only if the supporting hyperplane through q' passes on or above $q - \varepsilon$, or equivalently

$$v - \varepsilon \;\leq\; \langle a', u \rangle - b' \qquad \Longleftrightarrow \qquad b' \;\leq\; \langle u, a' \rangle - v + \varepsilon.$$

From the dual perspective, this is equivalent to saying that $p' \in H + \varepsilon$. Therefore, $q' \in D_q(U)$ if and only if $p' \in C_p(U^*)$, as desired. $\square$

Next, we explore how caps and dual caps are related through the polarity. We first introduce two related "flat" structures. Given $q \in \overline{\partial} U$, let p be its dual counterpart. By Lemma 6(i), the supporting hyperplane for U at q is p^*, and the supporting hyperplane for U^* at p is q^*. Given the dual cap $D_q(U)$, define its *base*, denoted $\Delta_q(U)$, to be $p^* \cap \mathrm{conv}(U \cup (q - \varepsilon))$ (see Figure 11(a)). Define the *base* of the cap $C_p(U^*)$, denoted $\Gamma_p(U^*)$, to be $U^* \cap (q^* + \varepsilon)$ (see Figure 11(b)).

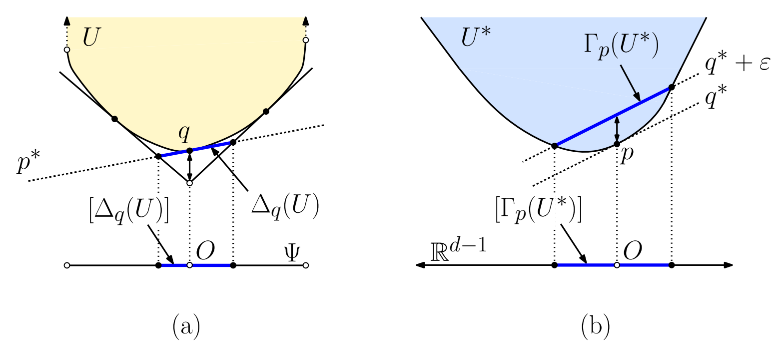

Figure 11: (a) The base of q's dual cap and its projection and (b) the base of p's cap and its projection. (Rather than showing the translation to the origin, we have shown the origin O at the projected center point.)

The relationship between bases is most clearly established through their projections. Given $\Delta_q(U)$ and $\Gamma_p(U^*)$ defined above, define $[\Delta_q(U)] = \Delta_q(U)^{\downarrow} - q^{\downarrow}$ and $[\Gamma_p(U^*)] = \Gamma_p(U^*)^{\downarrow} - p^{\downarrow}$. Intuitively, these take the vertical projections of the bases and then translate them so the origin coincides with the vertical projection of the defining point (see Figure 11). The next lemma shows that these two bodies are polars of each other, subject to a scale factor of ε.

Lemma 8. *Given a positive real ε, a U-shaped set U, and a point $q \in \overline{\partial}U$, let $p \in \overline{\partial}U^*$ be q's dual counterpart. Then $[\Delta_q(U)] = \varepsilon[\Gamma_p(U^*)]^\circ$.*

Proof. To simplify notation, let $\Delta_q = \Delta_q(U)$ and $\Gamma_p = \Gamma_p(U^*)$. Let $q = (u; v)$ and $p = (a; b)$, so that $q^\downarrow = u$ and $p^\downarrow = a$ (see Figure 12). Thus, $[\Delta_q(U)] = \Delta_q^\downarrow - u$ and $[\Gamma_p(U^*)] = \Gamma_p^\downarrow - a$. It suffices to show that $\Delta_q^\downarrow - u = \varepsilon(\Gamma_p^\downarrow - a)^\circ$.

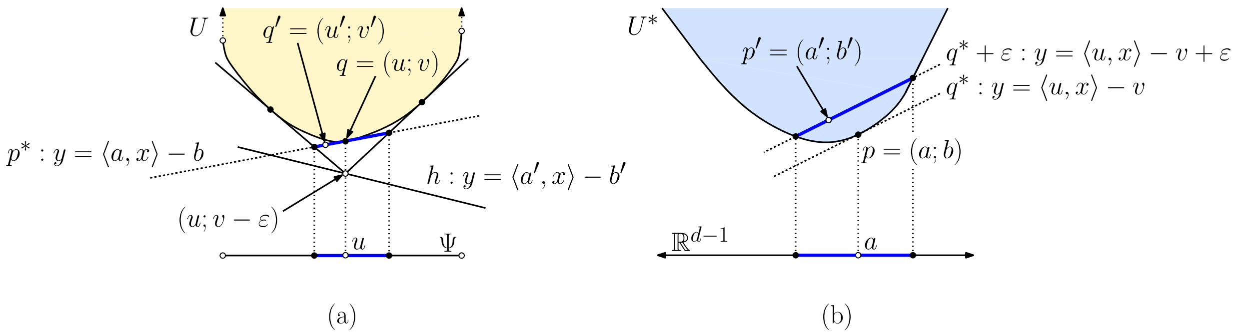

Figure 12: Proof of Lemma 8.

Recall the definition of the polar from Section 3.1. Because the polar transform is an involution, it follows that $w \in K$ if and only if $\langle w, z\rangle \leq 1$, for all $z \in K^\circ$. Thus, to show that $\Delta_q^\downarrow - u = \varepsilon(\Gamma_p^\downarrow - a)^\circ$, it suffices to show that $q' \in \Delta_q$ if and only if $\langle q'^\downarrow - u, p'^\downarrow - a\rangle \leq \varepsilon$, for all $p' \in \Gamma_p$. Letting $q' = (u'; v')$ and $p' = (a'; b')$, this is equivalent to

$$q' \in \Delta_q \iff \langle u' - u, a' - a\rangle \;\leq\; \varepsilon, \quad \forall p' \in \Gamma_p. \tag{3}$$

The remainder of the proof is devoted to establishing this assertion.

Define H_q to be the set of all hyperplanes h that pass through $q - \varepsilon$ such that U lies on or above h. Let H_q^* denote the duals of the hyperplanes of H_q. We assert that $H_q^* = \Gamma_p$. To see this, for each $h \in H_q$, where $h : y = \langle a', x\rangle - b'$, let $p' = h^* = (a'; b')$ denote its dual point. First, observe that since every hyperplane $h \in H_q$ passes through $q - \varepsilon$, by incidence preservation, its dual p' lies on the dual hyperplane $(q - \varepsilon)^* = q^* + \varepsilon$. The statement that U lies on or above h is equivalent to saying that for all points $s = (c; d) \in \overline{\partial}U$, s lies on or above h. That is,

$$d \;\geq\; \langle a', c\rangle - b', \qquad \text{or equivalently} \qquad b' \;\geq\; \langle c, a'\rangle - d.$$

The latter is equivalent to saying that p' lies on or above the hyperplane $s^* : y = \langle c, x\rangle - d$. Since s ranges over all the points in $\overline{\partial}U$, Lemma 6(i) implies that s^* ranges over all the supporting hyperplanes of $\overline{\partial}U^*$. Thus, p' lies on or above all the supporting hyperplanes of $\overline{\partial}U^*$, which is equivalent to saying that $p' \in U^*$. In summary, we have shown that $h \in H_q$ if and only if its dual point p' is in $U^* \cap (q^* + \varepsilon) = \Gamma_p$. Hence, $H_q^* = \Gamma_p$, as desired.

A point $q' = (u'; v')$ lies on Δ_q if and only if (i) it lies on q's supporting hyperplane p^*, and (ii) it lies in $\mathrm{conv}(U \cup \{q - \varepsilon\})$. Given that $p^* : y = \langle a, x\rangle - b$, condition (i) is equivalent to

$$v' \;=\; \langle a, u'\rangle - b. \tag{4}$$

Since p^* is the supporting hyperplane at q, we have $v = \langle a, u\rangle - b$ or equivalently $b = \langle a, u\rangle - v$. Combining this with Eq. (4), gives

$$v' \;=\; \langle a, u'\rangle - (\langle a, u\rangle - v) \;=\; \langle a, u' - u\rangle + v. \tag{5}$$

Next, let's consider condition (ii). Clearly, $\mathrm{conv}(U \cup \{q - \varepsilon\})$ is equal to the intersection of all the upper halfspaces of the hyperplanes in H_q, and so condition (ii) is equivalent to saying that q' lies on or above all of the hyperplanes $h \in H_q$. Expressing h as $y = \langle a', x\rangle - b'$, this is equivalent to

$$v' \;\geq\; \langle a', u'\rangle - b', \quad \forall h \in H_q.$$

Letting $p' = h^* = (a'; b')$ and by the earlier observation that $H_q^* = \Gamma_p$, we can rewrite this as

$$v' \;\geq\; \langle a', u'\rangle - b', \quad \forall p' \in \Gamma_p. \tag{6}$$

Since p' lies on $q^* + \varepsilon$, $b' = \langle u, a'\rangle - v + \varepsilon$. Combining this with Eq. (6), yields

$$v' \;\geq\; \langle a', u'\rangle - (\langle u, a'\rangle - v + \varepsilon) \;=\; \langle a', u' - u\rangle + v - \varepsilon, \quad \forall p' \in \Gamma_p. \tag{7}$$

By combining Eqs. (5) and (7), both v and v' are eliminated, yielding

$$\varepsilon \;\geq\; \langle a', u' - u\rangle - \langle a, u' - u\rangle \;=\; \langle u' - u, a' - a\rangle, \quad \forall p' \in \Gamma_p.$$

Thus, we have

$$q' \in \Delta_q \implies \langle u' - u, a' - a\rangle \;\leq\; \varepsilon, \quad \forall p' \in \Gamma_p.$$

Conversely, suppose that $\langle u' - u, a' - a\rangle \leq \varepsilon$, for all $p' = (a'; b') \in \Gamma_p$. Any point $u' \in \Delta_q^{\downarrow}$ arose from the projection of some point $q' = (u'; v')$ on the supporting hyperplane p^*, and hence $v' = \langle a, u'\rangle - b$. Such a point satisfies condition (i) above for being in Δ_q. As observed earlier, $v = \langle a, u\rangle - b$, and subtracting these yields $v' - v = \langle a, u' - u\rangle$. Therefore,

$$\begin{aligned} \varepsilon \;\geq\; \langle u' - u, a' - a\rangle \;&=\; \langle a' - a, u' - u\rangle \;=\; \langle a', u' - u\rangle - \langle a, u' - u\rangle \\ &=\; \langle a', u' - u\rangle - (v' - v), \quad \forall p' \in \Gamma_p, \end{aligned}$$

or equivalently

$$v' \;\geq\; \langle a', u' - u\rangle + v - \varepsilon, \quad \forall p' \in \Gamma_p.$$

This implies that q' also satisfies condition (ii), and hence $q' \in \Delta_q$ as desired. This establishes Eq. (3) and completes the proof. □

3.3 Additional Properties

Before describing our approximation constructions, in this section we present a few additional properties of caps and dual caps, which will be applied later in Section 4. The first two lemmas involve the sizes of useful caps and dual caps. In order to deal with dual caps whose whose defining point is very close to Ψ's boundary, we introduced the notion of usefulness. Recall that a point $q = (u; v) \in \overline{\partial} U$ is *useful* if $u \in \Psi$ and is at distance at least ε from $\partial\Psi$. Also, recall that a U-shaped set is (δ, γ)-*slope restricted* if for all $x \in \Psi \ominus \delta$, $\|\nabla f(x)\| \leq \gamma$. The following lemma shows that if a point $q \in \overline{\partial} U$ is useful, then the projected dual base $[\Delta_q(U)]$ contains a ball whose radius is, up to constant factors, at least ε.

Lemma 9. *Given positive reals δ, ε, and γ where $\delta \ll \varepsilon$, let U be a (δ, γ)-slope restricted U-shaped set over a domain Ψ, and let q be a useful point on $\overline{\partial} U$. Then $[\Delta_q(U)]$ contains a ball of radius $\varepsilon/2(1 + \gamma)$ centered at the origin.*

Proof. To simplify notation, we will drop reference to U when referring to dual caps and their bases. Let $q = (u; v)$, and let f denote the function that defines the lower boundary of U. Consider any point $s = (u_s; v_s)$ on $\partial\Delta_q$ (see Figure 13(a)). Since s lies on Δ_q, a ray shot from $q - \varepsilon$ through s hits $\overline{\partial}U$ at some point $q' = (u'; v')$, where $u' \in \Psi$.

Let $\Psi^- = \Psi \ominus \delta$. Clearly, $u \in \Psi^-$. If u' is also in Ψ^-, let $q'' = q'$. Otherwise, let $u'' = \partial\Psi^- \cap \overline{uu'}$ be the point where the segment $\overline{uu'}$ leaves Ψ^-, and let $q'' = (u''; v'')$ where $v'' = f(u'')$.

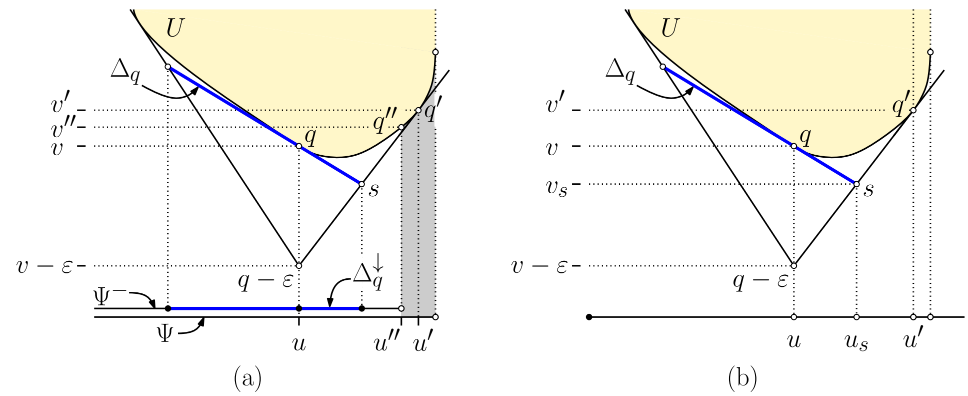

Figure 13: Proof of Lemma 9.

We first bound $|v'' - v|$. We can compute the vertical variation between q'' and q by integrating the directional derivative of f from u to u''.

$$v'' - v \;=\; \int_u^{u''} \langle \nabla f(x), dx \rangle.$$

By slope restriction, for all $x \in \Psi^-$, $\|\nabla f(x)\| \leq \gamma$. Therefore, by the Cauchy-Schwarz inequality we have

$$|v'' - v| \;\leq\; \left(\sup_{x \in \Psi^-} \|\nabla f(x)\| \right) \cdot \|u'' - u\| \;\leq\; \gamma \cdot \|u'' - u\|.$$

Next, we bound $\langle \nabla f(u'), u_s - u \rangle$. We consider two cases. First, if $u' \in \Psi^-$, then $\|\nabla f(u')\| \leq \gamma$, and again by Cauchy-Schwarz we have $\langle \nabla f(u'), u_s - u \rangle \leq \gamma \cdot \|u_s - u\|$.

Otherwise $u'' \in \partial\Psi^-$ and by hypothesis, $\|u'' - u\| \geq \varepsilon - \delta$. The supporting hyperplane through q' passes through $q - \varepsilon = (u; v - \varepsilon)$, and thus its defining equation is $y = \langle \nabla f(u'), x - u \rangle + v - \varepsilon$. Since s lies on this hyperplane, we have

$$v_s \;=\; \langle \nabla f(u'), u_s - u \rangle + v - \varepsilon. \tag{8}$$

Also, since u'' lies between u and u', q'' lies on or above this hyperplane, and thus,

$$v'' \;\geq\; \langle \nabla f(u'), u'' - u \rangle + v - \varepsilon.$$

Combining this with our bound on $|v'' - v|$ and the fact that $\|u'' - u\| \geq \varepsilon - \delta \geq \varepsilon/2$, we have

$$\begin{aligned} \langle \nabla f(u'), u'' - u \rangle \;&\leq\; (v'' - v) + \varepsilon \;\leq\; |v'' - v| + \varepsilon \;\leq\; \gamma \cdot \|u'' - u\| + \varepsilon \\ &\leq\; \gamma \cdot \|u'' - u\| + 2 \cdot \|u'' - u\| \;=\; \big(2 + \gamma\big) \cdot \|u'' - u\|. \end{aligned}$$

Since s lies along the line between $q - \varepsilon$ and q', we can scale both sides of the above inequality by $\|u_s - u\|/\|u'' - u\|$ to obtain

$$\langle \nabla f(u'), u_s - u \rangle \;\leq\; \big(2 + \gamma\big) \cdot \|u_s - u\|. \tag{9}$$

Observe that this bound applies in both cases.

The supporting hyperplane at q has the defining equation $y = \langle \nabla f(u), x - u \rangle + v$, and since it passes through s, we have

$$v_s \;=\; \langle \nabla f(u), u_s - u \rangle + v.$$

By combining this with Eq. (8), we obtain

$$\varepsilon \;=\; \langle \nabla f(u'), u_s - u \rangle - \langle \nabla f(u), u_s - u \rangle.$$

Since $u \in \Psi^-$, $\|\nabla f(u)\| \leq \gamma$, and together with the Cauchy-Schwarz inequality, we have $\langle \nabla f(u), u_s - u \rangle \leq \gamma \cdot \|u_s - u\|$. The triangle inequality and Eq. (9) yield

$$\varepsilon \;\leq\; (2 + \gamma) \cdot \|u_s - u\| + \gamma \cdot \|u_s - u\| \;\leq\; 2(1 + \gamma) \cdot \|u_s - u\|.$$

We conclude that

$$\|u_s - u\| \;\geq\; \frac{\varepsilon}{2(1 + \gamma)}.$$

Since this applies to every point $s \in \partial\Delta_q(U)$, we conclude that $\Delta_q(U)^{\downarrow}$ contains a ball of radius $\varepsilon/2(1 + \gamma)$ centered at $q^{\downarrow}$, as desired. $\square$

By the polar relationship between caps and dual cap projections, we obtain a complementary result for the projections of useful ε-caps. Recall that a cap is useful if its defining point $p \in \overline{\partial} U^*$ is the dual counterpart of a useful point $q \in \overline{\partial} U$.

Lemma 10. *Given positive reals δ, ε, and γ where $\delta \ll \varepsilon$, let U be a (δ, γ)-slope restricted U-shaped set over a domain Ψ, and let p be a useful point on $\overline{\partial} U^*$. Then $[\Gamma_p(U^*)]$ is contained within a ball of radius $2(1 + \gamma)$ centered at the origin.*

Proof. Consider any point $p \in \overline{\partial} U^*$ that induces a useful cap. By definition of usefulness, p is the dual counterpart of some useful point $q = (u; v) \in \overline{\partial} U$. Lemma 9 implies that $[\Delta_q(U)]$ contains a ball of radius $\varepsilon/2(1 + \gamma)$ centered at the origin, implying that $[\Delta_q(U)]/\varepsilon$ contains a ball of radius $1/2(1 + \gamma)$ centered at the origin. By Lemma 8, $[\Delta_q(U)]/\varepsilon = [\Gamma_p(U^*)]^{\circ}$. Due to the reciprocal nature of the polar transformation, it follows that $[\Gamma_p(U^*)]$ is contained within a ball of radius $2(1 + \gamma)$ centered at the origin. $\square$

The next result of this section shows that for any dual cap of a U-shaped set U, there is an ellipsoid of large volume that lies above the dual cap's base and close to U's lower boundary. More precisely, consider any point $q \in \overline{\partial} U$ and its associated dual cap. Let H^- denote the lower halfspace bounded by the supporting hyperplane at q (see Figure 14). The set $U \cap (H^- + \varepsilon)$ is convex, and hence by John's Theorem [37], it contains an ellipsoid whose volume is within a constant factor of $\mathrm{vol}(U \cap (H^- + \varepsilon))$. The following lemma strengthens this by showing that there is such an ellipsoid whose vertical projection is contained within the vertical projection of the base of q's dual cap. Recall from Section 1.1 that "$\gtrsim$" ignores constant factors.

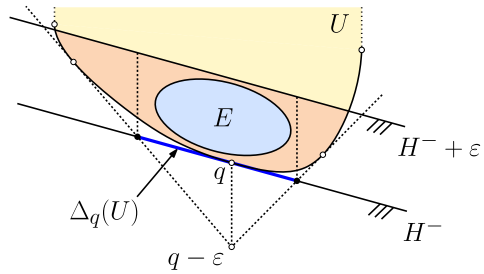

Figure 14: Lemma 11.

Lemma 11. *Given a positive real* ε*, a U-shaped set* U*, and a point* $q \in \overline{\partial} U$*, let* H^- *be the lower halfspace bounded by the supporting hyperplane to* U *at* q*. There exists an ellipsoid* $E \subseteq U \cap (H^- + \varepsilon)$ *such that* $\mathrm{vol}(E) \gtrsim \mathrm{vol}(U \cap (H^- + \varepsilon)) \gtrsim \varepsilon \cdot \mathrm{area}(\Delta_q(U)^{\downarrow})$ *and* $E^{\downarrow} \subseteq \Delta_q(U)^{\downarrow}$*.*

Proof. To simplify notation, let $\Delta_q = \Delta_q(U)$. The proof is based on the following construction. Let $T = \mathrm{conv}(\Delta_q \cap \{q - \varepsilon\})$ or equivalently $\mathrm{conv}(U \cup \{q - \varepsilon\}) \cap H^-$, and let T' be a scaling of T by a factor of 2 about $q - \varepsilon$ (see Figure 15(a)). We will prove the following facts about this construction.

(*i*) $T + \varepsilon \ \subseteq \ U \cap (H^- + \varepsilon) \ \subseteq \ T'$.

(*ii*) $\mathrm{vol}(U \cap (H^- + \varepsilon)) \ \gtrsim \ \varepsilon \cdot \mathrm{area}(\Delta_q^{\downarrow})$.

(*iii*) $T^{\downarrow} = \Delta_q^{\downarrow}$.

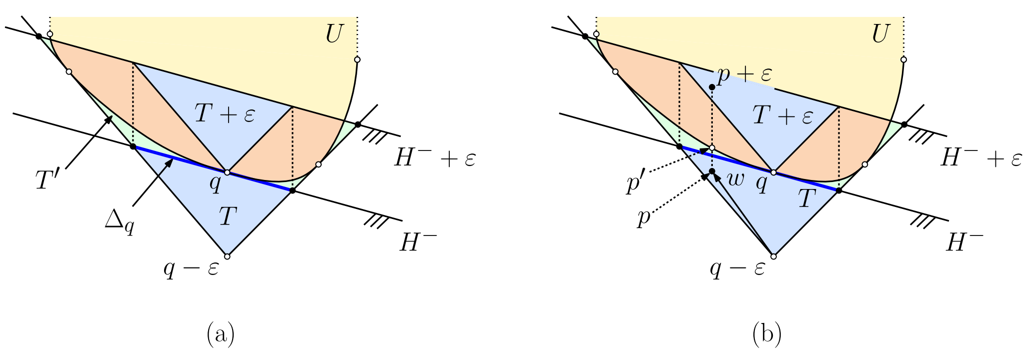

Figure 15: Proof of Lemma 11.

We begin by showing the first containment of (i). Since $T \subseteq H^-$, we have $T + \varepsilon \subseteq H^- + \varepsilon$. Thus, it suffices to show that $T + \varepsilon \subseteq U$. We will show that for all $p \in T$, $p + \varepsilon \in U$ (see Figure 15(b)). Consider the first point p' where a vertical ray shot up from p intersects U, that is, $p' = \min_{\beta} \{p + \beta\varepsilon \in U\}$. (Here, we are treating ε as if it were a vertical vector of length ε.) Letting β denote this minimum value, we will show that $\beta < 1$, implying that $p + \varepsilon \in U$.

Clearly, $p' \in \overline{\partial} U$. Since q is also on $\overline{\partial} U$, it follows from strict convexity that if we shoot a ray from q through p', every point beyond p' on this ray is external to U. That is, for all $\gamma > 1$, $q + \gamma(p' - q) \notin U$. Letting $w = p - (q - \varepsilon)$, for all $\gamma > 1$, U does not contain the point

$$q + \gamma((p + \beta\varepsilon) - q) \ = \ q + \gamma(p - (q - \varepsilon) + (\beta - 1)\varepsilon) \ = \ (q - \varepsilon) + \gamma w + ((1 - \gamma) + \gamma\beta)\varepsilon. \quad (10)$$

By definition of T, the extension of every ray from $q - \varepsilon$ through p intersects U. Thus, there exists a real $\alpha > 1$ such that $(q - \varepsilon) + \alpha w \in U$. Since $\alpha > 1$, it follows from Eq. (10) that U does not contain the point $(q - \varepsilon) + \alpha w + ((1 - \alpha) + \alpha\beta)\varepsilon$. But, as $(q - \varepsilon) + \alpha w \in U$, we conclude that $(1 - \alpha) + \alpha\beta < 0$, which implies that $\beta < (\alpha - 1)/\alpha < 1$, as desired.

To prove the second containment of (i), observe that T' is the intersection of $H^- + \varepsilon$ and the generalized cone defined as the intersection of the upper halfspaces of the supporting hyperplanes of U that pass through the apex $q - \varepsilon$. Clearly, this contains $U \cap (H^- + \varepsilon)$.

To prove (ii), observe that by the first containment of (i), we have $\mathrm{vol}(U \cap (H^- + \varepsilon)) \geq \mathrm{vol}(T)$ and by basic geometry, $\mathrm{vol}(T)$ is proportional to the product of the vertical distance between the vertex $q - \varepsilon$ and the triangle's base, which is ε, and the area of the vertical projection of its base $\Delta_q^{\downarrow}$.

Finally, (iii) follows from the fact that $T = \mathrm{conv}(\Delta_q \cup \{q - \varepsilon\})$. Since $q \in \Delta_q$, $(q - \varepsilon)$ lies vertically below the base, which implies that $T^{\downarrow} = \Delta_q^{\downarrow}$.

Returning to the proof, let E be the maximum volume ellipsoid contained in $T + \varepsilon$. By (i), $E \subseteq T + \varepsilon \subseteq U \cap (H^- + \varepsilon)$. By John's Theorem [37], $\mathrm{vol}(E) \geq \mathrm{vol}(T)/d^d$. By (i), $U \cap (H^- + \varepsilon) \subseteq T'$. Since T' is a factor-2 scaling of T, $\mathrm{vol}(T) = \mathrm{vol}(T')/2^d$. Therefore, by setting $c = 1/(2d)^d$, we have

$$\mathrm{vol}(E) \;\geq\; \frac{\mathrm{vol}(T)}{d^d} \;=\; \frac{\mathrm{vol}(T')}{(2d)^d} \;\geq\; \frac{\mathrm{vol}(U \cap (H^- + \varepsilon))}{(2d)^d} \;=\; c \cdot \mathrm{vol}(U \cap (H^- + \varepsilon)).$$

Also, combining the fact that $E \subseteq T + \varepsilon$ and (iii) implies that $E^{\downarrow} \subseteq \Delta_q^{\downarrow}$. □

4 Hitting Sets and Approximation

Armed with the tools developed in the previous section, in this section we will present the constructions to establish Theorem 2. Let us start with a high level description of how our construction works. Recall that we are given a U-shaped set U and $\varepsilon > 0$. U is the epigraph of a smooth, strictly convex function f defined over a bounded, open, convex domain Ψ. We assume that f is (δ, γ)-slope constrained, where $\delta > 0$ can be made arbitrarily small and γ is a constant. By Lemma 4, it suffices to show the existence of a set of points on $\overline{\partial} U$ that hit all the useful ε-dual caps of U. Recall that a dual cap is *useful* if it is induced by a point $q = (u; v) \in \overline{\partial} U$ such that $u \in \Psi$ and is at distance at least ε from $\partial\Psi$.

How do we select a small set of points to hit all these dual caps for each U? We will employ a classical structure from the study of convex bodies, called Macbeath regions (presented in Section 4.2 below). To generate the hitting set, we construct Macbeath regions along the lower boundary of U, select a small number of points from each region, and project each point vertically onto U's lower boundary.

The Macbeath-based approach will be efficient for hitting dual caps that are "large," meaning that the projected area of their base is sufficiently large. More precisely, these are the dual caps defined by points $q \in \overline{\partial} U$ such that $\mathrm{area}([\Delta_q(U)])$ is larger than some given threshold parameter. This will be presented in Section 4.3 below.

To handle the remaining "small" dual caps, we will exploit the polar connection between caps and dual caps in the primal and dual settings. Let p denote q's dual counterpart in U^*. Lemma 8 states that $[\Delta_q(U)]$ is related through the polar to the projected base of p's cap $[\Gamma_p(U^*)]$. Now, the Mahler volume bound of Lemma 5 can be applied to show that if q induces a small dual cap, then

p induces a relatively large cap (in Section 4.1 below). This means that we can efficiently apply the Macbeath-based sampling in the dual setting, and transfer the sampled points back to the primal setting. This construction will be presented in Section 4.4 below.

4.1 Hitting the Large and the Small

We begin by showing that small dual caps correspond to large caps in the dual. Let U be as defined above over the domain Ψ. Given any useful point $q \in \overline{\partial} U$, recall that $D_q(U)$ denotes q's ε-dual cap, $\Delta_q(U)$ denotes its base, and $[\Delta_q(U)]$ denotes the projection of the base, centered so the origin coincides with $q^{\downarrow}$. Define $t_0 = \sqrt{\mathrm{area}(\Psi)} \cdot \varepsilon^{(d-1)/2}$. We say that $D_q(U)$ is *large* if $\mathrm{area}([\Delta_q(U)]) \geq t_0$, and otherwise it is *small*.

By Lemma 6, each q has a unique dual counterpart $p \in \overline{\partial} U^*$. We say that the associated ε-cap, $C_p(U^*)$ is *large* if $\mathrm{area}([\Gamma_p(U^*)]) \geq \mu_{d-1} \varepsilon^{d-1}/t_0$. (Recall that μ_d is the dimension-dependent constant from Lemma 5.) Our next lemma establishes this size relationship between dual caps and caps.

Lemma 12. *Given a positive real ε, a U-shaped set U in $\mathbb{R}^d$, and a point $q \in \overline{\partial} U$ that induces a small ε-dual cap, then q's dual counterpart p induces a large ε-cap in U^*.*

Proof. Let Δ_q denote the base of q's dual cap in U, and let Γ_p denote the base of p's cap in U^*. By Lemma 8, $[\Delta_q] = \varepsilon [\Gamma_p]^\circ$. Noting that these are $(d-1)$-dimensional sets, scaling by ε alters the area by a factor of ε^{d-1}. Thus,

$$\mathrm{area}([\Delta_q]) \;=\; \mathrm{area}(\varepsilon \cdot [\Gamma_p]^\circ) \;=\; \varepsilon^{d-1} \cdot \mathrm{area}([\Gamma_p]^\circ).$$

Since D_q is small, $\mathrm{area}([\Delta_q]) \leq t_0$, and hence $\mathrm{area}([\Gamma_p]^\circ) \leq t_0/\varepsilon^{d-1}$. By Lemma 5,

$$\mathrm{area}([\Gamma_p]) \cdot \mathrm{area}([\Gamma_p]^\circ) \;\geq\; \mu_{d-1},$$

and therefore $\mathrm{area}([\Gamma_p]) \geq \mu_{d-1} \varepsilon^{d-1}/t_0$, as desired. ☐

We will describe the process for building hitting sets in Sections 4.3 and 4.4 below, but assuming these results for now, let us see how we use them to prove Theorem 2.

In Lemma 14 below, we will show that for any $t > 0$, there exists a set of $O(\mathrm{area}(\Psi)/t)$ points that hits all dual caps $D_q(U)$ where $\mathrm{area}([\Delta_q(U)]) \geq t$. By setting $t = t_0$ and observing that $\Delta_{d-1}(\Psi) \approx \mathrm{area}(\Psi)^{1/(d-1)}$, it follows that the resulting set $Q \subset \overline{\partial} U$ hits all the large dual caps, where

$$|Q| \;=\; O\left(\frac{\mathrm{area}(\Psi)}{t_0}\right) \;=\; O\left(\frac{\sqrt{\mathrm{area}(\Psi)}}{\varepsilon^{(d-1)/2}}\right) \;=\; O\left(\frac{\Delta_{d-1}(\Psi)}{\varepsilon}\right)^{\frac{d-1}{2}}.$$

Next, to hit the small dual caps, we will prove in Lemma 15 below that for any $t^* > 0$, there exists a set of $O(1/t^*)$ points that hits all useful ε-caps of U^* for which $\mathrm{area}([\Gamma_p(U^*)]) \geq t^*$. Setting $t^* = \mu_{d-1} \varepsilon^{d-1}/t_0$, it follows that the resulting set hits all the large, useful ε-caps in U^*. Let $Q' \subset \overline{\partial} U$ denote the dual counterparts of the points in this set. By Lemmas 7 and 12, each large cap in U^* is dually equivalent to a small dual cap in U. It follows that Q' hits all the small, useful ε-dual caps in U and has size

$$|Q'| \;=\; O\left(\frac{t_0}{\mu_{d-1}\varepsilon^{d-1}}\right) \;=\; O\left(\frac{\sqrt{\mathrm{area}(\Psi)}\cdot\varepsilon^{(d-1)/2}}{\varepsilon^{d-1}}\right) \;=\; O\left(\frac{\sqrt{\mathrm{area}(\Psi)}}{\varepsilon^{(d-1)/2}}\right)$$
$$\;=\; O\left(\frac{\Delta_{d-1}(\Psi)}{\varepsilon}\right)^{\frac{d-1}{2}}.$$

Therefore, the union $Q \cup Q'$ is a set of size $O((\Delta_{d-1}(\Psi)/\varepsilon)^{(d-1)/2})$ that hits all the useful ε-dual caps of U. By Lemma 4, the intersection of the upper supporting halfspaces bounded by the supporting hyperplanes at these points yields an ε-approximation to f. This completes the proof of Theorem 2. All that remains is to present the constructions of the hitting sets, which will be given in the subsequent sections.

4.2 Macbeath Regions

As discussed previously, our construction of hitting sets will employ a classical concept from the theory of convex sets called *Macbeath regions* (or *M-regions*). Given a convex body K, a point $x \in K$, and a real parameter $\lambda \geq 0$, the Macbeath region $M_K^\lambda(x)$ (also called an M-region) is defined as

$$M_K^\lambda(x) \;=\; x + \lambda((K - x) \cap (x - K)).$$

It is evident that $M_K^1(x)$ represents the largest centrally symmetric body centered at x and contained within K (see Figure 16(a)). Alternatively, it can be characterized as the intersection of K with the body obtained by reflecting K about x. When used in covering and packing applications, it is common to apply a constant scaling factor about the region's center point (see Figure 16(a)). When the body K is clear from context, we will omit explicit reference to it, defining $M^\lambda(x) \equiv M_K^\lambda(x)$.

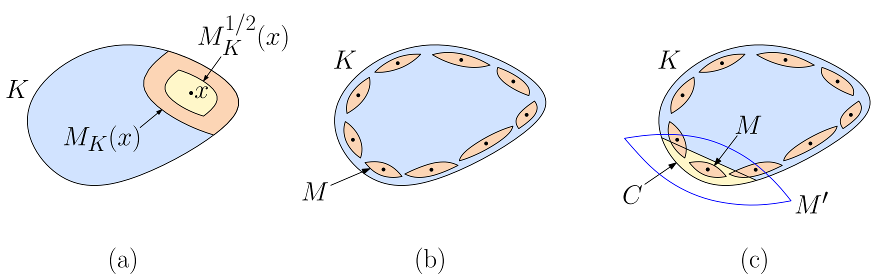

Figure 16: Macbeath regions.

This concept was introduced by Macbeath [40]. Macbeath regions have found numerous uses in the theory of convex sets and the geometry of numbers [17], and they have been applied to several problems in the field of computational geometry, including lower bounds [11, 13, 21], combinatorial complexity [4, 8, 10, 32, 46], approximate nearest neighbor searching [9], and computing the diameter and ε-kernels [7]. One of these applications involves the packing and covering of a collection of caps. (In contrast to our earlier usage, we are using "cap" in its traditional form as the intersection of a halfspace and a convex body.) This has been extensively explored in the works of Ewald, Larman, and Rogers [33], Bárány and Larman [18], Bárány [15, 16], Brönnimann *et al.* [21] and

Arya *et al.* [8]. We will need a variant of the covering lemma, which is presented next. The proof is a straightforward adaptation of Lemma 3.1 in [8], and for the sake of completeness, it is presented in Section 6. Recall that the notation "$\approx$" denotes equality up to constant factors, which may depend on d but not on K.

Lemma 13 (Cap Covering)**.** *Given a convex body $K \subset \mathbb{R}^d$ and any collection $\mathcal{C}$ of caps of K, there exist two collections of convex bodies, $\mathcal{M}$ and $\mathcal{M}'$, such that the bodies of $\mathcal{M}$ are contained within K and are pairwise disjoint (see Figure 16(b)). Each $M \in \mathcal{M}$ is associated with a corresponding body in $\mathcal{M}'$, denoted M', such that $M \subseteq M'$. M' is called M's* expanded body*. These sets satisfy the following:*

- (i) *For all $M \in \mathcal{M}$, $\mathrm{vol}(M') \approx \mathrm{vol}(M)$.*
- (ii) *For any cap $C \in \mathcal{C}$, there exists $M \in \mathcal{M}$ such that $M \subseteq C \subseteq M'$, where M' is M's expanded body (see Figure 16(c)).*

4.3 Hitting Large Dual Caps

In this section, we will explain how to apply the Macbeath region machinery to hit all the large dual caps. Broadly speaking, we will invoke Lemma 13 to construct a collection of regions covering the lower boundary of U, sample a constant number of points from each region, and then project these points onto U's lower boundary. We will show that these points hit all large dual caps. When combined with the points that hit all the small dual caps, Lemma 4 implies that the supporting hyperplanes at the resulting combined set of points will form the basis of the final approximation.

We bound the number of points needed by bounding the number of Macbeath regions. We will do this by recalling Lemma 11, which established the existence of an ellipsoid close to the boundary of U whose projection lies within the shadow of the dual cap's base. This will be combined with a classical sampling technique from computational geometry, based on the concept of ϵ-nets (described below).

Our next lemma is the main result of this section. It states that there exists a small set of points that hits all sufficiently large ε-dual caps of U. The notion of "large" is based on the projected area of the dual cap's base and a threshold parameter t.

Lemma 14. *Given positive reals ε and t and a U-shaped set U in $\mathbb{R}^d$ over a domain Ψ, there exists a set of $O(\mathrm{area}(\Psi)/t)$ points $Q \subset \overline{\partial} U$ that hits all ε-dual caps $D_q(U)$, where $\mathrm{area}([\Delta_q(U)]) \geq t$.*

The rest of this section is devoted to giving the proof. To simplify notation, we will omit references to U when specifying dual caps and their bases. Since U is unbounded, we cap it on top by intersecting it with a horizontal halfspace that is sufficiently high to contain all of the points within vertical distance ε of $\overline{\partial} U$. (The existence of such a halfspace is guaranteed by the fact that U-shaped functions have bounded domains and ranges.) Let U' denote the resulting convex body (see Figure 17(a)). We first apply Lemma 13 to U', where $\mathcal{C}$ is the collection of ε-caps induced by any point $q \in \overline{\partial} U$, such that $\mathrm{area}([\Delta_q]) \geq t$. Let $\mathcal{M}$ and $\mathcal{M}'$ denote the resulting collections of bodies. For each $M \in \mathcal{M}$, if M does not lie entirely within vertical distance ε of $\overline{\partial} U' = \overline{\partial} U$, then discard M from further consideration. For all the surviving regions M, let $M' \in \mathcal{M}'$ denote M's expanded body as described in Lemma 13(ii).

Before proceeding, let us recall some classical results in the theory of sampling and ϵ-nets (see, e.g., [41, 45]). A *set system* is a pair $(X, \mathcal{F})$, where X is a (possibly infinite) set and $\mathcal{F}$ is a

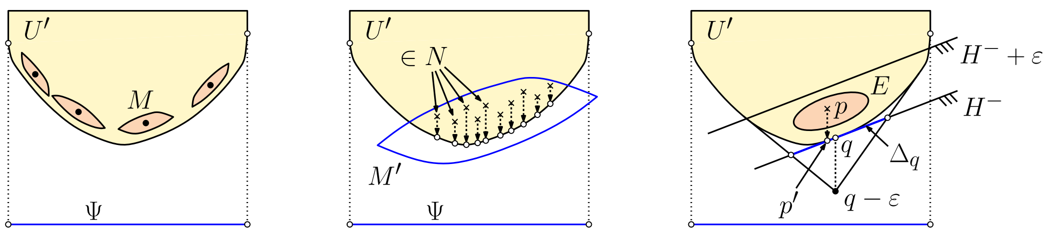

(a) (b) (c)

Figure 17: Proof of Lemma 14.

collection of subsets of X. Let μ be a measure over X. Given a set system $(X, \mathcal{F})$ and a parameter $\epsilon > 0$ (not to be confused with the ε used for approximation), a set $N \subseteq X$ is an *ϵ-net* of $(X, \mathcal{F})$ if for each $F \in \mathcal{F}$ with $\mu(F) \geq \epsilon \cdot \mu(X)$, F contains at least one point of N. The complexity of a set system can be described by a quantity called its *VC-dimension* [3]. We need only two facts regarding this concept. First, any set system of constant VC-dimension has an ϵ-net of size $O\big(\frac{1}{\epsilon}\log\frac{1}{\epsilon}\big)$ [3, 41], and second, the set system $(\Omega, \mathcal{E})$, where Ω is a bounded convex body in $\mathbb{R}^d$ endowed with the Lebesgue measure and $\mathcal{E}$ is the set of ellipsoids contained in Ω has VC-dimension at most $\binom{d+2}{d} = O(d^2)$ [41, Proposition 10.3.2].

Let c'' denote a constant whose value will be specified later. For each $M' \in \mathcal{M}'$, let $N \subset M'$ be an ϵ-net [45] chosen such that any ellipsoid contained within M' of volume at least $c'' \cdot \mathrm{vol}(M')$ contains at least one point of the net (see Figure 17(b)). By the above remarks, the number of points in the net is a constant. Project each point of N vertically downward onto $\overline{\partial}U$ (see Figure 17(b)), and add the resulting point to Q. Repeating this process for all $M \in \mathcal{M}$ yields the set Q.

To establish correctness, consider any ε-dual cap D_q such that $\mathrm{area}([\Delta_q]) \geq t$. Let H^- denote the lower halfspace bounded by the supporting hyperplane at q (see Figure 17(c)). Recall that "$\gtrsim$" denotes "$\geq$" with constant factors ignored. By applying Lemma 11 there exists an ellipsoid $E \subseteq U \cap (H^- + \varepsilon)$ such that

$$\mathrm{vol}(E) \;\gtrsim\; \mathrm{vol}(U \cap (H^- + \varepsilon)) \;\gtrsim\; \varepsilon \cdot \mathrm{area}(\Delta_q^{\downarrow}),$$

and $E^{\downarrow} \subseteq \Delta_q^{\downarrow}$. By Lemma 13(ii), there exists $M \in \mathcal{M}$ such that $M \subseteq U \cap (H^- + \varepsilon) \subseteq M'$. Note that M could not have been discarded in the construction process, since it lies within $H^- + \varepsilon$ and hence is entirely within vertical distance ε of $\overline{\partial}U$. By the above inclusions and Lemma 13(i), we have

$$\mathrm{vol}(E) \;\gtrsim\; \mathrm{vol}(U \cap (H^- + \varepsilon)) \;\gtrsim\; \mathrm{vol}(M) \;\gtrsim\; \mathrm{vol}(M'),$$

and so by choosing the constant c'' in the net construction appropriately, there exists a point p of the net that lies in the interior of E. Thus, the vertical projection p' of this point onto $\overline{\partial}U$ will be included in Q. Since $p^{\downarrow} \in E^{\downarrow} \subseteq \Delta_q^{\downarrow}$, p' hits D_q, as desired.

Finally, we bound the number of points in Q. To do this, consider the set of points of U that lie within vertical distance ε of its lower boundary. The volume of this region is clearly $\varepsilon \cdot \mathrm{area}(\Psi)$. Also, by the above inclusions and Lemma 11, we have

$$\mathrm{vol}(M) \;\gtrsim\; \mathrm{vol}(M') \;\gtrsim\; \mathrm{vol}(U \cap (H^- + \varepsilon)) \;\gtrsim\; \varepsilon \cdot \mathrm{area}(\Delta_q^{\downarrow}) \;=\; \varepsilon \cdot \mathrm{area}([\Delta_q]) \;\geq\; \varepsilon t.$$

The bodies of $\mathcal{M}$ are pairwise disjoint, and (after the discarding process) each of them lies within vertical distance ε of $\overline{\partial}U$. Therefore, by a simple packing argument, the number of such bodies is

$$O\left(\frac{\varepsilon \cdot \mathrm{area}(\Psi)}{\varepsilon t}\right) \;=\; O\left(\frac{\mathrm{area}(\Psi)}{t}\right),$$

as desired.

4.4 Hitting Large Caps in the Dual

In this section we consider the task of bounding the size of hitting sets for large ε-caps in the dual U^*. The construction is similar to the one given in the previous section, but since $(U^*)^{\downarrow}$ covers all of $\mathbb{R}^{d-1}$ some additional effort is required to show that the region of interest can be bounded.

Recall that a point $q = (u; v) \in \overline{\partial}U$ is *useful* if $u \in \Psi$ and $\mathrm{dist}(u, \partial\Psi) \geq \varepsilon$, and a point $p \in \overline{\partial}U^*$ and its associated ε-cap $C_p(U^*)$ are *useful* if p is the dual counterpart of such a point q. The following lemma shows how to apply the Macbeath-region machinery to hit all the useful ε-caps of U^* whose bases have sufficiently large area.

Lemma 15. *Given positive reals δ, ε, and γ where $\delta \ll \varepsilon$ and $\gamma \geq 1$, let U be a (δ, γ)-slope restricted U-shaped set over a domain Ψ. For any $t^* > 0$, there exists a set of $O(\gamma^{d-1}/t^*)$ points $P \subset \overline{\partial}U^*$ that hits all useful ε-caps $C_p(U^*)$ where* $\mathrm{area}([\Gamma_p(U^*)]) \geq t^*$.

Proof. Let Ψ denote U's domain. Let $p = (a; b) \in \overline{\partial}U^*$ be any useful point. To simplify notation, we will omit explicit reference to U^* when discussing caps and cap bases. By Lemma 10, $[\Gamma_p]$ is contained within a ball of radius $2(1 + \gamma)$ centered at the origin. Since p is useful, it is the dual counterpart of some useful point $q = (u; v) \in \overline{\partial}U$. By definition of usefulness, u is at distance at least ε from the boundary of Ψ. Since $\delta < \varepsilon$, slope restriction implies that $\|\nabla f(u)\| \leq \gamma$. Since p is the dual counterpart of q, by Lemma 6(ii) we have $a = \nabla f(u)$. This implies that a lies within distance γ of the origin. Since $[\Gamma_p] = \Gamma_p^{\downarrow} - a$, it follows from the triangle inequality that $\Gamma_p^{\downarrow}$ lies within a ball of radius $2(1 + \gamma) + \gamma \leq 5\gamma$ of the origin, which we denote by B. Since p is a generic useful point, it follows that the vertical projection of all useful ε-caps of U^* lie within B.

Therefore, for the purposes of approximation, we may restrict attention to the portion of U^* that lies above B. We can convert the unbounded set U^* into a convex body by intersecting it with the vertical cylinder whose cross section is B and then bounding it on top by any horizontal hyperplane that is sufficiently high that it does not intersect $\overline{\partial}U^*$. (The existence of such a hyperplane follows from the facts that we need only cover points $(a; b) \in \overline{\partial}U^*$, where $a \in B$ together with the fact that $\|\nabla f^*(a)\| = \|u\|$, which is bounded since $u \in \Psi$.)

We apply Lemma 13 to this restriction of U^*, where $\mathcal{C}$ is the collection of useful caps described in the statement of the lemma. Let $\mathcal{M}$ and $\mathcal{M}'$ denote the resulting collections of bodies. For each $M \in \mathcal{M}$, if M does not lie entirely within vertical distance ε of $\overline{\partial}U^*$, then discard M from further consideration (see Figure 18(a)). For each surviving body M, select an arbitrary point from it (say its center), project this point vertically downward onto $\overline{\partial}U^*$, and add the resulting point to P. Repeating this process for all $M \in \mathcal{M}$ yields the set P.

To establish correctness, consider any useful ε-cap C_p such that $\mathrm{area}([\Gamma_p]) \geq t^*$. Let H^- denote the lower halfspace bounded by p's supporting hyperplane (see Figure 18(b)). By Lemma 13(ii), there exists $M \in \mathcal{M}$ such that $M \subseteq U^* \cap (H^- + \varepsilon) \subseteq M'$. We claim that M could not have been discarded in the construction process. First, it lies within $H^- + \varepsilon$ and hence is entirely within

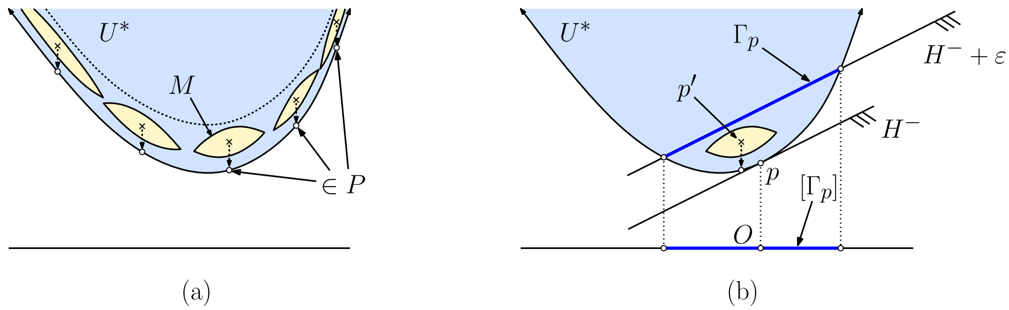

Figure 18: Proof of Lemma 15.

vertical distance ε of $\overline{\partial}U^*$. Second, by Lemma 10, the vertical projection of any useful ε-cap is contained within the ball B, and so this applies to M as well. Since M lies within $U^* \cap (H^- + \varepsilon)$, the vertical projection of any point $p' \in M$ onto $\overline{\partial}U^*$ hits C_p, as desired.

In order to bound the number of points in P, consider the set of points of $U^* \cap B$ that lie within vertical distance ε of $\overline{\partial}U^*$. Ignoring constant factors, the volume of this region is clearly $\varepsilon \cdot \mathrm{area}(B) \lesssim \varepsilon\gamma^{d-1}$. Also, for each body M, by the above inclusions and Lemma 13(i), we have

$$\mathrm{vol}(M) \ \gtrsim\ \mathrm{vol}(M') \ \geq\ \mathrm{vol}(U^* \cap (H^- + \varepsilon)).$$

The set $U^* \cap (H^- + \varepsilon)$ lies within the intersection of an infinite vertical cylinder whose horizontal cross section is $\Gamma_p^{\downarrow}$ and the slab $(H^- + \varepsilon) \setminus H^-$ of vertical height ε. This set's volume is $\varepsilon \cdot \mathrm{area}(\Gamma_p^{\downarrow})$. Thus, we have

$$\mathrm{vol}(M) \ \gtrsim\ \varepsilon \cdot \mathrm{area}(\Gamma_p^{\downarrow}) \ =\ \varepsilon \cdot \mathrm{area}([\Gamma_p]) \ \geq\ \varepsilon t^*.$$

The bodies of $\mathcal{M}$ are pairwise disjoint, and (after discarding) each of them lies within vertical distance ε of $\overline{\partial}U^*$. Therefore, by a simple packing argument, the number of such bodies is

$$O\left(\frac{\varepsilon\gamma^{d-1}}{\varepsilon t^*}\right) \ =\ O\left(\frac{\gamma^{d-1}}{t^*}\right),$$

as desired. $\square$

Having established Lemmas 14 and 15, the proof of Theorem 2 given in Section 2.2 is now complete.

5 Nonuniform Area-Based Bounds

In this section we note that a nonuniform bound very similar to ours can be derived from a result due to Gruber [34], who showed that if K is a strictly convex body and ∂K is twice differentiable (C^2 continuous), then there exists a constant k_d (depending only on the dimension d) and ε_0 depending on K, such that for any $0 < \varepsilon \leq \varepsilon_0$, the number of bounding halfspaces needed to achieve an ε-approximation to K is at most

$$k_d \left(\frac{1}{\varepsilon}\right)^{\frac{d-1}{2}} \int_{\partial K} \kappa(x)^{\frac{1}{2}} d\sigma(x), \tag{11}$$

where κ and σ denote the Gaussian curvature of K and ordinary surface area measure, respectively. (Böröczky showed that the requirement that K be "strictly" convex can be eliminated [14].) Because the square root function is concave and $\int_{\partial K} d\sigma(x) = \text{area}(K)$, we may apply Jensen's inequality to obtain

$$\frac{1}{\text{area}(K)} \int_{\partial K} \kappa(x)^{\frac{1}{2}} d\sigma(x) \;\leq\; \left(\frac{1}{\text{area}(K)} \int_{\partial K} \kappa(x) d\sigma(x)\right)^{\frac{1}{2}}.$$

Thus,

$$\int_{\partial K} \kappa(x)^{\frac{1}{2}} d\sigma(x) \;\leq\; \left(\text{area}(K) \int_{\partial K} \kappa(x) d\sigma(x)\right)^{\frac{1}{2}}.$$

By the Gauss-Bonnet theorem [30], the total Gaussian curvature of K is bounded by some quantity ζ_d, depending only on d. Also, letting α_{d-1} denote the surface area of the unit $(d-1)$-sphere, we have $\text{area}(K) = \alpha_{d-1}(\Delta_{d-1}(K)/2)^{d-1}$. Therefore,

$$\int_{\partial K} \kappa(x)^{\frac{1}{2}} d\sigma(x) \;\leq\; (\zeta_d \cdot \text{area}(K))^{\frac{1}{2}} \;=\; \left(\frac{\zeta_d \alpha_{d-1}}{2^{d-1}}\right)^{\frac{1}{2}} (\Delta_{d-1}(K))^{\frac{d-1}{2}}.$$

Substituting the above quantity into Eq.(11) and setting c_d to the constant $k_d(\zeta_d \alpha_{d-1}/2^{d-1})^{1/2}$, we obtain the following.

Theorem 3. *For any integer $d \geq 2$, there exists a constant c_d (depending on d) such that for any convex body $K \subseteq \mathbb{R}^d$ whose boundary is C^2 smooth, there exists ε_0 depending on K, such that for any $0 < \varepsilon \leq \varepsilon_0$, there exists an ε-approximating polytope P whose number of facets is at most*

$$c_d \left(\frac{\Delta_{d-1}(K)}{\varepsilon}\right)^{\frac{d-1}{2}}.$$

where c_d is a constant (depending on d) and $\Delta_{d-1}(K)$ is K's surface diameter.

Note that the bound in this theorem matches the uniform bound of Theorem 1. However, this approach cannot be used to produce a uniform bound. To see why, suppose to the contrary that such a bound existed, even in $\mathbb{R}^2$. That is, there exists a constant k_2 and positive ε_0 such that for all $\varepsilon \leq \varepsilon_0$ and all convex bodies K (of width at least ε in every direction) in $\mathbb{R}^2$, there exists an ε-approximating polygon whose number of sides satisfies Eq. (11). Consider any $\varepsilon \leq \min(\varepsilon_0, 1/9)$, and let $0 < \delta \leq \varepsilon$ be a sufficiently small value (chosen below). Set $m = \lfloor 1/\sqrt{\delta} \rfloor$, and define K_δ to be the Minkowski sum of a regular m-gon inscribed in a unit circle and a Euclidean ball of radius δ (see Figure 19(a)). Observe that since $m \geq 3$ and $\delta \leq 1/9$, K_δ satisfies the minimum width requirements. It consists of m straight edges, each of length $\Theta(\sqrt{\delta})$, connected by m circular arcs, each of radius δ and subtending an angle of $2\pi/m$. Since $\delta \leq \varepsilon$, it is straightforward to show that any convex polygon K_ε that ε-approximates K_δ requires $\Omega(1/\sqrt{\varepsilon})$ sides (see Figure 19(b)). (As δ decreases relative to ε, K_δ approaches a unit disk, and it is easy to show that in order to maintain a distance of at most ε, each side can have of length at most $c\sqrt{\varepsilon}$, for some constant c.)

Boundary points along the flat sides of K_δ have zero curvature, and boundary points within each each circular arc have curvature $1/\delta$. Since the circular arcs together cover a distance of $2\pi\delta$ of the boundary, it follows that

$$\int_{\partial K_\delta} \kappa(x)^{1/2} d\sigma(x) \;=\; \frac{2\pi\delta}{\sqrt{\delta}} \;=\; \Theta\left(\sqrt{\delta}\right).$$

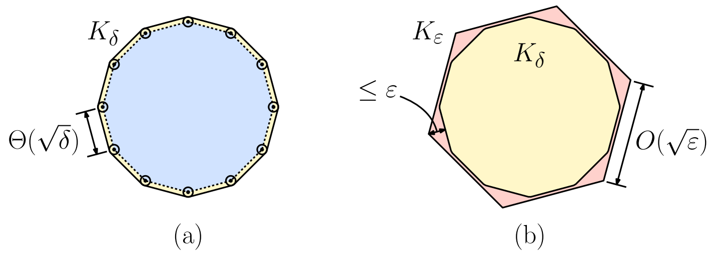

Figure 19: Why Theorem 3 cannot be used to generate a uniform bound.

Therefore, the hypothesized uniform bound would imply the existence of an ε-approximating polygon with $O\big(\sqrt{\delta/\varepsilon}\big)$ sides, contradicting the lower bound of $\Omega(1/\sqrt{\varepsilon})$ for all sufficiently small δ.[1]

6 Proof of the Cap-Covering Lemma

In this section, we present a proof of Lemma 13 on cap covering from Section 4.2. Our proof is similar in spirit to the proofs of related covering lemmas by Bárány and Larman [18], Bárány [16], and Arya *et al.* [8]. Before proceeding with the proof, we recall some standard definitions. Given a *cap* C defined as the nonempty intersection of K with a halfspace H, let h denote the hyperplane bounding H. Define C's *base* to be $h \cap K$, and define its *width* to be the distance between h and C's opposing parallel supporting hyperplane (see Figure 20(a)). Given any cap C of width w and a real parameter $\lambda \geq 0$, 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.

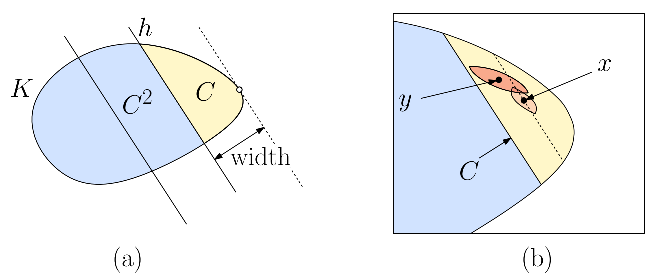

Figure 20: Proof of the Cap-Covering lemma.

Throughout this section, we assume that K is a convex body scaled to have unit volume, and ν_0 is a sufficiently small constant depending only on d. We begin with some important properties of Macbeath regions which were proved in [33] or [21].

The following lemma summarizes the basic properties of Macbeath regions that will be used in our proof. Claim (i) is a variant of Lemma 1 of [33] and was established by Brönnimann, Chazelle,

[1] Note that we cannot apply Gruber's or Böröczky's theorems directly to K_δ, since its boundary is not twice differentiable. In particular, the second derivative is discontinuous at the joints where each edge meets a circular arc. We can easily fix this by creating a sufficiently small gap at each joint and introducing a smooth polynomial spline of constant degree to fill the gap. Although the resulting body is not strictly convex, Böröczky showed that this assumption is not necessary for the bound to hold.

and Pach [21, Lemma 2.5]. Claim (ii) is a straightforward adaptation of Lemma 2.8 in [8] and is based on ideas from [33, Lemma 2] and [21, Lemma 2.6]. Claim (iii) is an immediate consequence of the definition of Macbeath regions.

Lemma 16. *Given a convex body K in $\mathbb{R}^d$ and a cap C of K:*

(*i*) *For $x, y \in K$, if $M^{1/5}(x) \cap M^{1/5}(y) \neq \emptyset$, then $M^{1/5}(x) \subseteq M(y)$.*
(*ii*) *If C has volume at most ν_0, then for any $\lambda \geq 1$, $C^\lambda \subseteq M^{3d(2\lambda-1)}(x)$, where x is the centroid of the base of C.*
(*iii*) *For any $x \in C$, $M(x) \subseteq C^2$.*

Our proof of the cap covering lemma is based on the following lemma, which establishes a relationship between scalings of Macbeath regions and caps (see Figure 20(b)).

Lemma 17. *Let C be a cap of K with volume at most ν_0. Let x denote the centroid of the base of the cap $C^{1/2}$. For any point $y \in K$ such that $M^{1/5}(x) \cap M^{1/5}(y) \neq \emptyset$,*

$$M^{1/5}(y) \;\subseteq\; C \;\subseteq\; M^{45d}(y).$$

Proof. Since $M^{1/5}(x) \cap M^{1/5}(y) \neq \emptyset$, Lemma 16(i) implies that $M^{1/5}(y) \subseteq M(x)$ and $M^{1/5}(x) \subseteq M(y)$. Also, by Lemma 16(iii), $M(x) \subseteq (C^{1/2})^2 = C$. Combining these, we obtain

$$M^{1/5}(y) \;\subseteq\; M(x) \subseteq C,$$

which establishes the first inclusion.

To prove the second inclusion, we apply Lemma 16(ii) setting C to $C^{1/2}$ and λ to 2. We obtain

$$C \;=\; (C^{1/2})^2 \;\subseteq\; M^{3d(2\lambda-1)}(x) \;=\; M^{9d}(x).$$

Recall that $M^{1/5}(x) \subseteq M(y)$. Scaling both of these centrally symmetric bodies by any positive factor about their respective centers preserves the inclusion (see, e.g., Bárány [15]), and hence

$$M^{9d}(x) \;=\; (M^{1/5}(x))^{45d} \;\subseteq\; M^{45d}(y).$$

Putting these together, we obtain $C \subseteq M^{9d}(x) \subseteq M^{45d}(y)$, as desired. ☐

We are now ready to present the proof of Lemma 13. Before considering the general case, let us assume that all the caps in $\mathcal{C}$ have volume at most ν_0. Let $\mathcal{M}$ be any maximal set of disjoint Macbeath regions of the form $M^{1/5}(x)$, where x is the centroid of the base of the cap $C^{1/2}$ for some $C \in \mathcal{C}$. For each Macbeath region $M^{1/5}(x)$, define its expanded body M' to be $M^{45d}(x)$. We will show that $\mathcal{M}$ and $\mathcal{M}'$ satisfy the properties given in the lemma. Property (i) is straightforward since M' is related to M by a constant scaling factor of $45d/(1/5) = 225d$. To show Property (ii), consider any cap $C \in \mathcal{C}$. Let x denote the centroid of the base of $C^{1/2}$. By the maximality of $\mathcal{M}$, there is a Macbeath region $M^{1/5}(y) \in \mathcal{M}$ such that $M^{1/5}(x) \cap M^{1/5}(y) \neq \emptyset$. Applying Lemma 17, it follows that $M^{1/5}(y) \subseteq C \subseteq M^{45d}(y)$. This establishes Property (ii) and thus proves the lemma for the special case when all caps have volume at most ν_0.

We now discuss the modifications required for handling the general case. For each cap $C \in \mathcal{C}$ whose volume exceeds ν_0, we replace it by the cap C^λ, where $\lambda < 1$ is chosen such that the volume of C^λ is exactly ν_0. Otherwise we retain the original cap C. Let $\mathcal{C}'$ represent the resulting set of caps. We construct the sets $\mathcal{M}$ and $\mathcal{M}'$ for the set $\mathcal{C}'$ exactly as described in the special case above. Finally, for each expanded body $M' \in \mathcal{M}'$, if its volume is at least ν_0, we replace it by the convex body K. Otherwise, we retain the same body M'. Let $\mathcal{M}''$ denote the resulting set of expanded bodies. We claim that the sets $\mathcal{M}$ and $\mathcal{M}''$ satisfy the properties given in the lemma for the set $\mathcal{C}$.

First, note that the argument given for the special case implies that the sets $\mathcal{M}$ and $\mathcal{M}'$ satisfy these properties for $\mathcal{C}'$. As we replace the expanded body only if its volume is at least constant ν_0, it follows that Property (ii) holds for $\mathcal{M}$ and $\mathcal{M}''$ (the ratio of the volume of M's expanded body to the volume of M increases by a factor of at most $1/\nu_0$).

To establish Property (ii), consider a cap $C \in \mathcal{C}$. If the volume of C is at most ν_0, then the argument given in the special case shows that there exists a body $M \in \mathcal{M}$ such that $M \subseteq C \subseteq M'$, where M' is M's expanded body in $\mathcal{M}'$. Regardless of whether M' is retained or replaced by K in constructing $\mathcal{M}''$, this property continues to hold. On the other hand, if the volume of C exceeds ν_0, then recall that it is replaced by a cap C^λ, where $\lambda < 1$ is chosen such that the volume of C^λ is exactly ν_0. The argument given in the special case shows that there exists a body $M \in \mathcal{M}$ such that $M \subseteq C^\lambda \subseteq M'$, where M' is M's expanded body in $\mathcal{M}'$. It follows that $\mathrm{vol}(M') \geq \mathrm{vol}(C^\lambda) = \nu_0$. Thus M' must be replaced by K in constructing $\mathcal{M}''$. In other words, M's expanded body in $\mathcal{M}''$ is K. Clearly, Property (ii) holds since $M \subseteq C \subseteq K$. This completes the proof of Lemma 13.

7 Concluding Remarks

In this paper, we have proved the existence of an ε-approximation to a convex body K in $\mathbb{R}^d$, whose size is sensitive to the shape of the body expressed in terms of its surface diameter. Our result yields a uniform bound, meaning that the result holds for all $\varepsilon \leq \varepsilon_0$, where ε_0 does not depend on K. We have shown that this bound is tight up to constant factors (depending on the dimension) as a function of surface diameter. Our approach is based on first decomposing K's boundary into a collection of surfaces, each of which is represented by a smooth, strictly convex function, and then exploiting properties of the Legendre transform and Macbeath regions to guide a sampling process and analyze its complexity. The connection between approximations convex bodies and approximations of convex functions in a computational context has been observed elsewhere [1, 9, 36, 50], and the techniques developed here may be of interest to future applications.

There are a number of interesting additional questions raised by our work. The surface diameter, Δ_{d-1}, (or alternatively the $(d-1)$st intrinsic volume) is but one way of introducing a measure of shape-sensitivity. As mentioned earlier, in another work, we have demonstrated a bound based on the volume diameter Δ_d, which is asymptotically superior to the bound presented here [12]. The approach presented here is very different from that one, and it may be applicable in contexts where the volume-sensitive approach is not, for example, approximating convex surface patches. An important question is what are the most robust measures of a body's ease of approximation, and which approximation techniques are most widely applicable. The ultimate goal would be a construction that yields the polytope of minimum combinatorial complexity that approximates a given body. Unfortunately, existing hardness results suggest that this may not be solvable in polynomial time [29].

The issue of how to compute the approximation has not been a focus of this paper, but we remark that all the elements of our construction can be implemented. Assuming that K is represented as the intersection of n halfspaces, we conjecture that the approximation can be constructed by a randomized algorithm in expected time $O\big(n + 1/\varepsilon^{O(d)}\big)$. The idea is to first reduce the size of K by computing an $\frac{\varepsilon}{2}$ approximation to K with $O\big(1/\varepsilon^{O(d)}\big)$ facets. Letting K' denote this approximation, we then compute an $\frac{\varepsilon}{2}$ approximation to K' using the construction presented in this paper. The polytope K' can be computed in time $O\big(n+1/\varepsilon^{O(d)}\big)$ by standard methods [7,25]. We believe that the elements of our construction (support functions, polar bodies, Macbeath regions, ϵ-nets) can all be computed (or approximated to the necessary precision) in time that is a polynomial function of ε and the number of facets of K'. Note that computing the ϵ-net involves random sampling. We conjecture that, given K', the entire construction can be performed in time $O\big(1/\varepsilon^{O(d)}\big)$. We leave open the question of whether there is a construction whose running time is sensitive to K's surface diameter.

8 Acknowledgments

We thank Quentin Merigót for pointing out the relationship to Gruber's result in Section 5. We also thank the anonymous reviewers of an earlier draft of this paper, for making a number of valuable suggestions regarding the paper's organization and approach.

References

[1] A. Abdelkader, S. Arya, G. D. da Fonseca, and D. M. Mount. Approximate nearest neighbor searching with non-Euclidean and weighted distances. In *Proc. 30th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 355–372, 2019. `doi:10.1137/1.9781611975482.23`.

[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] N. Alon, D. Haussler, and E. Welzl. Partitioning and geometric embedding of range spaces of finite Vapnik-Chervonenkis dimension. In *Proc. Third Annu. Sympos. Comput. Geom.*, pages 331–340, 1987. `doi:10.1145/41958.41994`.

[4] R. Arya, S. Arya, G. D. da Fonseca, and D. M. Mount. Optimal bound on the combinatorial complexity of approximating polytopes. *ACM Trans. Algorithms*, 18:1–29, 2022. `doi:10.1145/3559106`.

[5] S. Arya and T. 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. `doi:10.1145/2582112.2582161`.

[6] S. Arya, G. D. da Fonseca, and D. M. Mount. Polytope approximation and the Mahler volume. In *Proc. 23rd Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 29–42, 2012. `doi:10.1137/1.9781611973099.3`.

[7] S. Arya, G. D. da Fonseca, and D. M. Mount. Near-optimal ε-kernel construction and related problems. In *Proc. 33rd Internat. Sympos. Comput. Geom.*, pages 10:1–15, 2017. URL: `https://arxiv.org/abs/1703.10868`, `doi:10.4230/LIPIcs.SoCG.2017.10`.

[8] S. Arya, G. D. da Fonseca, and D. M. Mount. On the combinatorial complexity of approximating polytopes. *Discrete Comput. Geom.*, 58(4):849–870, 2017. doi:10.1007/s00454-016-9856-5.

[9] 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. doi:10.1137/1.9781611974782.18.

[10] S. Arya, G. D. da Fonseca, and D. M. Mount. Economical convex coverings and applications. In *Proc. 34th Annu. ACM-SIAM Sympos. Discrete Algorithms*, pages 1834–1861, 2023. doi:10.1137/1.9781611977554.ch70.

[11] 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. doi:10.1007/s00454-009-9140-z.

[12] S. Arya and D. M. Mount. Optimal volume-sensitive bounds for polytope approximation. In *Proc. 39th Internat. Sympos. Comput. Geom.*, pages 9:1–9:16, 2023. doi:10.4230/LIPIcs.SoCG.2023.9.

[13] 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.

[14] K. Böröczky Jr. Approximation of general smooth convex bodies. *Adv. Math.*, 153:325–341, 2000. doi:10.1006/aima.1999.1904.

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

[16] I. Bárány. The technique of M-regions and cap-coverings: A survey. *Rend. Circ. Mat. Palermo*, 65:21–38, 2000. URL: https://users.renyi.hu/~barany/.

[17] I. Bárány. Random polytopes, convex bodies, and approximation. In W. Weil, editor, *Stochastic Geometry*, volume 1892 of *Lecture Notes in Mathematics*, pages 77–118. Springer, 2007. doi:10.1007/978-3-540-38175-4_2.

[18] I. Bárány and D. G. Larman. Convex bodies, economic cap coverings, random polytopes. *Mathematika*, 35:274–291, 1988. doi:10.1112/S0025579300015266.

[19] G. Bonnet. Polytopal approximation of elongated convex bodies. *Advances in Geometry*, 18:105–114, 2018. doi:10.1515/advgeom-2017-0038.

[20] J. Bourgain and V. D. Milman. New volume ratio properties for convex symmetric bodies. *Invent. Math.*, 88:319–340, 1987. doi:10.1007/BF01388911.

[21] H. Brönnimann, B. Chazelle, and J. Pach. How hard is halfspace range searching? *Discrete Comput. Geom.*, 10:143–155, 1993. doi:10.1007/BF02573971.

[22] H. Brönnimann and M. T. Goodrich. Almost optimal set covers in finite VC-dimension,. *Discrete Comput. Geom.*, 14:463–479, 1995. doi:10.1007/BF02570718.

[23] E. M. Bronshteyn and L. D. Ivanov. The approximation of convex sets by polyhedra. *Siberian Math. J.*, 16:852–853, 1976. doi:10.1007/BF00967115.

[24] E. M. Bronstein. Approximation of convex sets by polytopes. *J. Math. Sci.*, 153(6):727–762, 2008. doi:10.1007/s10958-008-9144-x.

[25] 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.

[26] B. Chazelle, L. J. Guibas, and D. T. Lee. The power of geometric duality. *BIT*, 25:76–90, 1985. doi:10.1007/BF01934990.

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

[28] K. L. Clarkson. Building triangulations using ε-nets. In *Proc. 38th Annu. ACM Sympos. Theory Comput.*, pages 326–335, 2006. doi:10.1145/1132516.1132564.

[29] G. Das and D. Joseph. The complexity of minimum convex nested polyhedra. In *Proc. Second Canad. Conf. Comput. Geom.*, pages 296–301, 1990.

[30] M. P. do Carmo. *Differential Geometry of Curves and Surfaces.* Prentice Hall, 1976.

[31] R. M. Dudley. Metric entropy of some classes of sets with differentiable boundaries. *J. Approx. Theory*, 10(3):227–236, 1974. doi:10.1016/0021-9045(74)90120-8.

[32] K. Dutta, A. Ghosh, B. Jartoux, and N. H. Mustafa. Shallow packings, semialgebraic set systems, Macbeath regions and polynomial partitioning. *Discrete Comput. Geom.*, 61:756–777, 2019. doi:10.1007/s00454-019-00075-0.

[33] 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. doi:10.1112/S0025579300002655.

[34] P. M. Gruber. Aspects of approximation of convex bodies. In P. M. Gruber and J. M. Wills, editors, *Handbook of Convex Geometry*, chapter 1.10, pages 319–345. North-Holland, 1993. doi:10.1016/B978-0-444-89596-7.50015-8.

[35] B. Grünbaum. Measures of symmetry for convex sets. In *Proc. Sympos. Pure Math.*, volume VII, pages 233–270, 1963. doi:10.1090/pspum/007/0156259.

[36] S. Har-Peled and N. Kumar. Approximating minimization diagrams and generalized proximity search. *SIAM J. Comput.*, 44:944–974, 2015. doi:10.1137/140959067.

[37] 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.

[38] V. Klee. Some new results on smoothness and rotundity in normed linear spaces. *Math. Ann.*, 139:51–63, 1959. doi:10.1007/BF01459822.

[39] G. Kuperberg. From the Mahler conjecture to Gauss linking integrals. *Geom. Funct. Anal.*, 18:870–892, 2008. doi:10.1007/s00039-008-0669-4.

[40] A. M. Macbeath. A theorem on non-homogeneous lattices. *Ann. of Math.*, 56:269–293, 1952. doi:10.2307/1969800.

[41] J. Matoušek. *Lectures on Discrete Geometry.* Spring-Verlag, 2002. doi:10.1007/978-1-4613-0039-7.

[42] D. E. McClure and R. A. Vitalie. Polygonal approximation of plane convex bodies. *J. Math. Anal. Appl.*, 51:326–358, 1975. doi:10.1016/0022-247X(75)90125-0.

[43] P. McMullen. Non-linear angle-sum relations for polyhedral cones and polytopes. *Math. Proc. Cambridge Philos. Soc*, 78(2):247–261, 1975. doi:10.1017/S0305004100051665.

[44] P. McMullen. Inequalities between intrinsic volumes. *Monatshefte für Mathematik*, 111:47–53, 1991. doi:10.1007/BF01299276.

[45] N. H. Mustafa. *Sampling in Combinatorial and Geometric Set Systems*, volume 265 of *Mathematical Surveys and Monographs.* AMS, 2022. doi:10.1090/surv/265.

[46] 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. doi:10.4230/LIPIcs.STACS.2014.578.

[47] M. Naszódi, F. Nazarov, and D. Ryabogin. Fine approximation of convex bodies by polytopes. *Amer. J. Math*, 142:809–820, 2020. doi:10.1353/ajm.2020.0018.

[48] F. Nazarov. The Hörmander proof of the Bourgain-Milman theorem. In *Geometric Aspects of Functional Analysis*, pages 335–343. Springer, 2012. doi:10.1007/978-3-642-29849-3_20.

[49] R. T. Rockafellar. Conjugates and Legendre transforms of convex functions. *Canad. J. Math.*, pages 200–205, 1967. doi:10.4153/CJM-1967-012-4.

[50] G. Rote. The convergence rate of the sandwich algorithm for approximating convex functions. *Computing*, 48:337–361, 1992. doi:10.1007/BF02238642.

[51] R. Schneider. Polyhedral approximation of smooth convex bodies. *J. Math. Anal. Appl.*, 128:470–474, 1987. doi:10.1016/0022-247X(87)90197-1.

[52] L. F. Toth. Approximation by polygons and polyhedra. *Bull. Amer. Math. Soc.*, 54:431–438, 1948. doi:10.1090/S0002-9904-1948-09022-X.