



HAL
open science

Computable bounds for the reach and r -convexity of subsets of R^d

Ryan Cotsakis

► **To cite this version:**

Ryan Cotsakis. Computable bounds for the reach and r -convexity of subsets of R^d . 2023. hal-03854575v2

HAL Id: hal-03854575

<https://hal.science/hal-03854575v2>

Preprint submitted on 20 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Computable bounds for the reach and r -convexity of subsets of \mathbb{R}^d

Ryan Cotsakis

Laboratoire J.A. Dieudonné, Université Côte d'Azur
28, Avenue Valrose, 06108 Nice Cedex 2, France

ryan.cotsakis@univ-cotedazur.fr

Abstract

The convexity of a set can be generalized to the two weaker notions of reach and r -convexity; both describe the regularity of a set's boundary. For any compact subset of \mathbb{R}^d , we provide methods for computing upper bounds on these quantities from point cloud data. The bounds converge to the respective quantities as the point cloud becomes dense in the set, and the rate of convergence for the bound on the reach is given under a weak regularity condition. We also introduce the β -reach, a generalization of the reach that excludes small-scale features of size less than a parameter $\beta \in [0, \infty)$. Numerical studies suggest how the β -reach can be used in high-dimension to infer the reach and other geometric properties of smooth submanifolds.

Keywords: double offset, beta-reach, medial axis, high-dimensional, point clouds, geometric inference

Acknowledgements

I extend my gratitude to Elena Di Bernardino and Thomas Opitz for their attentive and indispensable guidance. Thank you to the anonymous reviewers, who significantly helped to improve the communication of ideas in this manuscript. This work has been supported by the French government, through the 3IA Côte d'Azur Investments in the Future project managed by the National Research Agency (ANR) with the reference number ANR-19-P3IA-0002.

1 Introduction

A number of concepts from convex geometry generalize from convex sets to much larger classes of sets. A classic example from [39] is the extension of kinematic formulas (in particular, Steiner’s formula) for convex sets to sets with *positive reach* (see Definition 2). Another example is the notion of the convex hull of a set, which can be weakened to the *r -convex hull*, for $r > 0$. This weak notion of a convex hull, instead of being expressed as the intersection of half-spaces, is expressed in terms of intersections of the complements of open balls of radius r . The resulting intersection is said to be *r -convex* (see Definition 4), which differs subtly from the notion of reach, and these differences have been studied in [27] and [32] for example. Both the reach and r -convexity—ranging from 0 to ∞ inclusive—can be seen as measures of the degree to which a set is convex. This paper introduces methods of computing upper bounds for these quantities from point cloud data in several, general settings.

There is a vast literature on properties of r -convex sets, dating from [46]. Their relations to other classes of sets that generalize convexity have been studied in [27, 32, 50, 53]. The use of r -convexity in image smoothing has been suggested in [31, 50, 53] as well as for set estimation in [4, 30–32, 43, 45]. In other literature, the r -convex hull of a set is referred to as its *double offset* [19, 20, 40]. In [34], the authors use a rolling-type condition (weaker than r -convexity) to improve the rate of convergence of their set boundary estimator. Efforts towards estimating the largest r such that a set is r -convex have been made in [49]; the authors test the uniformity of a point cloud on its r -convex hull using the statistical test of uniformity proposed in [11].

The reach is a popular measure of convexity largely due to its link with the Steiner formula; for the larger the reach of a set, the larger the interval on which the volume of a dilation of the set is polynomial in the dilation radius [39, Theorem 5.6]. The positivity of the reach implies a number of desirable properties, and so positive reach is a common hypothesis to guarantee convergence rates of statistical estimators of geometric quantities [13, 23, 29, 30, 48, 52]. In the field of geometric data analysis, the reach constitutes a commonly used measure of regularity of a set’s boundary, and is hence also referred to as the *condition number* [44]. In topological data analysis, a set’s reach is shown to quantify the ability to infer its homology from point cloud data [22, 40, 41, 44]. The statistical estimation of the reach has been of particular interest in recent literature. [2] and [3] suggest estimating the reach of a smooth submanifold of \mathbb{R}^d by measuring distances between it and its tangent spaces—a strategy inspired by the formulation of the reach in [39, Theorem 4.18]. These works obtain bounds on the minimax rate of convergence for this estimator when the manifold is C^3 -smooth, and when its tangent spaces are known. Following up to these works, [1] establishes an optimal convergence rate for minimax estimators of the reach of C^k -smooth submanifolds of \mathbb{R}^d without boundary. The rates that they establish adapt to the smoothness of the manifold, and to the type of phenomenon that limits the reach (see Remark 8). [10] looks to

the *convexity defect function* introduced in [7] and suggests an estimator for the reach of C^k -smooth manifolds by establishing a link between the convexity defect function and the reach. Convergence rates of their estimator are given for $k \geq 3$. In [25], the authors introduce a complete and tractable method for estimating the reach from point cloud data involving the computation of graph distances in a spatial network defined over the point cloud. Their approach is based on the formulation of the reach introduced in [16, Theorem 1].

Other strategies for estimating the reach involve a prior estimation of the *medial axis* (see Definition 3) [35, 36]. The λ -medial axis [21], μ -medial axis [18], and (λ, α) -medial axis [42] are all generalizations of the medial axis, and the reach can conceivably be approximated by the minimal distance from a set to the estimate of the medial axis of its complement. Such a strategy is suggested in [33] for the λ -medial axis, where the authors heavily rely on the notion of r -convexity in their construction.

A number of other authors have taken an interest in the mathematical properties of the reach. [47] identifies the sets of reach r with the r -proximally smooth sets introduced in [26], implying that the reach can be characterized by the gradients of the distance-to-set function. [27] makes a number of insightful connections between the reach and other geometrical properties of sets. [6] studies Vietoris–Rips complexes via the reach, proving that the reach of a set can only increase if intersected by sufficiently small balls. An alternative characterization of the reach, involving pairwise geodesic distances, is provided in [16]. The authors also study the relationship between midpoints of pairs of points in a set, and the set’s reach [16, Lemma 1]. We base the construction of our bound for the reach on this result (see Theorem 2).

This paper makes steps towards providing computationally tractable methods for bounding the reach and r -convexity of subsets of \mathbb{R}^d given point cloud data that represent the underlying sets.

Firstly, we establish some facts about the reach of closed subsets of \mathbb{R}^d . We prove that the r -convexity and reach are equivalent for compact subsets of \mathbb{R}^d whose topological boundary is a C^1 -smooth, $(d - 1)$ -dimensional manifold without boundary (see Theorem 1). In addition, for closed subsets of \mathbb{R}^d , we introduce the β -reach (see Definition 6), a quantity that loosely represents the reach of a set when features of size less than $\beta \in [0, \infty)$ are ignored. Indeed, the β -reach is identified with the reach for $\beta = 0$ (see Theorem 2).

These ideas are used to create methods of inferring bounds on the reach and r -convexity of sets from point cloud data. For general, closed subsets of \mathbb{R}^d (possibly having finite d -volume, and having *no smoothness conditions* on their topological boundaries), we provide an upper bound of the r -convexity of the set based on samples of the set and its complement at sampling locations that extend over \mathbb{R}^d . We show that, as the sampling locations become dense in \mathbb{R}^d , this bound converges to the largest r such that the set is r -convex (see Theorem 3). An example on real data in 3 dimensions shows that this method identifies regions where the underlying set is not locally r -convex, for $r > 0$, with a test specificity of 100%. Similarly, for any closed set $A \subseteq \mathbb{R}^d$, we define

an upper bound on the reach based on a set of points known to reside in A . As the set of points converges in the Hausdorff metric to A , we show that the bound converges to the reach of A , and provide the rate of convergence in terms of the Hausdorff distance between A and the sample points (see Theorem 4). A weak regularity condition on the β -reach of A , for β near 0, is used to show the convergence of our upper bound with a rate. In practice, both the β -reach of a point cloud and the upper bound on the reach can be computed efficiently in high-dimension. The computational complexity of the method increases only linearly with the dimension of the ambient space d .

The organization of the paper is as follows. Section 2 introduces the notation that we use throughout the document, explores the relationships between the reach and r -convexity, and introduces the β -reach. Section 3 describes the three methods that we propose for inferring bounds and approximations for the reach and r -convexity of general compact subsets of \mathbb{R}^d from point cloud data. The bounds for the r -convexity and reach are studied in Sections 3.1 and 3.2 respectively. Section 3.3 elaborates on how the β -reach of point clouds can be used to approximate the β -reach of the sets that they represent. In Section 4, we provide numerical studies that underline the computability of our methods. An application of the methods on real data is given in Section 4.1. In Section 4.2, the methods are tested against simulated data for which the reach and r -convexity are known, and empirical rates of convergence are provided. Some technical proofs and auxiliary results are postponed to Section 5.

2 Definitions and important notions

The sets that we study in this paper are subsets of \mathbb{R}^d , endowed with the Euclidean metric $\|\cdot\|_2$. For a set $S \subseteq \mathbb{R}^d$, let ∂S denote its topological boundary, let $\text{cl}(S) := S \cup \partial S$ denote its closure, and let S^c denote its complement in \mathbb{R}^d . Denote the closed ball with radius $r \in \mathbb{R}^+$ centered at $s \in \mathbb{R}^d$ by $B(s, r) := \{t \in \mathbb{R}^d : \|t - s\|_2 \leq r\}$. For $t \in \mathbb{R}^d$, denote the distance between t and a non-empty set S by $\delta_S(t) := \inf\{\|t - s\|_2 : s \in S\}$.

2.1 Set dilation, set erosion, and combinations of the two

Definition 1 (Operations on subsets of \mathbb{R}^d). We recall the Minkowski addition of two sets $A, B \subseteq \mathbb{R}^d$,

$$A \oplus B := \{x + y : x \in A, y \in B\}.$$

The Minkowski difference is given by

$$A \ominus B := \bigcap_{y \in B} (A \oplus \{y\}) = (A^c \oplus B)^c,$$

where A^c denotes the complement of A in \mathbb{R}^d . For $r \in \mathbb{R}$, let

$$A_r := \begin{cases} A \oplus B(0, r), & \text{for } r \geq 0, \\ A \ominus B(0, -r), & \text{for } r < 0, \end{cases}$$

denote either the dilation or erosion of a set A , depending on the sign of r . Finally, define $A_{\bullet r} := (A_r)_{-r}$ to be the closing of A by $B(0, r)$ if $r \geq 0$ and the opening of A by $B(0, -r)$ if $r < 0$.

For $r \geq 0$ and $A \subseteq \mathbb{R}^d$ closed, it follows from Definition 1 that A_r (also known as the r -offset of A in, e.g., [19]) denotes all the points in \mathbb{R}^d within a distance r of the set A . The set A_{-r} denotes all the points in A a distance of at least r from ∂A .

With these notions established, the Hausdorff distance between two closed sets $A, B \subseteq \mathbb{R}^d$ is defined as $d_H(A, B) := \inf\{r \in \mathbb{R}^+ : A \subseteq B_r, B \subseteq A_r\}$.

Lemma 1. *Let $r, s > 0$, and let $A \subseteq \mathbb{R}^d$. The following identities hold:*

- (a) $(A_r)^c = (A^c)_{-r}$,
- (b) $A \subseteq A_{\bullet r}$,
- (c) $(A_r)_s = A_{r+s}$,
- (d) $(A_r)_{-s} \supseteq A_{r-s}$,
- (e) $(A_{-r})_s \subseteq A_{s-r}$.

Proof of Lemma 1 Fix $r, s > 0$. The identity in (a) follows directly from Definition 1. Item (b) is proved by contradiction. Let $a \in A$ and suppose $a \in (A_{\bullet r})^c = (A_r)^c \oplus B(0, r)$. Then there is a $p \in (A_r)^c$ such that $a \in B(p, r) \Leftrightarrow \|p - a\|_2 \leq r \Leftrightarrow p \in B(a, r) \Leftrightarrow p \in A_r$; thus, a contradiction. To prove (c), remark that $B(0, r) \oplus B(0, s) = B(0, r+s)$ and that Minkowski addition is associative. To prove (d), consider the case where $r \geq s$, then $(A_r)_{-s} = (A_{r-s+s})_{-s} = ((A_{r-s})_s)_{-s} = (A_{r-s})_{\bullet s} \supseteq A_{r-s}$. For the case $r < s$, write $(A_r)_{-s} = (A_r)_{r-s-r} = ((A_r)_{-r})_{r-s} = (A_{\bullet r})_{r-s} \supseteq A_{r-s}$. To show (e), consider the complements of the sets in (d) and apply (a) repeatedly. \square

2.2 The reach and related concepts

Definition 2 (The reach). Recall from [39] that the *reach* of a set $A \subseteq \mathbb{R}^d$ is given by

$$\text{reach}(A) := \sup \{r \in \mathbb{R}^+ : \forall y \in A_r \exists! x \in A \text{ nearest to } y\}.$$

If $\text{reach}(A) > 0$, then A is said to have *positive reach*.

A useful notion related to the reach of a closed set A is the *medial axis* of A^c , originally proposed in [14].

Definition 3 (The medial axis). Let $O \subseteq \mathbb{R}^d$ be open. Its *medial axis* $\mathcal{M}(O)$ is the set of points in O with at least two closest points in ∂O .

The reach of a closed set A can be alternatively expressed as

$$\text{reach}(A) = \inf\{\|a - x\|_2 : a \in A, x \in \mathcal{M}(A^c)\}. \quad (1)$$

2.2.1 Connections to r -convexity

Definition 4 (r -convexity). A set $A \subseteq \mathbb{R}^d$ is said to be r -convex for $r \in \mathbb{R}^+$ if it is closed and $A_{\bullet s} = A$ for all $s \in (0, r)$ (see, e.g., [46]). Define the quantity $\text{rconv}(A) := \sup\{r \in \mathbb{R} : A_{\bullet r} = A\}$.

An equivalent definition of r -convexity is as follows: a set $A \subseteq \mathbb{R}^d$ is r -convex if and only if it can be expressed as the complement of a union of open balls of radius r .

Theorem 1. *Let A be closed in \mathbb{R}^d . It holds that*

$$\text{reach}(A) \leq \text{rconv}(A). \quad (2)$$

Moreover, if ∂A is a C^1 -smooth $(d-1)$ -dimensional manifold without boundary, then

$$\text{reach}(A) = \text{rconv}(A). \quad (3)$$

Equation (2) is proven in Cuevas et al. [32, Proposition 1] for compact A . Nonetheless, we reprove the statement for closed A in the proof of Theorem 1, which we postpone to Section 5. The novelty in Theorem 1 is that it provides a class of subsets of \mathbb{R}^d for which the reach and r -convexity are equal.

To see that the r -convexity and the reach of a set are indeed distinct notions for general subsets of \mathbb{R}^d , Figure 1 provides an example of a closed set A for which $\text{reach}(A) < \text{rconv}(A)$. Consider also the following remark.

Remark 1. Any closed subset of \mathbb{R}^d contained in a $(d-1)$ -dimensional affine linear subspace is r -convex for all $r > 0$. This is easy to see since the complement of the set in the $(d-1)$ -dimensional affine linear subspace is open, and is the union of open $(d-1)$ -balls of radius less than r . Each $(d-1)$ -ball can be expressed as the intersection of the affine linear subspace with an open ball of radius r in \mathbb{R}^d ; therefore, the closed subset in question can be expressed as the complement of a union of open balls of radius r and is hence r -convex.

Recently, a counterexample to Borsuk's conjecture that r -convex sets are locally contractible, was published in [24]. The conjecture is easily seen to be false by mapping a closed set in \mathbb{R}^{d-1} that is not locally contractible to a $(d-1)$ -dimensional hyperplane in \mathbb{R}^d under an isometry.




Fig. 1: The closed set A (in grey) has $\text{rconv}(A) > \text{reach}(A)$. The set A can be expressed as the complement of a union of open balls of radius $\text{rconv}(A)$ whose centers lie outside the dotted black curve. The dotted red line is the medial axis $\mathcal{M}(A^c)$.

Here, we provide some corollaries to Theorem 1 that provide alternative sets of sufficient conditions for the equality of the reach and the r -convexity. The first applies to sets in Serra's regular model [50, p. 144].

Corollary 1. *Let $A \subset \mathbb{R}^d$ be non-empty, compact, and path-connected. If $\text{rconv}(\text{cl}(A^c)) > 0$, then $\text{reach}(A) = \text{rconv}(A)$.*

The proof of Corollary 1, which relies heavily on Theorem 1 in [53], is postponed to Section 5. Likewise, we prove the following result in Section 5.

Corollary 2. *Let A be a closed set in \mathbb{R}^d and let $\epsilon \in \mathbb{R}^+$. Then $\text{reach}(A_\epsilon) = \text{rconv}(A_\epsilon)$.*

2.2.2 The β -reach

We show in Theorem 2 below, that the reach of a set can be formulated in terms of pairs of points in the set, and the distance to the set from their midpoints. We define a parametrized version of the reach by restricting to pairs of points whose midpoints are sufficiently far from the set. The so-called β -reach is constructed from the following family of functions.

Definition 5 (Spherical cap geometry). Define for $\alpha \in [0, \infty)$ and $x \in [0, \alpha/2]$,

$$g_\alpha(x) := \begin{cases} \frac{\alpha^2}{8x} + \frac{x}{2}, & x > 0, \\ \infty, & x = 0, \end{cases} \quad (4)$$

and its inverse for $r \geq \alpha/2$,

$$g_\alpha^{-1}(r) := r - \sqrt{r^2 - \frac{\alpha^2}{4}}. \quad (5)$$




Fig. 2: $g_\alpha(x)$ in Definition 5 plotted for several values of α .




Fig. 3: A geometric interpretation of $g_\alpha(x)$ in Definition 5.

See Figure 3 for a geometric interpretation of the function in (4) and its inverse in (5). Evidently from the figure, this function is derived from the height of a spherical cap; it is written about in [6, 7, 10, 16, 37] in reference to the reach.

Definition 6 (The β -reach). For a closed set $A \subseteq \mathbb{R}^d$ and $\beta \in [0, \infty)$, let the β -reach of A be defined as

$$\text{reach}_\beta(A) := \inf \left\{ g_{\|a_2 - a_1\|_2}(x) : a_1, a_2 \in A, x = \delta_A \left(\frac{a_1 + a_2}{2} \right) \geq \beta \right\},$$

where $g_{\|a_2 - a_1\|_2}(x)$ is defined in (4). Recall that $\delta_A : \mathbb{R}^d \rightarrow \mathbb{R}$ maps each point in \mathbb{R}^d to its distance from A . See Figure 4 for a visual aid.

By restricting to pairs of points $a_1, a_2 \in A$ in Figure 4 that yield $x \geq \beta$, the β -reach of A is the largest lower bound of the resulting values of $g_{\|a_2 - a_1\|_2}(x)$.




Fig. 4: For $\beta \in [0, \infty)$, the β -reach of A (see Definition 6) is the largest lower bound of $g_{\|a_2-a_1\|_2}(x)$ (see Definition 5), for the pairs $a_1, a_2 \in A$ satisfying $x \geq \beta$.




Fig. 5: The interior of the large ball centered at p does not intersect A , so the distance from the midpoint $(a_1+a_2)/2$ to the set A is at least $g_{\|a_2-a_1\|_2}^{-1}(\delta_A(p))$. This construction is used in the proof of Theorem 2.

If one does not restrict the size of x (i.e., for $\beta = 0$), then this largest lower bound is precisely $\text{reach}(A)$. This is formalized by the following theorem.

Theorem 2. *Let A be closed in \mathbb{R}^d . The map $\beta \mapsto \text{reach}_\beta(A)$ for $\beta \in \mathbb{R}^+$ is non-decreasing; moreover,*

$$\lim_{\beta \searrow 0} \text{reach}_\beta(A) = \text{reach}_0(A) = \text{reach}(A). \quad (6)$$

Proof of Theorem 2 The non-decreasing property is seen immediately via the inclusion

$$\left\{ (a_1, a_2) \in A^2 : \delta_A \left(\frac{a_1 + a_2}{2} \right) \geq \beta_2 \right\} \subseteq \left\{ (a_1, a_2) \in A^2 : \delta_A \left(\frac{a_1 + a_2}{2} \right) \geq \beta_1 \right\},$$

for all $\beta_1, \beta_2 \in \mathbb{R}$ satisfying $\beta_1 < \beta_2$.

Now, we start by proving the second equality in (6). The proof is largely supplied by Lemma 1 of [16] which provides

$$\text{reach}(A) \leq g_{\|a_2 - a_1\|_2} \circ \delta_A \left(\frac{a_1 + a_2}{2} \right), \quad (7)$$

for all $a_1, a_2 \in A$, since $g_\alpha(x)$ is non-increasing in x . What remains to show is that $\text{reach}(A)$ is the largest lower bound in (7); *i.e.*, if $r > \text{reach}(A)$ then there exist $a_1, a_2 \in A$ such that r exceeds the right-hand side of (7). Let $r > \text{reach}(A)$ and let $\tilde{r} \in (\text{reach}(A), r)$. By the definition of $\text{reach}(A)$ (Definition 2), $\exists p \in A_{\tilde{r}}$ and $a_1, a_2 \in A$ such that $\|a_1 - p\|_2 = \|a_2 - p\|_2 = \delta_A(p) \leq \tilde{r}$. Since the interior of $B(p, \delta_A(p))$ does not intersect A , we have

$$\delta_A \left(\frac{a_1 + a_2}{2} \right) \geq g_{\|a_2 - a_1\|_2}^{-1}(\delta_A(p)),$$

(see Figure 5) and

$$g_{\|a_2 - a_1\|_2} \circ \delta_A \left(\frac{a_1 + a_2}{2} \right) \leq \delta_A(p) \leq \tilde{r} < r.$$

This proves that $\text{reach}_0(A) = \text{reach}(A)$.

Finally, we show the first equality in (6). The non-decreasing property gives $\lim_{\beta \searrow 0} \text{reach}_\beta(A) \geq \text{reach}_0(A)$. Now, it suffices to show the reverse inequality. For any $\epsilon > 0$, there exists $a_1, a_2 \in A$ satisfying

$$g_{\|a_2 - a_1\|_2} \circ \delta_A \left(\frac{a_1 + a_2}{2} \right) < \text{reach}_0(A) + \epsilon.$$

Any such pair (a_1, a_2) must satisfy $\delta_A \left(\frac{a_1 + a_2}{2} \right) > 0$, and so for $\beta \in (0, \delta_A \left(\frac{a_1 + a_2}{2} \right))$, it holds that $\text{reach}_\beta(A) < \text{reach}_0(A) + \epsilon$. Thus, $\lim_{\beta \searrow 0} \text{reach}_\beta(A) \leq \text{reach}_0(A)$. \square

Remark 2. For any $\alpha \in [0, \infty)$ and $x \in [0, \alpha/2]$, one has $g_\alpha(x) \geq x$ (see Figure 2). Thus, for any closed set A and $\beta \in \mathbb{R}^+$,

$$\text{reach}_\beta(A) \geq \beta.$$

Intuitively, this means that the β -reach excludes small, reach-limiting features of the set that have scale less than β .

The values of β for which $\text{reach}_\beta(A) = \beta$ are the distances from the critical points of the generalized gradient function [18] to the set A .

Remark 3. Other generalizations of the reach are constructed similarly to the β -reach, in that they formulate the reach as an infimum or supremum over some set, and add or remove elements in the set using some real parameter (in our case, β). For example, [16, Theorem 1] expresses the reach of a set $A \subset \mathbb{R}^d$ as the supremum over a subset of \mathbb{R}^+ that satisfies a certain condition relating to A . By “continuously” weakening the condition with a parameter $\delta \in [0, \infty)$, the authors in [1] introduce the *spherical distortion radius* as the supremum of the larger subset of \mathbb{R}^+ satisfying the weaker condition parametrized by δ . The spherical distortion radius is identified with the reach for $\delta = 0$. In an earlier work, [18] introduces the μ -reach by considering the shortest distance from an element in the set to the μ -medial axis, a filtered version of the medial

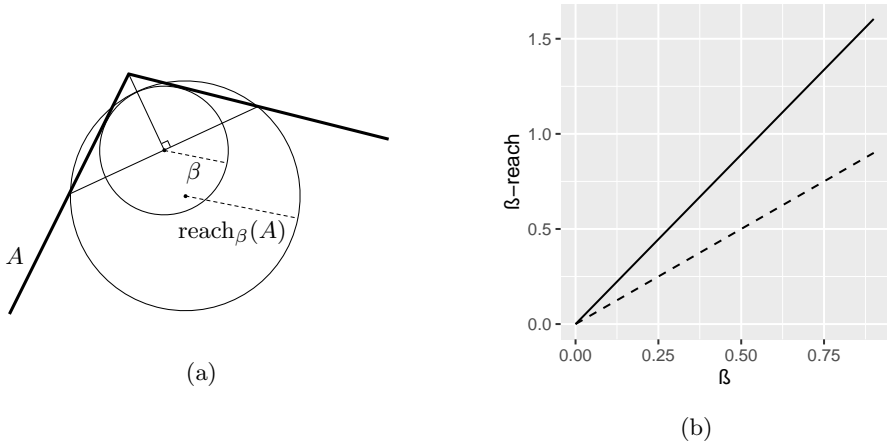


Fig. 6: **(a):** The construction of $\text{reach}_\beta(A)$ for $\beta > 0$ with A , the union of two rays in \mathbb{R}^d that originate from the same point. **(b):** The β -reach profile of A shown as a solid line. The lower bound with slope 1 is shown as a dashed line.

axis by considering regions where the generalized gradient function [41] is less than some $\mu \in (0, 1]$. The β -reach, the μ -reach, and the spherical distortion radius, are all generalizations of the reach that exclude “small-scale” features as decided by the corresponding parameter β , μ , or δ .

One significant advantage of the β -reach is its computability for high-dimensional point cloud data (see Section 3.3). In this setting, the formulation of the reach in terms of the β -reach for $\beta = 0$ can also be used to construct an upper bound for the reach of any compact $A \subset \mathbb{R}^d$ (see Section 3.2).

For a closed set $A \subseteq \mathbb{R}^d$, one can study the map $\beta \mapsto \text{reach}_\beta(A)$ for $\beta \in [0, \infty)$, which we refer to as the β -reach profile of the set A . We will see in Section 3 that a set’s β -reach profile provides pertinent information regarding the estimation of the reach from point cloud data, especially through its first-order approximation at $\beta = 0$.

Some exemplary sets and their β -reach profiles are considered in the several examples that follow.

Example 1 (The β -reach of a corner). For two line segments in \mathbb{R}^d with ends joined by an angle $\theta \in (0, \pi)$, the β -reach of their union is $\frac{\beta}{2} (1 + \sec^2(\frac{\theta}{2}))$, for sufficiently small $\beta > 0$. See Figure 6 for an illustration.

Example 2 (The β -reach of an arc). The β -reach of a circular arc of angle at most π is equal to the radius of the arc, for sufficiently small $\beta > 0$.

Example 3 (The β -reach of a bottleneck structure). If the reach of a set $A \subset \mathbb{R}^d$ is determined by a bottleneck structure like those described in [2],




Fig. 7: **(a):** The construction of $\text{reach}_\beta(A)$ for $\beta > 0$ with A , a paraboloid embedded in \mathbb{R}^d . **(b):** The β -reach profile of A shown as a solid line. The lower bound with slope 1 is shown as a dashed line.

then $\text{reach}_\beta(A) = \text{reach}(A)$ for $\beta \in [0, \text{reach}(A)]$. This is the case, for example, when A is a finite set of points in \mathbb{R}^d .

Example 4 (The β -reach of a C^2 -smooth curve). Let $h : \mathbb{R} \rightarrow \mathbb{R}$ be a C^2 -smooth function with $h'(0) > 0$. Suppose that the graph of the function $f : [-1, 1] \rightarrow \mathbb{R}$ defined by $f(x) := h(x^2)$ obtains its maximal curvature at $x = 0$. Then, the graph $A := \{(x, f(x)) : -1 \leq x \leq 1\} \subset \mathbb{R}^2$ satisfies

$$\text{reach}(A) = \frac{1}{2h'(0)} \quad (8)$$

and

$$\text{reach}_\beta(A) = \frac{1}{2h'(0)} + \left(\frac{1}{2} - \frac{h''(0)}{4h'(0)^3} \right) \beta + o(\beta), \quad (9)$$

for sufficiently small $\beta > 0$. Justifications for Equations (8) and (9) are provided at the end of Section 5. Figure 7 depicts a special case of this example with $h(x) = x/2$.

3 Methods for point cloud data

In practice, one might be interested in identifying bounds on the reach and r -convexity of a set $A \subseteq \mathbb{R}^d$ given a discrete set of points that are known to reside in A . The goal of this section is to provide computational methods for bounding the reach and the r -convexity from above, and producing diagnostics for possible approximations. The three main settings that we consider are:

- (a) One has access to a point cloud that extends over \mathbb{R}^d in the sense that \mathbb{R}^d can be covered by balls of fixed radius ϵ centered at each point. Moreover, one knows the partition of the points that lie in $A \subseteq \mathbb{R}^d$, and those that lie in A^c . See section 3.1 for a treatment of this setting.
- (b) One has access to a set of points for which it is known that $A \subseteq \mathbb{R}^d$ can be covered by balls of fixed radius ϵ centered at each point. Here, ϵ is known. See section 3.2.
- (c) The set A is a submanifold of \mathbb{R}^d of dimension $m < d$. One has access to a set of points that is known to be contained in A . See section 3.3.

The mathematical results that we present in this section hold in arbitrary dimension d . Nevertheless, the computational complexity of the method described in Section 3.1 for setting (a) increases quite drastically as higher dimensions are considered.

For the method that we present for setting (b) described in Section 3.2, its computation time depends on the ambient dimension only through the computation of distances in \mathbb{R}^d , which is linear in d . The algorithm is largely dependent on the number of points used, which may be large in high dimension to ensure small ϵ .

The computational complexity of visualizing the β -reach profile of the manifold in setting (c) is also linear in the dimension of the ambient space, and so the algorithm that we describe in Section 3.3 can adapt to large values of d . Contrast this with existing methods that aim to approximate the reach by first approximating the medial axis [18, 21, 33, 35, 36, 42], where in high dimension, accessing the medial axis becomes computationally challenging. The relationship between d , m , n , and our method's performance is elaborated on in Remark 10 in Section 3.3.

3.1 An upper bound for the r -convexity of a set

In this section, we introduce a method for identifying when sets are not r -convex, for $r > 0$, with a true negative rate of 100%. That is, given a set $A \subset \mathbb{R}^d$ that is observed over some grid of points (possibly lacking structure), we show how to correctly identify for which values of $r > 0$ the set A is *certainly not* r -convex. The smallest of these values of r provides an upper bound for $\text{rconv}(A)$, and by Equation (2), an upper bound for $\text{reach}(A)$. Moreover, if the sampling becomes dense in \mathbb{R}^d , we show that this upper bound converges to $\text{rconv}(A)$.

Here, we define operations analogous to dilation and erosion, for discrete sets of points.

Definition 7. Let Φ be a point cloud in \mathbb{R}^d , *i.e.*, a countable subset of \mathbb{R}^d . For a set $\hat{A} \subseteq \Phi$, we make a slight abuse of notation and denote for $r \in \mathbb{R}$,

$$\hat{A}_r := \begin{cases} \{p \in \Phi : \delta_{\hat{A}}(p) \leq r\}, & \text{for } r \geq 0, \\ \{a \in \hat{A} : \delta_{\Phi \setminus \hat{A}}(a) > -r\}, & \text{for } r < 0. \end{cases}$$




Fig. 8: An example where $(\widehat{A}_r)_{-r}$ in (c) is strictly larger than \widehat{A} in (a), while the underlying set A is r -convex. The boundary ∂A is shown as a dashed line and r is three times the pixel width.

Remark 4. We emphasize that for a point cloud Φ , a subset $\widehat{A} \subseteq \Phi$, and a real number $r \in \mathbb{R}$, Definition 7 implies that $\widehat{A}_r \subseteq \Phi$. Contrast this with the set $\widehat{A} \oplus B(0, |r|)$, which is not contained in Φ for $r > 0$.

Let us illustrate the operations in Definition 7 using the example in Figure 8. The set $A \subset \mathbb{R}^2$ occupies the bottom left corner of the domain, just until the dashed line. When sampled on a square lattice Φ , one obtains the image in panel (a); each pixel is centered on a point in Φ , and grey if the point lies in A . Dilating the grey pixels \widehat{A} by $r = 3 \times (\text{pixel width})$, one obtains the grey pixels in panel (b), \widehat{A}_r . If one then erodes \widehat{A}_r by r (equivalent to dilating $\Phi \setminus \widehat{A}_r$ by r and taking complements), one obtains the grey pixels in panel (c), $(\widehat{A}_r)_{-r}$. Notice, however, that there is a point in $(\widehat{A}_r)_{-r}$ that is not in \widehat{A} . This might be surprising since we chose a set A such that $A_{\bullet, r} = A$. This illustration shows that naively testing for r -convexity using the discrete dilation and erosion operations *mistakenly classifies sets as not r -convex* when indeed they are. The following theorem shows how to correctly identify sets as not being r -convex, and gives an upper bound for the r -convexity of a set that is tight in some sense.

Theorem 3. For $n \in \mathbb{N}^+$, let $\Phi^{(n)}$ be a point cloud in \mathbb{R}^d , and suppose that

$$\epsilon_n := \sup \{ \delta_{\Phi^{(n)}}(q) : q \in \mathbb{R}^d \} \quad (10)$$

is finite. Let A be a compact subset of \mathbb{R}^d , and for $n \in \mathbb{N}^+$, define $\widehat{A}^{(n)} := A \cap \Phi^{(n)}$ and the corresponding bound,

$$\widehat{r}^{(\epsilon_n)}(\widehat{A}^{(n)}) := \inf \left\{ r > \epsilon_n : (\widehat{A}_{r-\epsilon_n}^{(n)})_{-(r+\epsilon_n)} \not\subseteq \widehat{A}^{(n)} \right\}. \quad (11)$$




Fig. 9: (a): A set A is shown as a pixelated image in grey. The set of grey pixel centers is $\widehat{A}^{(n)}$. (b): The set $(\widehat{A}^{(n)})_{r-\epsilon_n}-(r+\epsilon_n)$ contains the dark grey pixels, which are not contained in $\widehat{A}^{(n)}$. Thus, $r \geq \text{rconv}(A)$. Here, r is nine times the pixel width.

It holds that

$$\inf_{n \in \mathbb{N}^+} \widehat{r}^{(\epsilon_n)}(\widehat{A}^{(n)}) \geq \text{rconv}(A). \quad (12)$$

Moreover, if $\epsilon_n \rightarrow 0$ and $d_H(\widehat{A}^{(n)}, A) \rightarrow 0$ as $n \rightarrow \infty$, then

$$\lim_{n \rightarrow \infty} \widehat{r}^{(\epsilon_n)}(\widehat{A}^{(n)}) = \text{rconv}(A). \quad (13)$$

Theorem 3 provides a method for correctly identifying which subsets of \mathbb{R}^d are not r -convex. If a dilation of the the discretization $\widehat{A}^{(n)}$ by $r - \epsilon_n$ followed by an erosion of $r + \epsilon_n$ produces a set that is not contained in $\widehat{A}^{(n)}$, then $\text{rconv}(A) \leq r$. See Figure 9 for an example of a set A for which $(\widehat{A}^{(n)})_{r-\epsilon_n}-(r+\epsilon_n) \not\subseteq \widehat{A}^{(n)}$. The precise regions where r -convexity does not hold are highlighted by the method.

Remark 5. Note that $\widehat{r}^{(\epsilon_n)}(\widehat{A}^{(n)})$ in (11) can be computed using entirely available information, since ϵ_n in (10) is a feature of the point cloud $\Phi^{(n)}$ and not of the unknown set A . In particular, there is no need to estimate $d_H(\widehat{A}^{(n)}, A)$ for the construction of $\widehat{r}^{(\epsilon_n)}(\widehat{A}^{(n)})$. A binary search algorithm, along with the discrete dilation operations in Definition 7, are sufficient for the numerical calculation of $\widehat{r}^{(\epsilon_n)}(\widehat{A}^{(n)})$.

Remark 6. The requirement that $\Phi^{(n)}$ extends over all of \mathbb{R}^d (see Equation (10)) allows for a mathematical simplification and is not needed in practice. In real applications, one only requires that $d_H(\Phi^{(n)}, T) = \epsilon_n$ for some compact $T \subset \mathbb{R}^d$ that contains $A_{\text{rconv}(A)+2\epsilon_n}$. In the case where A is lower dimensional than the embedding space \mathbb{R}^d , one can add any finite number of points uniformly distributed on T to $\Phi^{(n)}$, and they will almost surely not belong to A .




Fig. 10: A discrete dilation of the set $\widehat{A}^{(n)}$ (black points) by the length of the green arrows ($\text{rconv}(A) - \epsilon_n$) followed by a discrete erosion by the length of the longer blue arrow ($\text{rconv}(A) + \epsilon_n$) leaves all of the white points out of the resulting set. The boundary of $A_{\text{rconv}(A) - \epsilon_n}$ is shown as a dashed line.

The proof of Theorem 3 is quite technical, and so we postpone it to Section 5. Nevertheless, Figure 10 provides an intuitive illustration that helps understand Equation (12). This is an example of a worst-case scenario, in that a dilation of $\widehat{A}^{(n)}$ by any more than $\text{rconv}(A) - \epsilon_n$ results in all of $\Phi^{(n)}$ being consumed. The dilation radius $\text{rconv}(A) - \epsilon_n$ is maximal in the sense that, by dilating any less, there is guaranteed to be at least one element of $\Phi^{(n)}$ that is not in the dilation. The distance between this one remaining element of $\Phi^{(n)}$ and the other points in $\Phi^{(n)} \setminus \widehat{A}^{(n)}$ might be anywhere up to $r + \epsilon_n$, where $r - \epsilon_n \in (0, \text{rconv}(A) - \epsilon_n)$ is the dilation radius. Therefore, an erosion by at least $r + \epsilon_n$ is necessary to guarantee that the resulting set is a subset of $\widehat{A}^{(n)}$. In this sense, the bound in (12) is tight.

The following result from [49], while interesting on its own, is instrumental in the proof of Equation (13).

Proposition 1 (Lemma 8.3 in [49]). *Let $r \in \mathbb{R}^+$ and let $A \subset \mathbb{R}^d$ be a closed set satisfying $\text{rconv}(A) < r$. The set $A_{\bullet, r} \setminus A$ contains an open subset of \mathbb{R}^d .*

The relationship between Proposition 1 and Equation (13) is that, for $r > \text{rconv}(A)$, as the point cloud $\Phi^{(n)}$ becomes more dense in \mathbb{R}^d , the open subset in $A_{\bullet, r} \setminus A$ fills with points that remain in $(\widehat{A}_{r - \epsilon_n}^{(n)})_{-(r + \epsilon_n)}$ for ϵ_n sufficiently small. Thus,

$$\lim_{n \rightarrow \infty} \widehat{r}^{(\epsilon_n)}(\widehat{A}^{(n)}) < r,$$

and since $r \in (\text{rconv}(A), \infty)$ is arbitrary, the limit is at most $\text{rconv}(A)$. Equality then holds by Equation (12). The full proof of Theorem 3 and an alternative proof of Proposition 1 are provided in Section 5.

Remark 7. Equation (13) is provided without a rate of convergence. This is due to the fact that Proposition 1 provides no guarantees on the size of the open subset that can be found in $A_{\bullet, r} \setminus A$ for $r > \text{rconv}(A)$. To provide a rate of

convergence, a deeper analysis is needed to understand the rate at which the size of the largest open ball in $A_{\bullet r} \setminus A$ decreases as $r \searrow \text{rconv}(A)$. Conversely, the bound that we construct for the reach in Section 3.2 converges to the reach at a known rate (see Theorem 4 below).

3.2 An upper bound for the reach of a set

We have already seen that $\widehat{\text{rch}}^{(\epsilon_n)}(\widehat{A}^{(n)})$ in (11) is an upper bound for the reach by Equations (2) and (12). In this section, we introduce another computable bound for the reach based on the expression for the β -reach in Definition 6. The context in which one can apply this bound was introduced as setting (b) at the start of Section 3: a countable set of points $\widehat{A}^{(n)}$ is known to be included in a compact set $A \subset \mathbb{R}^d$, and its Hausdorff distance to A is known to be at most ϵ_n .

A very weak regularity condition is imposed on the set A in terms of its β -reach profile near 0; it is used to prove the rate at which our upper bound converges to $\text{reach}(A)$ as $\widehat{A}^{(n)} \rightarrow A$ in the Hausdorff metric.

Assumption 1. Suppose that the set $A \subset \mathbb{R}^d$ is compact, and that there exists $\delta > 0$ such that the map $\beta \mapsto \text{reach}_\beta(A)$ for $\beta \in \mathbb{R}$ is either constant or strictly increasing on $[0, \delta]$. In addition, suppose that

$$K_A := \lim_{\beta \searrow 0} \frac{\text{reach}_\beta(A) - \text{reach}(A)}{\beta} \quad (14)$$

exists and is finite.

Assumption 1 is imposed on the β -reach of A , as opposed to the regularity of ∂A , since the β -reach lends very naturally to the proof of Theorem 4. Moreover, Assumption 1 allows for most compact sets in \mathbb{R}^d , excluding some pathological counterexamples such as hypothetical sets admitting the β -reach profiles in Figure 11. Assumption 1 is seen to hold in all of Examples 1 through 4 in Section 2.2.2.

Theorem 4. Let $A \subseteq \mathbb{R}^d$ be closed. For each $n \in \mathbb{N}^+$, let $\widehat{A}^{(n)}$ be a countable subset of A . Suppose that $\widehat{A}^{(n)} \rightarrow A$ in the Hausdorff metric. i.e., there is a sequence $(\epsilon_n)_{n \geq 1}$ in \mathbb{R}^+ tending to 0 as $n \rightarrow \infty$ such that $\epsilon_n \geq d_H(\widehat{A}^{(n)}, A)$ for all $n \in \mathbb{N}^+$. For each $n \in \mathbb{N}^+$, define the corresponding bound on $\text{reach}(A)$ by

$$\widehat{\text{rch}}^{(\epsilon_n)}(\widehat{A}^{(n)}) := \inf \left\{ g_{\|a_2 - a_1\|_2}(x - \epsilon_n) : a_1, a_2 \in \widehat{A}^{(n)}, \right. \quad (15)$$

$$\left. x = \delta_{\widehat{A}^{(n)}} \left(\frac{a_1 + a_2}{2} \right) \geq \epsilon_n \right\},$$




Fig. 11: Hypothetical β -reach profiles of sets that would *not* satisfy Assumption 1. **(a):** The β -reach profile is neither constant nor strictly increasing on any neighbourhood of 0. **(b):** The limit in (14) is infinite.

where $g_{\|a_2 - a_1\|_2}$ is defined in (4). Then,

$$\widehat{\text{rch}}^{(\epsilon_n)}(\widehat{A}^{(n)}) \geq \text{reach}(A), \quad (16)$$

for $n \in \mathbb{N}^+$. Furthermore, if A satisfies Assumption 1 and is not convex, then there exists $n_0 \in \mathbb{N}^+$ such that

$$\widehat{\text{rch}}^{(\epsilon_n)}(\widehat{A}^{(n)}) - \text{reach}(A) \leq (2 \text{reach}(A) + K_A) \sqrt{\epsilon_n}, \quad (17)$$

for all $n \geq n_0$, where K_A is defined as in (14).

Proof of Theorem 4 We begin by showing (16) in the case where A is not convex (for A convex, the result is trivial since $\widehat{\text{rch}}^{(\epsilon_n)}(\widehat{A}^{(n)}) = \text{reach}(A) = \infty$ [39]). Suppose that for some fixed $n \in \mathbb{N}^+$, we have $\epsilon_n \geq d_H(\widehat{A}^{(n)}, A)$. We write,

$$\begin{aligned} \text{reach}_0(A) &= \inf \left\{ g_{\|a_2 - a_1\|_2} \circ \delta_A \left(\frac{a_1 + a_2}{2} \right) : a_1, a_2 \in A \right\} \\ &\leq \inf \left\{ g_{\|a_2 - a_1\|_2} \circ \delta_A \left(\frac{a_1 + a_2}{2} \right) : a_1, a_2 \in \widehat{A}^{(n)} \right. \\ &\quad \left. \delta_{\widehat{A}^{(n)}} \left(\frac{a_1 + a_2}{2} \right) \geq \epsilon_n \right\}. \end{aligned} \quad (18)$$

The equality in (18) is an application of Definition 6, and the inequality holds since the infimum is taken over a smaller subset. For any $a_1, a_2 \in A$, there exists a projection of $\frac{a_1 + a_2}{2}$ onto A , namely $a^\pi \in A$, satisfying

$$\left\| a^\pi - \frac{a_1 + a_2}{2} \right\|_2 = \delta_A \left(\frac{a_1 + a_2}{2} \right).$$




Fig. 12: Visual aid for Equation (21). The real numbers $\delta_{\hat{A}^{(n)}}(\frac{p_1+p_2}{2})$ and $\delta_A(\frac{p_1+p_2}{2})$ are contained in their respective intervals, positioned relative to $\delta_A(\frac{a_1+a_2}{2})$. Ticks have a spacing of ϵ_n .

By the triangle inequality and the fact that $\epsilon_n \geq d_H(\hat{A}^{(n)}, A)$,

$$\delta_{\hat{A}^{(n)}}\left(\frac{a_1+a_2}{2}\right) \leq \delta_{\hat{A}^{(n)}}(a^\pi) + \left\|a^\pi - \frac{a_1+a_2}{2}\right\|_2 \leq \epsilon_n + \delta_A\left(\frac{a_1+a_2}{2}\right).$$

Rearranging gives,

$$\delta_{\hat{A}^{(n)}}\left(\frac{a_1+a_2}{2}\right) - \epsilon_n \leq \delta_A\left(\frac{a_1+a_2}{2}\right). \quad (19)$$

Since $g_{\|a_2-a_1\|_2}$ is non-increasing, the rightmost expression in (18) cannot decrease when $\delta_A(\frac{a_1+a_2}{2})$ is replaced by $\delta_{\hat{A}^{(n)}}(\frac{a_1+a_2}{2}) - \epsilon_n$. Thus,

$$\widehat{\text{rch}}^{(\epsilon_n)}(\hat{A}^{(n)}) \geq \text{reach}_0(A) = \text{reach}(A),$$

where the latter equality is an application of Theorem 2.

Now, we proceed to show (17). Recall the equality in (18), and split the analysis into two cases.

Case 1: There exists $a_1, a_2 \in A$ such that $a_1 \neq a_2$ and $\text{reach}(A) = g_{\|a_2-a_1\|_2} \circ \delta_A(\frac{a_1+a_2}{2})$.

For this pair $a_1, a_2 \in A$ there exists $p_1, p_2 \in \hat{A}^{(n)}$ satisfying $\|p_2 - a_2\|_2, \|p_1 - a_1\|_2 \leq \epsilon_n$. This implies

$$\left\|\frac{p_1+p_2}{2} - \frac{a_1+a_2}{2}\right\|_2 \leq \epsilon_n. \quad (20)$$

Equation (20) tells us that the two midpoints are close, and so their distances to A cannot differ by more than the distance between them. *i.e.*,

$$\left|\delta_A\left(\frac{a_1+a_2}{2}\right) - \delta_A\left(\frac{p_1+p_2}{2}\right)\right| \leq \epsilon_n.$$

Since Equation (19) holds for $p_1, p_2 \in A$, and $\delta_A(\frac{p_1+p_2}{2}) \leq \delta_{\hat{A}^{(n)}}(\frac{p_1+p_2}{2})$, we also have

$$\delta_A\left(\frac{a_1+a_2}{2}\right) - \epsilon_n \leq \delta_{\hat{A}^{(n)}}\left(\frac{p_1+p_2}{2}\right) \leq \delta_A\left(\frac{a_1+a_2}{2}\right) + 2\epsilon_n, \quad (21)$$

by the triangle inequality (see Figure 12).

Now, suppose that n is sufficiently large such that $2\epsilon_n < y := \delta_A(\frac{a_1+a_2}{2})$. Then, by (21), $x := \delta_{\hat{A}^{(n)}}(\frac{p_1+p_2}{2}) - \epsilon_n > 0$ and so

$$\widehat{\text{rch}}^{(\epsilon_n)}(\hat{A}^{(n)}) \leq g_{\|p_2-p_1\|_2}(x) = \frac{\|p_2-p_1\|_2^2}{8x} + \frac{x}{2}$$

$$\begin{aligned}
 &\leq \frac{(\|a_2 - a_1\|_2 + 2\epsilon_n)^2}{8x} + \frac{x}{2} \\
 &\leq \frac{(\|a_2 - a_1\|_2 + 2\epsilon_n)^2}{8(y - 2\epsilon_n)} + \frac{y + \epsilon_n}{2}.
 \end{aligned} \tag{22}$$

The second inequality in (22) holds since $\|p_2 - p_1\|_2 \leq \|a_2 - a_1\|_2 + 2\epsilon_n$, and the final inequality holds by Equation (21).

Recall that $\text{reach}(A) = \frac{\|a_2 - a_1\|_2^2}{8y} + \frac{y}{2}$. The difference between this and the final expression in (22) is of the order $\mathcal{O}(\epsilon_n)$, which is stronger than required.

Case 2: $\text{reach}(A) < g_{\|a_2 - a_1\|_2} \circ \delta_A \left(\frac{a_1 + a_2}{2} \right)$ for all distinct $a_1, a_2 \in A$.

By Assumption 1, for all sufficiently large n , the function $\text{reach}_\beta(A)$ is strictly increasing for β in a neighbourhood of $\beta_n := \sqrt{\epsilon_n} + 2\epsilon_n$. For each of these values of n , define

$$\Psi_n := \{(a_1, a_2) \in A^2 : \delta_A \left(\frac{a_1 + a_2}{2} \right) \geq \beta_n\},$$

which is compact by Assumption 1, and the continuous map $f_n : \Psi_n \rightarrow \mathbb{R}^+$ by

$$f_n((\psi_1, \psi_2)) = g_{\|\psi_2 - \psi_1\|_2} \circ \delta_A \left(\frac{\psi_1 + \psi_2}{2} \right).$$

Remark that $\text{reach}_{\beta_n}(A) = \inf_{\psi \in \Psi_n} f_n(\psi)$, and so there is an element $\psi_n^* := (\psi_{1n}^*, \psi_{2n}^*) \in \Psi_n$ for which

$$\text{reach}_{\beta_n}(A) = f_n(\psi_n^*).$$

By Definition 6, $\text{reach}_\beta(A) \leq f_n(\psi_n^*)$ for $\beta = \delta_A \left(\frac{\psi_{1n}^* + \psi_{2n}^*}{2} \right)$. The β -reach profile of A is thus constant on the interval $[\beta_n, \delta_A \left(\frac{\psi_{1n}^* + \psi_{2n}^*}{2} \right)]$ by Theorem 2, but by Assumption 1, any such interval must have length 0. Therefore,

$$\beta_n = \delta_A \left(\frac{\psi_{1n}^* + \psi_{2n}^*}{2} \right).$$

Again by Assumption 1,

$$\text{reach}_{\beta_n}(A) = \frac{\|\psi_{2n}^* - \psi_{1n}^*\|_2^2}{8\beta_n} + \frac{\beta_n}{2} = \text{reach}(A) + K_A \beta_n + o(\beta_n).$$

Therefore,

$$\begin{aligned}
 \|\psi_{2n}^* - \psi_{1n}^*\|_2^2 &= 8\beta_n \text{reach}(A) + (8K_A - 4)\beta_n^2 + o(\beta_n^2) \\
 &= 8\sqrt{\epsilon_n} \text{reach}(A) + (16 \text{reach}(A) + 8K_A - 4)\epsilon_n + o(\epsilon_n).
 \end{aligned}$$

As seen in Case 1, Equation (22),

$$\begin{aligned}
 \widehat{\text{rch}}^{(\epsilon_n)}(\widehat{A}^{(n)}) &\leq \frac{(\|\psi_{2n}^* - \psi_{1n}^*\|_2 + 2\epsilon_n)^2}{8(\beta_n - 2\epsilon_n)} + \frac{\beta_n + \epsilon_n}{2} \\
 &= \text{reach}(A) + (2 \text{reach}(A) + K_A) \sqrt{\epsilon_n} + o(\sqrt{\epsilon_n}).
 \end{aligned}$$

□

Remark 8. The split of the proof of Theorem 4 into two cases is natural in the context of reach estimation. Recall Theorem 3.4 in [2]: For a compact set A with $0 < \text{reach}(A) < \infty$, there is either a bottleneck structure (two distinct points $a_1, a_2 \in A$ such that $\|a_2 - a_1\|_2 = 2\delta_A \left(\frac{a_1 + a_2}{2} \right) = 2 \text{reach}(A)$) or an arc-length parametrized geodesic of A with curvature $1/\text{reach}(A)$. The convergence rate developed in Case 1 applies to sets whose reach is decided by a bottleneck structure, and that of Case 2 applies to sets whose reach is decided by a region of high curvature.

Remark 9. In [1], the authors provide a minimax optimal rates of convergence for statistical estimators of the reach. The authors proceed by assuming that $A \subset \mathbb{R}^d$ is a C^k -smooth manifold without boundary, with $k \geq 3$. The minimax rates in [1, Theorem 6.6] are expressed in terms of k , the number of continuous derivatives of the underlying manifold.

It is not surprising that our bound on the reach—for sets satisfying the very general Assumption 1—does not converge to the reach at the optimal rates given in [1]. Nevertheless, we do observe a similar phenomenon in that the rate of convergence derived for a set whose reach is determined by a bottleneck structure ($\mathcal{O}(\epsilon_n)$ in Case 1 in the proof of Theorem 4) is faster than the rate when the reach is determined by a region of high curvature ($\mathcal{O}(\sqrt{\epsilon_n})$ in Case 2).

3.3 The β -reach profiles of high-dimensional point clouds

Here, we tackle the problem of approximating the reach of smooth manifolds of low dimension embedded in high-dimensional Euclidean space. As discussed in the introduction, this setting is the focus of many works. Given a C^1 -smooth manifold $M \subset \mathbb{R}^d$ of dimension $m < d$, we suppose that one has access to a set of points $\widehat{M} \subset M$ from which one would like to infer $\text{reach}(M)$.

Recall from Definition 6 that the β -reach of a point cloud is easily computable, and the computation time does not heavily depend on the dimension d of the ambient space. The β -reach profile of a point cloud can be obtained by considering pairs of points in the point cloud and computing the distance from their midpoint to the closest neighbour in the point cloud. For a fixed number of points, the computation time for calculating the *exact* β -reach profile of the point-cloud scales linearly in d . Of course, the *beta*-reach profile of \widehat{M} is not identical to that of M , however, we argue here that it provides a good approximation for the values of β away from 0.

Recall that the bounds developed in Sections 3.1 and 3.2 are for subsets of \mathbb{R}^d that possibly have positive d -volume. This means that a point cloud \widehat{A} contained in a set $A \subset \mathbb{R}^d$ can recede a distance $\epsilon := d_H(\widehat{A}, A)$ from the topological boundary of A . Thus, the distance from a point in A^c to A can be up to ϵ *shorter* than the distance to the nearest point in \widehat{A} . Conversely, for a smooth manifold M embedded in \mathbb{R}^d , projection of an element $p \in \mathbb{R}^d$ onto M is along a vector normal to M at the projection in M . Thus, the distance from p to its nearest point in \widehat{M} is more similar to the length of the projection vector, so in this setting, the bounds constructed previously are overly robust. Moreover, existing algorithms for computing a mesh from point cloud data can be leveraged to improve quality of the estimate of the distance to the nearest point in M , thus improving the estimate of the β -reach.

Example 5 (3-dimensional point cloud on a 2-dimensional manifold). Let $M \subset \mathbb{R}^3$ be a section of a two-dimensional paraboloid, where its reach is decided by a point of maximal curvature at the vertex ($\text{reach}(M) = 4$). Let $\widehat{M} \subset M$ be a realization of 1500 points uniformly distributed on M shown in

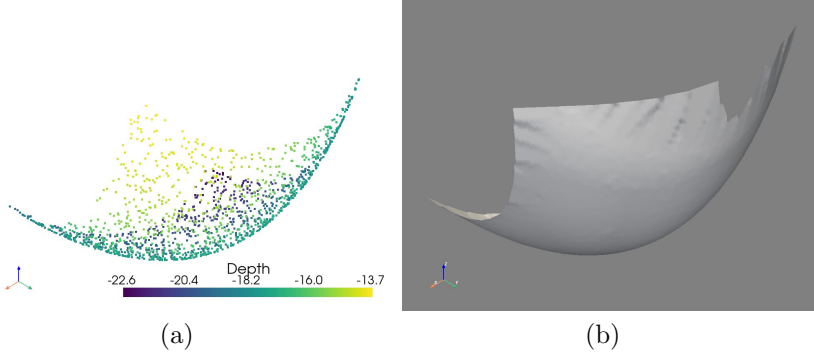


Fig. 13: (a): The point cloud \widehat{M} in Example 5. (b): The mesh $\mathcal{M}(\widehat{M})$ computed with the Python library `pyvista`.




Fig. 14: The exact β -reach profile (see Definition 6) of \widehat{M} in Example 5 and Figure 13 (a) is shown as the blue solid curve. Using the formulation in (23), and the mesh in Figure 13 (b), one may compute an improved approximation (green dashed curve) to the β -reach of the underlying manifold M (red dotted curve).

Figure 13 (a). Using the Python library `pyvista` [51], we construct a mesh $\mathcal{M}(\widehat{M})$ over the point cloud \widehat{M} (see Figure 13 (b)). This allows for an improved approximation of the β -reach profile of M ,

$$\text{reach}_{\beta}^{\mathcal{M}}(\widehat{M}) := \inf \left\{ g_{\|a_2 - a_1\|_2}(x) : a_1, a_2 \in \widehat{M}, \right. \quad (23)$$

$$\left. x = \delta_{\mathcal{M}(\widehat{M})} \left(\frac{a_1 + a_2}{2} \right) \geq \beta \right\}.$$




Fig. 15: The exact β -reach profile of \widehat{M} in Example 6 is shown as the blue solid curve. The β -reach of the underlying manifold M (a 3-dimensional paraboloid embedded in \mathbb{R}^d) is shown as a red dotted curve.

The β -reach of M can be computed exactly as $\text{reach}_\beta(M) = 4 + \frac{\beta}{2}$, for β between 0 and some positive constant. As seen in Figure 14, this is well approximated by $\text{reach}_\beta(\widehat{M})$ for $\beta > 1.3$, and by $\text{reach}_\beta^{\mathcal{M}}(\widehat{M})$ for $\beta > 0.4$; both approximations are efficiently computable from the point cloud \widehat{M} .

To estimate $\text{reach}(M)$ from the β -reach profiles in Example 5, one can perform a linear regression on $\text{reach}_\beta^{\mathcal{M}}(\widehat{M})$, for β in the range where the curve appears linear, and estimate the reach as the model intercept.

Although the manifold M in Example 5 is derived from a simple paraboloid, this example is representative of other two-dimensional manifolds whose reach is determined by a region of high curvature. The difference might be that the linear term in the first order approximation of the β -reach profile at $\beta = 0$ does not have a coefficient of $1/2$ (see Example 4).

The following two examples emphasize the applicability of this method in higher dimensions.

Example 6 (d -dimensional point cloud on a 3-dimensional manifold). Let $d \geq 4$, and let $M \subset \mathbb{R}^d$ be a section of the three-dimensional paraboloid defined by $x_1^2 + x_2^2 + x_3^2 - 12x_4 = 0$, such that its reach is decided by a point of maximal curvature at the origin. It can be shown that $\text{reach}_\beta(M) = 6 + \frac{\beta}{2}$, for β between 0 and some positive threshold.

Let $\widehat{M} \subset M$ be a realization of 3000 points, uniformly distributed on M . The Python package `pyvista` does not currently have methods for computing meshes in dimension higher than 3, so in this example, we do not compute $\text{reach}_\beta^{\mathcal{M}}(\widehat{M})$ in (23). Mesh generation in higher dimension is itself an active field of research [15, 17, 38]. Nonetheless, we plot the exact β -reach profile of \widehat{M} in Figure 15.

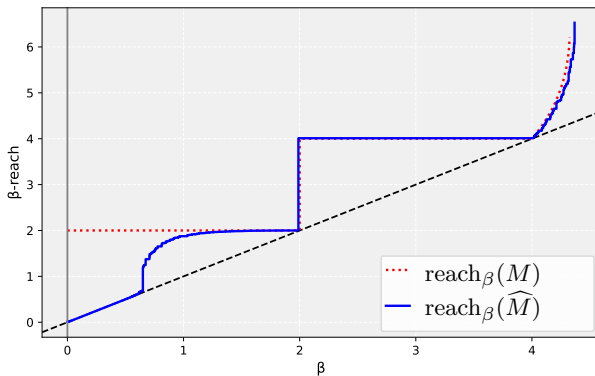


Fig. 16: The exact β -reach profile of \widehat{M} in Example 7. The β -reach of the underlying manifold M is shown as a red dotted curve. Here, M is the union of two 3-dimensional hyperspheres of radius 2 embedded in \mathbb{R}^d .

Remark 10. The dimension of the ambient space $d \geq 4$ in Example 6 does not affect the shape of the β -reach profile of the point cloud \widehat{M} , computed purely in terms of distances between the points, and distances to midpoints. These distances are preserved under isometries to higher dimensional spaces. However, the number of points in \widehat{M} plays a role in the shape of the β -reach profile through the Hausdorff distance between \widehat{M} and the underlying manifold M . The exact relationship between the number of points in \widehat{M} and the Hausdorff distance $d_H(\widehat{M}, M)$ cannot be made precise with no prior knowledge of how the points are distributed on M —which also plays a role in the shape of the β -reach profile of \widehat{M} . Nonetheless, the dimension of M determines to a large extent the number of points needed to ensure that $d_H(\widehat{M}, M)$ is less than a given tolerance.

Example 7 (Two 3-dimensional hyperspheres). Let $d \geq 4$, and let $M \subset \mathbb{R}^d$ be the union of two three-dimensional hyperspheres of radius 2 whose centers are 12 units apart. Let $\widehat{M} \subset M$ be a realization of 3000 points uniformly distributed on M . The β -reach profiles of M and a realization of \widehat{M} are shown in Figure 16.

Example 7 highlights a few nice features of the β -reach profile. That is, it adapts very well to situations where the reach is determined by a bottleneck structure, as is the case for a hypersphere. For $\beta \in (2, 4]$, the β -reach is no longer determined by the radius of the hypersphere, but by half the distance between the surfaces of the spheres (in this case, $(12 - 2 - 2)/2 = 4$ units). Thus, the β -reach profile also provides information about the large scale features of




Fig. 17: (a): The point cloud data \widehat{P} in Example 8 is used to guarantee that the r -convexity of the underlying surface of the plane P does not exceed the chosen test radius $r = 0.03$. The areas of the plane responsible for this restriction on the r -convexity are highlighted in red. (b): The `pyvista` mesh $\mathcal{M}(\widehat{N})$ constructed from $\widehat{N} \subset \widehat{P}$, the points at the nose of the plane.

the data. Notably, it gives the scales at which it becomes possible to distinguish these features from one another.

Finally, recall from Remark 2 that one can read from the plot information about the critical points of the generalized gradient function. In the case of Example 7, the distances from the three critical points to M are 2, 2, and 4.

4 Numerical studies

In this section, we test the methods for bounding the reach and r -convexity introduced in Section 3 against numerical data. First, in Section 4.1, we study the performance of our methods on real data. Then, in Section 4.2, we test the convergence results in Theorems 3 and 4.

4.1 Aircraft data

The real data in Example 8, below, is studied using the tools developed in this document. After a short description of the data, we perform several analyses using the tools developed in Section 3.

Example 8 (Aircraft data). [8] provides a point cloud dataset that is sampled over the surface of a commercial aircraft (the white points in Figure 17 (a)). The diameter of the raw point cloud data (which corresponds to the length of the plane) is 0.749 units. Denote the surface of the aircraft by P and the approximating point cloud by \widehat{P} .

We are also interested in the nose of the aircraft, $N \subset P$, the surface of the first 0.056 length units of the aircraft. A two-dimensional mesh approximating $\widehat{N} := N \cap \widehat{P}$ is shown in Figure 17 (b).




Fig. 18: The exact β -reach profile of \widehat{N} in Example 8 is shown as the blue solid curve. Using the formulation in (23), and the mesh in Figure 17 (b), one may compute an improved approximation (green dashed curve) to the β -reach of the underlying (unknown) manifold N .

It is possible that in an engineering practice, one may want to identify which regions of a point cloud are not r -convex for a predefined value of r . Panel (a) of Figure 17 highlights in red the regions of the plane that are *certainly* not r -convex for the choice of $r = 0.03$.

The surface of the plane P is known to be two-dimensional, and so a cubic lattice of points φ with lattice spacing $a = 0.004$ superimposed over \widehat{P} will almost surely not intersect P . The discrete dilation and erosion operations in Definition 7 are well defined using the larger point cloud $\Phi := \varphi \cup \widehat{P}$. The red points in Figure 17 (a) are the elements of the discrete set $(\widehat{P}_{r-\epsilon})_{-(r+\epsilon)} \cap \varphi$, with $\epsilon := \sqrt{3}a/2 = \sup\{\delta_{\Phi}(q) : q \in \mathbb{R}^3\}$. Since this set is not empty, one has that P is a proper subset of $P_{\bullet r}$. In addition, by (2), one has conclusive evidence that $\text{reach}(P) \leq r$.

This example illustrates that, with this method, one can identify the regions responsible for limiting the r -convexity (and thus the reach), with a test specificity of 100%. One can improve the sensitivity of the test by decreasing the lattice spacing a . Then, other regions that are not r -convex (such as the interiors of the horizontal and vertical stabilizers) would be identified as such.

Remark that the nose of the plane is marked in red, due of course to a region of high curvature. By considering only the smooth manifold N in Example 8, corresponding to the nose of the plane, we can approximate the β -reach profile of N by $\text{reach}_{\beta}(\widehat{N})$ in Definition 6, or by $\text{reach}_{\beta}^{\mathcal{M}}(\widehat{N})$ in Equation (23).

The resulting β -reach profiles are shown in Figure 18. From the figure, the reach is clearly seen to be much smaller than 0.03 as indicated by the r -convexity experiment. Assuming that the β -reach profile of N maintains a constant slope

near 0 as seen before in Figure 14, the approximation $\text{reach}(N) \approx 0.012$ can be read from Figure 18.

There are two main issues with using the reach bound in (15) on the aircraft data in Example 8. First, it is expected to produce a large overestimate of $\text{reach}(N)$ since N is a smooth manifold (see the discussion in Section 3.3). In addition, it is impossible to know the Hausdorff distance between N and its approximating point cloud \widehat{N} . Nevertheless, one can obtain a good approximation by considering the persistence diagram of growing balls centered at the points in \widehat{N} . The largest of the death-times of the topological features corresponds to the smallest radius for which the union of balls is homotopic to a point. This is likely to correspond closely to the Hausdorff distance $d_H(N, \widehat{N})$, and is thus a good choice for ϵ in (15). The persistence diagram, generated from the Python module `ripser` [9], gives that the largest death-time is $\epsilon = 0.0065$. From this, one calculates

$$\widehat{\text{rch}}^{(\epsilon)}(\widehat{N}) = 0.022.$$

Although the reach bound in (15) is less appropriate to use in this setting, we will see in the following section that it (and the r -convexity bound in (11)) is highly applicable when studying binary images.

4.2 Numerical convergence of reach and r -convexity bounds

In stochastic geometry literature, the *excursion set* of a random field is the subset of its domain on which the random field surpasses a predefined threshold (see [5] for a comprehensive reference). The set in Figure 9, for example, is one realization of the excursion set of a stationary, isotropic Gaussian random field sampled on a square lattice.

The following two examples are meant to imitate the discretized excursion set of C^2 continuous random fields on square lattices (see, e.g., [12, 28, 29]). A key feature of both examples, is that the reach and r -convexity of the sets are known, and so we can study the convergence of the bounds in Theorems 3 and 4 as the grid of sampling points becomes dense in \mathbb{R}^2 . Each example illustrates one of the two cases mentioned in Remark 8 concerning the relationship between the reach, regions of high curvature, and bottleneck structures.

Example 9 (Reach determined by curvature). Let $U := \{(x, y) \in \mathbb{R}^2 : y \leq x^2/2\} \cap B(0, 10)$. It is easy to check that $\text{reach}(U) = \text{rconv}(U) = 1$. Moreover, the reach is determined by a point of maximal curvature at $(0, 0) \in U$.

Let $(\Phi^{(n)})_{n \geq 1}$ be a sequence of square lattices over \mathbb{R}^2 with lattice spacing $a_n = 0.7/n$, and independent, uniformly random position and orientation. For $n \in \mathbb{N}^+$, let $\widehat{U}^{(n)} := U \cap \Phi^{(n)}$ be the set U sampled on $\Phi^{(n)}$. Figure 19 depicts square subsets of realizations of $\widehat{U}^{(2)}$ and $\widehat{U}^{(10)}$ shown as binary images.




Fig. 19: Realizations of the random set $\widehat{U}^{(n)}$ in Example 9. (a): $n = 2$. (b): $n = 10$.




Fig. 20: Realizations of the random set $\widehat{W}^{(n)}$ in Example 10. (a): $n = 2$. (b): $n = 10$.

Example 10 (Reach determined by a bottleneck structure). Let $W := \{(x, y) \in \mathbb{R}^2 : |y| \geq x^2/2 + 1\} \cap B(0, 10)$. One has $\text{reach}(W) = \text{rconv}(W) = 1$, which is half the distance between its two connected components. For the sequence of square lattices $(\Phi^{(n)})_{n \geq 1}$ in Example 9, let $\widehat{W}^{(n)} := W \cap \Phi^{(n)}$. Figure 20 depicts square subsets of realizations of $\widehat{W}^{(2)}$ and $\widehat{W}^{(10)}$ shown as binary images.

Remark 11. The choice to use sets with a reach of 1 in Examples 9 and 10 does not limit their generality. The importance lies in the scale of the sampling lattice $\Phi^{(n)}$ relative to the reach, and the ratio of the two tends to 0 as $n \rightarrow \infty$ in both examples.

The following observation makes the generality of these examples even clearer: to consider the lattice of sampling points at various scales and orientations is equivalent to considering a fixed lattice and various scales and orientations of the underlying sets U and W .

For each $n \in \{2, \dots, 12\}$, we compute 50 independent realizations of $\widehat{U}^{(n)}$ and 50 independent realizations of $\widehat{W}^{(n)}$. For each of the resulting discrete sets, we




Fig. 21: (a): The mean and 95% confidence interval for the values of $\widehat{r}^{(a_n/\sqrt{2})}(\widehat{U}^{(n)})$ (squares) and $\widehat{rch}^{(\sqrt{1.25}a_n)}(\widehat{U}^{(n)})$ (circles) for 50 independent realizations of $\widehat{U}^{(n)}$ in Example 9 (b): The mean and 95% confidence interval for the values of $\widehat{r}^{(a_n/\sqrt{2})}(\widehat{W}^{(n)})$ (squares) and $\widehat{rch}^{(\sqrt{1.25}a_n)}(\widehat{W}^{(n)})$ (circles) for 50 independent realizations of $\widehat{W}^{(n)}$ in Example 10.




Fig. 22: The grey points and the white points can be separated by a curve with arbitrarily small curvature (arbitrarily large reach). The midpoint of two grey points may be up to a distance of $\sqrt{1.25}a_n$ from another grey point, while remaining arbitrarily close to the separating curve.

compute the bounds of r -convexity and reach in (11) and (15) respectively. The means of the bounds are plotted in Figure 21 with empirical 95% confidence intervals. Panels (a) and (b) correspond to the results for the replications of $\widehat{U}^{(n)}$ and $\widehat{W}^{(n)}$ respectively.

For the r -convexity bound, $\epsilon_n = a_n/\sqrt{2}$ is chosen since $d_H(\Phi^{(n)}, \mathbb{R}^2) = a_n/\sqrt{2}$, for $n \in \mathbb{N}^+$.

For the reach bound, we set $\epsilon_n = \sqrt{1.25}a_n$ for reasons that are more complicated. Even though this choice of ϵ_n leads to $d_H(\widehat{U}^{(n)}, U) > \epsilon_n$, or $d_H(\widehat{W}^{(n)}, W) > \epsilon_n$ with positive probability, it is sufficient for the result of Theorem 4 in this case. We provide Figure 22 for some intuition, but a formal proof is omitted.




Fig. 23: The quantities used in the proof of Theorem 1. Sending $\epsilon \rightarrow 0$ leads to the largest circle intersecting ∂A if it is to always intersect $B(a_1, \epsilon)$.

A linear regression on the log-log plot of the sample means of the data in Figure 21 (a) shows that

$$\mathbb{E}[\widehat{r}^{(a_n/\sqrt{2})}(\widehat{U}^{(n)})] \approx 1 + 1.71n^{-0.59} \quad \text{and} \quad \mathbb{E}[\widehat{\text{rch}}^{(\sqrt{1.25}a_n)}(\widehat{U}^{(n)})] \approx 1 + 1.54n^{-0.62}.$$

These empirical convergence rates — $\mathcal{O}(\epsilon_n^{0.59})$ for the bound on $\text{rch}(U)$ and $\mathcal{O}(\epsilon_n^{0.62})$ for the bound on $\text{reach}(U)$ — are both slightly faster than the predicted rate of $\mathcal{O}(\sqrt{\epsilon_n})$ established in (17) for the bound on the reach. Since the reach of U is determined by a region of maximal curvature, the convergence rate is governed by the analysis in Case 2 in the proof of Theorem 4.

For the data in panel (b) of Figure 21,

$$\mathbb{E}[\widehat{r}^{(a_n/\sqrt{2})}(\widehat{W}^{(n)})] \approx 1 + 2.81n^{-1.20} \quad \text{and} \quad \mathbb{E}[\widehat{\text{rch}}^{(\sqrt{1.25}a_n)}(\widehat{W}^{(n)})] \approx 1 + 0.65n^{-1.50}.$$

Again, these empirical convergence rates are faster than the anticipated rate of $\mathcal{O}(\epsilon_n)$ established for the bound on the reach in Case 1 (*i.e.*, when the reach is determined by a bottleneck structure) in the proof of Theorem 4.

5 Proofs and technical results

Here, we provide proofs and auxiliary lemmas that support the results in Sections 2 and 3.

Proof of Theorem 1 We begin by showing (2). Suppose that $r \in (0, \text{reach}(A))$. By Lemma 1, $A \subseteq A_{\bullet, r}$, so it suffices to show that $A^c \subseteq (A_{\bullet, r})^c$. Let $x \in A^c$. If $x \in (A_r)^c$,

then clearly, $x \in (A_r)^c \oplus B(0, r) = (A_{\bullet r})^c$. If $x \in A_r \setminus A$, then by Federer [39, Corollary 4.9],

$$\delta_{(A_r)^c}(x) = r - \delta_A(x) < r,$$

and so $x \in (A_r)^c \oplus B(0, r) = (A_{\bullet r})^c$, which proves (2). What remains to be shown is that if ∂A is C^1 -smooth and $(d-1)$ -dimensional, then $\text{rconv}(A) \leq \text{reach}(A)$. This inequality is shown via proof by contradiction. Suppose that

- (i) $r > \text{reach}(A)$ and
- (ii) $A_{\bullet r} = A$.

By (i), there exists $p \in A_{\frac{r+\text{reach}(A)}{2}}$ with no unique point in A closest to p . In particular, since A is closed, there exists two non-identical points $a_1, a_2 \in \partial A \subseteq A$ such that $\|p - a_1\|_2 = \|p - a_2\|_2 = \delta_A(p) < r$. Let n_1 be the unit normal vector to ∂A at a_1 , pointing towards p . Since a_1 is a limit point of A , one has $A^c \cap B(a_1, \epsilon) \neq \emptyset$ for all $\epsilon > 0$. By (ii), $A^c = \bigcup_{x \in (A_r)^c} B(x, r)$, and so for all $\epsilon > 0$, there exists $x \in (A_r)^c$ such that

- (iii) $B(x, r) \cap B(a_1, \epsilon) \neq \emptyset$ and
- (iv) $B(x, r) \cap A = \emptyset$.

The boundary ∂A is C^1 -smooth at a_1 , so it is easily checked that for ϵ close to 0, the locations x that satisfy both (iii) and (iv) are necessarily contained in a small neighbourhood around $a_1 + rn_1$ (see Figure 23). That is, there exists a mapping $\theta : \epsilon \mapsto \theta(\epsilon) \in \mathbb{R}^+$ that tends to 0 as $\epsilon \rightarrow 0$, such that if $x \in (A_r)^c$ and $\epsilon \in \mathbb{R}^+$ satisfy (iii) and (iv), then $x \in B(a_1 + rn_1, \theta(\epsilon))$. Note that $\|(a_1 + rn_1) - a_2\|_2 < r$ since $a_2 \in B(p, \delta_A(p)) \subset B(a_1 + rn_1, r)$. Choose ϵ such that $\theta(\epsilon) < r - \|(a_1 + rn_1) - a_2\|_2$. Then for any $x \in (A_r)^c \cap B(a_1 + rn_1, \theta(\epsilon))$,

$$\begin{aligned} \|x - a_2\|_2 &\leq \|x - (a_1 + rn_1)\|_2 + \|(a_1 + rn_1) - a_2\|_2 \\ &\leq \theta(\epsilon) + \|(a_1 + rn_1) - a_2\|_2 < r. \end{aligned}$$

This contradicts (iv). □

Lemma 2. *Let U be an open set in \mathbb{R}^d , and let $r \geq 0$. Then,*

$$\text{cl}(U)_{\bullet r} = \text{cl}(U) \quad \implies \quad U_{\bullet r} = U.$$

Proof of Lemma 2 By Lemma 1, item (b), $U \subseteq U_{\bullet r} \subseteq \text{cl}(U)_{\bullet r}$. Note also that $U_{\bullet r}$ is open. Suppose that $\text{cl}(U)_{\bullet r} = \text{cl}(U)$. Then $U_{\bullet r}$ is an open subset of $\text{cl}(U)$ that contains all of the interior points of $\text{cl}(U)$. Therefore, $U_{\bullet r} = U$. □

Lemma 3. *Let $A \subseteq \mathbb{R}^d$ be closed. It holds that*

$$\text{rconv}(\text{cl}(A^c)) = \sup\{r \in \mathbb{R} : A_{\bullet -r} = A\}$$

Proof of Lemma 3 By Lemma 2 and item (a) in Lemma 1, $\{r \in \mathbb{R} : \text{cl}(A^c)_{\bullet r} = \text{cl}(A^c)\} \subseteq \{r \in \mathbb{R} : (A^c)_{\bullet r} = A^c\} = \{r \in \mathbb{R} : A_{\bullet -r} = A\}$. Therefore, $\text{rconv}(\text{cl}(A^c)) \leq \sup\{r \in \mathbb{R} : A_{\bullet -r} = A\}$. Now we show the reverse inequality. Let $r, \tilde{r} \in \mathbb{R}^+$ satisfy $r > \tilde{r} > \text{rconv}(\text{cl}(A^c))$. By Proposition 1, $\text{cl}(A^c)_{\bullet \tilde{r}} \setminus \text{cl}(A^c)$ contains a ball of radius

ϵ for some $\epsilon \in (0, r)$. Now, let $\delta = \min(\epsilon, 2(r - \tilde{r}))$. Note that $(A^c)_{\delta/2} \supseteq \text{cl}(A^c)$, and so by Lemma 1, $((A^c)_{\delta/2})_{\bullet, r-\delta/2} = ((A^c)_r)_{-(r-\delta/2)} \supseteq \text{cl}(A^c)_{\bullet, r-\delta/2} \supseteq \text{cl}(A^c)_{\bullet, \tilde{r}}$, which contains the ball of radius ϵ . Now, eroding by $\delta/2$ preserves a ball of radius $\epsilon - \delta/2 \geq \epsilon/2$. That is, $((A^c)_{\delta/2})_{\bullet, r-\delta/2} -_{\delta/2} = (A^c)_{\bullet, r}$ contains a ball of radius $\epsilon/2$ that does not intersect $\text{cl}(A^c)$, and so $(A^c)_{\bullet, r} \neq A^c$ which implies the desired $A_{\bullet, -r} \neq A$ by Lemma 1, item (a). \square

Proof of Corollary 1 If $\text{rconv}(A) = 0$, then by Theorem 1, $\text{reach}(A) = 0$ and we are done. Now assume $\text{rconv}(A) > 0$ and $\text{rconv}(\text{cl}(A^c)) > 0$, then by Lemma 3, there exists $\delta > 0$ such that $A_{\bullet, r} = A$ for all $r \in (-\delta, \delta)$. Theorem 1 in [53] states that ∂A is $(d-1)$ -dimensional and C^1 -smooth, which implies Equation 3 in our Theorem 1. \square

Proof of Corollary 2 By Theorem 1, the statement holds if $\partial(A_\epsilon)$ is C^1 -smooth and $(d-1)$ -dimensional. Otherwise, if there is indeed a point c on the boundary of $A_\epsilon = \bigcup_{a \in A} B(a, \epsilon)$ that does not have a continuous derivative, then c is an intersection point of two distinct d -spheres centered in A of radius ϵ . In other words, c is at cusp that points inwards towards the interior of A_ϵ , making $\text{reach}(A_\epsilon) = \text{rconv}(A_\epsilon) = 0$. \square

Remark 12. Remark that $(A_\epsilon)_{\bullet, -\epsilon} = A_\epsilon$. Therefore, if the hypotheses of Corollary 2 are strengthened to those of Corollary 1, then the proof of Corollary 2 holds by applying Theorem 1 in [53] followed by our Theorem 1.

Here, we present an auxiliary lemma, and use it in our proof of Proposition 1 in Section 3.1.

Lemma 4. *For closed $A \subseteq \mathbb{R}^d$ and $r \in \mathbb{R}^+$, it holds that*

$$A_{\bullet, r} = \{p \in \mathbb{R}^d : \forall x \in B(0, r), B(p+x, r) \cap A \neq \emptyset\}. \quad (24)$$

Proof of Lemma 4 By manipulating the expressions in Definition 1, one obtains

$$A_{\bullet, r} = \left(\bigcup_{y \in (A \oplus B(0, r))^c} B(y, r) \right)^c, \quad (25)$$

which reads: *the elements of $A_{\bullet, r}$ are the $p \in \mathbb{R}^d$ such that p is not contained in any closed ball of radius r that does not intersect A . This statement is equivalent to: $p \notin A_{\bullet, r}$ if and only if there exists a closed ball of radius r that contains p but does not intersect A . Thus, we have shown that $(A_{\bullet, r})^c$ is equal to the complement of the RHS of (24), which proves the desired result. \square*

Proof of Proposition 1 Let $\tilde{r} \in (\text{rconv}(A), r)$ and fix $p \in A_{\bullet, \tilde{r}} \setminus A$. Since A is closed, there is an open neighbourhood containing p that does not intersect A . Choose $\epsilon \in \mathbb{R}^+$ such that $B(p, \epsilon) \cap A = \emptyset$. There exists a sufficiently small $\delta \in (0, \epsilon)$ such that for every $y \in B(p, \delta+r)$, there exists an $x \in B(0, \tilde{r})$ that satisfies $B(p+x, \tilde{r}) \setminus B(p, \epsilon) \subset B(y, r)$.

Since $p \in A_{\bullet, \tilde{r}}$, one has that $B(p+x, \tilde{r}) \setminus B(p, \epsilon)$ contains an element of A by Lemma 4. By inclusion, $B(y, r)$ contains a point in A as well. Let $z \in B(p, \delta)$. We have shown that z is in the right-hand side of (24), since, for all $x \in B(0, r)$, one has $y := z+x \in B(z, \delta+r)$ by the triangle inequality, and so by previous arguments, $B(z+x, r)$ contains an element of A . Therefore, by Lemma 4, $z \in A_{\bullet, r}$. But z is not in A since $B(p, \delta) \subseteq B(p, \epsilon) \subseteq A^c$. Therefore, $z \in A_{\bullet, r} \setminus A$ and so $B(p, \delta) \subseteq A_{\bullet, r} \setminus A$. \square

Proof of Theorem 3 We begin by showing (12). Let $n \in \mathbb{N}^+$ and fix $p \in \Phi^{(n)} \setminus \widehat{A}^{(n)}$. If $\epsilon_n \geq \text{rconv}(A)$, then (12) holds trivially. Now, let $r \in \mathbb{R}^+$ be such that $\epsilon_n < r < \text{rconv}(A)$ so that $A_{\bullet, r} = A$. We aim to show that $p \notin (\widehat{A}_{r-\epsilon_n}^{(n)})_{-(r+\epsilon_n)}$. Indeed, by the r -convexity of A , there exists $x \in (A_r)^c$ such that $\|x-p\|_2 < r$. In addition, $B(x, \epsilon_n) \cap A_{r-\epsilon_n} = \emptyset$, so there exists $q \in \Phi^{(n)} \setminus \widehat{A}_{r-\epsilon_n}^{(n)}$ such that $\|q-x\|_2 \leq \epsilon_n$. Therefore, by the triangle inequality, $\|q-p\|_2 \leq r+\epsilon_n$, which implies $\delta_{\Phi^{(n)} \setminus \widehat{A}_{r-\epsilon_n}^{(n)}}(p) \leq r+\epsilon_n$. Thus, $p \notin (\widehat{A}_{r-\epsilon_n}^{(n)})_{-(r+\epsilon_n)}$ as desired.

Now, to prove (13), first fix $r > \text{rconv}(A)$. To simplify notation, let $\delta_n := d_H(\widehat{A}^{(n)}, A)$. By Proposition 1, there exists an open subset $O \subseteq A_{\bullet, r} \setminus A$ that contains a closed ball of radius $\delta_{n_0} + 3\epsilon_{n_0}$ for sufficiently large $n_0 \in \mathbb{N}^+$. Let $n \geq n_0$. The point cloud $\Phi^{(n)}$ is sufficiently dense such that there exists $q \in O_{-(\delta_n + 2\epsilon_n)} \cap \Phi^{(n)}$. Importantly, this implies that $q \in (A_r)_{-(r+\delta_n+2\epsilon_n)} \cap \Phi^{(n)}$. By Lemma 1, for all $s > r$, we have $(A_r)_{-(r+\delta_n+2\epsilon_n)} \subseteq ((A_r)_{\bullet, (s-r)})_{-(r+\delta_n+2\epsilon_n)} = (A_s)_{-(s+\delta_n+2\epsilon_n)}$, and therefore $q \in (A_s)_{-(s+\delta_n+2\epsilon_n)} \cap \Phi^{(n)}$. Notice that $A_s \cap \Phi^{(n)} \subseteq \widehat{A}_{\delta_n+s}^{(n)}$, which implies $(A_s)^c \supseteq \Phi^{(n)} \setminus \widehat{A}_{s+\delta_n}^{(n)}$ and thus $(A_s)_{-(s+\delta_n+2\epsilon_n)} \cap \Phi^{(n)} \subseteq (\widehat{A}_{\delta_n+s}^{(n)})_{-(s+\delta_n+2\epsilon_n)}$. Summarizing, we have shown that there is a point q in $(\widehat{A}_{\delta_n+s}^{(n)})_{-(s+\delta_n+2\epsilon_n)}$ that is not in $\widehat{A}^{(n)}$. By the change of variables $\tilde{s} := s + \delta_n + \epsilon_n$, it follows that $(\widehat{A}_{\tilde{s}-\epsilon_n}^{(n)})_{-(\tilde{s}+\epsilon_n)} \setminus \widehat{A}^{(n)} \neq \emptyset$ for all $\tilde{s} > r + \delta_n + \epsilon_n$, which implies $\widehat{r}^{(\epsilon_n)}(\widehat{A}^{(n)}) \leq r + \delta_n + \epsilon_n$. Sending $n \rightarrow \infty$ yields $\lim_{n \rightarrow \infty} \widehat{r}^{(\epsilon_n)}(\widehat{A}^{(n)}) \leq r$, and since $r \in (\text{rconv}(A), \infty)$ was chosen freely, $\lim_{n \rightarrow \infty} \widehat{r}^{(\epsilon_n)}(\widehat{A}^{(n)}) \leq \text{rconv}(A)$. This result, along with (12), gives the convergence in (13). \square

Justification for Equations (8) and (9).

The reach of A is equal to the inverse of the curvature of f at $x=0$ [2, Theorem 3.4]. Equation (8) holds since the curvature of f at $x=0$ is $f''(0) = 2h'(0)$.

Now we show (9). Without loss of generality, suppose $h(0) = f(0) = 0$. For x in a neighbourhood of 0, consider the symmetric points $a_1 = (x, f(x))$ and $a_2 = (-x, f(x))$ in A , and remark that

$$g_{\|a_2 - a_1\|_2} \circ \delta_A \left(\frac{a_1 + a_2}{2} \right) = \frac{x^2}{2f(x)} + \frac{f(x)}{2}, \quad (26)$$

since $\delta_A \left(\frac{a_1 + a_2}{2} \right) = f(x)$. We claim without proof that for all $\tilde{a}_1, \tilde{a}_2 \in A$ satisfying $\delta_A \left(\frac{\tilde{a}_1 + \tilde{a}_2}{2} \right) \geq f(x)$, it holds that $g_{\|\tilde{a}_2 - \tilde{a}_1\|_2} \circ \delta_A \left(\frac{\tilde{a}_1 + \tilde{a}_2}{2} \right) \geq g_{\|a_2 - a_1\|_2} \circ$

$\delta_A \left(\frac{a_1 + a_2}{2} \right)$. In the language of the β -reach, this is equivalent to

$$\text{reach}_{f(x)}(A) = g_{\|a_2 - a_1\|_2} \circ \delta_A \left(\frac{a_1 + a_2}{2} \right). \quad (27)$$

There is a $\delta > 0$ such that h has an inverse on $[0, \delta]$, and so for $\beta \in [0, \delta]$, choose x such that $f(x) = \beta$. Remark that

$$x^2 = h^{-1}(\beta) = \frac{\beta}{h'(0)} - \frac{\beta^2 h''(0)}{2h'(0)^3} + o(\beta^2),$$

by Taylor's theorem. Plugging into (26) and applying (27), one obtains (9).

References

- [1] Aamari, E., Berenfeld, C., and Levrard, C. (2022). Optimal reach estimation and metric learning. *Preprint arXiv:2207.06074*, pages 1–47.
- [2] Aamari, E., Kim, J., Chazal, F., Michel, B., Rinaldo, A., and Wasserman, L. (2019). Estimating the reach of a manifold. *Electronic Journal of Statistics*, 13(1):1359–1399.
- [3] Aamari, E. and Levrard, C. (2019). Nonasymptotic rates for manifold, tangent space and curvature estimation. *The Annals of Statistics*, 47(1):177–204.
- [4] Aaron, C., Cholaquidis, A., and Fraiman, R. (2022). Estimation of surface area. *Electronic Journal of Statistics*, 16(2):3751–3788.
- [5] Adler, R. J. and Taylor, J. E. (2007). *Random fields and geometry*. Springer Monographs in Mathematics. Springer, New York.
- [6] Attali, D. and Lieutier, A. (2015). Geometry-driven collapses for converting a Čech complex into a triangulation of a nicely triangulable shape. *Discrete & Computational Geometry*, 54:798–825.
- [7] Attali, D., Lieutier, A., and Salinas, D. (2013). Vietoris–rips complexes also provide topologically correct reconstructions of sampled shapes. *Computational Geometry*, 46(4):448–465. 27th Annual Symposium on Computational Geometry (SoCG 2011).
- [8] Baorui, M., Yu-Shen, L., Matthias, Z., and Zhizhong, H. (2022). Surface reconstruction from point clouds by learning predictive context priors. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- [9] Bauer, U. (2021). Ripser: efficient computation of Vietoris–Rips persistence barcodes. *J. Appl. Comput. Topol.*, 5(3):391–423.
- [10] Berenfeld, C., Harvey, J., Hoffmann, M., and Shankar, K. (2022). Estimating the Reach of a Manifold via its Convexity Defect Function. *Discrete & Computational Geometry*, 67:403–438.

- [11] Berrendero, J. R., Cuevas, A., and Pateiro López, B. (2012). A multivariate uniformity test for the case of unknown support. *Statistics and Computing*, 22:259–271.
- [12] Biermé, H. and Desolneux, A. (2021). The effect of discretization on the mean geometry of a 2D random field. *Annales Henri Lebesgue*, 4:1295–1345.
- [13] Biermé, H., Di Bernardino, E., Duval, C., and Estrade, A. (2019). Lipschitz-Killing curvatures of excursion sets for two-dimensional random fields. *Electronic Journal of Statistics*, 13(1):536–581.
- [14] Blum, H. (1967). A transformation for extracting new descriptions of shape. *Models for the perception of speech and visual form*, pages 362–380.
- [15] Boissonnat, J.-D., Devillers, O., and Hornus, S. (2009). Incremental construction of the delaunay triangulation and the delaunay graph in medium dimension. In *Proceedings of the twenty-fifth annual symposium on Computational geometry*. ACM.
- [16] Boissonnat, J.-D., Lieutier, A., and Wintraecken, M. (2019). The reach, metric distortion, geodesic convexity and the variation of tangent spaces. *Journal of Applied and Computational Topology*, 3(1):29–58.
- [17] Boissonnat, J.-D., Wormser, C., and Yvinec, M. (2008). Locally uniform anisotropic meshing. In *Proceedings of the twenty-fourth annual symposium on Computational geometry*. ACM.
- [18] Chazal, F., Cohen-Steiner, D., and Lieutier, A. (2009a). A Sampling Theory for Compact Sets in Euclidean Space. *Discrete & Computational Geometry*, 41:461–479.
- [19] Chazal, F., Cohen-Steiner, D., Lieutier, A., and Thibert, B. (2007). Shape Smoothing Using Double Offset. *Proceedings - SPM 2007: ACM Symposium on Solid and Physical Modeling*.
- [20] Chazal, F., Cohen-Steiner, D., Lieutier, A., and Thibert, B. (2009b). Stability of curvature measures. In *Computer Graphics Forum*, volume 28, pages 1485–1496. Wiley Online Library.
- [21] Chazal, F. and Lieutier, A. (2005a). The “ λ -medial axis”. *Graphical Models*, 67(4):304–331.
- [22] Chazal, F. and Lieutier, A. (2005b). Weak feature size and persistent homology: Computing homology of solids in \mathbb{R}^n from noisy data samples. In *Proceedings of the Twenty-First Annual Symposium on Computational Geometry*, SCG ’05, pages 255–262, New York, NY, USA. Association for Computing Machinery.
- [23] Chazal, F. and Lieutier, A. (2008). Smooth manifold reconstruction from noisy and non-uniform approximation with guarantees. *Computational Geometry*, 40(2):156–170.
- [24] Cholaquidis, A. (2023). A counter example to a borsuk conjecture. *Applied General Topology*, 24(1):125—128.
- [25] Cholaquidis, A., Fraiman, R., and Moreno, L. (2022). Universally consistent estimation of the reach. *Preprint arXiv:2110.12208v2*, pages 1–15.
- [26] Clarke, F. and Wolenski, P. (1995). Proximal Smoothness and the Lower- C^2 Property. *Journal of Convex Analysis*, 2:117–144.

- [27] Colesanti, A. and Manselli, P. (2010). Geometric and isoperimetric properties of sets of positive reach in \mathbf{E}^d . *Preprint*, pages 1–15.
- [28] Cotsakis, R., Di Bernardino, E., and Duval, C. (2022a). Surface area and volume of excursion sets observed on point cloud based polytopic tessellations. *Preprint arXiv:2209.10383*.
- [29] Cotsakis, R., Di Bernardino, E., and Opitz, T. (2022b). On the perimeter estimation of pixelated excursion sets of 2D anisotropic random fields. *Preprint hal-03582844v2*, pages 1–28.
- [30] Cuevas, A. (2009). Set estimation: Another bridge between statistics and geometry. *BEIO, Boletín de Estadística e Investigación Operativa*, 25:71–85.
- [31] Cuevas, A., Fraiman, R., and Casal, A. (2007). A nonparametric approach to the estimation of lengths and surface areas. *The Annals of Statistics*, 35:1031–1051.
- [32] Cuevas, A., Fraiman, R., and Pateiro-López, B. (2012). On Statistical Properties of Sets Fulfilling Rolling-Type Conditions. *Advances in Applied Probability*, 44(2):311–329.
- [33] Cuevas, A., Llop, P., and Pateiro-López, B. (2014). On the estimation of the medial axis and inner parallel body. *Journal of Multivariate Analysis*, 129:171–185.
- [34] Cuevas, A. and Rodríguez-Casal, A. (2004). On Boundary Estimation. *Advances in Applied Probability*, 36(2):340–354.
- [35] Dey, T. K. and Sun, J. (2006). Normal and feature approximations from noisy point clouds. In Arun-Kumar, S. and Garg, N., editors, *FSTTCS 2006: Foundations of Software Technology and Theoretical Computer Science*, pages 21–32, Berlin, Heidelberg. Springer Berlin Heidelberg.
- [36] Dey, T. K. and Zhao, W. (2003). Approximating the Medial Axis from the Voronoi Diagram with a Convergence Guarantee. *Algorithmica*, 38(1):179–200.
- [37] Divol, V. (2021). Minimax adaptive estimation in manifold inference. *Electronic Journal of Statistics*, 15:5888–5932.
- [38] Edelsbrunner, H. (2001). *Geometry and Topology for Mesh Generation*. Cambridge University Press.
- [39] Federer, H. (1959). Curvature measures. *Transactions of the American Mathematical Society*, 93(3):418–491.
- [40] Kim, J., Shin, J., Chazal, F., Rinaldo, A., and Wasserman, L. (2019). Homotopy reconstruction via the cech complex and the vietoris-rips complex. *Preprint arXiv:1903.06955*.
- [41] Lieutier, A. (2004). Any open bounded subset of \mathbb{R}^n has the same homotopy type as its medial axis. *Computer-Aided Design*, 36(11):1029–1046.
- [42] Lieutier, A. and Wintraecken, M. (2023). Hausdorff and gromov-hausdorff stable subsets of the medial axis. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023*, pages 1768—1776, New York, NY, USA. Association for Computing Machinery.
- [43] Mani-Levitska, P. (1993). Characterizations of Convex Sets. In *Handbook of Convex Geometry*, pages 19–41. North-Holland.

- [44] Niyogi, P., Smale, S., and Weinberger, S. (2008). Finding the Homology of Submanifolds with High Confidence from Random Samples. *Discrete & Computational Geometry*, 39:419–441.
- [45] Pateiro López, B. (2008). Set estimation under convexity type restrictions. Ph.D. thesis. [\[Link to thesis\]](#).
- [46] Perkal, J. (1956). Sur les ensembles ϵ -convexes. *Colloquium Mathematicum*, 4:1–10.
- [47] Poliquin, R. A., Rockafellar, R. T., and Thibault, L. (2000). Local Differentiability of Distance Functions. *Transactions of the American Mathematical Society*, 352(11):5231–5249.
- [48] Rataj, J. and Zähle, M. (2019). Sets with positive reach. In *Curvature Measures of Singular Sets*, pages 55–86. Springer.
- [49] Rodríguez Casal, A. and Saavedra-Nieves, P. (2016). A fully data-driven method for estimating the shape of a point cloud. *ESAIM: Probability and Statistics*, 20:332–348.
- [50] Serra, J. (1984). *Image Analysis and Mathematical Morphology*. Academic Press.
- [51] Sullivan, C. and Kaszynski, A. (2019). PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK). *Journal of Open Source Software*, 4(37):1450.
- [52] Thäle, C. (2008). 50 years sets with positive reach—a survey. *Surveys in Mathematics and its Applications*, 3:123–165.
- [53] Walther, G. (1999). On a Generalization of Blaschke’s Rolling Theorem and the Smoothing of Surfaces. *Mathematical Methods in the Applied Sciences*, 22:301–316.