Found Comput Math
DOI 10.1007/s10208-009-9056-2

FOUNDATIONS OF COMPUTATIONAL MATHEMATICS
The Journal of the Society for the Foundations of Computational Mathematics

Boundary Measures for Geometric Inference

Frédéric Chazal · David Cohen-Steiner · Quentin Mérigot

Received: 6 February 2009 / Revised: 29 October 2009 / Accepted: 29 October 2009

Abstract We study the *boundary measures* of compact subsets of the d-dimensional Euclidean space, which are closely related to Federer's *curvature measures*. We show that they can be computed efficiently for point clouds and suggest that these measures can be used for geometric inference. The main contribution of this work is the proof of a quantitative stability theorem for boundary measures using tools of convex analysis and geometric measure theory. As a corollary we obtain a stability result for Federer's curvature measures of a compact set, showing that they can be reliably estimated from point-cloud approximations.

Keywords Geometric inference · Curvature measures · Convex functions · Nearest neighbors · Point clouds

Mathematics Subject Classification (2000) 52A39 · 51K10 · 49Q15

1 Introduction

Recently there has been a growing interest in applying geometric methods to data analysis. This approach uses well-known geometric or topological properties such as

Communicated by Herbert Edelsbrunner.

F. Chazal
Geometrica, INRIA Saclay, Orsay, France

D. Cohen-Steiner · Q. Mérigot (✉)
Geometrica, INRIA Sophia-Antipolis, Sophia-Antipolis, France
e-mail: Quentin.Merigot@sophia.inria.fr

Published online: 08 December 2009

curvature and intrinsic dimension in order to describe and understand the structure of data represented by high dimensional point clouds.

The general problem of *geometric inference* can be stated as follows: given a noisy point-cloud approximation C of a compact set $K \subseteq \mathbb{R}^d$, how can we recover geometric and topological information about K, such as its curvature, sharp edges, boundaries, Betti numbers, or Euler–Poincaré characteristic, etc., knowing only the point cloud C?

Previous Work For smooth surfaces embedded in $\mathbb{R}^3$, inference problems are by now well mastered. In particular, several algorithms (e.g. [2]) allow us to reconstruct a topologically correct and geometrically close piecewise-linear approximation from a sufficiently dense point cloud. From these reconstructions, differential quantities such as curvature tensors can be reliably estimated [10]. Most of these surface reconstruction algorithms are based on the study of the shape of Voronoi cells for points sampled on smooth surfaces. For smooth submanifolds in higher dimensional spaces, similar ideas lead to provable local dimension estimation algorithms [12]. It has also been shown that appropriate offsets, or equivalently α-shapes [13], provided reconstructions with the correct homotopy type [16] under sampling conditions similar to the one used in 3 dimensional surface reconstruction. This last result was recently extended to compact sets more general than smooth submanifolds [6], using a weaker sampling condition. Unfortunately, all these approaches are impractical in high dimensions since they require computing the Voronoi diagram, which has exponential complexity with respect to ambient dimension. Finally, it has been shown that persistent homology can be used to reliably estimate Betti numbers of a wide class of compact sets under an even weaker sampling condition [8, 9]. Computing persistent homology used to require the computation of a Delaunay triangulation, which is impractical in high dimension. However, approximate computations are made possible by using witness complexes [11]. This approach proved useful in particular in the study of the space of natural images [5].

As described above, substantial progress was recently made for inferring topological invariants of possibly nonsmooth sampled objects embedded in arbitrary dimensional spaces. However, very little is known on how to infer more geometric invariants, such as curvatures or singularities, for such objects. In this paper, we address this problem from a particular angle. Given a compact set $K \subseteq \mathbb{R}^d$, our approach consists in exploiting the geometric information contained in the growth of the volume of its offsets $K^r = \{x \in \mathbb{R}^d ;\, \mathrm{d}(x, K) \leqslant r\}$. The well-known tube formula states that for smooth [21] or convex [19] objects, this volume is a polynomial in r provided r is small enough. The coefficients of this polynomial, called intrinsic volumes, Lipschitz–Killing curvatures, Minkowski functionals, or Quermassintegrale in the literature, encode important geometric information about K, such as dimension, curvatures, angles of sharp features, or even Euler characteristic. Federer [14] later showed that such a polynomial behavior actually always holds as long as r does not exceed the *reach* of K, which is well known in computational geometry as the (minimum of) the *local feature size*, that is, the minimum distance between K and its medial axis. Also, he introduced the important notion of *curvature measure*, which, loosely speaking, details the contribution of each part of K to the intrinsic volumes, and hence gives local information about curvatures, dimensions, and sharp feature angles.

Contributions Curvature measures have been studied extensively in differential geometry [21] and geometric measure theory [14]. In this paper, we show that they can also prove useful from the perspective of geometric inference.

Given a compact set $K \subseteq \mathbb{R}^d$, we define the *boundary measure* of K at scale r as the mass distribution μ_{K,K^r} on K such that $\mu_{K,K^r}(B)$ is the contribution of a region $B \subset K$ to the volume of K^r. The tube formula states that, when K has positive reach, the curvature measures of K can be retrieved if one knows the boundary measures at several scales. The usability of boundary measures for geometric inference depends on the following two questions:

1. is it practically feasible to compute the boundary measure of a point cloud $C \subseteq \mathbb{R}^d$?
2. if C is a good approximation of K (i.e. dense enough and without too much noise), does the boundary measure μ_{C,C^r} carry approximately the same geometric information as μ_{K,K^r}?

The answer to the first question is given in Sect. 5 in the form of a very simple Monte Carlo algorithm allowing us to compute the boundary measure of a point cloud C embedded in the space $\mathbb{R}^d$. Standard arguments show that if C has n points, an ε-approximation of μ_{C,C^r} can be obtained with high probability (e.g. 99%) in time $\mathrm{O}(dn^2 \ln(1/\varepsilon)/\varepsilon^2)$ without using any sophisticated data structure. A more careful analysis shows that the n^2 behavior can be replaced by n times an appropriate covering number of C, which indicates that the cost is linear both in n and d for low entropy point clouds. Hence this algorithm is practical, at least for moderate size point clouds in high dimension.

The main contribution of this article is the proof of a stability theorem for boundary measures, which also gives a positive and quantitative answer to the second question. The following statement is a simplified version of the theorem. The set of compact subsets of $\mathbb{R}^d$ is endowed with the Hausdorff distance $\mathrm{d_H}$, and the set of mass distributions on $\mathbb{R}^d$ with the bounded-Lipschitz distance $\mathrm{d_{bL}}$ (which will be defined in Sect. 4).

Theorem 1.1 *For every compact set $K \subseteq \mathbb{R}^d$, there exists some constant $C(K)$ such that*

$$\mathrm{d_{bL}}(\mu_{K,K^r}, \mu_{C,C^r}) \leqslant C(K)\mathrm{d_H}(C, K)^{1/2}$$

as soon as $\mathrm{d_H}(C, K)$ is small enough.

In the sequel we will make this statement more precise by giving explicit constants. A similar stability result for a generalization of curvature measures is deduced from this theorem. At the heart of these two stability results is an L^1 stability property for (closest point) projections on compact sets. The proof of the projection stability theorem involves a new inequality in convex analysis, which may also be interesting in its own right. Recall that a subset S of $\mathbb{R}^d$ is called $(d-1)$-*rectifiable* if it can be written as a countable union of patches of Lipschitz hypersurfaces (see e.g. [15]).

Theorem 1.2 *Let E be an open subset of $\mathbb{R}^d$ with $(d-1)$-rectifiable boundary, and f, g be two convex functions such that* $\operatorname{diam}(\nabla f(E) \cup \nabla g(E)) \leqslant k$. *Then there exists a constant $C(d, E, k)$ such that for $\|f-g\|_\infty$ small enough,*

$$\|\nabla f - \nabla g\|_{\mathrm{L}^1(E)} \leqslant C(d, E, k)\|f-g\|_\infty^{1/2}.$$

Outline In Sect. 2 we introduce the mathematical objects involved in this article, in particular Hausdorff measures, boundary, and curvature measures. In Sect. 3 we prove the main technical result of this paper concerning the gradient of convex functions and show how this leads to the projection stability theorem as well as to the stability of a variant of boundary measures. Section 4 presents the stability result for boundary measures, from which one can deduce that curvature measures are also stable for the Hausdorff distance. In Sect. 5 we show how Monte Carlo methods can be applied to the computation of boundary measures of point clouds.

2 Mathematical Preliminaries

In this section, we introduce the mathematical objects which will be used throughout the article.

Measures A (nonnegative) *measure* μ associates to any (Borel) subset B of $\mathbb{R}^d$ a nonnegative number $\mu(B)$. It should also enjoy the following additivity property: if (B_i) is a countable family of disjoint subsets, then $\mu(\bigcup B_i) = \sum_i \mu(B_i)$. A measure on $\mathbb{R}^d$ should be seen as a mass distribution on $\mathbb{R}^d$, which one can probe using Borel sets: $\mu(B)$ is the mass of μ contained in B. The *restriction* of a measure μ to a subset $C \subseteq \mathbb{R}^d$ is the measure $\mu_{|C}$ defined by $\mu_{|C}(B) = \mu(B \cap C)$. The simplest examples of measures are Dirac masses which, given a point $x \in \mathbb{R}^d$, are defined by $\delta_x(B) = 1$ if and only if x is in B. In what follows, we will also use the k-dimensional *Hausdorff measures* $\mathcal{H}^k$ which, loosely speaking, associate to any k-dimensional subset of $\mathbb{R}^d$ its k-dimensional area $(0 \leqslant k \leqslant d)$. We refer to [15] for more details. For example, if $S \subset \mathbb{R}^d$ is k-dimensional, $\mathcal{H}^k_{|S}$ models a mass distribution uniformly distributed on S.

Boundary and Curvature Measures of a Compact Set We begin with some well-known facts about the distance function and projections on a compact set $K \subseteq \mathbb{R}^d$. The distance to K is defined as $\mathrm{d}_K(x) = \min_{y \in K}\|x-y\|$. The *medial axis* of K, denoted by $\mathcal{M}(K)$, is the set of points of $\mathbb{R}^d \setminus K$ which admit more than one closest point on K (see Fig. 1). The *projection* on K maps any point x outside the medial axis to its unique closest point on K, $\mathrm{p}_K(x)$. Since medial axes always have zero $\mathcal{H}^d$ measure, projections are defined almost everywhere. The *r-offset* of a subset $K \subseteq \mathbb{R}^d$ is the set of points at distance at most r from K, and is denoted by K^r.

As mentioned in the introduction, there is a lot of geometric information lying in the intensity of the contribution of a part of K to the volume of K^r. To make this statement precise, we introduce the notion of boundary measure.

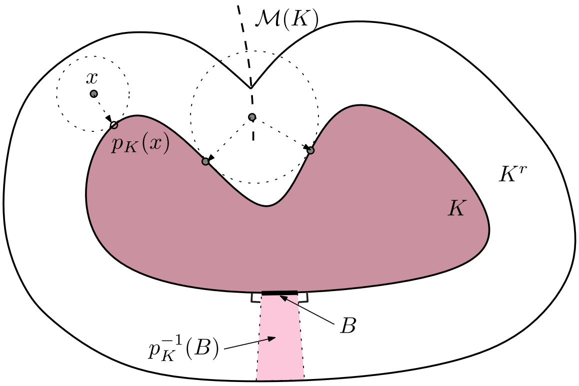

Fig. 1 Boundary measure of $K \subset \mathbb{R}^d$

Definition If K is a compact subset and E a domain of $\mathbb{R}^d$, the *boundary measure* $\mu_{K,E}$ is defined as follows: for any subset $B \subseteq \mathbb{R}^d$, $\mu_{K,E}(B)$ is the d-volume of the set of points of E whose projection on K is in B, i.e.

$$\mu_{K,E}(B) = \mathcal{H}^d\big(\mathrm{p}_K^{-1}(B \cap K) \cap E\big).$$

We will be particularly interested in the boundary measure μ_{K,K^r}, see Fig. 1. While the above definition makes sense for any compact $K \subseteq \mathbb{R}^d$, boundary measures have been mostly studied in the convex case and in the smooth case. Let us first give two examples in the convex case. Let S be a unit-length segment in the plane with endpoints a and b. The set S^r is the union of a rectangle of dimension $1 \times 2r$ whose points project on the segment and two half-disks of radius r whose points are projected on a and b. It follows that

$$\mu_{S,S^r} = 2r\mathcal{H}^1\big|_S + \frac{\pi}{2}r^2\delta_a + \frac{\pi}{2}r^2\delta_b.$$

If P is a convex solid polyhedron of $\mathbb{R}^3$, F its faces, E its edges, and V its vertices, then one can see that

$$\mu_{P,P^r} = \mathcal{H}^3\big|_P + r\sum_{f\in F}\mathcal{H}^2\big|_f + r^2\sum_{e\in E}K(e)\mathcal{H}^1\big|_e + r^3\sum_{v\in V}K(v)\delta_v$$

where $K(e)$ is the angle between the normals of the faces adjacent to the edge e and $K(v)$ the solid angle formed by the normals of the faces adjacent to v. As shown by Steiner and Minkowski, for general convex polyhedra the measure μ_{K,K^r} can be written as a sum of weighted Hausdorff measures supported on the i-skeleton of K, and whose local density is the local external dihedral angle.

Weyl [21] proved that the polynomial behavior for $r \mapsto \mu_{K,K^r}$ shown above is also true for small values of r when K is a compact and smooth submanifold of $\mathbb{R}^d$. Moreover he also proved that the coefficients of this polynomial can be computed from the second fundamental form of K. For example, if K is a d-dimensional subset with smooth boundary, then we have for sufficiently small $r > 0$ and for all $B \subset K$,

$$\mu_{K,K^r}(B) = \sum_{i=0}^{d} \omega_i \, \Phi_K^{d-i}(B)\, r^i \tag{2.1}$$

where $\Phi_K^d(B)$ is (half of) the d-volume of B and Φ_K^k $(k < d)$ are the (signed) measures with density given by symmetric functions of the principal curvatures of ∂K:

$$\Phi_K^k(B) = \mathrm{const}(d,k) \int_B \left[\sum_{i(1)<\cdots<i(k)} \kappa_{i(1)}(p) \ldots \kappa_{i(k)}(p) \right] \mathrm{d}p. \tag{2.2}$$

Hence the principal curvatures of K can be retrieved from the curvature measures Φ_K^k—at least in principle. Indeed, for every $p \in \partial K$ one has

$$\sigma_k(p) := \sum_{i(1)<\cdots<i(k)} \kappa_{i(1)}(p) \ldots \kappa_{i(k)}(p) = \mathrm{const}(d,k) \lim_{r\to 0} \frac{\Phi_K^k(\mathrm{B}(p,r))}{\mathcal{H}^{d-1}(\partial K \cap \mathrm{B}(p,r))}.$$

Knowing all the elementary symmetric functions $\sigma_1(p), \ldots, \sigma_{d-1}(p)$ of the principal curvatures at p is enough to recover them (up to a permutation). Formulas similar to (2.1)–(2.2) exist for higher codimension submanifolds (see e.g. [21]). In general, if K has intrinsic dimension n, then the measures Φ_K^k vanish identically when $k > n$. Hence, these measures also encode the dimension of K.

In [14], Federer generalized Steiner's and Weyl's tube formula to the class of compact sets with positive *reach*. A compact set $K \subseteq \mathbb{R}^d$ is said to have a reach greater than R if the minimum distance between K and its medial axis is greater than R, i.e. if the projection p_K is well defined everywhere in the interior of K^R. What Federer showed is that for any compact set K with reach at least R, there exists a unique set of signed measures Φ_K^k, which he called *curvature measures*, such that (2.1) holds. From the discussion above, we see that curvature measures encode local dimensions, angles of sharp features of various dimensions, and principal curvatures. The following generalization of the Gauss–Bonnet theorem shows that they also contain topological information about K.

Theorem *Given any compact set K with positive reach, $\Phi_K^0(K)$ is equal to the Euler–Poincaré characteristic of K.*

The curvature measures can be retrieved from the knowledge of boundary measures μ_{K,K^r} by polynomial fitting (cf. Sect. 4). Hence all the information contained in curvature measures is also contained in boundary measures, which is our main motivation for studying them.

3 Stability of Boundary Measures

In this section, we suppose that E is a fixed open set with rectifiable boundary, and we obtain a quantitative stability theorem for the map $K \mapsto \mu_{K,E}$. What we mean by *stable* is that if two compact sets K and K' are close, then the measures $\mu_{K,E}$ and $\mu_{K',E}$ are also close. In order to be able to formulate a precise statement, we need to choose a notion of distance on the space of compact subsets of $\mathbb{R}^d$ and on the set of measures on $\mathbb{R}^d$.

To measure the distance between two compact subsets K and K' of $\mathbb{R}^d$, we will use the Hausdorff distance: $\mathrm{d_H}(K, K')$ is by definition the smallest positive constant η such that both $K' \subseteq K^\eta$ and $K \subseteq K'^\eta$. It is also the uniform distance between the two distance functions d_K and $\mathrm{d}_{K'}$: $\mathrm{d_H}(K, K') = \sup_{x\in\mathbb{R}^d} |\mathrm{d}_K(x) - \mathrm{d}_{K'}(x)|$. The next paragraph describes the distance we use to compare measures.

Wasserstein Distance The *Wasserstein distance* (of exponent 1) between two measures μ and ν on $\mathbb{R}^d$ having the same total mass $\mu(\mathbb{R}^d) = \nu(\mathbb{R}^d)$ is a nonnegative number which quantifies the cost of the optimal transportation from the mass distribution defined by μ to the mass distribution defined by μ (cf. [20]). It is denoted by $\mathrm{d_W}(\mu, \nu)$. More precisely, it is defined as

$$\mathrm{d_W}(\mu, \nu) = \inf_{X,Y} \mathbb{E}\big[d(X, Y)\big]$$

where the infimum is taken on all pairs of $\mathbb{R}^d$-valued random variables X and Y whose law are μ and ν respectively. Notice that when λ is a finite measure on a space X whose mass is not one, the expectation $\mathbb{E}_\lambda(f)$ should be interpreted as the unnormalized mean $\int_X f(x)\,\mathrm{d}\lambda(x)$.

This distance is also known as the earth-mover distance, and has been used in vision [17] and image retrieval [18]. One of the interesting properties of the Wasserstein distance is the Kantorovich–Rubinstein duality theorem. Recall that a function $f : \mathbb{R}^d \to \mathbb{R}$ is 1-Lipschitz if for every choice of x and y, $|f(x) - f(y)| \leqslant \|x - y\|$.

Theorem (Kantorovich–Rubinstein) *If μ and ν are two probability measures on $\mathbb{R}^d$, then*

$$\mathrm{d_W}(\mu, \nu) = \sup_f \left| \int f\,\mathrm{d}\mu - \int f\,\mathrm{d}\nu \right|$$

where the supremum is taken on all 1-Lipschitz function in $\mathbb{R}^d$.

The Lipschitz function f in the theorem can be thought of as a way of probing the measure μ. For example, if f is a tent function, e.g. $f(y) = (1 - \|x - y\|)_+$, then $\int f\,\mathrm{d}\mu$ gives an information about the local density of μ near x. The Kantorovich–Rubinstein theorem asserts that if two measures μ and ν are Wasserstein-close, then one can control the probing error between μ and ν by Lipschitz functions.

The following proposition reduces the stability result of the map $K \mapsto \mu_{K,E}$ with respect to the Wasserstein distance to a stability result for the map $K \mapsto \mathrm{p}_K|_E$.

Proposition 3.1 *If E is a subset of $\mathbb{R}^d$, and K and K' two compact sets, then*

$$\mathrm{d_W}(\mu_{K,E}, \mu_{K',E}) \leqslant \int_E \big\|\mathrm{p}_K(x) - \mathrm{p}_{K'}(x)\big\|\,\mathrm{d}x.$$

Proof Let Z be a random variable whose law is $\mathcal{H}^d|_E$. Then, $X = \mathrm{p}_K \circ Z$ and $Y = \mathrm{p}_{K'} \circ Z$ have law $\mu_{K,E}$ and $\mu_{K',E}$ respectively. Hence by definition, $\mathrm{d_W}(\mu_{K,E}, \mu_{K',E}) \leqslant \mathbb{E}(\mathrm{d}(\mathrm{p}_K \circ Z, \mathrm{p}_{K'} \circ Z))$, which is the desired bound. □

A L^1 *Stability Theorem for Projections* From now on, if E is a subset of $\mathbb{R}^d$, f an integrable function on E, and g a continuous function on E, we define the two following norms:

$$\|f\|_{\mathrm{L}^1(E)} = \int_E \|f(x)\| \,\mathrm{d}x \quad \text{and} \quad \|g\|_{\infty,E} = \sup_{x\in E} \|g(x)\|.$$

The key result of this paper is the following upper bound for the L^1-norm $\|\mathrm{p}_K - \mathrm{p}_{K'}\|_{\mathrm{L}^1(E)}$, when K and K' are two Hausdorff-close compact sets.

Theorem 3.2 (Projection Stability) *Let E be an open set in $\mathbb{R}^d$ with rectifiable boundary, and K and K' be two compact subsets of $\mathbb{R}^d$ and $R_K = \|d_K\|_{\infty,E}$. Then, there is a constant $C_1(d)$ such that*

$$\|\mathrm{p}_K - \mathrm{p}_{K'}\|_{\mathrm{L}^1(E)} \leqslant C_1(d)\big[\mathcal{H}^d(E) + \mathrm{diam}(K)\mathcal{H}^{d-1}(\partial E)\big] \times \sqrt{R_K \mathrm{d_H}(K, K')}$$

assuming $\mathrm{d_H}(K, K') \leqslant \min(R_K, \mathrm{diam}(K), \mathrm{diam}(K)^2/R_K)$.

The question of whether projection maps are stable has already drawn attention in the past. In particular, Federer proved the following result in [14].

Theorem *Let R be a positive number, and $K_n \subseteq \mathbb{R}^d$ be a sequence of compact sets whose reach is greater than R, which Hausdorff-converges to some compact $K \subseteq \mathbb{R}^d$ with* $\mathrm{reach}(K) > R$. *Then p_{K_n} converges to p_K uniformly on K^R as n grows to infinity.*

A drawback of this theorem is that it does not say anything about the speed of convergence. But more importantly, the very strong assumptions that all compact sets in the sequence have their reach bounded from below makes it completely unusable in the setting of geometric inference: indeed, since the reach of a point cloud is the minimum distance between any two of its points, if a sequence of point clouds C_n converges to some nondiscrete compact set K (e.g. a segment), then necessarily $\lim_n \mathrm{reach}(C_n) = 0$. In fact, if K is the union of two distinct points x and y, and $K_n = \{x + \frac{1}{n}(x - y), y\}$, p_{K_n} does not converge uniformly to p_K near the medial hyperplane of x and y. So one cannot hope for a generalization of the above theorem to a uniform convergence result of p_{K_n} to p_K on an arbitrary open set E.

From the stability of the gradient of the distance function (see [6]) one can deduce that the projections p_K and p_{K_n} can differ dramatically only near the medial axis of K. This makes it reasonable to expect a L^1 convergence property of the projections, i.e. if K_n converges to K, then

$$\lim_n \int_E \|\mathrm{p}_{K_n}(x) - \mathrm{p}_K(x)\| \,\mathrm{d}x = 0.$$

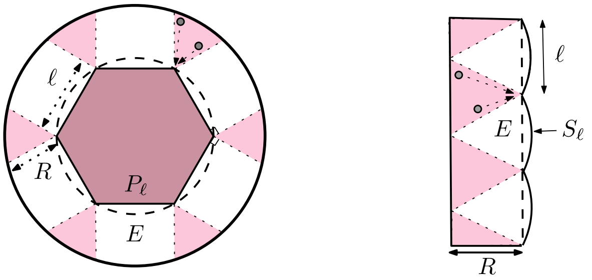

Fig. 2 Optimality of the stability theorem

Unfortunately, for a generic compact set, the medial axis $\mathcal{M}(K)$ is dense in $\mathbb{R}^d$ (see [22] for a proof). In this case, every point in $\mathbb{R}^d$ is (arbitrarily) close to a point in $\mathcal{M}(K)$, and the above remark is of little help. The actual proof of the Projection Stability Theorem relies on a new theorem in convex analysis, and is postponed to the next section.

Let us now comment on the optimality of this Projection Stability Theorem. First, the speed of convergence of $\mu_{K',E}$ to $\mu_{K,E}$ cannot be (in general) faster than $\mathrm{O}(\mathrm{d_H}(K, K')^{1/2})$. Indeed, if D is the closed unit disk in the plane, and P_ℓ is a regular polygon inscribed in D with sidelength ℓ, then $\mathrm{d_H}(D, P_\ell) \simeq \ell^2$. Now we let E be the disk of radius $1 + R$. Then, a constant fraction of the mass of E will be projected onto the vertices of P_ℓ by the projection p_{P_ℓ} (lightly shaded area in Fig. 2). Now, the cost of spreading out the mass concentrated on these vertices to get a uniform measure on the circle is proportional to the distance between consecutive vertices, so $\mathrm{d_W}(\mu_{D,E}, \mu_{P_\ell,E}) = \Omega(\ell)$. Hence, by Proposition 3.1, $\|\mathrm{p}_D - \mathrm{p}_{P_\ell}\|_{\mathrm{L}^1(E)} = \Omega(\ell) = \Omega(\mathrm{d_H}(D, P_\ell)^{1/2})$. Note that this $\mathrm{O}(\mathrm{d_H}(K, K')^{1/2})$ behavior does not come from the curvature of the disk, since one can also find an example of a family of compacts S_ℓ made of small circle arcs converging to the unit segment S such that $\|\mathrm{p}_S - \mathrm{p}_{S_\ell}\|_{\mathrm{L}^1(E)} = \Omega(\mathrm{d_H}(S, S_\ell)^{1/2})$ (see Fig. 2).

The second remark concerning the optimality of the theorem is that the second term of the bound involving $\mathcal{H}^{d-1}(\partial E)$ cannot be avoided. Indeed, let us suppose that a bound $\|\mathrm{p}_K - \mathrm{p}_{K'}\|_{\mathrm{L}^1(E)} \leqslant C(K)\mathcal{H}^d(E)\sqrt{\varepsilon}$ were true around K, where ε is the Hausdorff distance between K and K'. Now let K be the union of two parallel hyperplanes at distance R intersected with a big sphere centered at a point x of their medial hyperplane M. Let E_ε be a ball of radius $\varepsilon/2$ tangent to M at x and K_ε be the translate by ε of K along the common normal of the hyperplanes such that the ball E_ε lies in the slab bounded by the medial hyperplanes of K and K_ε. Then, $\|\mathrm{p}_K - \mathrm{p}_{K'}\|_{\mathrm{L}^1(E_\varepsilon)} \simeq R \times \mathcal{H}^d(E_\varepsilon)$, which exceeds the assumed bound for a small enough ε.

Proof of the Projection Stability Theorem The Projection Stability Theorem will follow from a more general theorem on the L^1 norm of the difference between the gradients of two convex functions defined on some open set E with rectifiable boundary. The connection between projections and convex analysis is that any projection p_K derives from a convex potential v_K.

Lemma 3.3 *The function* $v_K : \mathbb{R}^d \to \mathbb{R}$ *defined by* $v_K(x) = \|x\|^2 - \mathrm{d}_K(x)^2$ *is convex with gradient* $\nabla v_K = 2\mathrm{p}_K$ *almost everywhere.*

Proof By definition, $v_K(x) = \sup_{y\in K} \|x\|^2 - \|x - y\|^2 = \sup_{y\in K} v_{K,y}(x)$ with $v_{K,y}(x) = 2\langle x|y\rangle - \|y\|^2$. Hence v_K is convex as a supremum of affine functions, and is differentiable almost everywhere. Then, $\nabla v_K(x) = 2\mathrm{d}_K(x)\nabla_x \mathrm{d}_K - 2x$. The equality $\nabla \mathrm{d}_K(x) = (\mathrm{p}_K(x) - x)/\mathrm{d}_K(x)$, valid when x is not in the medial axis, concludes the proof. $\square$

Hence if K, K' are two compact subsets of $\mathbb{R}^d$, we have

$$\|\mathrm{p}_K - \mathrm{p}_{K'}\|_{\mathrm{L}^1(E)} = \frac{1}{2}\|\nabla v_K - \nabla v_{K'}\|_{\mathrm{L}^1(E)}.$$

Moreover, denoting $R_K = \|\mathrm{d}_K\|_{\infty,E}$, the two following properties for the functions v_K and $v_{K'}$ can be deduced from a simple calculation.

Lemma 3.4 *If* $\mathrm{d_H}(K, K') \leqslant \min(R_K, \mathrm{diam}(K))$, *then*

$$\|v_K - v_{K'}\|_{\infty,E} := \sup_{x\in E}\big|\mathrm{d}_K(x)^2 - \mathrm{d}_{K'}(x)^2\big| \leqslant 3\mathrm{d_H}(K, K')R_K,$$

$$\mathrm{diam}\big(\nabla v_K(E) \cup \nabla v_{K'}(E)\big) \leqslant 3\,\mathrm{diam}(K).$$

Let us now state the stability result for gradients of convex functions.

Theorem 3.5 *Let E be an open subset of $\mathbb{R}^d$ with rectifiable boundary, and f, g be two locally convex functions on E such that* $\mathrm{diam}(\nabla f(E) \cup \nabla g(E)) \leqslant k$. *Then,*

$$\begin{aligned}&\|\nabla f - \nabla g\|_{\mathrm{L}^1(E)}\\ &\quad\leqslant C_2(n) \times \big(\mathcal{H}^d(E) + \big(k + \|f - g\|_{\infty,E}^{1/2}\big)\mathcal{H}^{d-1}(\partial E)\big)\|f - g\|_{\infty,E}^{1/2}.\end{aligned}$$

We note that this result may be viewed as the converse of classical inequalities (*Poincaré inequalities*) stating that if the gradients of two functions are close, then the two functions are also (up to an additive constant) close. Such converse results (*reverse Poincaré inequalities*) have been the subject of intense research in functional analysis (e.g. [3, 4]), but in a rather different setting.

The Projection Stability Theorem is easily deduced from Theorem 3.5 and Lemmas 3.3 and 3.4. We now turn to the proof of the 1-dimensional case of Theorem 3.5. The general case will follow using an argument of *integral geometry*—i.e., we will integrate the 1-dimensional inequality over the set of lines of $\mathbb{R}^d$ and use the Cauchy–Crofton formulas (3.1) and (3.2) below to get the d-dimensional inequality.

Cauchy–Crofton formulas give a way to compute the volume of a set E in terms of the expectation of the length of $E \cap \ell$ where ℓ is a random line in $\mathbb{R}^d$. More precisely, if one denotes by $\mathcal{L}^d$ the set of oriented lines of $\mathbb{R}^d$, then denoting by $\mathrm{d}\mathcal{L}^d$ the properly normalized rigid motion invariant measure on $\mathcal{L}^d$, we have

$$\mathcal{H}^d(E) = \frac{1}{\omega_{d-1}} \int_{\ell\in\mathcal{L}^d} \mathrm{length}(\ell \cap E)\,\mathrm{d}\mathcal{L}^d \tag{3.1}$$

where ω_{d-1} is the $(d-1)$-volume of the unit sphere in $\mathbb{R}^d$. If S is a rectifiable hypersurface of $\mathbb{R}^d$—i.e., a countable union of patches of Lipschitz hypersurfaces,

Fig. 3 Proof of the 1-dimensional inequality

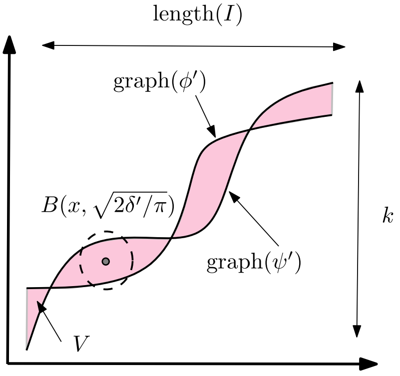

see [15]—then

$$\mathcal{H}^{d-1}(S) = \frac{1}{2\beta_{d-1}} \int_{\ell \in \mathcal{L}^d} \#(\ell \cap S)\, \mathrm{d}\mathcal{L}^d \tag{3.2}$$

where β_{d-1} is the $(d-1)$-volume of the unit ball in $\mathbb{R}^{d-1}$.

Proposition 3.6 *If I is an interval, and $\varphi : I \to \mathbb{R}$ and $\psi : I \to \mathbb{R}$ are two convex functions such that* $\operatorname{diam}(\varphi'(I) \cup \psi'(I)) \leqslant k$, *then letting* $\delta = \|\varphi - \psi\|_{\mathrm{L}^\infty(I)}$,

$$\int_I |\varphi' - \psi'| \leqslant 8\pi\big(\operatorname{length}(I) + k + \delta^{1/2}\big)\delta^{1/2}.$$

Proof Since φ and ψ are convex, their derivatives are nondecreasing. Let V be the closure of the set of points (x, y) in $I \times \mathbb{R}$ such that y is in the segment $[\varphi'(x), \psi'(x)]$ (or $[\psi'(x), \varphi'(x)]$ if $\varphi'(x) \geqslant \psi'(x)$). By the definition of V, $\int_I |\varphi' - \psi'| = \mathcal{H}^2(V)$.

If D is a disk included in V and $[x_0, x_1] \subset I$ is the projection of D on the x-axis, then the sign of the derivative of the difference $\kappa = \varphi - \psi$ does not change on $[x_0, x_1]$. Assuming w.l.o.g. that κ is nondecreasing on $[x_0, x_1]$, we have $|\kappa(x_0) - \kappa(x_1)| = \int_{x_0}^{x_1} |\kappa'| \geqslant \mathcal{H}^2(D)$.

But since $\|\kappa\|_\infty = \delta$, the area of D cannot be greater than 2δ. Thus, if p is any point of V, for any $\delta' > \delta$ the disk $B(p, \sqrt{2\delta'/\pi})$, whose area is $2\delta'$, necessarily intersects the boundary ∂V (see Fig. 3). This proves that V is contained in the offset $(\partial V)^{\sqrt{2\delta/\pi}}$. It follows that

$$\int_I |\varphi' - \psi'| \leqslant \mathcal{H}^2\big((\partial V)^{\sqrt{2\delta/\pi}}\big). \tag{3.3}$$

Now, ∂V can be written as the union of two xy-monotone curves Φ and Ψ joining the lower left corner of V and the upper right corner of V so that $(\partial V)^r \subseteq \Phi^r \cup \Psi^r$.

We now find a bound for $\mathcal{H}^2(\Phi^r)$ (the same bound will of course apply to Ψ). Since the curve Φ is xy-monotone, we have $\operatorname{length}(\Phi) \leqslant \operatorname{length}(I) + k$. Thus, for any $r > 0$ there exists a subset $X \subseteq \Phi$ of $N = \lceil(\operatorname{length}(I) + k)/r\rceil$ points such that $\Phi \subseteq X^r$, implying

$$\mathcal{H}^2\big(\Phi^r\big) \leqslant \mathcal{H}^2\big(X^{2r}\big) \leqslant 4\pi r^2 N \leqslant 4\pi r\big(\operatorname{length}(I) + k + r\big). \tag{3.4}$$

Using (3.3) and (3.4), and $\sqrt{2\delta/\pi} \leqslant \sqrt{\delta}$, one finally obtains

$$\int_I |\varphi' - \psi'| \leqslant \mathcal{H}^2\big(\Phi^{\sqrt{2\delta/\pi}} \cup \Psi^{\sqrt{2\delta/\pi}}\big) \leqslant 8\pi\big(\mathrm{length}(I) + k + \sqrt{\delta}\big)\sqrt{\delta}. \qquad \square$$

Proof of Theorem 3.5 The 1-dimensional case follows directly from Proposition 3.6: in this case, E is a countable union of intervals on which f and g satisfy the hypothesis of the proposition. Summing the inequalities gives the result with $C_2(1) = 8\pi$.

We now turn to the general case. Given any L^1 vector field X one has

$$\int_E \|X\|\,\mathrm{d}x = \frac{d}{2\omega_{d-2}} \int_{\ell\in\mathcal{L}^d} \int_{y\in\ell\cap E} \big|\big\langle X(y)|u(\ell)\big\rangle\big|\,\mathrm{d}y\,\mathrm{d}\ell$$

where $u(\ell)$ is a unit directing vector for ℓ (see Lemma III.4 in [7] for a proof of this formula). Letting $X = \nabla f - \nabla g$, $f_\ell = f|_{\ell\cap E}$ and $g_\ell = g|_{\ell\cap E}$, one gets, with $D(d) = d/(2\omega_{d-2})$,

$$\begin{aligned}\int_E \|\nabla f - \nabla g\| &= D(d) \int_{\ell\in\mathcal{L}^d} \int_{y\in\ell\cap E} \big|\big\langle \nabla f - \nabla g|u(\ell)\big\rangle\big|\,\mathrm{d}y\,\mathrm{d}\ell \\ &= D(d) \int_{\ell\in\mathcal{L}^d} \int_{y\in\ell\cap E} |f'_\ell - g'_\ell|\,\mathrm{d}y\,\mathrm{d}\ell.\end{aligned}$$

The functions f_ℓ and g_ℓ satisfy the hypothesis of the 1-dimensional case, so that for each choice of ℓ, and with $\delta = \|f - g\|_{\mathrm{L}^\infty(E)}$,

$$\int_{y\in\ell\cap E} |f'_\ell - g'_\ell|\,\mathrm{d}y \leqslant 8\pi\delta^{1/2} \times \big(\mathcal{H}^1(E\cap\ell) + \big(k + \delta^{1/2}\big)\mathcal{H}^0(\partial E\cap\ell)\big).$$

It follows by integration on $\mathcal{L}^d$ that

$$\begin{aligned}&\int_E \|\nabla f - \nabla g\| \\ &\quad\leqslant 8\pi D(d)\delta^{1/2} \times \left(\int_{\mathcal{L}^d} \mathcal{H}^1(E\cap\ell)\,\mathrm{d}\mathcal{L}^d + \big(k + \delta^{1/2}\big)\int_{\mathcal{L}^d} \mathcal{H}^0(\partial E\cap\ell)\,\mathrm{d}\mathcal{L}^d\right).\end{aligned}$$

The formulas (3.1) and (3.2) show that the first integral in the second term is equal (up to a constant) to the volume of E and the second one to the $(d-1)$-measure of ∂E. This proves the theorem with $C_2(d) = 8\pi D(d)(\omega_{d-1} + 2\beta_{d-1})$. $\square$

3.1 Stability of the Pushforward of a Measure by a Projection

The boundary measures defined above are a special case of *pushforward* of a measure by a function. The pushforward of a measure μ on $\mathbb{R}^d$ by the projection p_K is another measure, denoted by $\mathrm{p}_{K\#}\mu$, concentrated on K and defined by the formula $\mathrm{p}_{K\#}\mu(B) = \mu(\mathrm{p}_K^{-1}(B))$.

The stability results for the boundary measures $K \mapsto \mu_{K,E}$ can be generalized to prove the stability of the map $K \mapsto \mathrm{p}_{K\#}\mu$ where μ has a density $u : \mathbb{R}^d \to \mathbb{R}^+$, which means that $\mu(B) = \int_B u(x)\,\mathrm{d}x$. We need the measure μ to be finite, which is the same as asking that the function u belong to the space $\mathrm{L}^1(\mathbb{R}^d)$ of integrable functions.

We also need the function $u \in \mathrm{L}^1(\mathbb{R}^d)$ to have *bounded variation*. We recall the following basic facts of the theory of functions with bounded variation, taken from [1]. If u is an integrable function on $\mathbb{R}^d$, the *total variation* of u is

$$\operatorname{var}(u) = \sup\left\{\int_{\mathbb{R}^d} u \operatorname{div}\varphi;\, \varphi \in \mathcal{C}_c^1\big(\mathbb{R}^d\big),\, \|\varphi\|_\infty \leqslant 1\right\}.$$

A function $u \in \mathrm{L}^1(\mathbb{R}^d)$ has bounded variation if $\operatorname{var}(u) < +\infty$. The set of functions of bounded variation on $\mathbb{R}^d$ is denoted by $\mathrm{BV}(\mathbb{R}^d)$. We also mention that if u is Lipschitz, then $\operatorname{var}(u) = \|\nabla u\|_{\mathrm{L}^1(\mathbb{R}^d)}$.

Theorem 3.7 *Let μ be a measure with density $u \in \mathrm{BV}(\mathbb{R}^d)$ with respect to the Lebesgue measure, and K be a compact subset of $\mathbb{R}^d$. We suppose that the support of u is contained in the offset K^R. Then, if $\mathrm{d}_H(K, K')$ is small enough,*

$$\begin{aligned}\mathrm{d_W}(\mathrm{p}_{K\#}\mu, \mathrm{p}_{K'\#}\mu) \leqslant C_2(n)\big(\|u\|_{\mathrm{L}^1(K^R)} + \operatorname{diam}(K)\operatorname{var}(u)\big)\\ \times \sqrt{R}\,\mathrm{d_H}(K, K')^{1/2}.\end{aligned}$$

Proof We begin with the additional assumption that u has class $\mathcal{C}^\infty$. The function u can be written as an integral over $t \in \mathbb{R}$ of the characteristic functions of its superlevel sets $E_t = \{u > t\}$, i.e. $u(x) = \int_0^\infty \chi_{E_t}(x)\,\mathrm{d}t$. Fubini's theorem then ensures that for any 1-Lipschitz function f defined on $\mathbb{R}^d$ with $\|f\|_\infty \leqslant 1$,

$$\begin{aligned}\mathrm{p}_{K'\#}\mu(f) &= \int_{\mathbb{R}^d} f \circ \mathrm{p}_{K'}(x)u(x)\,\mathrm{d}x\\ &= \int_{\mathbb{R}}\int_{\mathbb{R}^d} f \circ \mathrm{p}_{K'}(x)\chi_{\{u \geqslant t\}}(x)\,\mathrm{d}x\,\mathrm{d}t.\end{aligned}$$

By Sard's theorem, for almost any t, $\partial E_t = u^{-1}(t)$ is an $(n-1)$-rectifiable subset of $\mathbb{R}^d$. Thus, for those t the Projection Stability Theorem implies, for $\varepsilon = \mathrm{d_H}(K, K') \leqslant \varepsilon_0 = \min(R, \operatorname{diam}(K), \operatorname{diam}(K)^2/R_K)$,

$$\begin{aligned}\int_{E_t} \big|f \circ \mathrm{p}_K(x) - f \circ \mathrm{p}_{K'}(x)\big|\,\mathrm{d}x &\leqslant \|\mathrm{p}_K - \mathrm{p}_{K'}\|_{\mathrm{L}^1(E_t)}\\ &\leqslant C_2(n)\big[\mathcal{H}^d(E_t) + \operatorname{diam}(K)\mathcal{H}^{n-1}(\partial E_t)\big]\sqrt{R\varepsilon}.\end{aligned}$$

Putting this inequality into the last equality gives

$$\big|\mathrm{p}_{K\#}\mu(f) - \mathrm{p}_{K'\#}\mu(f)\big| \leqslant C_2(n)\bigg(\int_{\mathbb{R}} \mathcal{H}^d(E_t) + \operatorname{diam}(K)\mathcal{H}^{n-1}(\partial E_t)\,\mathrm{d}t\bigg)\sqrt{R\varepsilon}.$$

Using Fubini's theorem again and the coarea formula, one finally gets that

$$\big|\mathrm{p}_{K\#}\mu(f) - \mathrm{p}_{K'\#}\mu(f)\big| \leqslant C_2(n)\big(\|u\|_{\mathrm{L}^1(K^R)} + \operatorname{diam}(K)\operatorname{var}(u)\big)\sqrt{R\varepsilon}.$$

From the Kantorovich–Rubinstein theorem, this yields the desired inequality on $\mathrm{d_W}(p_{K\#}\mu(f), p_{K'\#}\mu(f))$, and concludes the proof of the theorem in the case of a $\mathcal{C}^\infty$ function u. To get the general case, one has to approximate the bounded variation function u by a sequence of $\mathcal{C}^\infty$ functions (u_n) such that both $\|u - u_n\|_{\mathrm{L}^1(K^R)}$ and $|\operatorname{var}(u) - \operatorname{var}(u_n)|$ converge to zero. This is always possible thanks to Theorem 3.9 in [1]. □

4 Stability of Curvature Measures

The definition of Wasserstein distance assumes that both measures are positive and have the same mass. While this is true for $\mu_{K,E}$ and $\mu_{K',E}$ (whose mass is the volume of E), this is not the case anymore when considering μ_{K,K^r} and μ_{K',K'^r} whose masses are respectively $\mathcal{H}^d(K^r)$ and $\mathcal{H}^d(K'^r)$. We thus need to introduce another distance on the space of (signed) measures.

Distance Between Two Boundary Measures The Kantorovich–Rubinstein theorem makes it natural to introduce the *bounded Lipschitz distance* between two measures μ and ν as follows:

$$\mathrm{d_{bL}}(\mu,\nu) = \sup_{f\in \mathrm{BL}_1(\mathbb{R}^d)} \left| \int f\,\mathrm{d}\mu - \int f\,\mathrm{d}\nu \right|$$

the supremum being taken on the space of 1-Lipschitz functions f on $\mathbb{R}^d$ such that $\sup_{\mathbb{R}^d} |f| \leqslant 1$. With this definition, one gets the following.

Proposition 4.1 *If K, K' are compact subsets of $\mathbb{R}^d$,*

$$\mathrm{d_{bL}}(\mu_{K,K^r}, \mu_{K',K'^r}) \leqslant \int_{K^r\cap K'^r} \|\mathrm{p}_K(x) - \mathrm{p}_{K'}(x)\|\,\mathrm{d}x + \mathcal{H}^d\big(K^r \Delta K'^r\big),$$

where $K^r \Delta K'^r$ is the symmetric difference between the offsets K^r and K'^r.

Proof Let φ be a 1-Lipschitz function on $\mathbb{R}^d$ bounded by 1. Using the change-of-variable formula one has

$$\begin{aligned}
&\left| \int \varphi(x)\,\mathrm{d}\mu_{K,K^r} - \int \varphi(x)\,\mathrm{d}\mu_{K',K'^r} \right| \\
&\quad = \left| \int_{K^r} \varphi\circ \mathrm{p}_K(x)\,\mathrm{d}x - \int_{K'^r} \varphi\circ \mathrm{p}_{K'}(x)\,\mathrm{d}x \right| \\
&\quad \leqslant \int_{K'^r\cap K^r} \big|\varphi\circ \mathrm{p}_K(x) - \varphi\circ \mathrm{p}_{K'}(x)\big|\,\mathrm{d}x + \int_{K\Delta K'} \big|\varphi(x)\big|\,\mathrm{d}x.
\end{aligned}$$

By the Lipschitz condition, $\|\varphi\circ \mathrm{p}_K(x) - \varphi\circ \mathrm{p}_{K'}(x)\| \leqslant |\mathrm{p}_K(x) - \mathrm{p}_{K'}(x)|$, thus giving the desired inequality. □

The new term appearing in this proposition involves the volume of the symmetric difference $K^r \Delta K'^r$. In order to get a result similar to the Projection Stability Theorem but for the map $K \mapsto \mu_{K,K^r}$, we need to study how fast this symmetric difference vanishes as K' converges to K. It is not hard to see that if $\mathrm{d_H}(K, K')$ is smaller than ε, then $K^r \Delta K'^r$ is contained in $K^{r+\varepsilon} \setminus K^{r-\varepsilon}$. Assuming that $\mathrm{d_H}(K, K') < \varepsilon$, using the coarea formula (see [15]), we can bound the volume of this annulus around K as follows:

$$\mathcal{H}^d\big(K^r \Delta K'^r\big) \leqslant \int_{r-\varepsilon}^{r+\varepsilon} \mathcal{H}^{d-1}\big(\partial K^s\big)\,\mathrm{d}s.$$

In the next paragraph we give a bound on the area of the boundary of the offset K^r, which we will then use to obtain a stability result for boundary measures μ_{K,K^r} and curvature measures.

Area of Offset Boundaries The next proposition gives a bound for the measure of the r-level set ∂K^r of a compact set $K \subseteq \mathbb{R}^d$ depending only on its covering number $\mathcal{N}(K,r)$. The covering number $\mathcal{N}(K,r)$ is defined as the minimal number of closed balls of radius r needed to cover K and is a way to measure the complexity of K—for instance, if K can be embedded in $\mathbb{R}^d$, then $\mathcal{N}(K,r) = \mathrm{O}(r^{-d})$. Precisely, we prove the following proposition, denoting by $\omega_{d-1}(r)$ the volume of the $(d-1)$-dimensional sphere of radius r in $\mathbb{R}^d$:

Proposition 4.2 *If K is a compact set in $\mathbb{R}^d$, for every positive r, ∂K^r is rectifiable and*

$$\mathcal{H}^{d-1}\big(\partial K^r\big) \leqslant \mathcal{N}(\partial K, r) \times \omega_{d-1}(2r).$$

We first prove Proposition 4.2 in the special case of *r-flowers*. An r-flower F is the boundary of the r-offset of a compact set contained in a ball $B(x,r)$, i.e. $F = \partial K^r$ where $K \subseteq B(x,r)$. The difference from the general case is that if $K \subseteq B(x,r)$, then K^r is star-shaped with respect to x (this will be established in the proof of Lemma 4.3). Thus we can define a ray-shooting map $s_K : \mathcal{S}^{d-1} \to \partial K^r$, which sends any $v \in \mathcal{S}^{d-1}$ to the intersection of the ray emanating from x with direction v with ∂K^r.

Lemma 4.3 *If K is a compact set contained in a ball $B(x,r)$, the ray-shooting map s_K defined above is $2r$-Lipschitz, so that $\mathcal{H}^{d-1}(\partial K^r) \leqslant \omega_{d-1}(2r)$.*

Proof Since $\partial K^r = s_K(B(0,1))$, assuming that s_K is $2r$-Lipschitz, we will indeed have $\mathcal{H}^{d-1}(K^r) \leqslant (2r)^{d-1}\mathcal{H}^{d-1}(B(0,1)) = \omega_{d-1}(2r)$. Let us now compute the Lipschitz constant of the ray-shooting map s_K.

If we let $t_K(v)$ be the distance between x and $s_K(v)$, we have $t_K(v) = \sup_{e\in K} t_e(v)$. Since t_K is the supremum of all the t_e, in order to prove that s_K is $2r$-Lipschitz, we only need to prove that each s_e is $2r$-Lipschitz. Without loss of generality, we will suppose that e is at the origin, $x \in \mathrm{B}(0,r)$, and let $s = s_e$. Solving the equation $\|x + t(v)v\| = r$ with $t \geqslant 0$ gives

$$t(v) = \sqrt{\langle x|v\rangle^2 + r^2 - \|x\|^2} - \langle x|v\rangle.$$

This gives the following expression for the derivative of $s(v) = x + t(v)v$:

$$\begin{aligned}
\mathrm{d}_v s(w) &= \left\|s(v) - x\right\| w + \mathrm{d}_v t(w) v \\
&= \left\langle s(v) - x | v \right\rangle w + \left(\frac{\langle x|v\rangle \langle x|w\rangle}{\langle s(v)|v\rangle} - \langle x|w\rangle \right) v \\
&= \left\langle s(v) - x | v \right\rangle \frac{\langle s(v)|v\rangle w - \langle x|w\rangle v}{\langle s(v)|v\rangle}.
\end{aligned}$$

If w is orthogonal to x, then $\|\mathrm{d}_v s(w)\| \leqslant \|s(v) - x\| \|w\| \leqslant 2r\|w\|$ and we are done. We now suppose that w is contained in the plane spanned by x and v. Since w is tangent to the sphere at v, it is also orthogonal to v. Hence, $\langle s(v) - x|w\rangle = 0$, and $\langle x|w\rangle = \langle s(v)|w\rangle$.

$$\mathrm{d}_v s(w) = \left\|s(v) - x\right\|^2 \frac{\langle s(v)|v\rangle w - \langle s(v)|w\rangle v}{\langle s(v)|s(v) - x\rangle}.$$

If we suppose that both v and w are unit vectors, $\langle s(v)|v\rangle w - \langle s(v)|w\rangle v$ is the rotation of $s(v)$ by an angle of $\pi/4$. By linearity we get $\|\langle s(v)|v\rangle w - \langle s(v)|w\rangle v\| = \|w\| \|s(v)\|$ (we still have $\|v\| = 1$). Now let us remark that

$$\left\langle s(v)|s(v) - x\right\rangle = \frac{1}{2}\left(\left\|x - s(v)\right\|^2 + \left\|s(v)\right\|^2 - \|x\|^2\right) \geqslant \frac{1}{2}\left\|x - s(v)\right\|^2.$$

Using this we deduce the following:

$$\left\|\mathrm{d}_v s(w)\right\| \leqslant \left\|s(v) - x\right\|^2 \frac{\|w\| \|s(v)\|}{\frac{1}{2}\|x - s(v)\|^2} = 2\left\|s(v)\right\| \|w\| \leqslant 2r\|w\|.$$

We have just proved that for any w tangent to the sphere at v, $\|\mathrm{d}_v s(w)\| \leqslant 2r\|w\|$, from which we can conclude that s is $2r$-Lipschitz. □

Proof of Proposition 4.2 By definition of the covering number, there exists a finite family of points $x_1, \ldots, x_n$, with $n = \mathcal{N}(K, r)$, such that union of the open balls $B(x_i)$ covers ∂K. If one denotes by K_i the intersection of ∂K with $\overline{B}(x_i, r)$, the boundary ∂K^r is contained in the union $\bigcup_i \partial K_i^r$. Hence its Hausdorff measure does not exceed the sum $\sum_i \mathcal{H}^{d-1}(\partial K_i^r)$. Since, for each i, ∂K_i^r is a flower, one concludes by applying the preceding lemma. □

From the discussion above we easily get the following.

Corollary 4.4 *For any compact sets $K, K' \subseteq \mathbb{R}^d$, with $\mathrm{d_H}(K, K') \leqslant r/2$,*

$$\mathcal{H}^d\left(K^r \Delta K'^r\right) \leqslant 2\mathcal{N}(K, r/2)\omega_{d-1}(3r) \times \mathrm{d_H}(K, K').$$

Stability of Approximate Curvature Measures Combining the results of Proposition 4.1, Corollary 4.4, and the Projection Stability Theorem, one obtains the following stability result for boundary measures μ_{K,K^r}.

Theorem 4.5 *If K and K' are two compact sets of $\mathbb{R}^d$,*

$$\mathrm{d_{bL}}(\mu_{K,K^r},\mu_{K',K'^r}) \leqslant C_3(d)\mathcal{N}(K,r/2)r^d\big[r+\mathrm{diam}(K)\big]\sqrt{\frac{\mathrm{d_H}(K,K')}{r}}$$

provided that $\mathrm{d_H}(K,K') \leqslant \min(\mathrm{diam}\, K, r/2, r^2/\mathrm{diam}\, K)$.

To define the *approximate curvature measures*, let us fix a sequence (r_i) of $d+1$ distinct numbers $0 < r_0 < \cdots < r_d$. For any compact set K and Borel subset $B \subset K$, we let $[\Phi_{K,i}^{(r)}(B)]_i$ be the solutions of the linear system

$$\forall i \quad \text{s.t} \quad 0 \leqslant i \leqslant d, \qquad \sum_{j=0}^{d} \omega_{d-j}\Phi_{K,j}^{(r)}(B) r_i^{d-j} = \mu_{K,K^{r_i}}(B).$$

We call $\Phi_{K,j}^{(r)}$ the (r)-approximate curvature measure. Since this is a linear system, the functions $\Phi_{K,i}^{(r)}$ also are additive. Hence the (r)-approximate curvature measure $\Phi_{K,i}^{(r)}$ defines a signed measure on $\mathbb{R}^d$. We also note that if K has a reach greater than r_d, then the measures $\Phi_{K,i}^{(r)}$ coincide with Federer's curvature measures of K, as introduced in Sect. 2. From these remarks and from Theorem 4.5, we have the following.

Corollary 4.6 *For each compact set K whose reach is greater than r_d, there exists a constant $C_4(K,(r),d)$ depending on K, (r) and d such that for any $K' \subseteq \mathbb{R}^d$ close enough to K,*

$$\mathrm{d_{bL}}\big(\Phi_{K',i}^{(r)}, \Phi_K^i\big) \leqslant C_4\big(K,(r),d\big)\mathrm{d_H}(K,K')^{1/2}.$$

This corollary gives a way to approximate the curvature measures of a compact set K with positive reach from the (r)-approximate curvature measures of any point cloud close to K.

5 Computing Boundary Measures

If $C = \{p_i; 1 \leqslant i \leqslant n\}$ is a point cloud, that is, a finite set of points of $\mathbb{R}^d$, then μ_{C,C^r} is a sum of weighted Dirac measures: letting $\mathrm{Vor}_C(p_i)$ denote the Voronoi cell of p_i, we have

$$\mu_{C,C^r} = \sum_{i=1}^{n} \mathcal{H}^d\big(\mathrm{Vor}_C(p_i) \cap C^r\big)\delta_{p_i}.$$

Hence, computing boundary measures amounts to finding the volume of intersections of Voronoi cells with balls. This method is practical in dimension 3, but in higher dimensions it becomes prohibitive due to the exponential cost of Voronoi diagram computations. We instead compute approximations of boundary measures using a Monte Carlo method. Let us first recall some standard facts about these approximations.

Approximation by Empirical Measures If μ is a probability measure on $\mathbb{R}^d$, one can define another measure as follows: let $X_1, \ldots, X_N$ be a family of independent random vectors in $\mathbb{R}^d$ whose law is μ, and let $\overline{\mu}_N$ be the sum of Dirac $\frac{1}{N}\sum_i \delta_{X_i}$. Convergence results of the *empirical measure* $\overline{\mu}_N$ to μ are known as *the uniform law of large numbers*. Using standard arguments based on Hoeffding's inequality, covering numbers of spaces of Lipschitz functions and the union bound, it can be shown [7] that if μ supported in $K \subseteq \mathbb{R}^d$, then the following estimate on the bounded-Lipschitz distance between the empirical and the real measure holds:

$$\mathbb{P}\big[\mathrm{d_{bL}}(\mu_N, \mu) \geqslant \varepsilon\big] \leqslant 2\exp\big(\ln(16/\varepsilon)\mathcal{N}(K, \varepsilon/16) - N\varepsilon^2/2\big).$$

In particular, if μ is supported on a point cloud C, with $\#C = n$, then $\mathcal{N}(C, \varepsilon/16) \leqslant n$. This shows that computing an ε-approximation of the measure μ with high probability (e.g. 99%) can always be done with $N = \mathrm{O}(n\ln(1/\varepsilon)/\varepsilon^2)$. However, if C is sampled near a k-dimensional object, then for ε in an appropriate range we have $\mathcal{N}(C, \varepsilon/16) \leqslant \text{const}\ \varepsilon^{-k}$, in which case N is of the order of $-\ln(\varepsilon)\varepsilon^{-k-2}$.

Monte Carlo Approximation of Boundary Measures Let $C = \{p_1, \ldots, p_n\}$ be a point cloud. Applying the ideas of the previous paragraph to the probability measure $\beta_{C,C^r} = \frac{\mu_{C,C^r}}{\mathcal{H}^d(C^r)}$, we get the approximation algorithm (Algorithm 1).

To simulate the uniform measure on C^r in step I, one cannot simply generate points in a bounding box of C^r, keeping only those that are actually in C^r since the probability of finding a point in C^r might decrease exponentially with the ambient dimension.

Luckily, there is a simple algorithm to generate points according to this law which relies on picking a random point x_i in the cloud C and then a point X in $B(x_i, r)$—taking into account the overlap of the balls $B(x, r)$ where $x \in C$ (Algorithm 2). Instead of completely rejecting a point if it lies in k balls with probability $1 - 1/k$, one can instead modify Algorithm 1 to attribute a weight $1/k$ to the Dirac mass added at this point. Step III requires an estimate of $\mathcal{H}^d(C^r)$. Using the same empirical measure convergence argument, one can prove that if *TL* is the total number of times the loop of Algorithm 2 was run, and *TN* is the total number of points generated, then $TN/TL \times n\mathcal{H}^d(B(0, r))$ converges to $\mathcal{H}^d(C^r)$.

In Fig. 4, we show the result of this computation for a point-cloud approximation of a mechanical part. Each point x_i of the point cloud C is represented by a sphere

Algorithm 1 Monte Carlo algorithm to approximate μ_{C,C^r}

Input: a point cloud C, a scalar r, a number N
Output: an approximation of μ_{C,C^r} in the form $\frac{1}{N}\sum n(p_i)\delta_{p_i}$
while $k \leqslant N$ **do**
  [**I.**] Choose a random point X with probability distribution $\frac{1}{\mathcal{H}^d(C^r)}\mathcal{H}^d|_{C^r}$
  [**II.**] Find its closest point p_i in the cloud C, add 1 to $n(p_i)$
end while
[**III.**] Multiply each $n(p_i)$ by $\mathcal{H}^d(C^r)$.

Algorithm 2 Simulating the uniform measure in C^r

Input: a point cloud $C = \{p_i\}$, a scalar r
Output: a random point in C^r whose law is $\mathcal{H}^d|_{C^r}$
repeat
  Pick a random point p_i in the point cloud C
  Pick a random point X in the ball $B(p_i, r)$
  Count the number k of points $p_j \in C$ at distance at most r from X
  Pick a random integer d between 1 and k
until $d = 1$
return X.

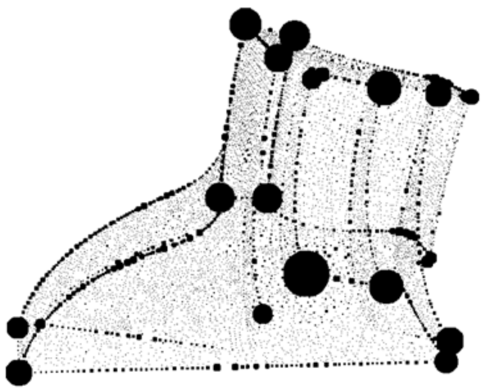

Fig. 4 Boundary measure for a sampled mechanical part

whose radius is proportional to the convolved value $\sum_j \chi(x_i - x_j)\mu_{C,C^r}(\{x_j\})$, where χ is a tent function of appropriate radius. As expected, the relevant features of the shape are visually highlighted.

6 Discussion

In this article, we introduced the notion of boundary measure. We showed how to compute these measures efficiently for point clouds using a simple Monte Carlo algorithm. More importantly, we proved that they depend continuously on the compact set. That is, the boundary measure of a point-cloud approximation of a compact set K retains the geometric information contained in the boundary measure of K, which is, as we know, very rich. We also obtain the first quantitative stability result for Federer's curvature measures, which shows their usefulness for geometric inference. However, several questions are still to be investigated.

On the algorithmic side, finding approximate nearest neighbors is usually much cheaper than finding exact nearest neighbors. It is thus tempting to replace nearest neighbor projections by approximate nearest neighbors in the definition of boundary measures. From the perspective of inference, it would be interesting to know whether the obtained measure still enjoys comparable stability properties.

Also, we know that if K has reach greater than R, then μ_{K,K^r} is polynomial on $[0, R]$. Because of the stability property proved in this article, if C is a point cloud approximating K, then μ_{C,C^r} is almost polynomial on $[0, R]$. But is the converse also true? If so, our approach would provide a test for knowing whether a point cloud is close to a compact of positive reach.

Finally, the boundary measures as defined here are sensitive to outliers: a single outlier may get a significant fraction of the mass of the boundary measure. While this may be seen as an interesting feature for outlier detection, it would be useful to modify the general definition of boundary measures so as to make it also robust to this kind of noise.

Acknowledgements The authors are grateful to C. Villani and F. Morgan for valuable discussions and suggestions. This work was partly supported by the ANR Grant GeoTopAl JC05_41922.

References

1. L. Ambrosio, N. Fusco, D. Pallara, *Functions of Bounded Variation and Free Discontinuity Problems*. Oxford Mathematical Monographs (Oxford University Press, London, 2000).
2. N. Amenta, S. Choi, T.K. Dey, N. Leekha, A simple algorithm for homeomorphic surface reconstruction, *Int. J. Comput. Geom. Appl.* **12**(1–2), 125–141 (2002).
3. S. Bobkov, C. Houdré, Converse Poincaré-type inequalities for convex functions, *Stat. Probab. Lett.* **34**, 37–42 (1997).
4. S. Bobkov, C. Houdré, A converse Gaussian Poincaré-type inequality for convex functions, *Stat. Probab. Lett.* **44**, 281–290 (1999).
5. G. Carlsson, T. Ishkhanov, V. De Silva, A. Zomorodian, On the local behavior of spaces of natural images, *Int. J. Comput. Vis.* **76**(1), 1–12 (2008).
6. F. Chazal, D. Cohen-Steiner, A. Lieutier, A sampling theory for compacts in Euclidean space, in *Proceedings of the 22th Annual Symposium on Computational Geometry* (2006).
7. F. Chazal, D. Cohen-Steiner, Q. Mérigot, Stability of boundary measures (2007). arXiv:0706.2153.
8. F. Chazal, A. Lieutier, Stability and computation of topological invariants of solids in $\mathbb{R}^n$, *Discrete Comput. Geom.* **37**(4), 601–617 (2007).
9. D. Cohen-Steiner, H. Edelsbrunner, J. Harer, Stability of persistence diagrams, *Discrete Comput. Geom.* **37**, 103–120 (2007).
10. D. Cohen-Steiner, J. Morvan, Restricted Delaunay triangulations and normal cycle, in *Proceedings of the Nineteenth Annual Symposium on Computational Geometry* (2003), pp. 312–321.
11. V. de Silva, G. Carlsson, Topological estimation using witness complexes, in *Proceedings of the First IEEE/Eurographics Symposium on Point-based Graphics* (2004), pp. 157–166.
12. T. Dey, J. Giesen, S. Goswami, W. Zhao, Shape dimension and approximation from samples, *Discrete Comput. Geom.* **29**(3), 419–434 (2003).
13. H. Edelsbrunner, The union of balls and its dual shape, *Discrete Comput. Geom.* **13**, 415–440 (1995).
14. H. Federer, Curvature measures, *Trans. Am. Math. Soc.* **93**(3), 418–491 (1959).
15. F. Morgan, *Geometric Measure Theory: A Beginner's Guide* (Academic Press, San Diego, 1988).
16. P. Niyogi, S. Smale, S. Weinberger, Finding the homology of submanifolds with confidence from random samples, *Discrete Comput. Geom.* (2006).
17. S. Peleg, M. Werman, H. Rom, A unified approach to the change of resolution: space and gray-level, *IEEE Trans. Pattern Anal. Mach. Intell.* **11**(7), 739–742 (1989).
18. Y. Rubner, C. Tomasi, L. Guibas, The earth mover's distance as a metric for image retrieval, *Int. J. Comput. Vis.* **40**(2), 99–121 (2000).
19. J. Steiner, Über parallele Flächen, in *Monatsbericht die Akademie der Wissenschaften zu Berlin* (1840), pp. 114–118.
20. C. Villani, *Topics in Optimal Transportation* (American Mathematical Society, Providence, 2003).
21. H. Weyl, On the volume of tubes, *Am. J. Math.* **61**(2), 461–472 (1939).
22. T. Zamfirescu, On the cut locus in Alexandrov spaces and applications to convex surfaces, *Pac. J. Math.* **217**, 375–386 (2004).