



HAL
open science

Voronoi-Based Curvature and Feature Estimation from Point Clouds

Quentin Mérigot, Maks Ovsjanikov, Leonidas J. Guibas

► **To cite this version:**

Quentin Mérigot, Maks Ovsjanikov, Leonidas J. Guibas. Voronoi-Based Curvature and Feature Estimation from Point Clouds. IEEE Transactions on Visualization and Computer Graphics, 2011, 17 (6), pp.743 - 756. <10.1109/TVCG.2010.261>. <inria-00406575v2>

HAL Id: inria-00406575

<https://inria.hal.science/inria-00406575v2>

Submitted on 12 Jan 2016

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.



HAL Authorization

VORONOI-BASED CURVATURE AND FEATURE ESTIMATION FROM POINT CLOUDS

QUENTIN MÉRIGOT, MAKS OVSJANIKOV, AND LEONIDAS GUIBAS

ABSTRACT. We present an efficient and robust method for extracting curvature information, sharp features and normal directions of a piecewise smooth surface from its point cloud sampling in a unified framework. Our method is integral in nature and uses convolved covariance matrices of Voronoi cells of the point cloud which makes it provably robust in the presence of noise. We show that these matrices contain information related to curvature in the smooth parts of the surface, and information about the directions and angles of sharp edges around the features of a piecewise-smooth surface. Our method is applicable in both two and three dimensions, and can be easily parallelized, making it possible to process arbitrarily large point clouds, which was a challenge for Voronoi-based methods. In addition, we describe a Monte-Carlo version of our method, which is applicable in any dimension. We illustrate the correctness of both principal curvature information and feature extraction in the presence of varying levels of noise and sampling density on a variety of models. As a sample application, we use our feature detection method to segment point cloud samplings of piecewise-smooth surfaces.

INTRODUCTION

Estimating surface normals, principal curvatures and sharp edges from a noisy point cloud sampling has many applications in computer graphics, geometry processing and reverse engineering. Principal curvatures are rotation-invariant local descriptors, which together with principal curvature directions have proven useful in detecting structural regularity [PMW⁺08], global matching [AMCO08], modeling and rendering of point-based surfaces [GTE⁺06], and anisotropic smoothing [LP05] to name just a few. In these applications, various notions of curvature serve as local descriptors that encode second order variations of the surface. The location of sharp edges and highly curved areas of the surface is a precious piece of information in settings that include feature-aware reconstruction [JWS08], non photorealistic rendering [PKG03], and industrial metrology.

In practice, it is often interesting to recover this information when the input is an unstructured collection of point coordinates, obtained by a range scanner, before attempting surface reconstruction. These point clouds can be noisy, and can exhibit strong sampling bias. The ability to reliably estimate surface normals, principal curvatures, and curvature directions as well as sharp features directly on such point clouds can be used in both geometry processing algorithms and in surface reconstruction to improve the quality of the resulting mesh.

Devising robust local descriptors, which can handle both non-uniform noise, sampling bias and sharp edges is a challenging task. This is mainly because we are

This work is published in *IEEE Trans. Vis. Comp. Graph.* Vol 17, Issue 6, p743–756.

trying to estimate differential quantities which by nature are very sensitive to local perturbations. The lack of a natural parametrization of point cloud data introduces another challenge by making it difficult to estimate angles and areas on the surface. Finally, devising a method with theoretical guarantees on the approximation quality is not easy in the absence of a definition of curvature and sharp features that that would incorporate both point clouds and piecewise-smooth surfaces.

In this paper, we address some of these challenges by presenting a method for robustly estimating the location and direction of sharp edges as well as highly curved area from a possibly noisy point cloud sampling. We also show that the same method can be used to recover curvature information of the underlying surface, when the sampling is dense enough. The sharp edge estimation technique comes with theoretical guarantees on the quality of the results as a function of the Hausdorff distance between the point cloud and the underlying surface. We also address a certain class of outliers.

Prior work on feature and curvature estimation. The questions of curvature estimation and sharp feature detection are tightly related: sharp edges and corners can be thought as parts of the surface with infinite concentration of curvature (mean and Gaussian, respectively). Surprisingly, however, these research topics have known very different developments; we review the existing results separately.

Curvature estimation on meshes. Estimating the curvature of a smooth surface from a mesh is an important question in discrete differential geometry has been studied for many years and is now a well understood topic (see [MSR07] for an extensive comparison of methods to estimate Gaussian and mean curvatures). The approach commonly used is to define a notion of curvature for meshes, and to study how well it approximates the real curvature of the underlying surface. The usual requirement is that (i) the Hausdorff distance between the mesh and the surface is low, and (ii) the normals to the mesh do not deviate from those of the surface too much.

Normal cycle. Cohen-Steiner and Morvan [CSM03] used the normal cycle from geometric measure theory in order to define a closed-form second fundamental measure of a mesh. They also explicitly bound the error between the second fundamental measure of the mesh and the underlying surface, supposing (i) and (ii). This typically applies to the case of a mesh obtained by a Delaunay-based reconstruction of the surface from an ε -sample. (It should also be noted that the second fundamental measure can be used to compute reliably approximate the curvature of *offsets* of the underlying surface from a point cloud approximation [CCSLT09].

Integral methods. A very different approach estimates the curvature and sharp features location from a mesh approximation of a surface S that bounds a domain D , by considering the volume and covariance matrices of intersections of small balls with the domain D [Con86, PWHY09]. This approach can be used both for curvature estimation and sharp features detection. Another interesting aspect using an approximation D' of the domain D yields an error that depends on the volume of the symmetric difference $D\Delta D'$ only. This makes these integral methods resilient to noise in the normal directions. However, they require a knowledge of both the surface and the interior domain, and cannot be applied to point clouds straightforwardly.




FIGURE 1. Features (in red) computed by our algorithm from a point cloud sampling of the surface in yellow.

Curvature estimation from points clouds. Only recently several methods have been proposed for computing curvature and feature information on the point cloud directly. One class of methods for computing curvature information from a point cloud sampling work by estimating directional curvatures in the neighborhood of a point and using them to estimate the principal curvatures and curvature directions: [LP05, HM02]. Robust methods such as least squares [HM02] and local re-weighting [LP05] help these methods cope with noise and irregular sampling. Berkmann et al. [BC94] propose a method that rely on computing the covariance matrices of normal vectors to approximate the shape operator. More recently, Kalogerakis et al. [KNSS09] use robust statistical methods to fit the curvature tensor of a point cloud by considering normal variation in a neighborhood of a given point. Yang and Qian [YQ07] derive analytic expressions for computing principal curvatures based on the implicit definition of the moving least squares surface by Amenta and Kil [AK04].

These algorithms start by estimating normals to the surface, or assume that they are given. Any error made in this estimation is therefore only aggravated in the computation of principal curvatures. A possible way to compute the normal and the second fundamental form at the same time is to fit local jets (see Cazals and Pouget [CP05]). Least-square jet-fitting is very fast, and its convergence properties are well understood. Unfortunately, a necessary condition for this method to give correct estimations is that the point cloud should be uniformly sampled. This can be a problem for some point clouds such as laser scans which are virtually never free of noise and often exhibit oversampled clusters of points along horizontal lines.

Sharp features estimation from point clouds. Although extracting sharp edges and corners is closely related to curvature estimation, research in these areas has been relatively independent. Fleishman et al. [FCOS05] detect sharp edges and corners by segmenting neighborhoods of points into regions corresponding to the same part of the surface. They achieve robustness by using a forward search technique which finds reference planes, corresponding to each segment. This work is extended by Daniels et al. [DIOHS08] to extract feature-curves, by locally projecting onto feature points, and growing smooth polylines through the projected cloud. Lipman and colleagues [LCOL07] extract sharp edges within the MLS projection framework by examining the error of the MLS surface approximation. Jenke et al. [JWS08] detect sharp features by robustly fitting local surface patches and computing intersections of nearby patches with dissimilar normals. In a similar spirit, Oztireli et al. [OGG09]

define a feature-preserving MLS projection operator by noting that angles between normal vectors, rather than point coordinates can be used to discard points from unrelated parts of a piecewise-smooth surface.

To sum things up, there is currently no provably-correct algorithm for estimating the sharp features of a piecewise smooth surface from an unoriented point cloud (i.e. without normals).

Contributions. In this paper, we associate to every point cloud C (or surface S) its *Voronoi covariance measure* $\mathcal{V}_{C,R}$ (resp. $\mathcal{V}_{S,R}$), from which one can extract information on the feature and curvature of the underlying surface. This measure is defined, following the Voronoi-based normal estimation technique of Alliez et al [ACSTD07], by considering the covariance matrices of Voronoi cells. However, instead of intersecting them with a large bounding box of the point cloud C , we intersect them with an offset C^R , thus obtaining more *local* information about the variation in shape and size of the Voronoi cell, which is crucial for feature estimation. We present two algorithms for computing the Voronoi covariance measure: a Monte Carlo method, and a method based on tessellation (§2).

In order to gain robustness to Hausdorff noise, we have to consider convolved versions of the Voronoi covariance measure. For any positive convolution radius r , this defines a function $\mathcal{V}_{C,R} * \chi_r$ that maps points of \mathbb{R}^d to symmetric matrices. The Stability Theorem (Theorem 4.1), asserts that replacing the convolved VCM of a surface S by the convolved VCM of a discrete approximation C yields an error of $O(d_H(C, S)^{1/2})$. We show in §5 that the eigenvectors of the convolved VCM of a smooth surface are the normal to the surface followed by the principal curvature directions; and experimentally, this information remains in the VCM of point clouds that approximate S densely and without noise. Much more interestingly, it appears that the principal curvature directions computed this way at *highly curved* parts of the surface remain good even when the sampling is very anisotropic or when Hausdorff noise is added to the point cloud.

Our main contribution is an algorithm that estimates the location and direction of sharp edges of a piecewise-smooth surface S from a point cloud approximation (§6) from the eigendecomposition of the convolved VCM. The algorithm requires as input a (rough) lower bound on the one-sided reach of the surface, a notion defined in §6.1. The existence of such a lower bound is similar to the positive local feature size (lfs) assumption needed for most provable surface reconstruction techniques such as [AB99]. The algorithm uses a threshold parameter, which can be specified by the user and is explicitly related to the minimum expected dihedral angle between smooth parts of the surface.

The Stability Theorem ensures that the sharp feature detection algorithm recovers the correct location of edges and is robust to Hausdorff noise. We experimentally check this fact on different models, under various levels of noise. Finally, as a sample application, we use this feature detection method to segment point cloud samplings of piecewise-smooth surfaces.

1. OVERVIEW OF THE METHOD

1.1. Geometric intuition. The method we propose is motivated by the same intuition as Voronoi-based normal estimation techniques such as the work by Amenta and Bern [AB99] and Alliez et al. [ACSTD07]. Recall that the Voronoi cell of a

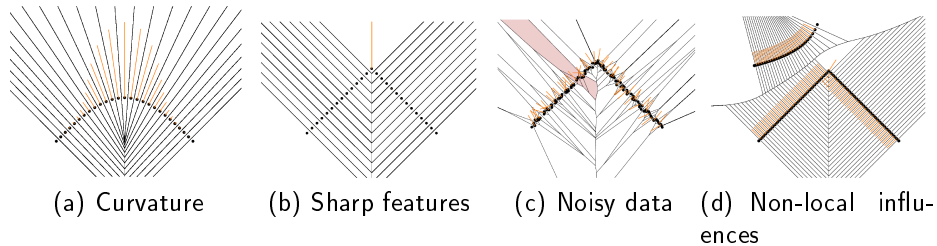


FIGURE 2. The ratio of the two eigenvalues of the covariance matrix of a Voronoi cell (length of the direction in orange) contains curvature information (a). In the limit, Voronoi cells of points sampled from sharp corners are close to being isotropic (b). Noisy data (c) and non-local parts of the surface (d) can perturb the Voronoi cells, thus preventing correct curvature and sharp features estimations. Our method recovers from these effects using convolution and intersection with an offset.

point p in a point cloud C is the set of points of the space which are closer to p than to any other point in C . Denote this Voronoi cell by $\text{Vor}_C(p)$.

Amenta and Bern remarked in [AB99] that in the absence of noise, if a point cloud samples a smooth surface densely enough, its Voronoi cells will be elongated in the direction of the normal to the surface 2(a). This allows to approximate the normal direction at a sample point p_0 by considering the shape of its Voronoi cell, by either computing its extremal points (poles) [AB99] or by performing principal component analysis [ACSTD07]. Note, however, that the shapes of Voronoi cells provide more information than just the normal direction. Intuitively, the Voronoi cells at points sampled from highly curved areas will deviate from straight pencil shapes (which is well-approximated by a skinny cylinder), and will be V-shaped with the angle related to the curvature of the underlying surface. Fig. 2(a) illustrates this effect on points sampled from a smooth planar curve. The lengths of the principal directions of each Voronoi cell, depicted in orange, are drawn proportional to the ratio $\lambda_{\min}/\lambda_{\max}$, of the two eigenvalues of the covariance matrix of the corresponding Voronoi cell, intersected with a large bounding box. Note that this ratio becomes larger for points sampled from highly curved areas of the shape, and therefore, it should not be surprising that the shape of the Voronoi cells, which can be captured by its covariance matrix, is related to the curvature of the underlying surface.

The deviation of Voronoi cells from being pencil shaped is most pronounced at points sampled from sharp corners and sharp edges of the surface, as shown in Fig. 2(b) and Fig. 6. Note that when the underlying surface is not smooth, some of its points will have normal cones rather than single normal directions. Nevertheless, even in this case, the shapes of Voronoi cells accurately reflect the shapes of the underlying normal cones, that can be characterized by the principal components of their covariance matrices. This observation is the key to allowing us to study both the smooth and piecewise-smooth surfaces in a unified framework.

One of the principal drawbacks of using Voronoi-based methods directly, is their sensitivity to noise. As shown in Fig. 2(c), even moderate noise can change the

shapes of Voronoi cells dramatically. A beautiful insight made by Alliez et al. [ACSTD07] is that although individual Voronoi cells can become ill shaped, by considering the union of Voronoi cells of a few neighbors, one can recover the correct normal directions. In their work, Alliez et al. used covariance matrices of the union of Voronoi cells to combat noisy sampling. One drawback of this approach is that taking the union unnecessarily produces more isotropic (fat) cells. A way to overcome this effect is to take a weighted average of covariance matrices of a few neighboring Voronoi cells instead of taking the covariance matrix of their union. In the continuous case, this average directly translates to convolution. In this paper we show that these averaged covariance matrices capture information about the shapes of normal cones of the underlying surface.

Because of the global nature of the Voronoi diagram, points sampled from remote parts of the shape can influence the shapes of Voronoi cells of points even in absence of noise, thus leading to erroneous estimations of the underlying normal cones. Fig. 2(d) illustrates this effect. Note that the Voronoi cell of the point at the tip of the sharp corner is “invaded” by Voronoi cells of points from a different curve, which adversely influences the computation of the normal direction at the point. Note that this effect can occur even due to remote points of the same curve, is very global in nature and therefore difficult to rule out locally. To overcome this problem, we compute the covariance matrix of the Voronoi cell of each point p intersected with a small ball around p . This allows us to obtain purely local information about the variation of the surface around p , which is crucial for curvature and feature estimation.

1.2. Voronoi covariance of a point cloud. As described in the previous section, our algorithm aims to extract information about curvature and sharp features of the underlying surface from its point cloud sampling C , by analyzing the shapes of Voronoi cells of C intersected with small balls around each point. A way to get information about the shape of a domain E of \mathbb{R}^d with finite volume is through its covariance matrix with respect to a base point p , denoted by $\text{cov}(E, p)$. Specifically, the eigenvectors of $\text{cov}(E, p)$ capture the principal axes of E while the ratio of the eigenvalues gives information about the anisotropy of E . The covariance matrix of a domain E can be computed through the following formula: $\text{cov}(E, p) := \int_E (x - p)(x - p)^T dx$.

1.2.1. Definition. Given a point cloud C , and an offset radius R , we define the Voronoi covariance measure (VCM) of a point $p \in C$ at scale R to be:

$$(1.1) \quad \mathcal{V}_{C,R}(\{p\}) := \text{cov}(\text{Vor}_C(p) \cap \text{B}(p, R), p)$$

The VCM $\mathcal{V}_{C,R}(\{p\})$ of any point p in C is a symmetric matrix that captures information about the shape of the Voronoi cell of p . As mentioned earlier, however, Voronoi cells and as a result VCM can be unstable under noisy sampling. Thus, we define the *convolved Voronoi covariance measure* at a point x (not necessarily in C) by summing the covariance matrices defined above among all points $p \in C$ in a neighborhood of x :

$$(1.2) \quad \mathcal{V}_{C,R} * \chi_r(x) := \sum_{p \in \text{B}(x,r) \cap C} \mathcal{V}_{C,R}(\{p\})$$

Intuitively, the convolved VCM of a point near the underlying surface S provides information about the normals and curvature of S .

1.3. Outline. In this paper, we demonstrate the practical implications of the Voronoi covariance measure. First, we indicate how to compute it efficiently from three-dimensional point clouds (§2). In §3 we give a general definition of the VCM suitable for any compact set, and prove a stability theorem of VCM under arbitrary Hausdorff noise (§4). Using the stability of the VCM, we devise robust curvature estimation and feature extraction algorithms (§5,6). Finally, we apply our feature extraction method to segment samplings of piece-wise smooth surfaces. (§7).

2. COMPUTING VCM

In this section, we describe two algorithms for computing the VCM of a point cloud in practice. The first method is easy to implement and is applicable in any ambient dimension. The second algorithm is much faster, and is the one we used for all of our results. We then describe a straightforward way to convolve the VCM.

Given a point cloud $C = \{p_1, \dots, p_n\}$, computing the VCM of any point p_i at scale R amounts to evaluating the integral $\text{cov}(B_i, p_i) = \int_{B_i} (x-p)(x-p)^T dx$ where B_i is the intersection $B(p_i, R) \cap \text{Vor}(p_i)$ between a ball of radius R centered at p_i and the Voronoi cell of p_i . No simple closed-form expression seems to exist for the covariance matrix of the intersection B_i , and we have to resort to approximation to compute it.

2.1. Monte-Carlo approximation of the VCM. The Voronoi covariance measure is a modification of the *boundary measure* introduced in [CCSM10]. In this work, instead of considering the covariance matrix $\text{cov}(B_i)$, the authors only considered its volume. The Monte-Carlo algorithm for computing the boundary measures can be easily adapted to compute the VCM.

The input of Algorithm 3 is a point cloud C , an offset radius R and a number of steps N . For every point p in C it maintains a matrix $C(p)$ which approximates the covariance matrix $\text{cov}(\text{Vor}(p) \cap B(p, R), p)$ as the number of steps N grows to infinity. The arguments presented in [CCSM10, p.6] show that the output of this algorithm converge to the real VCM with high probability: computing an ε -approximation of the VCM with high probability (e.g. 99%) can be done in $N = O(|C| \ln(1/\varepsilon)/\varepsilon^2)$ steps.

This algorithm only uses nearest-neighbor queries and is easy to implement in any ambient dimension. However it is too slow in practice for computing the VCM of point clouds with hundreds of thousands of points. In the following, we describe a deterministic approach that can be used to improve the computation speed in low dimensions.

FIGURE 3. Monte-Carlo algorithm for VCM

1: For all $p \in C$, set $C(p)$ to the zero matrix, set $M \leftarrow 0$.

2: Repeat N times :

2a: Select a random point p in C and generate a random point x in the ball $B(p, R)$.

2b: $k \leftarrow$ number of points of C in the ball $B(x, R)$.

2c: Find the nearest neighbor p' of x in C , and set $C(p') \leftarrow C(p') + \frac{1}{k}(x - p')(x - p')^T$, $M \leftarrow M + \frac{1}{k}$.

3: For all $p \in C$, set $C(p) \leftarrow \frac{C(p)}{M}$.

2.2. Approximating the VCM in 3D by tessellation. We base our second method on the fact that the covariance matrix of a tetrahedron can be computed analytically [ACSTD07]. Therefore, using the additivity of the integral, in order to compute the covariance matrix of the intersection of a Voronoi cell with a ball, it is sufficient to approximate it with a union of tetrahedra. For this, we triangulate the boundary of this intersection and build tetrahedra by connecting each of these triangles to the center of the Voronoi cell — this can be done because the intersection is convex.

We start with an arbitrary triangulation of the boundary of the Voronoi cell $\partial(\text{Vor}(p_i))$. Our goal is to subdivide each triangle so that its projection onto the ball $B(p_i, R)$ is a sufficiently good approximation of the corresponding spherical triangle. For this, we process each triangle t in the original triangulation according to three simple rules (Algorithm 4).

The output of this algorithm yields a tetrahedrization of the intersection of the Voronoi cell with a ball of a given radius, centered at the same point. We then use the closed-form formulas of [ACSTD07] to compute the covariance matrix of this intersection.

Convolution of VCM. Computing the convolved VCM of a point x simply amounts to summing the VCM matrices of all points p of C lying in the ball $B(x, r)$. The points belonging to this intersection $B(x, r) \cap C$ are the k nearest neighbors of x in C for a suitable value of k . The value of k and the corresponding neighbors can be determined by a binary search, using a structure adapted to k -NN queries (such as a k D-tree).

2.3. Implementation. We implemented the two algorithms described above. The tessellation of the Voronoi cells was done using the 3D Delaunay Triangulation package of CGAL [cga]. Its running times of this algorithm on some standard point clouds is reported in Table 1. The convolution step is implemented using the ANN library [MA], which includes a query for finding the set of points of a point cloud contained in a given ball. The time of the convolution step depends on the radius of convolution, but stays within 10 seconds for most models.

FIGURE 4. Tetrahedrization of $\text{Vor}_C(p_i) \cap B(p_i, R)$
1: $\mathcal{T} \leftarrow$ triangulation of the boundary of $\text{Vor}_C(p_i)$.

$\mathcal{T}' \leftarrow \emptyset$

2: For all triangle $t \in \mathcal{T}$:

A: If t is completely outside the ball $B(p_i, R)$:

A1: Recursively subdivide t into a family of smaller triangles $\{t_k\}$.

A2: $\mathcal{T}' \leftarrow \mathcal{T}' \cup \{t_k\}$

B: If t is completely inside the ball: $\mathcal{T}' \leftarrow \mathcal{T}' \cup \{t\}$

C: If t crosses the sphere $\partial B(p_i, R)$:

C1: Subdivide t by adding points along the circular arc of intersection $t \cap \partial B(p_i, R)$

C2: Apply A or B to the constructed triangles.

3: Return the tetrahedra joining p_i and triangles $t \in \mathcal{T}'$

TABLE 1. Computation times for VCM (in sec., on a 3GHz CPU)

Model	Delaunay	Tessellation	Total
Blade (195k)	23.73	90.82	114.55
Bimba (502k)	79.04	305.42	384.46
Nicolò (947k)	95.08	465.53	560.61

Parallelization. Note that both of the algorithms above can be easily parallelized. The Monte Carlo method relies on the ability to answer nearest neighbor queries, and therefore each of the N iterations of the method can be performed on a separate processor once a single shared nearest neighbor data structure is computed. This data structure can also be used for convolution, which, again relies on nothing but nearest neighbor computations.

Note also that although the Voronoi diagram is a global object, whose computation is known to be difficult to parallelize, our tessellation-based method only requires *local* information, and can be parallelized as follows: when computing the VCM of a point p , rather than building a global Voronoi diagram, compute the Voronoi diagram of all points $q \in C$, s.t. $d(p, q) \leq 2R$. Since we are only interested in the intersection of the Voronoi cell of p with a ball of radius R , this set of points is sufficient to compute the VCM of p at scale R (the correctness of this observation follows from the proof of Lemma 4.2 below.) This means that we can limit the amount of memory necessary for computation by only processing batches of the point cloud that sample contiguous regions of the surface.

3. VORONOI COVARIANCE MEASURE

In this section, we review the mathematical background necessary for the analysis of our method. We then define the Voronoi covariance measure of a general compact set, which will allow us to analyze both the robustness of the VCM with respect to noise, and its convergence properties in a single Stability Theorem (Th. 4.1).

3.1. General definition. In order to generalize the definition of VCM to any compact subset of \mathbb{R}^d , we make use of the notion of *projection function*. Given a compact set K of \mathbb{R}^d , the projection function maps any point x of \mathbb{R}^d to its only¹ closest point in K . Denote this function by $p_K : \mathbb{R}^d \rightarrow \mathbb{R}^d$. If C is a finite set of points $\{p_1, \dots, p_n\}$, the projection on C maps any $x \in \mathbb{R}^d$ to one of the points p_i in C whose Voronoi cells contain x . In particular, notice that $p_C^{-1}(p_i)$ is exactly the Voronoi cell of p_i . By analogy, if K is any compact set, we will refer to the inverse image $p_K^{-1}(B)$ of a subset B of K by the projection function p_K as the *Voronoi cell of B* (see Fig. 5). Likewise, we will call $p_K^{-1}(\{p\})$ the *infinitesimal Voronoi cell* of the point p in K .

3.1.1. Definition. The Voronoi covariance measure of K with respect to K^R is a tensor-valued measure² denoted by $\mathcal{V}_{K,R}$. It maps every subset $B \subseteq \mathbb{R}^d$ to a positive

¹In reality, the projection function is uniquely defined everywhere but on the medial axis. Since the d -volume of the medial axis is zero, this has no consequences on our definition.

²Measure has to be understood as in Lebesgue measure theory, as a mass distribution on the space.

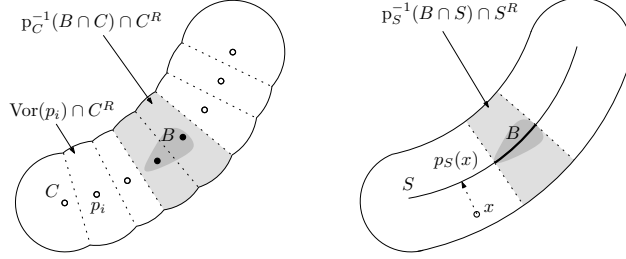


FIGURE 5. Voronoi covariance measures $\mathcal{V}_{C,R}(B)$ of a 2D point cloud C and of the underlying curve S , with respect to the same probing set B .

semi-definite symmetric matrix $\mathcal{V}_{K,R}(B)$ defined by:

$$(3.3) \quad \mathcal{V}_{K,R}(B) = \int_{K^R \cap p_K^{-1}(B \cap K)} (x - p_K(x))(x - p_K(x))^T dx$$

The domain of integration in the definition of $\mathcal{V}_{K,R}(B)$ is the Voronoi cell $p_K^{-1}(B \cap K)$ of the subset B intersected with the chosen offset of K (see Fig. 5). In general, the matrix defined by (3.3) can be thought of as the covariance matrix of the domain of integration with a *varying base point*. Note that for a point cloud, this definition coincides with the definition given in §1.2.

3.1.2. Normal Cones. There is a strong relation between the notion of infinitesimal Voronoi cells and normal cones as defined in many fields of geometry, that is captured by the Voronoi covariance measure. For any p in K , let $\mathcal{N}_{p,R}K$ be the set of vectors defined by $\mathcal{N}_{p,R}K = \{x - p; x \in p_K^{-1}(\{p\}) \cap K^R\}$. This set consists of all vectors that are normal to K at p , at a scale defined by the offset parameter R .

Remark that if S is a smooth hypersurface, and R is small enough, then $\mathcal{N}_{p,R}S$ is the segment $[-R\mathbf{n}(p), R\mathbf{n}(p)]$ (cf Fig. 5). Let B be a small neighborhood of the point p in S such as the ball $B(p, r) \cap S$. Then, by definition, the matrix $\mathcal{V}_{K,R}(B)$ is the integral of all the covariance matrices of the normal cones $\mathcal{N}_{p,R}K$, with p in B . Therefore, this matrix captures the variation of the normals to the surface S around p , which is related to the curvature at p .

If, on the other hand, K is a convex polyhedron of \mathbb{R}^3 , the set $\mathcal{N}_{p,R}K$ coincides with the normal cone at p as defined in convex geometry, intersected with the ball $B(p, R)$. In particular, this set is two-dimensional if the point p lies on an edge of K and three-dimensional when p is a vertex of K (cf Fig. 6). The Voronoi covariance measure of K will reflect this fact, leading to the feature detection algorithm of presented in §6.

3.2. Convolved Voronoi covariance measure. The Voronoi covariance measure can be convolved by any *convolution kernel* $\chi : \mathbb{R}^d \rightarrow \mathbb{R}^+$, which turns it into a function denoted by $\mathcal{V}_{K,R} * \chi$. This operation smoothes the information contained in the covariance matrices, which is particularly useful when dealing with point clouds that can be noisy or that are anisotropically sampled. The value of the

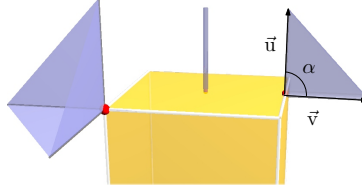


FIGURE 6. The infinitesimal Voronoi cell $p_C^{-1}(x)$ of a point x on a cube is pencil, triangle or cone-shaped depending on the dimension of the normal cone.

convolved VCM at a point p of \mathbb{R}^d is the symmetric matrix defined by the formula:

$$(3.4) \quad \mathcal{V}_{K,R} * \chi(p) := \int_{K^R} (x - p_K(x))(x - p_K(x))^T \chi(p_K(x) - p) dx$$

If χ_r is the *indicator function* of the ball $B(0, r)$ i.e. $\chi_r(p) = 1$ if x belongs to $B(0, r)$ and $\chi_r(p) = 0$ otherwise, then:

$$(3.5) \quad \mathcal{V}_{K,R} * \chi_r(p) = \mathcal{V}_{K,R}(B(p, r))$$

In this work, the convolution kernel is always chosen to be a Lipschitz approximation of such an indicator function, like the “hat function” $\chi(p) = \max(0, r - \|p\|)^2$. We also often use the indicator function itself for which we get good results even though the theoretical guarantees of convergence do not apply directly.

4. ROBUSTNESS OF THE CONVOLVED VCM

Throughout our analysis, we will be using the notion of Hausdorff distance, which can be defined for any two compact sets, K and K' :

$$d_H(K, K') = \inf_{\varepsilon} \text{ s.t. } K \subseteq K'^{\varepsilon} \text{ and } K' \subseteq K^{\varepsilon}$$

where $K^{\varepsilon} = \{x \in \mathbb{R}, d(x, K) < \varepsilon\}$.

For instance, if K is a point cloud sampling of a smooth surface K' , then $d_H(K, K')$ will be small provided that the sampling is dense enough and that there are no outliers. Note that this notion of distance is purely geometric and does not make any assumptions on the uniformity of the sampling, or on the sample points lying on the surface.

This section is devoted to the proof of robustness of the VCM. By robustness we mean that if two compact sets K and K' are very close in the Hausdorff sense, then their convolved VCM are also close. In applications, K will be a sampled surface, and K' will be a point cloud; however this result is very general and does not depend on the nature of these compact sets.

We will suppose that the convolution kernel χ is k -Lipschitz and bounded by a constant M . This means that $|\chi(x) - \chi(y)| \leq k \|x - y\|$ for every pair of points x, y of \mathbb{R}^d , and $|\chi(x)| \leq M$ for any point x . Under this assumption, it is possible to show that the convolved Voronoi covariance measure $\mathcal{V}_{K,R} * \chi$ depends continuously on the compact set K .

THEOREM 4.1 (Stability Theorem for VCM). *Let χ be a bounded k -Lipschitz kernel. For every compact subset K of \mathbb{R}^d and $R > 0$, there exists a constant $\text{cst}(K, R)$*

such that for every point x of \mathbb{R}^d ,

$$\|\mathcal{V}_{K,R} * \chi(p) - \mathcal{V}_{K',R} * \chi(p)\|_{\text{op}} \leq \text{cst}(K, R) d_H(K, K')^{1/2}$$

where the operator norm $\|\cdot\|_{\text{op}}$ is the standard matrix norm induced by the Euclidean metric.

Proof. Let E be the intersection between the offsets K^R and K'^R . Thanks to [CCSM10, Corollary 4.4], we know that the volume of $K^R \setminus E$ behaves as $O(d_H(K, K'))$. Hence, letting $\delta(x) = x - \mathfrak{p}_K(x)$, it is safe to change the integration domain in the definition of VCM and make the approximation,

$$\mathcal{V}_{K,R} * \chi(p) \simeq \int_E \delta(x) \delta(x)^\mathbf{T} \chi(\mathfrak{p}_K(x) - p) dx$$

The goal is then to bound the operator norm of the difference: $M = \int_E P(x) - P'(x) dx$, where P (resp. P') is defined by $P(x) = \delta(x) \delta(x)^\mathbf{T} \chi(\mathfrak{p}_K(x) - p)$. We have

$$\begin{aligned} (4.6) \quad & P(x) - P'(x) \\ &= \chi(\mathfrak{p}_K(x) - p) (\delta(x) \delta(x)^\mathbf{T} - \delta'(x) \delta'(x)^\mathbf{T}) \\ & \quad + (\chi(\mathfrak{p}_K(x) - p) - \chi(\mathfrak{p}_{K'}(x) - p)) \delta'(x) \delta'(x)^\mathbf{T} \end{aligned}$$

Using the definition of δ , one can factor the first term of this sum:

$$\begin{aligned} (4.7) \quad & \chi(\mathfrak{p}_K(x) - p) \left((\mathfrak{p}_K(x) - x) (\mathfrak{p}_K(x) - \mathfrak{p}_{K'}(x))^\mathbf{T} \right. \\ & \quad \left. + (\mathfrak{p}_{K'}(x) - \mathfrak{p}_K(x)) (x - \mathfrak{p}_{K'}(x))^\mathbf{T} \right) \end{aligned}$$

Recall that the kernel χ is bounded by M and that $\|\mathfrak{p}_K(x) - x\| \leq R$ for every $x \in E$. Using the triangle inequality and these facts, the operator norm of (4.7) is bounded by $2RM \|\mathfrak{p}_{K'}(x) - \mathfrak{p}_K(x)\|$.

Thanks to the k -Lipschitz property for the kernel χ , the norm of the second term of the sum in (4.6) can be bounded by

$$\begin{aligned} & |\chi(\mathfrak{p}_K(x) - p) - \chi(\mathfrak{p}_{K'}(x) - p)| \|\delta'(x) \delta'(x)^\mathbf{T}\| \\ & \leq k \|\mathfrak{p}_K(x) - \mathfrak{p}_{K'}(x)\| R^2 \end{aligned}$$

Combining these two bounds in (4.6),

$$\|P(x) - P'(x)\|_{\text{op}} \leq (2RM + kR^2) \|\mathfrak{p}_K(x) - \mathfrak{p}_{K'}(x)\|$$

Integrating this inequality on the set E yields

$$\|M\|_{\text{op}} \leq (2RM + kR^2) \int_E \|\mathfrak{p}_K(x) - \mathfrak{p}_{K'}(x)\| dx$$

The Projection Stability Theorem [CCSM10, Th. 3.2] asserts that $\int_E \|\mathfrak{p}_K(x) - \mathfrak{p}_{K'}(x)\| dx \leq C(d, K, E) d_H(K, K')^{1/2}$. This is enough to conclude. \square

4.1. Consequences of the Stability Theorem. We now let χ_r be a Lipschitz approximation of the characteristic function of the ball $B(0, r)$. If S is a piecewise smooth surface, p a point of S and C_n a sequence of point clouds converging to S in the Hausdorff sense, the Stability Theorem asserts that the convolved VCM $\mathcal{V}_{C_n, R} * \chi_r(p)$ of the point cloud C_n converges to the convolved VCM $\mathcal{V}_{S, R} * \chi_r(p)$ of S with respect to the operator norm. It also quantifies the speed of convergence.

Using the standard results of matrix perturbation theory (see e.g. [SS90]), one then obtains the convergence of the eigenvalues and eigenvectors of $\mathcal{V}_{C_n, R} * \chi_r(p)$ to

those of $\mathcal{V}_{S,R} * \chi_r(p)$, provided that the eigenvalues have multiplicity one. The speed of this convergence depends on the eigengap. As we will see in §5, at a smooth point p of the surface S the eigenvalues of $\mathcal{V}_{S,R} * \chi_r$ are proportional to 1 , $\kappa_1^2(p) \frac{r^2}{4}$ and $\kappa_2^2(p) \frac{r^2}{4}$ where $\kappa_i(p)$ is the i th principal curvature of S at p . One can then expect a faster convergence rate for the estimated normal, and a faster convergence of the estimated principal curvature directions at points where the principal curvatures are very different (i.e. very non-umbilical points).

4.2. Robustness to some outliers. One limitation of the Hausdorff distance in the bound above, is its sensitivity to outliers. Indeed, even under controlled noise, outliers can contaminate the point cloud, and influence the shapes of the Voronoi cells. Nevertheless, as pointed out earlier, intersecting the Voronoi cells with an offset allows us to obtain local information which is unaffected by a certain class of outliers. The following Lemma shows that if the outliers are sufficiently far away from the point cloud, the convolved VCM will remain the same.

Lemma 4.2. *Let C be a point cloud and O a set of outliers with $d(o, C) \geq 2R$ for any o in O . Then the convolved VCM $\mathcal{V}_{C,R} * \chi(p)$ and $\mathcal{V}_{C',R} * \chi(p)$ (where $C' = C \cup O$) agree on C provided that the support of χ is contained in a ball of radius smaller than R .*

Proof. See [MOG09, Lemma 3.3]. □

5. VORONOI COVARIANCE AND CURVATURE

Let S be a compact smooth surface in \mathbb{R}^3 and \mathbf{n} an oriented normal vector field on the surface. As before, let S^R denotes the R -offset of S . The relation between the Voronoi covariance measure and curvature relies on the following two facts:

- (i) For any point x in the ambient space, $x = p_S(x) \pm d_S(x)\mathbf{n}(x)$, where $d_S(x)$ is the distance from x to S .
- (ii) For any point p on the surface, let $\kappa_i(p)$ and $P_i(p)$, $i = 1, 2$, denote the two principal curvatures and principal curvature directions at p . Then the two eigenpairs of the derivative of the normal map $d\mathbf{n}(p)$ are $(\kappa_i(p), P_i(p))$ for $i = 1, 2$.

Combining (i) and the change-of-variable formula, one can prove the following estimation:

$$\begin{aligned}
 & \mathcal{V}_{S,R}(\mathbb{B}(p, r)) \\
 &= \int_{S^R \cap \mathbb{P}_S^{-1}(\mathbb{B}(p, r))} [d_S(x)\mathbf{n}(p_S(x))] [d_S(x)\mathbf{n}(p_S(x))]^T dx \\
 (5.8) \quad & \simeq \int_{q \in \mathbb{B}(p, r) \cap S} \int_{-R}^R [t\mathbf{n}(q)] [t\mathbf{n}(q)]^T dt dq \\
 & \simeq \frac{2}{3} R^3 \int_{q \in \mathbb{B}(p, r) \cap S} \mathbf{n}(q)\mathbf{n}(q)^T dq
 \end{aligned}$$

Putting (ii) in the first-order Taylor expansion of \mathbf{n} gives

$$\begin{aligned}
 & \mathbf{n}(q) \simeq \mathbf{n}(p) + d\mathbf{n}(p)[q - p] \\
 (5.9) \quad & \simeq \mathbf{n}(p) + \sum_{i=1}^2 \kappa_i(p) \langle P_i(p) | q - p \rangle P_i(p)
 \end{aligned}$$

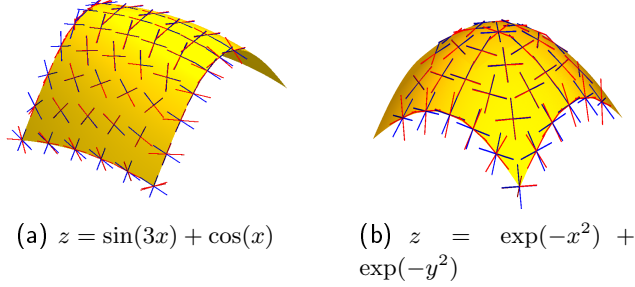


FIGURE 7. Parametric surfaces with exact (in blue) and computed (in red) principal curvature directions. At the boundary computed directions follow the edges of the surface.

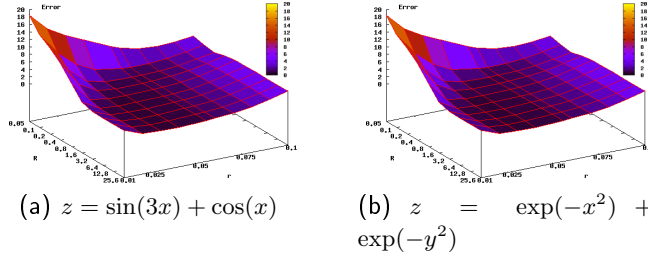


FIGURE 8. Average deviation (in degrees) of computed principal curvature directions from exact ones for different values of parameters R and r .

Then, using (5.8) and (5.9) one obtains the expansion

$$(5.10) \quad \begin{aligned} & \mathcal{V}_{S,R}(B(p,r)) \\ & \simeq \frac{2\pi}{3} R^3 r^2 \left[\mathbf{n}(p) \mathbf{n}(p)^\top + \frac{r^2}{4} \sum_{i=1}^2 \kappa_i^2(p) P_i(p) P_i(p)^\top \right] \end{aligned}$$

More details on these computations are available in [MOG09]. It follows from equation (5.10) that if the offset and convolution radii R and r are chosen small enough, the eigenvectors of the VCM near a smooth point are close to the principal curvature directions at that point.

5.1. Parametric Surfaces. To validate the theoretical guarantees presented above, we tested our method on parametric surfaces for which principal curvatures and principal curvature directions can be computed exactly.

We sampled two functions $z = \sin(3x) + \cos(y)$, with (x, y) in $[0, 1]^2$ and $z = \exp(-x^2) + \exp(-y^2)$, with (x, y) in $[-\frac{1}{2}, \frac{1}{2}]^2$. In both examples we used 100,000 points that were chosen uniformly at random within the domain. Fig. 7 shows these two surfaces with the exact and computed principal curvature directions, using $R = 1$ and $r = 0.055$. As expected, away from the boundary the computed and the exact directions match very closely, except possibly at umbilical points (tip of the second surface). Near the boundary of the domain, the principal directions



FIGURE 9. Mean absolute curvature plot for the Horse and Max Planck models.

computed using our method follow the edges of the surface, which forms the basis of our feature detection method.

To measure the dependence of our method on the offset and convolution radii, we computed the average deviation (in degrees) of principal directions from exact ones. To minimize the effect of the points close to the boundary, where the directions are not expected to match, we only considered 90 percent of the data-set that is farthest from the boundary of the domain. Fig. 8 shows the average deviation for these two parametric surfaces. As can be seen, the results are quite stable for different choices of the parameters.

To illustrate the performance of our curvature estimation algorithm on complex shapes, we also plotted the mean absolute curvature: $|\kappa_1| + |\kappa_2|$ for the standard Max Planck and horse models (Figure 9), where the colors range from yellow to red and pink to signify low to high mean absolute curvature.

5.2. Comparison with Polynomial Fitting.

5.2.1. Sampling Conditions. As mentioned in the introduction, the most common method of estimating principal curvatures and principal curvature directions on point clouds is to locally fit a polynomial of a given degree, and to analytically compute principal curvatures of this fitted polynomial [CP05].

Although these methods work well on noiseless and regularly sampled surfaces, polynomial fitting performs poorly if the data has strong noise and sampling bias. Our method, however, is oblivious to the sampling density, as a consequence of Hausdorff-robustness proved in section 4. To illustrate this, we added 50,000 points in a small band along the diagonal (in the parameter space) to the sampling of the surface shown in Fig. 7(a). The results obtained with our method and with the state of the art polynomial jet-fitting algorithm implemented in CGAL [CP05] are reported in Fig. 10. We used second order polynomial fitting with different neighborhood sizes k , which gave satisfactory results for the original sampling. As can be seen, the extra points do not affect the accuracy of our method. The results obtained with jet-fitting, however, become unreliable and strongly biased in the direction of the oversampling.

5.2.2. Robustness. Other areas that are challenging for polynomial fitting algorithms include parts of the shape with high curvature, and regions where separate parts of the shape come in close contact, thus adversely influencing the quality of

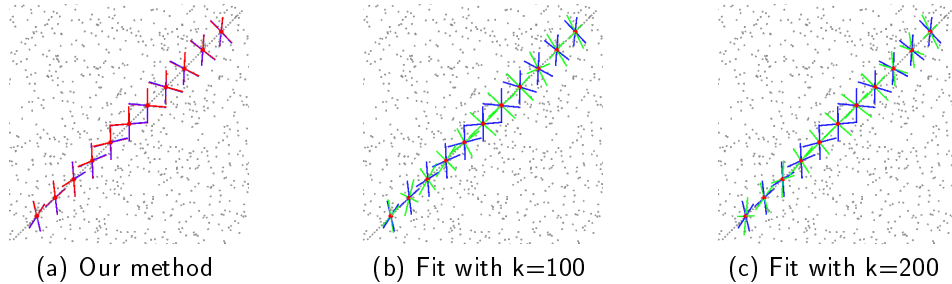


FIGURE 10. Principal curvature directions on a biased dataset. Jet fitting (*b-c*) produces unreliable directions (in green) following oversampled areas.




FIGURE 11. Principal curvature directions computed with our method (in red) are stable under noise. Directions computed by jet-fitting (in blue) are unreliable, especially when points from separate parts of the shape begin to mix.

the fit. While the first problem can be addressed by fitting higher order polynomials, both of these settings are severely aggravated in the presence of noise. To illustrate this, we sampled a surface made by smoothly joining a small cylinder lying parallel to the z axis, with two planes on either side of the x axis. We used 0.1 as the radius of the cylinder, so the curvature at points along the z axis equals 10. Fig. 11 shows the principal curvature directions obtained on this model with varying levels of noise, using our method and using jet-fitting with 200 nearest neighbors and second order polynomial fitting. Note that when points from different parts of the shape mix, even the robust low order polynomial fitting fails, while our method preserves robustness

6. DETECTING SHARP EDGES USING VCM

As remarked by Dey et al in [DGGZ03], the shape of the Voronoi cell of a point p can be used to estimate the dimension of the part of the underlying shape this point belongs to. For instance, consider a point p lying on a sharp edge of a piecewise smooth surface S . As Fig. 6 shows, the infinitesimal Voronoi cell of p is triangle shaped, and the normal to the triangle is aligned with the direction of the sharp edge. Under good sampling conditions, the corresponding Voronoi cell will be very thin, and almost two-dimensional. Here, we show how to use the Voronoi covariance measure to translate this remark into a sharp feature detection algorithm, whose

sampling conditions only involves the Hausdorff distance between the point cloud and the underlying surface.

Our main ingredient for this is to estimate the dimension of the infinitesimal Voronoi cell at a point p of the input point cloud by the number of “small” eigenvalues s of the convolved VCM of the cloud. This number can then be used to distinguish between a point lying on a face ($s = 2$), on a sharp edge ($s = 1$) and on a corner ($s = 0$). The actual meaning of “small” is made more precise in Proposition 6.1.

6.1. Theoretical derivation. We assume that the piecewise-smooth surface S we consider has a one-sided reach at least R , which we denote by $\text{reach}_1(S) \geq R$. This means the following: for any point p on the surface S , and any unit vector v in the normal cone³ to p , the projection of either $p + Rv$ or $p - Rv$ on the surface S should lie at p .

When a point p belongs to an edge of the surface S , denote by $\mathbf{u}_\pm(p)$ the leftmost and rightmost normals to the surface S at P . Also let $\mathbf{e}_1(p)$ be the mean of these two extreme normals, $e_1(p) = \frac{1}{2}(\mathbf{u}_+(p) + \mathbf{u}_-(p))$, and $\mathbf{e}_2(p)$ be orthogonal to $e_1(p)$ in the plane spanned by $\mathbf{u}_\pm(p)$, i.e $e_2(p) = \frac{1}{2}(\mathbf{u}_+(p) - \mathbf{u}_-(p))$. Finally, let $\alpha(p)$ be the positive angle between the two normals $\mathbf{u}_\pm(p)$ in the range $[0, \pi]$, and $\mathbf{t}(p)$ be the tangent direction to the edge.

PROPOSITION 6.1. *Let S be a polyhedron of \mathbb{R}^3 whose 1-reach is lower bounded by R , and let p lie on a sharp edge of S . Then, the three eigenvalues of $\mathcal{V}_{S,R}(\mathbb{B}(p, r))$ are*

$$\begin{aligned}\lambda_1(p) &= \frac{R^4 r}{8} [\alpha(p) + \sin \alpha(p) + O(r/R)]; \\ \lambda_2(p) &= \frac{R^4 r}{8} [\alpha(p) - \sin \alpha(p) + O(r/R)]; \\ \lambda_3(p) &= 0.\end{aligned}$$

The corresponding eigenvectors are $\mathbf{e}_1(p)$, $\mathbf{e}_2(p)$ and $\mathbf{t}(p)$.

Proof. Suppose that r is small enough so that the only parts of the surface intersected by the ball $\mathbb{B}(p, r)$ are the edge E containing p and the two adjacent faces F_+ and F_- . Thanks to the additivity property of the VCM, one can write:

$$(6.11) \quad \begin{aligned}\mathcal{V}_{S,R}(\mathbb{B}(p, r)) &= \mathcal{V}_{S,R}(\mathbb{B}(p, r) \cap E) + \mathcal{V}_{S,R}(\mathbb{B}(p, r) \cap F_+) \\ &\quad + \mathcal{V}_{S,R}(\mathbb{B}(p, r) \cap F_-)\end{aligned}$$

Using Lemma 6.2 below, we show that the term of (6.11) involving the edge, $\mathcal{V}_{S,R}(\mathbb{B}(p, r) \cap E)$, yields the constant parts of the eigenvalues $\lambda_1(p)$ and $\lambda_2(p)$. In order to get the $O(R/r)$ approximation error, one just needs to bound the two terms of equation (6.11) involving the faces F_\pm :

$$\begin{aligned}\|\mathcal{V}_{S,R}(\mathbb{B}(p, r) \cap F_\pm)\| &\leq \int_{F_\pm \cap \mathbb{B}(p, r)} \int_{-R}^R \|(\mathbf{t}\mathbf{n}_\pm)(\mathbf{t}\mathbf{n}_\pm)^\mathbf{T}\| dt dx \\ &= \pi \frac{r^2 R^3}{3} = \frac{R^4 r}{8} O(r/R)\end{aligned}$$

where $\mathbf{n}_\pm = \mathbf{u}_\pm(p)$ is the unit normal to the face F_\pm .

³Recall that a vector v belongs to the normal cone to a point p of the surface S iff for some $\lambda > 0$, the projection of $p + \lambda v$ onto S is p .

□

Lemma 6.2. *Let \mathbf{u}_\pm be two unit vectors in the plane, and Q denote the positive cone spanned by them intersected with the ball $B(0, R)$, i.e $Q = \{a\vec{u} + b\vec{v}; a, b > 0\} \cap B(0, R)$. The eigenpairs of the covariance matrix $\text{cov}(Q, 0)$ are:*

$$\begin{aligned}\lambda_1 &= \frac{R^4}{8} (\alpha + \sin(\alpha)), \quad \mathbf{e}_1 = \mathbf{u}_+ + \mathbf{u}_- \quad \text{and} \\ \lambda_2 &= \frac{R^4}{8} (\alpha - \sin(\alpha)), \quad \mathbf{e}_2 = \mathbf{u}_+ - \mathbf{u}_-\end{aligned}$$

where α is the angle between \mathbf{u}_- and \mathbf{u}_+ in the range $[0, \pi]$.

Proof. Let $M = \text{cov}(Q, 0)$, and consider the unit vector $\mathbf{d}(\beta)$ making an angle β with the vector \mathbf{u}_- . Then,

$$\begin{aligned}(6.12) \quad \mathbf{d}^\mathbf{T}(\beta)M\mathbf{d}(\beta) &= \int_{\theta=0}^{\alpha} \int_{r=0}^R r^2 \cos^2(\theta - \beta) r dr d\theta \\ &= \frac{R^4}{8} \int_0^{\alpha} \cos(2(\theta - \beta)) + 1 d\theta \\ &= \frac{R^4}{8} (\sin(2(\alpha - \beta)) + \sin(2\beta) + \alpha)\end{aligned}$$

By symmetry of Q , the vector $\mathbf{d}(\alpha/2)$ is necessarily an eigenvector of M . The second eigenvector is then $\mathbf{d}(\pi/2 + \alpha/2)$. To compute the eigenvalues associated with these eigenvectors, we plug in the values into (6.12). □

From Proposition 6.1, we deduce that the eigenvector corresponding to the smallest eigenvalue is the tangent direction to the feature, with an eigengap of $O(R/r)$.

6.1.1. Sharp corner of a surface. Similar calculations can be done if p is a vertex of the surface whose normal cone is isotropic. In this case, the two smallest eigenvalues will be equal, and contained in the plane orthogonal to the direction of the cone. The case where the normal cone isn't isotropic is more difficult to analyze precisely because the computations will depend on the sharp edge graph at p .

6.2. Description of our edge-detection method. The method needs three parameters: the offset radius R , the convolution radius r and a threshold T . It can be summarized as follows:

- 1:** Compute the Voronoi covariance measure $\mathcal{V}_{C,R}$.
- 2:** For every point $p \in C$:
 - 2a:** Compute $\mathcal{V}_{C,R}(B(p, r))$ and diagonalize it.
 - 2b:** Sort the eigenpairs $(\lambda_i(p), \mathbf{e}_i(p))_{i \in \{1,2,3\}}$ by decreasing order of eigenvalues.
 - 2c:** Compute the ratio $r(p) = \lambda_2(p) / (\sum_i \lambda_i(p))$.
 - The point p is on a sharp edge if $r(p) \geq T$
 - The edge direction is given by $e_3(p)$.

6.2.1. Choice of parameters. The choice of the **offset radius** R should be done by considering the expected geometry of the underlying set. As Proposition 6.1 indicates, if the offset radius is less than the one-sided reach of the object, all of the sharp edges will be recovered by the method above. A too large offset parameter R can decrease the anisotropy of the Voronoi cell of a point that lies on a sharp feature (see Fig. 12), thus making it invisible to the method. From Theorem 4.1, one also

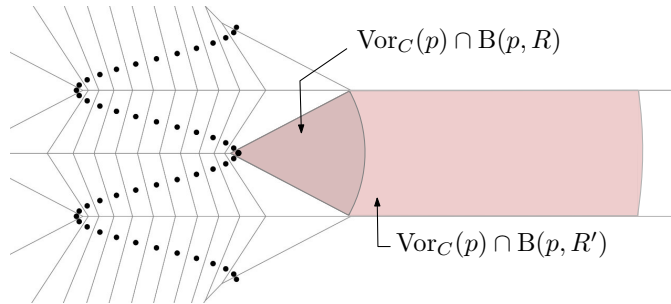


FIGURE 12. Effect of the choice of offset radius on the anisotropy of the Voronoi of a point lying on a sharp feature.

sees that a higher value of the offset radius R will make the computations more robust. In conclusion, the offset radius should be chosen as the largest possible lower bound on the one-sided reach of the underlying surface.

The **convolution radius** r will be chosen depending on the expected density and amount of noise in the point cloud: it should be large enough so that every ball $B(p, r)$ with p in C contains at least a few tens of points (eg. 50), and should definitely be larger than the amount of noise in the normal direction to the surface. However, a compromise should be made: the larger r , the less localized the detected sharp features are. We explore the dependence of the results on the choice of the convolution parameter in section 6.3.

We consider a point as a feature if the ratio of the second smallest eigenvalue to the sum of the three eigenvalues of the convolved Voronoi covariance measure is greater than some **threshold parameter**. This ratio provide a way to measures the thickness of the infinitesimal Voronoi cell. If p belongs to a sharp edge, Proposition 6.1 shows that

$$r(p) \simeq \frac{1}{2} \left(\frac{\alpha - \sin(\alpha)}{\alpha} \right) \simeq \frac{\alpha^2}{12}$$

Thus, given a threshold parameter T , a sharp edge with external angle α will be selected if $\alpha \geq 2\sqrt{3}T^{1/2}$.

Note that sometimes it may be necessary to detect only sharp edges filtering out corners as well as smooth parts of the shape. In this case, we use the ratio of the second smallest to the smallest eigenvalue. This ratio will be large for points around sharp edges, and close to one for points in the smooth areas as well as around the corners of the surface.

6.3. Experimental results. In order to assess the quality of the set of sharp edges detected by the method, we compare it to the set of sharp edges of the underlying surface.

6.3.1. Resilience to noise. We evaluated the sharp edge estimation method and its resilience to noise on a unit icosahedron. We sampled 100k points randomly on it, ran the computations with $R = 20$ and various convolution parameters r . In order to test the resilience of our method to noise, the original 100k point cloud on the icosahedron is perturbed by adding to each point a random vector uniformly chosen in a ball of given radius (the radius of the noise). Fig. 13 shows the estimated feature directions on the icosahedron, for different noise and convolution radii.

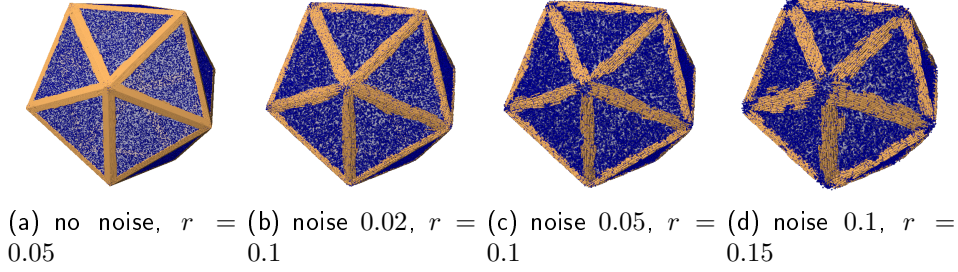


FIGURE 13. Estimated feature directions on a unit icosahedron, with various convolution radii r and noise values.

TABLE 2. Distances between the detected features (D) and real features (R) of an icosahedron, with varying noise radius and convolution radius. The distances δ_∞ , δ_1 and α_1 are between D and R while δ'_∞ and δ'_1 are between R and D

Noise	r	δ_∞	δ_1	α_1	δ'_∞	δ_1
0.0	0.05	0.35	0.037	3.25	0.076	0.011
0.0	0.1	0.118	0.051	0.33	0.124	0.016
0.02	0.1	0.226	0.049	1.65	0.139	0.020
0.05	0.1	0.220	0.050	2.82	0.155	0.025
0.1	0.15	0.271	0.069	3.12	0.178	0.036

We call *oriented set* a finite subset C of \mathbb{R}^d , each of whose points is endowed with a tangent direction $\mathbf{t}(p)$. We assess the similarity of two such sets C and C' using the following metrics. The *Half-Hausdorff distance* $\delta_\infty(C, C')$ (resp. average distance $\delta_1(C, C')$) is the maximum (resp. average) distance between a point in C and its nearest neighbor in C' . Finally, the *average angular deviation*: $\alpha_1(C, C')$ is the average angle (in degrees) between the direction $\mathbf{t}(p)$ of a point p in C and the direction $\mathbf{t}(p')$ of its nearest neighbor p' in C' .

Table 2 summarizes the values of the distances between the oriented set D of detected features, and the set R of real features, for various noise and convolution radii. The distance $\delta_\infty := \delta_\infty(D, R)$ measures the presence of isolated outliers, while $\delta_1 := \delta_1(D, R)$ measures the spread of the estimated features. One can see that increasing the convolution radius removes isolated outliers, but increases the spread of feature. The distances with primes (δ'_∞ and δ'_1) correspond to the same metrics when swapping the role of D and R . Both δ'_∞ and δ'_1 measure how “complete” the set of features is, i.e. how far points lying on actual sharp features are to the set of detected feature D . Most of the error here comes from the corners, that are discarded based on the ratio between the second and third eigenvalues.

6.3.2. Sharp edges with low external angle. In order to understand the effect of sharpness of the edge on the quality of the feature detection, we sampled a surface

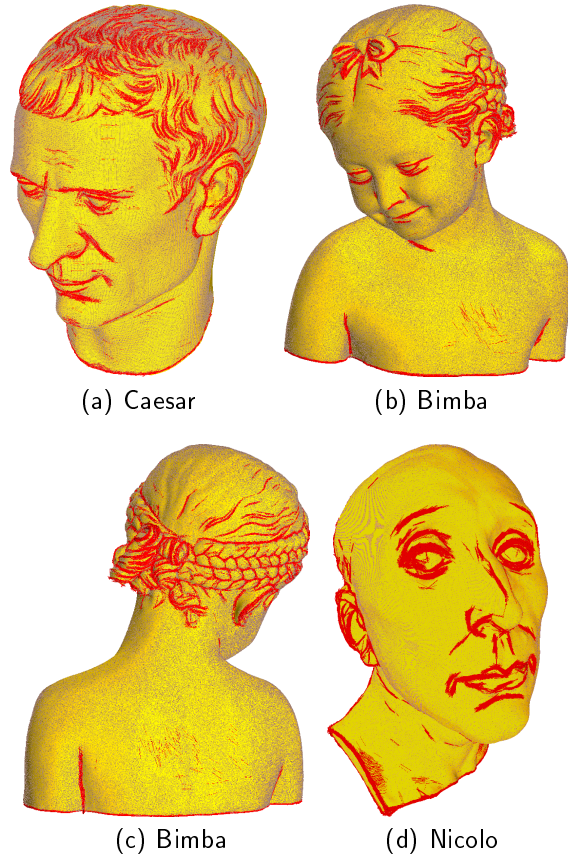


FIGURE 14. Results of the feature detection algorithm on the Julius Caesar (a), Bimba (b-c) and Nicolò (d) point clouds.

made of two planar rectangular patches joined by a common edge and whose normals differ by an angle of 2α . As shown in Fig. 15, the feature estimation method described above is able to reliably detect sharp edges whose external dihedral angle is as small as 2° ; all the results were produced using the same convolution radius and threshold.

6.3.3. Results on more complex point clouds. On larger range scan point clouds, we were able to decrease both the offset and the convolution radii while keeping the detected features almost noiseless. This enables the algorithm to capture very small and non-sharp features, like the hair of Caesar or the braiding of Bimba, as shown in Fig. 14(a),14(b),14(c).

7. SURFACE SEGMENTATION

As a sample application to our feature detection method, we use it to segment point cloud samplings of piecewise smooth surfaces. Here, given a point cloud, our goal is to segment it into components sampled from contiguous regions of the surface separated by sharp edges and highly curved areas (see Fig. 16). The use

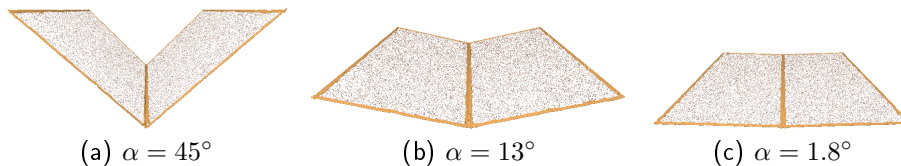


FIGURE 15. Estimated feature directions on a folded rectangle, for different values of the external angle α .

of our feature detection method for point cloud segmentation was suggested by Frederic Chazal and Steve Oudot, who also contributed to the work presented in this section.

Shape segmentation has applications in reverse engineering and CAD model generation as well as surface reconstruction [JWS08], rendering of point clouds [FCOS05, LCOL07], and object recognition and matching [DGG04]. As in the case of curvature and feature estimation, most existing methods rely on meshes (see [Sha08] for a survey). Nevertheless, a rich body of work also exists in the context of point cloud segmentation with various targeted applications (e.g. [GKF09, GG04]). Unfortunately most existing methods lack a sound theoretical foundation. In this section, we demonstrate that our provably correct feature detection method can be used in shape segmentation. Note that rather than opposing existing techniques, our objective is to illustrate the performance of our feature detection method for shape segmentation and to make a step towards a provably correct segmentation technique.

A recent method for segmenting point clouds with theoretical guarantees on the number of components was proposed by Chazal and colleagues [CGOS09], using ideas from topological persistence. Given a scalar function f defined over a point cloud C , and a neighborhood parameter μ , their method starts by building a Rips graph of C by introducing an edge between any two points of C whose Euclidean distance is less than μ . Then, each point is associated to a local minima of the function f through a simple gradient descent on the Rips graph. The corresponding attraction basins are merged according to a persistence threshold parameter α . Using methods from the theory of topological persistence, the authors are able to prove that the number of components thus obtained is stable under perturbations of the function and irregular point cloud sampling of the underlying space.

We apply the method of Chazal et al. [CGOS09] to segment point clouds into smooth regions. To define a function on the point cloud, we use a two step process: first we identify sharp edges and corners using the method described above. Then, we build a Rips graph on the point cloud using the same μ parameter as defined earlier. For each point p in the cloud, define $g(p)$ as the geodesic distance in the Rips graph from p to the closest feature point. Since the method of Chazal et al. identifies local minima rather than maxima, we apply their algorithm to the function $f(p) = -g(p)$.

This process is illustrated in Fig. 16. The input point cloud is a 200k sampling of the Fandisk model 16(a). The feature points identified by our method are shown in Fig. 16(b); the negative distance function to the features in Fig. 16(c), and the final point cloud segmentation in Fig. 16(d). In Fig. 17, we show the results of this

segmentation method on point cloud samplings of several other models. Note that one of the advantages of persistence-based clustering [CGOS09] is that this method is very global in nature. Therefore, even if spurious features are detected, e.g. due to a very noisy sampling, introducing additional local minima of the negative distance function, the segments associated with these minima will be merged with more persistent “true” local minima.

CONCLUSION AND FUTURE WORK

In this paper, we have described a method for detecting sharp features and feature directions on point cloud data in a unified fashion. We have provided theoretical guarantees on its robustness, and implemented it and tested it on various point cloud data with varying amounts of sampling bias and noise.

In the future, it would be interesting to combine the theoretical stability properties of persistence-based clustering [CGOS09] with our feature detection method to give explicit guarantees on the quality of the segmentation method for point cloud samplings of piecewise-smooth surfaces described in §7.

One drawback of the Voronoi covariance measure is its sensitivity to outliers that are at intermediate distance from the point cloud. A possible way to overcome this limitation would be to use the recently introduced notion of distance-to-measure [CCSM09].

ACKNOWLEDGMENTS

This work has been supported by INRIA/Stanford Associated Team TGDA, NSF grants FRG 0354543, TF 0634803 and FODAVA 808515, as well as by DARPA grant HR0011-05-1-0007. Some of the models are issued from the AIM@Shape Shape Repository. We would also like to thank the authors of [CGOS09] for sharing the implementation of their clustering algorithm.

REFERENCES

- [AB99] N. Amenta and M. Bern, *Surface reconstruction by voronoi filtering*, Discrete Comput. Geom **22** (1999), 481–504.
- [ACSTD07] P. Alliez, D. Cohen-Steiner, Y. Tong, and M. Desbrun, *Voronoi-based variational reconstruction of unoriented point sets*, Proc. Eurographics Symp. Geom. Proc., 2007, pp. 39–48.
- [AK04] N. Amenta and Y.J. Kil, *Defining point-set surfaces*, ACM Trans. Graph. **23** (2004), no. 3, 264–270.
- [AMCO08] D. Aiger, N. Mitra, and D. Cohen-Or, *4-points congruent sets for robust pairwise surface registration*, ACM Trans. Graph. **27** (2008), no. 3.
- [BC94] J. Berkman and T. Caelli, *Computation of surface geometry and segmentation using covariance techniques*, IEEE Trans. Pattern Anal. Mach. Intell. **16** (1994), no. 11, 1114–1116.
- [CCSLT09] F. Chazal, D. Cohen-Steiner, A. Lieutier, and B. Thibert, *Stability of Curvature Measures*, Proc. Eurographics Symp. Geom. Proc., 2009.
- [CCSM09] F. Chazal, D. Cohen-Steiner, and Q. Mériçot, *Geometric Inference for Measures based on Distance Functions*, Research Report RR-6930, INRIA, 2009.
- [CCSM10] ———, *Boundary measures for geometric inference*, Found. Comput. Math. **10** (2010), no. 2, 221–240.
- [cgal] CGAL, *Computational Geometry Algorithms Library*, <http://www.cgal.org>.
- [CGOS09] F. Chazal, L. Guibas, S. Y. Oudot, and P. Skraba, *Analysis of scalar fields over point cloud data*, Proc. ACM/SIAM Symp. Disc. Algo., 2009, pp. 1021–1030.
- [Con86] M.L. Connolly, *Measurement of protein surface shape by solid angles*, J. Mol. Graph. **4** (1986), no. 1, 3–6.

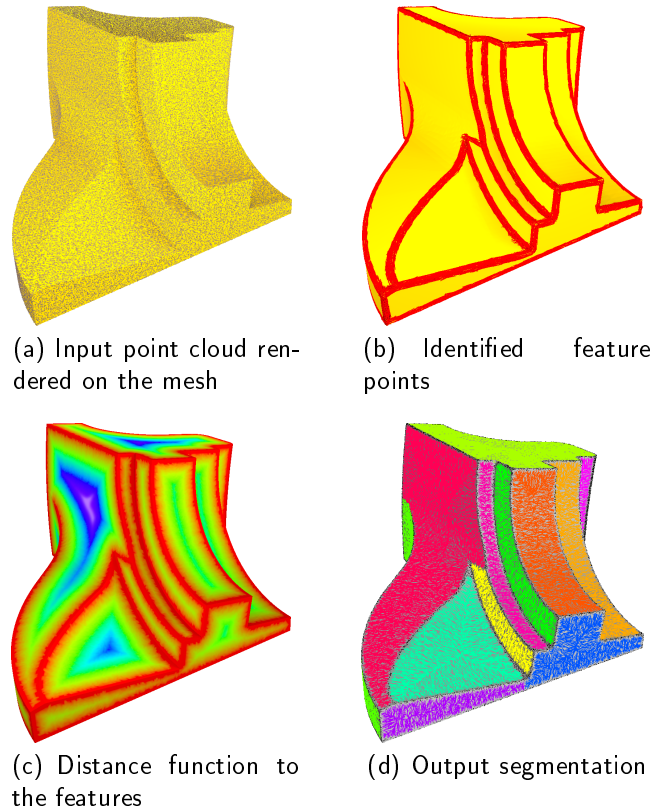


FIGURE 16. A segmentation of a point cloud sampling of the fan-disk model using our feature detection method and the clustering algorithm of Chazal et al [CGOS09].

- [CP05] F. Cazals and M. Pouget, *Estimating differential quantities using polynomial fitting of osculating jets*, *Comp. Aided Geom. D.*, 2005, pp. 121–146.
- [CSM03] D. Cohen-Steiner and J.M. Morvan, *Restricted delaunay triangulations and normal cycle*, *Proc. ACM Symp. Comput. Geom.*, 2003, pp. 312–321.
- [DGG04] T. K. Dey, J. Giesen, and S. Goswami, *Shape segmentation and matching from noisy point clouds*, *Proc. Eurographics Symp. Point-based Graph.*, 2004, pp. 193–199.
- [DGGZ03] T.K. Dey, J. Giesen, S. Goswami, and W. Zhao, *Shape dimension and approximation from samples*, *Discrete and Computational Geometry* **29** (2003), no. 3, 419–434.
- [DIOHS08] J. Daniels II, T. Ochotta, L.K. Ha, and C.T. Silva, *Spline-based feature curves from point-sampled geometry*, *The Visual Computer* **24** (2008), no. 6, 449–462.
- [FCOS05] S. Fleishman, D. Cohen-Or, and C. T. Silva, *Robust moving least-squares fitting with sharp features*, *ACM Trans. Graph.* **24** (2005), no. 3, 544–552.
- [GG04] Natasha Gelfand and Leonidas J. Guibas, *Shape segmentation using local slippage analysis*, *Proc. Eurographics Symp. Geom. Proc.*, 2004, pp. 214–223.
- [GKF09] Aleksey Golovinskiy, Vladimir G. Kim, and Thomas Funkhouser, *Shape-based recognition of 3D point clouds in urban environments*, *IEEE Int. Conf. on Comp. Vis.* (2009).
- [GTE⁺06] J. P. Gois, E. Tejada, T. Etienne, L. G. Nonato, A. Castelo, and T. Ertl, *Curvature-driven modeling and rendering of point-based surfaces*, *Braz. Symp. Comp. Graph. Imag. Proc.*, 2006, pp. 27–36.

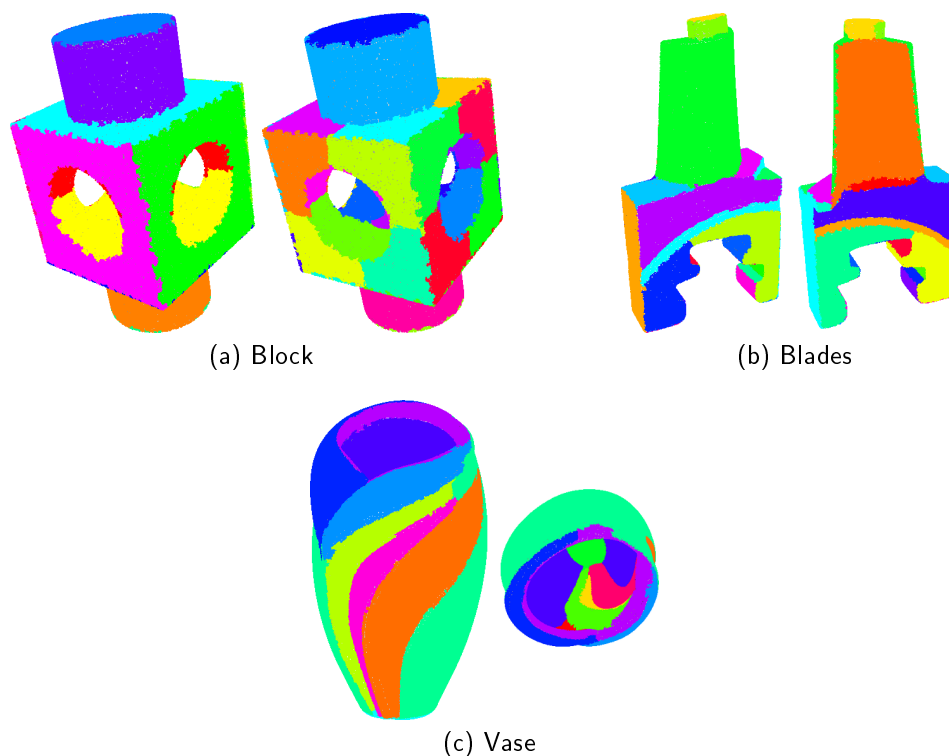


FIGURE 17. Segmentations of point cloud samplings of the different models using our method.

- [HM02] J. Huang and C. H. Menq, *Combinatorial manifold mesh reconstruction and optimization from unorganized points with arbitrary topology*, *Comp. Aided Geom. D.* **34** (2002), no. 2, 149–165.
- [JWS08] P. Jenke, M. Wand, and W. Straßer, *Patch-graph reconstruction for piecewise smooth surfaces*, *Proc. of Vis., Mod., Vis.*, 2008.
- [KNSS09] Evangelos Kalogerakis, Derek Nowrouzezahrai, Patricio Simari, and Karan Singh, *Extracting lines of curvature from noisy point clouds*, *Special Issue of the Elsevier Computer-Aided Design journal on Point-Based Computational Techniques* **41** (2009), no. 4, 282–292.
- [LCOL07] Y. Lipman, D. Cohen-Or, and D. Levin, *Data-dependent MLS for faithful surface approximation*, *Proc. Eurographics Symp. Geom. Proc.*, 2007, pp. 59–67.
- [LP05] C. Lange and K. Polthier, *Anisotropic smoothing of point sets*, *Comp. Aided Geom. D.* **22** (2005), no. 7, 680–692.
- [MA] D. Mount and S. Arya, *ANN: A library for approximate nearest neighbor searching*.
- [MOG09] Q. Mérigot, M. Ovsjanikov, and L. Guibas, *Robust Voronoi-based curvature and feature estimation*, *Proc. ACM Symp. Solid Phys. Modeling*, 2009, pp. 1–12.
- [MSR07] E. Magid, O. Soldea, and E. Rivlin, *A comparison of Gaussian and mean curvature estimation methods on triangular meshes of range image data*, *Comput. Vis. Image Und.* **107** (2007), no. 3, 139–159.
- [OGG09] C. Oztireli, G. Guennebaud, and M. Gross, *Feature preserving point set surfaces based on non-linear kernel regression*, *Eurographics*, 2009.
- [PKG03] M. Pauly, R. Keiser, and M. Gross, *Multi-scale Feature Extraction on Point-Sampled Surfaces*, *Comput. Graph. Forum*, 2003.

- [PMW⁺08] M. Pauly, N. Mitra, J. Wallner, H. Pottmann, and L. Guibas, *Discovering structural regularity in 3d geometry*, ACM Trans. Graph. **27** (2008), no. 3.
- [PWHY09] H. Pottmann, J. Wallner, Q. Huang, and Y.L. Yang, *Integral invariants for robust geometry processing*, Comp. Aided Geom. D. **26** (2009), 37–60.
- [Sha08] Ariel Shamir, *A survey on mesh segmentation techniques*, Comput. Graph. Forum **27** (2008), no. 6, 1539–1556.
- [SS90] G.W. Stewart and J. Sun, *Matrix perturbation theory*, Computer Science and Scientific Computing, 1990.
- [YQ07] P. Yang and X. Qian, *Direct computing of surface curvatures for point-set surfaces*, Proc. Eurographics Symp. Point-based Graph., 2007.

CNRS, LJK, F-38000 GRENOBLE, FRANCE

COMPUTER SCIENCE DEPARTMENT, STANFORD UNIVERSITY, STANFORD CA, CA 94305

COMPUTER SCIENCE DEPARTMENT, STANFORD UNIVERSITY, STANFORD CA, CA 94305