



HAL
open science

Growing Least Squares for the Continuous Analysis of Manifolds in Scale-Space

Nicolas Mellado, Pascal Barla, Gaël Guennebaud, Patrick Reuter, Christophe Schlick

► **To cite this version:**

Nicolas Mellado, Pascal Barla, Gaël Guennebaud, Patrick Reuter, Christophe Schlick. Growing Least Squares for the Continuous Analysis of Manifolds in Scale-Space. Computer Graphics Forum, 2012, Proceedings of Symposium on Geometry Processing 2012, p. hal-00713678v1

HAL Id: hal-00713678

<https://inria.hal.science/hal-00713678v1>

Submitted on 18 Jul 2012 (v1), last revised 24 Feb 2015 (v2)

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.

Growing Least Squares for the Analysis of Manifolds in Scale-Space

Nicolas Mellado[†], Gaël Guennebaud, Pascal Barla, Patrick Reuter, Christophe Schlick
Inria - Univ. Bordeaux - IOGS - CNRS

Abstract

We present a novel approach to the multi-scale analysis of point-sampled manifolds of co-dimension 1. It is based on a variant of Moving Least Squares, whereby the evolution of a geometric descriptor at increasing scales is used to locate pertinent locations in scale-space, hence the name “Growing Least Squares”. Compared to existing scale-space analysis methods, our approach is the first to provide a continuous solution in space and scale dimensions, without requiring any parametrization, connectivity or uniform sampling. An important implication is that we identify multiple pertinent scales for any point on a manifold, a property that had not yet been demonstrated in the literature. In practice, our approach exhibits an improved robustness to change of input, and is easily implemented in a parallel fashion on the GPU. We compare our method to state-of-the-art scale-space analysis techniques and illustrate its practical relevance in a few application scenarios.

1. Introduction

The use of multi-scale algorithms is ubiquitous in image and geometry processing, and may serve very diverse purposes. In computer graphics, *multi-resolution* representations have been widely used. They are based on the decomposition of a spatial signal into base and detail layers, which may be used for instance in signal processing (e.g., [BEA83]) or for compression purposes (e.g., [Mal08]). In both cases, the goal of a multi-resolution structure is to help *modify* a signal in ways that depend on the targeted application. In this paper, we are rather interested in *scale-space* techniques [Lin94, Rom09], where the goal is altogether different: it consists in analyzing a signal at different scales to *discover* its geometric structure, independently of any application. While multi-resolution methods strive to decompose the signal into independent layers, scale-space techniques critically rely upon redundant information between scales to characterize the structure of a spatial signal.

A classical approach consists in tracking signal extrema (0-crossings) at increasing scales, using the scale at which they get annihilated to characterize this structure. Unfortunately, this method raises important issues. First, the construction of a scale space requires a dense regular sampling in both spatial and scale dimensions. Most importantly, the tracking of extrema provides information only for a subset of scale-space (at signal extrema); it is not able to identify multiple pertinent scales per point; and it is sensitive to noise or

small changes in the input signal, as detailed in Section 2.1. Moreover, it is difficult to adapt scale-space theory to the analysis of point-sampled manifolds as explained in Section 2.2. To the best of our knowledge, there is no method that has proposed a continuous alternative to extrema tracking, and that avoids the need of a regular parametrization while remaining computationally tractable for manifolds of co-dimension 1 (e.g., curves in 2D, surfaces in 3D).

The key idea of this paper is to perform the scale-space analysis of manifolds by means of continuous algebraic fits. In our approach, a scale-space is built through least-square fits of a low-degree algebraic surface onto neighborhoods of continuously increasing sizes. In some sense, it can be seen as an adaptation of the Moving Least Squares formalism [Lev98a] to continuously varying scales, hence the name “Growing Least Squares”. The use of an algebraic surface ensures robust fits even at large scales and yields a rich geometric descriptor with only a few parameters. The continuity of the fitting process through scales provides for a stable and elegant analysis of geometric variations.

Our approach exhibits a number of advantages over previous scale-space analysis methods of manifolds:

- It provides a fully continuous alternative to extrema tracking, based on the derivative of a local geometric descriptor tailored to scale-space analysis;
- It does not require any connectivity or parametrization, and works from points equipped with normals;
- It is easy to implement on the GPU, and robust to the scale sampling and to small changes in input points.

[†] nicolas.mellado@inria.fr

2. Background

In this section, we present previous work relevant to the analysis of manifolds in scale-space. This excludes pyramidal methods [FMR11], spectral processing [LZ09], or topological persistence [ELZ02, CGOS11], since these are based on formalisms that do not permit to analyze redundancy.

2.1. Scale-space theory

Scale-space theory has first been applied to 1D signals [Wit83] and gray-level images [Lin94]. Its main idea is to represent the input signal as a one-parameter family of smoothed signals. The parameter that controls smoothing size is called the *scale*, and the intuition behind this approach is to ignore small variations at increasing scales. The most common smoothing operator is the Gaussian kernel, in which case we talk about linear scale-space. Contrary to multi-resolution approaches that mostly rely on pyramidal representations, scale-space makes use of stack representations: the resolution is kept unchanged, which creates a lot of redundancy from one scale to the other.

Such a redundancy is central to scale-space analysis, since it permits to study accurately how a signal evolves from scale to scale. The core of analysis techniques relies on *differential invariants* (e.g., gradient magnitude, curvature). The common approach consists in tracking 0-crossings of those invariants along the scale dimension, and to record locations where pairs of them annihilate. This analysis exhibits what has been called the *deep structure* (e.g., [Rom09], p.154): 0-crossings are considered more or less pertinent depending on the scale at which they get annihilated. Once identified, the deep structure may be put to use in a variety of applications: denoising, matching, simplification, etc.

The main benefit of scale-space theory is that the signal is analyzed independently of any potential application, fully automatically, without requiring the tuning of parameters. Unfortunately, there are some important limitations: the pairing of 0-crossings is not robust to small changes in the input (as shown in [Mal00], p.86), and spurious 0-crossings may appear at increasing scales in 2D [Lin94]. Practical algorithms have limitations of their own: the tracking of 0-crossings requires a dense sampling in the scale dimension; and the smoothing operator requires a regular uniform sampling. Alternative methods based on the tracking of level-crossings [WS90, DUM*11] have partly addressed theoretical limitations, but they are still limited by the aforementioned practical issues.

Most importantly, the adaptation of tracking-based methods to the analysis of manifolds is far from straightforward. As opposed to the image domain which offers a natural parametrization, manifolds are $(d-1)$ -dimensional signals in \mathbb{R}^d , and often lack parametrization, uniform sampling and even connectivity. Hence, specific solutions are required to analyze manifolds in scale-space.

2.2. Manifolds in scale-space

Differential geometry. One way to adapt the scale-space theory to manifolds is to compute differential invariants of point coordinates. This approach has first been applied to curves and surfaces by Mokhtarian et al. [MM86, MKY01]. Its main limitation is that it requires a parametrization and a regular sampling on extended neighborhoods, which are often difficult and/or costly to obtain. A simpler approach employs the Difference of Gaussians (DoG) that approximates the Laplacian, first introduced for images [MH80] and later adapted to meshes [ZBVH09]. It works in two stages: 1) smooth the manifold at two scales; 2) approximate curvature at a point by the distance between its two smoothed locations. This method does not require any parametrization.

The DoG method is preferable to methods based on parametrization when dealing with complex (e.g., scanned) geometry, although it still relies on a connected mesh in existing implementations. Moreover, it only provides local curvature information. Both approaches suffer from common scale-space issues: need of a dense sampling and lack of robustness of 0-crossings. Even though spurious 0-crossings could be discarded (e.g., using [ELZ02, CGOS11]), other practical issues will remain.

Numerical approximation. An alternative solution to avoid computing a parametrization is to perform measurements of manifold properties that are related to curvature. One example is direct curvature space smoothing [ZH97, ZH99], which uses a global spherical space to smooth curvature-like values at different scales. Heat kernel smoothing [SOG09] replaces the classic smoothing step by a geodesic Gaussian diffusion to compute a heat kernel value at each scale. Integral invariants [HFG*06, PWHY09] compute the volume occupied by the object in a ball of varying size, which plays the role of scale.

Like the DoG, these methods do not require any parametrization, and have thus been used in applications such as matching and registration of scanned surfaces (they also require a mesh reconstruction step when applied to point sets). Although they have proven to be efficient for extracting global characteristics of manifolds, they are not well adapted to scale-space analysis where we are interested in identifying precise geometric structures: the computed values are valid curvature approximations only at small scales.

Local regression. A third approach to analyze manifolds in scale-space consists in first locally fitting simple geometric kernels to input data, and then extracting properties related to differential invariants. Performing regression with a first-order kernel intuitively corresponds to the fitting of a plane to input points and normals. Surface anisotropy may then be extracted [YLHP06, LG05, IT11], as well as a covariance measure [PKG03]. Regression through 2nd-order kernels [BSF02, CPG09] provides curvature information, but requires successive fitting steps (first fit a plane, then fit a quadric to the residual height data). This raises issues when the surface “folds-over” itself.

The main benefit of local regression methods is that they work directly from point sets, while all other methods require a connected mesh and ignore normals altogether. As before, the differential invariants that may be extracted from fitted kernels are limited. Indeed, although curvature is computed with second order kernels, the use of two fitting steps greatly complexifies subsequent scale-space analysis: indeed, curvature values are not directly related when the support plane of the fitted quadric changes from scale to scale.

3. Our approach

The approach taken in this paper is based on local regression. It works directly with unorganized point sets equipped with normals and thus does not require any connectivity. Moreover, contrary to previous techniques, it provides a geometric descriptor that is directly amenable to *continuous* scale-space analysis.

Our approach draws inspiration from Moving Least Squares (MLS) [Lev98b], whereby local regression is performed in a globally consistent way, resulting in continuously differentiable surfaces. Our first main contribution is to adapt MLS to perform scale-space analysis, taking advantage of their continuity along the scale dimension. In a nutshell, instead of shifting in space a weight function of constant scale as with MLS, we study the evolution of the fitting results at a constant position with a growing weight function support. This leads us to our second key contribution: instead of considering simple curvature measures or eigenvalue ratios, we propose to exploit and analyze the entire low-degree algebraic surfaces resulting from the fits.

4. Growing least squares

We first briefly explain our choice of local regression in section 4.1, along with a reparametrization that makes it relevant for the analysis of manifolds in scale-space. We then extend the method in section 4.2 to provide an alternative to the tracking of 0-crossings, based on analytic derivatives along the scale dimension. Third, we show how to adapt our approach in section 4.3 to measure pair-wise geometric dissimilarities at arbitrary scales and positions.

4.1. Scale-space via local regression

The first step of our approach is to characterize any point \mathbf{p} of a manifold at any scale t by a low-degree algebraic surface that best approximates its neighborhood \mathcal{P}_t . In a discrete setting, our manifold is described by a set of points $\mathbf{q}_i \in \mathbb{R}^d$, with d being the dimension of the ambient space, and the neighborhood \mathcal{P}_t consists in the set of data points contained in a ball of radius t centered at \mathbf{p} : $\mathcal{P}_t = \{\mathbf{q}_i; \|\mathbf{q}_i - \mathbf{p}\| \leq t\}$. Inspired by recent work on MLS reconstruction [GG07], we use algebraic hyper-spheres which have the advantages of being easy to fit in a robust manner, while providing second-order information with a minimal number of parameters. We assume each point \mathbf{q}_i is equipped with a normal $\mathbf{n}_i \in \mathbb{R}^d$. In case normals are not provided we estimate them using covariance analysis, as done in previous work.




Figure 1: Fitting and reparametrization. The weight function (in green) around a point \mathbf{p} is defined in a neighborhood of size t (dashed green line). Points \mathbf{q}_i and normals \mathbf{n}_i that belong to this neighborhood are fitted by an algebraic hyper-sphere $s_{\mathbf{u}}$; its 0-isosurface is shown in black, and its scalar field with a color code. This sphere is reparametrized in terms of three geometric parameters: the mean curvature κ , the offset τ , and the gradient direction $\boldsymbol{\eta}$ at \mathbf{p} .

Fitting. An algebraic sphere is implicitly defined as the 0-isosurface of the following scalar field (see Figure 1):

$$s_{\mathbf{u}}(\mathbf{x}) = [1 \ \mathbf{x}^T \ \mathbf{x}^T \ \mathbf{x}] \cdot \mathbf{u}, \quad (1)$$

where $\mathbf{u} \in \mathbb{R}^{d+2}$, $\mathbf{u} = [u_c \ \mathbf{u}_\ell \ u_q]^T$ is the vector of (respectively constant, linear and quadratic) parameters. In order to fit such a sphere onto a set of neighborhood points \mathcal{P}_t , we employ the fast fitting technique of Guennebaud et al. [GGG08]. First, \mathbf{u}_ℓ and u_q are computed by minimizing $\sum_i w_i(t) \|\nabla s_{\mathbf{u}}(\mathbf{q}_i) - \mathbf{n}_i\|^2$, where $\mathbf{q}_i \in \mathcal{P}_t$ and w_i is a *scale-dependent* weight function:

$$w_i(t) = \left(\frac{\|\mathbf{q}_i - \mathbf{p}\|^2}{t^2} - 1 \right)^2. \quad (2)$$

Second, the constant coefficient u_c is obtained by minimizing in a least square sense the algebraic distance to the samples: $\sum_i w_i(t) \|s_{\mathbf{u}}(\mathbf{q}_i)\|^2$. These two minimizations yield closed-form formulas which are recalled in the Appendix.

Normalization. Contrary to [GG07], our goal is not to reconstruct a surface from a point set, but instead to analyze its shape at multiple scales. To this end, we want to assign a *unique* and *meaningful* geometric descriptor for any choice of point \mathbf{p} and scale t . A straightforward solution would be to use the center \mathbf{c} and radius r of the hyper-sphere. Unfortunately, this leads to degenerated cases when the surface is locally planar: in particular, \mathbf{c} becomes undefined.

We thus rather consider the scalar field itself $s_{\mathbf{u}}$ as a geometric descriptor. However, there exists an infinity of scalar fields (based on scalar multiples of \mathbf{u}) that correspond to the same hyper-sphere. To solve this issue and consistently pick a unique solution, we use Pratt's normalization [Pra87]: its basic idea is to constrain the scalar field to have a unitary gradient vector on the 0-isosurface, yielding:

$$\hat{\mathbf{u}} = [\hat{u}_c \ \hat{\mathbf{u}}_\ell \ \hat{u}_q]^T = \mathbf{u} / \sqrt{\|\mathbf{u}_\ell\|^2 - 4u_c u_q}. \quad (3)$$

This choice has the additional advantage to make algebraic distances near-Euclidean for points close to the 0-isosurface.

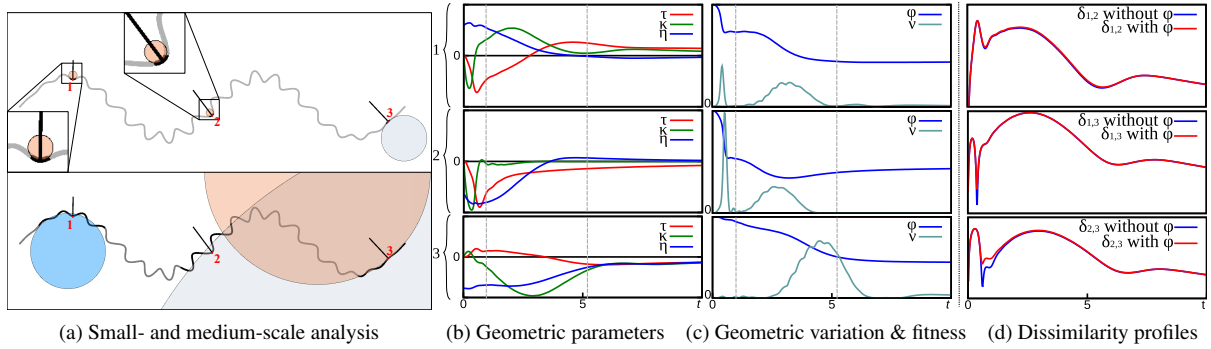


Figure 2: 2D analysis. The analysis of a synthetic 2D curve composed of two sinusoids of different frequencies is illustrated at 3 different points. In (a), we show their geometric descriptor at two scales, whose parameters are visualized in (b) for all scales, with one point per row. In (c), we display geometric variations and fitness: note that the 3rd point has a more stable structure at small scales since the magnitude of the high-frequency component is low at its location. In (d), we display dissimilarity measures for all pairs of points. Observe how the use of the fitness parameter helps disambiguate between the two types of inflexion points at intermediate scales. In all plots, scale sampling is quadratic.

Reparametrization. After normalization, we are left with a scalar field $s_{\hat{\mathbf{u}}}$ for which a geometric interpretation is far from evident. Firstly, $\hat{\mathbf{u}}_c$ and $\hat{\mathbf{u}}_\ell$ do not correspond to any measurable physical quantity. Secondly, all $d+2$ parameters are still interdependent, since the normalization binds them together with: $\|\hat{\mathbf{u}}_\ell\|^2 - 4\hat{u}_c\hat{u}_q = 1$. We propose an alternative parametrization of the scalar field, which is illustrated in Figure 1. Intuitively, its parameters consist of: the algebraic offset distance τ between the evaluation point \mathbf{p} and the 0-isosurface; the unit normal $\boldsymbol{\eta}$ of the scalar field at \mathbf{p} ; the signed curvature κ of the hyper-sphere. When the fitting degenerates to a plane, τ represents the distance from the origin to the plane, $\boldsymbol{\eta}$ its normal, and κ vanishes (see Figures 2 and 3 for 2D&3D illustrations, respectively).

Formally, the geometric parameters are given by:

$$\tau = s_{\hat{\mathbf{u}}}(\mathbf{p}); \quad \boldsymbol{\eta} = \frac{\nabla s_{\hat{\mathbf{u}}}(\mathbf{p})}{\|\nabla s_{\hat{\mathbf{u}}}(\mathbf{p})\|}; \quad \kappa = 2\hat{u}_q. \quad (4)$$

Thanks to Pratt's normalization, the offset τ provides a close approximation to the Euclidean distance between \mathbf{p} and the 0-isosurface. The normal parameter $\boldsymbol{\eta}$ gives the direction to the point on the hyper-sphere that is closest to \mathbf{p} . The curvature parameter simply corresponds to the inverse of the hyper-sphere radius r , and has the advantage of behaving continuously when passing through a locally planar surface, while r tends toward infinity. Note that with these parameters, the scalar field can no longer be expressed as a linear combination of monomials, as shown in the Appendix.

Fitness. Once reparametrized, the scalar field fitted from \mathcal{P}_t yields a univocal geometric descriptor invariant to rigid transformations. However, a given geometric descriptor can be associated to a space of generator neighborhoods. Considering the sphere as a whole instead of, for instance, local curvature only, already permits to significantly reduce the size of these spaces. It is interesting to remark that they can be further reduced by looking at the fitness ϕ that exhibits how close the \mathbf{q}_i are to the fitted scalar field $s_{\hat{\mathbf{u}}}$. We define this




Figure 4: The fitness parameter ϕ helps to disambiguate two identical fits, e.g., a smooth versus a bumped neighborhood (left), or a flat versus a saddle configuration (right).

additional parameter by $\phi = \sum_i w_i(t) \nabla s_{\hat{\mathbf{u}}}(\mathbf{q}_i) \cdot \mathbf{n}_i / \sum_i w_i(t)$, where w_i is the same weighting function as before. Using the fitting equations 7, it can be shown that for this fitting procedure, ϕ boils down to Pratt's norm for \mathbf{u} : $\phi = \|\mathbf{u}_\ell\|^2 - 4u_c u_q$. Hence, ϕ does not have to be measured or stored explicitly since it is embedded in the vector of parameters \mathbf{u} . Note that by construction, ϕ is dimension-less, scale-invariant, and varies in the $[0, 1]$ range, with $\phi = 1$ meaning a perfect alignment between the fitted scalar field and input normals. As illustrated in Figure 4, this typically permits to disambiguate between surfaces that locally have the same geometric description, but differs from a pure algebraic sphere.

The chain of operations described in this section augments an arbitrary point \mathbf{p} at an arbitrary scale t with a geometric descriptor that characterizes data points $\mathbf{q}_i \in \mathcal{P}_t$, as illustrated in Figure 2(a-b) for a 2D curve, and Figure 3(a-c) for a 3D surface. It thus describes an elegant method for building a fully continuous scale-space from sampled manifolds of co-dimension 1, providing meaningful surface information in the form of the $\tau, \boldsymbol{\eta}, \kappa$ geometric parameters, and a fitness parameter ϕ that further helps disambiguate similar descriptors. Another handy property of our geometric descriptor is that negating its parameters yields the *complement* descriptor, which is equivalent to the descriptor of the same surface with an opposite orientation. As detailed in Section 5, this is a significant improvement compared to existing methods. One may be tempted to use it to track 0-crossings of the curvature κ for instance, as in previous work. Instead, we show in the next section that scale-space analysis may also be performed in a fully continuous manner.

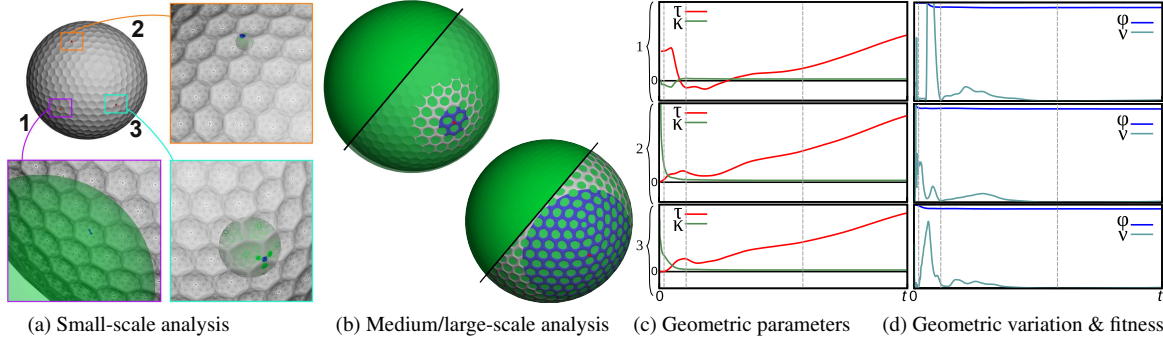


Figure 3: 3D analysis. The analysis of a 3D golf ball model is illustrated at 3 different points: a concavity, an edge and a junction. This is shown at a small scale in (a), where we see that the three points (with their neighborhoods in blue) have quite dissimilar geometric descriptors (in green). However, at medium and large scales (b), all 3 points converge to a same global sphere (only one point is shown for clarity). This is best observed in (c), where we display their geometric parameters (except for η), and in (d), where all 3 points converge together. In all plots, scale sampling is quadratic.

4.2. Continuous scale-space analysis

The main purpose of analyzing a spatial signal in scale-space is to track its variations at increasing scales to discover its geometric structure. A classical approach consists in tracking invariants in the form of 0-crossings of a *spatial* derivative of the signal; then find locations in scale-space where they get annihilated. As mentioned in Section 2, this leads to many shortcomings, especially when trying to deal with manifolds. Most importantly, it restricts the analysis to a subset of locations in scale-space while requiring a parametrization.

We propose a different approach to discover the multi-scale structure of a manifold. Our key insight is to observe that, in general, a pertinent scale for \mathbf{p} is one where its geometric descriptor exhibits *minimal variation* when the neighborhood size increases. This suggests that, at such scales, the parameters of our descriptor do not crucially depend on scale, but rather indicate stable geometric properties of the manifold. In this section we focus on the derivation of such a general geometric variation; we will show its relevance in Section 5 and how it can be exploited in Section 6.

One may think that the curvature parameter κ is the one mostly involved in geometric variations. Figure 5 shows counter examples where either τ or η have significant influence. We thus compute the variations of all 3 geometric parameters and combine them in a natural fashion to yield a geometric variation function $v(\mathbf{p}, t)$ that describes the scale-space structure of the input manifold.

Scale derivatives. The variation of geometric descriptor parameters we are looking for are simply given by their partial derivatives along the scale dimension at (\mathbf{p}, t) : $d\tau/dt$, $d\eta/dt$, and $d\kappa/dt$. We emphasize that we are interested in geometric variations only, hence the fitness ϕ does not play a role here. Since we make use of a local regression that is both continuous and given in closed-form, these derivatives are easily computed analytically, provided the weight functions themselves are differentiable. Their computation does not yield any difficulty, and simply involves differentiating the chain




Figure 5: Variations of τ and η . We show the hyper-spheres fitted for the red point with various support sizes (indicated by dark dots). On the left, both the offset τ and curvature κ vary, while on the right only the normal direction η varies.

of equations presented in Section 4.1, i.e. from top to bottom: weighting (Eq. 2), fitting (Eq. 7), normalization (Eq. 3) and reparametrization (Eq. 4).

Geometric variation. The geometric variation function $v(\mathbf{p}, t)$ is obtained by a weighted squared sum of these partial derivatives. In order not to introduce any bias, a special care has to be taken in the choice of these weights. In particular, we propose the following weighting scheme:

$$v(\mathbf{p}, t) = \left(\frac{d\tau}{dt} \right)^2 + \left(t \frac{d\eta}{dt} \right)^2 + \left(t^2 \frac{d\kappa}{dt} \right)^2, \quad (5)$$

which has the fundamental advantages to yield a dimensionless and scale-invariant measure, and to naturally give equal importance to each parameter. Indeed, let us for instance choose meters m for the unit of length. Thanks to our intuitive reparametrization of Section 4.1, we have τ in m , the unit-less η , and κ in m^{-1} . Moreover, by construction it is reasonable to expect to have τ to be mostly comprised in $[-t, t]$, κ in $[-\frac{1}{t}, \frac{1}{t}]$, while $\|\eta\| = 1$. Therefore, a reasonable choice is to scale the parameters (τ, η, κ) by $(1/t, 1, t)$ respectively in order to get scale-invariant and unit-less quantities of the same order of magnitude. Finally, in order to compensate for the differentiation over the scale t that is in m , it is natural to multiply by the scale t , thus leading to the scaling factors $(1, t, t^2)$ of Equation 5.

The function $v(\mathbf{p}, t)$ is one of the key contribution of this paper. It provides a continuous description of pertinent

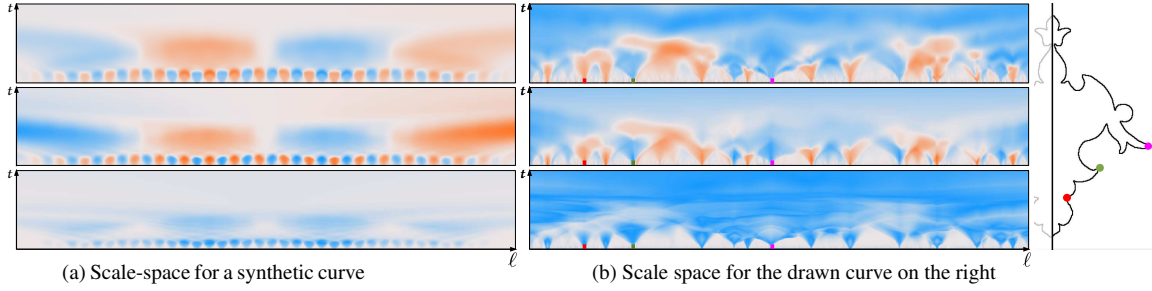


Figure 6: Curvature scale spaces (from top to bottom: DoG, ours and CA), with orange colors for concavities and blue colors for convexities. In (a), we show curvatures for the sinus-based function of Figure 2(a); in (b), we show curvatures for the symmetric curve shown on the right. Compared to DoG, our approach exhibits less noise at small scales and a better convergence at large scales. CA does not make a difference between convex and concave regions, which merge together.

scales for any point \mathbf{p} , and it is robust to small changes in the input. This is in contrast with previous approaches that rely on the annihilation of extremal points (0-crossings) and may lead to altogether different structures when the input changes only a little, as shown in Section 5. Most importantly, our approach is the first to identify *multiple* pertinent scales for individual points on manifolds, as exemplified in Figures 2(c) and 3(d) for 2D and 3D cases, respectively. This opens the door to many new applications, for which we sketch a few examples in Section 6.

4.3. Pairwise dissimilarity in scale-space

For a variety of applications, it is also interesting to compare a pair of arbitrary scale-space locations (\mathbf{p}_a, t_a) and (\mathbf{p}_b, t_b) . In this context, it is crucial to provide a measure invariant to similarity transformations. Invariance to translation is readily available because the geometric parameters are defined relative to their fitted points. Since our descriptor is isotropic, invariance to rotation is achieved by aligning the respective unit normals $\boldsymbol{\eta}_a$ and $\boldsymbol{\eta}_b$, which amounts to ignore these parameters. Finally, scale invariance is obtained using the weights derived in the previous paragraph, and we thus define the dimension-less dissimilarity function $\delta_{a,b}$ by:

$$\delta_{a,b} = \left(t_a^{-1} \tau_a - t_b^{-1} \tau_b \right)^2 + (t_a \kappa_a - t_b \kappa_b)^2 + (\varphi_a - \varphi_b)^2 \quad (6)$$

Here the fitting errors φ_a and φ_b help disambiguate between similar descriptors that may correspond to different surfaces. Recall that it is dimension-less, scale-invariant, and of the same order of magnitude than the other quantities since it varies between 0 and 1.

The dissimilarity measure is illustrated in Figure 2(d), where we compare three pairs of points at multiple scales. Other examples are presented in Section 6, where we combine geometric variation and dissimilarity measurements.

5. Comparisons with previous work

In this section, we demonstrate the benefits of our scale space construction and analysis techniques and compare them to standard approaches. We only consider methods

that, like ours, avoid fold-over issues, do not require any parametrization and provide valid approximations of curvature at all scales: DoG curvature scale space [ZBVH09] and covariance analysis (CA) [PKG03].

Nevertheless, this section mostly makes use of parametrized 2D curves to 1) visualize the different scale-spaces in a comprehensive manner, and 2) permit comparisons with DoG and 0-crossings.

Curvature scale space. Let us first focus on the scalar-valued scale-space obtained using our curvature measure κ . All our 2D visualizations use the arclength ℓ and scale t parameters for the horizontal and vertical axis, respectively. Figure 6(a) compares our approach to DoG and CA for the smooth sinus-like curve of Figure 2. On this example, all methods manage to identify the two signals, though CA cannot distinguish between concavities and convexities, and DoG fails to capture important variations at curve borders. On a more complex example, as shown in Figure 6(b), the different methods start to exhibit more different behaviors. In particular, DoG does not reach convergence at large scales: it keeps on introducing meaningless structures. In contrast, our curvature κ produces a much smoother and coherent scale space, even though it does not rely on input point connectivity.

Stability comparison. In Figure 7, we illustrate an important limitation of tracking-based scale-space analysis. When the input curve is modified so that one of its bumps is slightly more prominent, then the corresponding set of 0-crossings changes abruptly, suggesting that the intrinsic structure of the curve has changed in similar respects. In contrast, our geometric variation scale-space evolves continuously to reflect the more subtle change of structure implied by this slight modification. In particular, intermediate scales around the bump region are no longer considered persistent once the more prominent bump is set in place. Indeed, increasing the amplitude of the bump has the effect of “breaking the structure” of the slope on the left side of the curve.

Relationship to 0-crossings. Figure 8 shows a layering of 0-crossings on top of our geometric variation scale-space.




Figure 7: Stability comparison. Top row: a convex bump is added to the left side of a simple open curve. Middle row: the structure of 0-crossings changes abruptly even for small changes in the input (red and green lines show the 0-crossings of κ and of its curvilinear derivative κ' , respectively). Bottom row: the geometric variation v changes continuously to reflect the subtle change in the input curve (white means low variation). Note how the stable region at intermediate scales, which corresponds to the slope on the left side, is progressively filled with a more complex structure.




Figure 8: Relationship to 0-crossings. The geometric variation v for the second drawing in Figure 7 is displayed with 0-crossings of κ (in red) and κ' (in green) layed on top.

Two types of differential invariants are shown: 0-crossings of κ and its spatial derivative κ' . Let's write the corresponding scale-space points C and C' for convenience. First, note that all annihilation events for both C and C' occur in regions of low geometric variation, but there are other such regions which are never reached by either C or C' . This suggests that the tracking of extrema actually leads to a *subset* of pertinent scale-space locations according to v . It is also interesting to note that points in C (resp. C') seem to be attracted toward regions of high (resp. low) variation, a tendency we plan to study in future work.

Our second observation concerns the largest scale: by construction, only a few points from C' (and none from C) reach the top of scale-space, and they are related through tracking to a single point on the original manifold. However, there is no reason to identify a precise location as a representative of such a large scale, since it corresponds to an extremely smooth shape. Our approach reflects this convergence since most points have a low geometric variation in this case, and correspond to a *global* geometric descriptor.

6. Application scenarios

We have shown that our geometric variation scale-space provides a robust and continuous characterization of the geometric structure of a manifold. We now give examples of its usefulness through three application scenarios related to denoising, feature detection and matching/registration. A full

comparison with dedicated techniques in each of these application domains is out of the scope of this paper. Instead, we focus on the novelties brought by our scale-space analysis, essentially its continuity and its ability to detect multiple pertinent scales per point. Some implementation details are given at the end of this section.

Adaptive Bandwidth. When a point-sampled manifold is corrupted with spatially varying noise as shown in Figure 9, it is not appropriate to reconstruct the signal at a single global scale. Instead, we must find a spatially-varying scale locally adapted to the amplitude of noise, also called adaptive bandwidth [WSS09]. However, care must be taken not to over-estimate this minimum scale, otherwise pertinent manifold structures may be damaged.

Thanks to our geometric variation scale-space, it is possible to adaptively estimate a proper bandwidth: intuitively, we only have to find the smallest pertinent scale for each point \mathbf{p} . As a proof of concept, we propose a simple top-down heuristic that works well for smooth objects. We detect the noise in a coarse-to-fine fashion: for each point we set a bandwidth at scales where dv/dt is greater than a given threshold (we use 0.01 in this example). We then regularize the result with a spatial smoothing across data points. We will study the extension of this method to more complex objects in future work. The resulting adaptive bandwidth is used as a variable support size to reconstruct a smooth curve, shown in orange in Figure 9 for different kind of noise.

Continuous features. To deal with complex manifolds, many geometry processing applications rely on a preliminary feature extraction step, which identifies a subset of salient points to consider for further processing. Probably the most famous technique in this field is SIFT (e.g., [IT11]), which identifies local features that are related to locations where 0-crossings of κ' annihilate. Because such 0-crossings are sensitive to noise and small changes in the input, SIFT points may be filtered out in a number of ways to keep only the most salient locations.

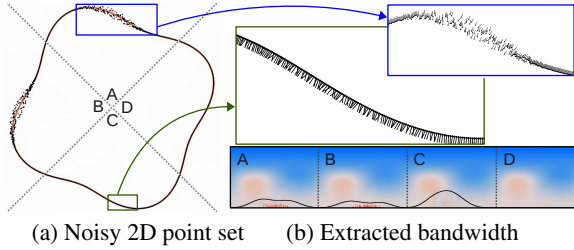


Figure 9: Adaptive bandwidth: a 2D point set (a) is corrupted with noise on: normals and positions (A), positions (B) and normals (C); D is without noise. The reconstructed curve (shown in orange) is obtained using an adaptive bandwidth (black curve) extracted from our continuous scale-space representation (b).

Instead of relying on individual points, we propose the continuous feature function $f(\mathbf{p}) = \int \bar{\nu}(\mathbf{p}, t) dt$, where $\bar{\nu}$ is a smooth remapping of ν that disregards too strong geometric variations (we use a *tanh* function to this end). The intuition behind this formula is that points for which f is small are subject to nearly no geometric variations across scales (e.g., $f = 0$ on a sphere); whereas points with a high f indicate that our geometric descriptor is pertinent only at a few scales.

This is shown on complex curves and surfaces in Figure 10 with a color gradient. Observe how the most pronounced features (in red) may correspond to widely different concave and/or convex regions; this is desirable because there is often no *a priori* on the sign of curvature of features. Also note that regions with repetitive small details are not considered as important, whereas isolated or more prominent shape details emerge as key features of the model.

Multi-scale dissimilarity. The ability to evaluate local surface similarities is the cornerstone of many matching or registration techniques. However, the notion of multi-scale similarity hides an important question: which scales should actually be taken into account? One may think that taking all scales at once is a natural solution. However, as shown in Figure 11(a), this might not always be the case, at least for some applications. The question becomes delicate as soon as multiple scales are nested in a same object: one may want to find similarities at the smallest of these scales, or perhaps in other situations at the largest scale.

Our approach permits to make this choice in a continuous manner, and we illustrate it with a simple picking tool: the user selects a point on a surface, and the system finds all similar points on the same object, given a rough scale prior. The basic idea is to combine the global prior with the local geometric variation to compute a per-scale dissimilarity, which is then integrated over the scales t_i to yield a multi-scale dissimilarity $\Delta_{\mathbf{p}, \mathbf{q}} = \sum_i \delta_{\mathbf{p}, \mathbf{q}}(t_i) h(t_i)$, where $\delta_{\mathbf{p}, \mathbf{q}}(t)$ is the dissimilarity measure from Equation 6, and h is a normalized weighting function over scales that defines the global prior. In our example, we use a simple box filter for h .

Figure 11(b-c) visualizes $\Delta_{\mathbf{p}, \mathbf{q}}$ as a function of \mathbf{q} (in blue) for a given point \mathbf{p} (in red) located both on a small ridge and

on the 'S' of the SGP acronym. By varying the global prior, our method identifies unambiguously either one detail layer or the other. Although our approach is based on an isotropic regression, this result demonstrates that anisotropic features at pertinent scales are properly extracted, even junctions and corners. Moreover, this is done irrespective of the shape of the base surface, since the letters SGP are extracted similarly on different locations of the torus.

Implementation details. In our system we use a kd-tree for the neighbor search, and the fitting at scale t is performed by collecting all the neighbors within the distance t of the current point. This makes the complexity of our algorithm quadratic with respect to t for a 3D surface. We implemented our approach both on the CPU (on all 8 cores of an Intel I7 3.40Ghz) and on the GPU (using CUDA on a GTX 580). As expected, the GPU implementation outperforms the CPU version; for the example of Figure 10(c), the analysis takes only 6.3s versus 221s. All reported timings correspond to our CUDA implementation. We believe a huge speed-up could be achieved by iteratively smoothing and simplifying the input point cloud to compute the larger scales [PKG06].

7. Discussion & Future work

We have presented a novel approach to scale-space analysis based on local regression, that automatically and robustly characterizes the stable structure of a manifold of co-dimension 1 in a continuous manner. A strength of this approach is that it is entirely independent of any targeted application. Indeed, our continuous geometric variation scale-space may be interpreted and processed differently by different types of applications, as demonstrated in Section 6.

Our solution considerably improves previous work: it is fully continuous in both space and scale, it is robust to noise (Figures 9 and 12), it does not require any parametrization or connectivity, and it naturally deals with manifold borders. The DoG method could be adapted to handle point sets and manifold borders using local regression for smoothing. However, it would still have to be applied to data points with geodesic neighborhoods, which are often difficult to obtain, and it would provide only curvature measurements. Our approach identifies a more complete geometric descriptor while only requiring an Euclidean neighborhood with a reasonable number of point samples (at least about 5 in 2D and 12 in 3D). Quadric fitting is another alternative to compute curvature, but it is not adapted to scale-space analysis since it requires a local planar parametrization which is not globally coherent and not robust to fold-overs. In contrast, our method avoids these shortcomings by employing a total-least square fitting procedure.

A limitation of our approach comes from the choice of an isotropic regression: surface anisotropy is not explicitly identified. However, as shown in Figure 11, our method is able to detect indirectly complex anisotropic structures. As shown in Figure 13, our scale-space analysis even permits to identify saddle-like shapes. The case of a perfect saddle




Figure 10: Continuous features are displayed using a color gradient (a), on a symmetric 2D closed curve (b), and on 3 variants of the Armadillo 3D model. Observe how features of varying shapes and sizes (in red) are properly selected in all cases.

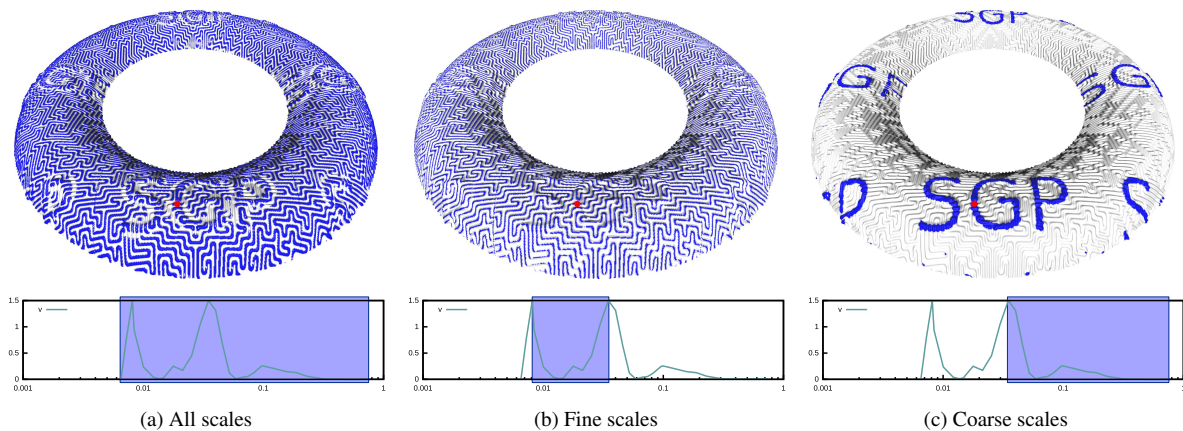


Figure 11: Multi-scale similarity. Top row: Points similar to the picked point (in red) are selected (in blue) via our dissimilarity measure. Bottom row: the type of selected feature depends on a user-controlled global prior (blue box), which is locally refined by our geometric variation. In (a), all scales are selected. In (b), only the fine displacement pattern emerges. In (c), the large-scale SGP letters are properly segmented (SGP torus: 500k pts, 20 scales, 42sec). We use a log remapping on the scale axis.

patch requires the take fitness into account, as shown in Figure 4. In future work, we also plan to study spatial variations of our geometric descriptor, in a way similar to what is classically done with MLS. It would permit to explicitly identify direction fields on manifolds and the scales at which these are pertinent. In this paper, we have only considered shape geometry and ignored additional information such as color attributes. In future work, we plan to extend our weight functions to reject neighbor points of dissimilar attributes, yielding a *non-linear* version of our continuous scale-space analysis. We have also observed oscillations in v for input manifolds with regular structures (see Figures 2(c) and 3(d)). For some applications, they might lead to false positives, which could be detected through a frequential analysis of v .

Our versatile and continuous geometric scale-space analysis has a great potential for a large variety of applications that we want to explore in future work. For example, in the domain of Cultural Heritage, where it is common to con-

front different hypotheses, our continuous scale-space could be used to design semi-automatic methods that guide archaeologists in shape matching tasks. Multi-resolution representations could also benefit from our analysis as a pre-process, for instance for the identification of the most meaningful scales at which details should be computed. This could be particularly interesting in the context of base/relief surface decomposition. Finally, the geometric descriptors that correspond to pertinent scales could be used for curve and surface abstraction purposes.

Acknowledgment

This work has been supported by the ANR SeARCH project (ANR-09-CORD-019), and the European Community's Seventh Framework Program [FP7 - 2007/2013] under the Grant Agreement 270404 (V-must.net). The model of figure 3 is provided courtesy of S. Rusinkiewicz, and figure 10 courtesy of AIM@SHAPE Shape Repository.

References

- [BEA83] BURT P. J., EDWARD, ADELSON E. H.: The laplacian pyramid as a compact image code. *IEEE Transactions on Communications* 31 (1983), 532–540. 1
- [BSF02] BOYER K. L., SRIKANTIAH R., FLYNN P. J.: Saliency sequential surface organization for free-form object recognition. *Comput. Vis. Image Underst.* 88, 3 (2002), 152–188. 2
- [CGOS11] CHAZAL F., GUIBAS L., OUDOT S., SKRABA P.: Scalar field analysis over point cloud data. *Discrete & Computational Geometry* 46 (2011), 743–775. 2
- [CPG09] CIPRIANO G., PHILLIPS G. N., GLEICHER M.: Multi-scale surface descriptors. *IEEE Transactions on Visualization and Computer Graphics* 15, 6 (2009), 1201–8. 2
- [DUM*11] DUANGGATE C., UYYANONVARA B., MAKHANOV S. S., BARMAN S., WILLIAMSON T.: Object detection with feature stability over scale space. *Journal of Visual Communication and Image Representation* 22, 4 (2011), 345–352. 2
- [ELZ02] EDELSBRUNNER H., LETSCHER D., ZOMORODIAN A.: Topological persistence and simplification. *Discrete & Computational Geometry* 28 (2002), 511–533. 2
- [FMR11] FIORIO C., MERCAT C., RIEUX F.: Adaptive Discrete Laplace Operator. In *Advances in Visual Computing* (2011), vol. 6939/2011, pp. 377–386. 2
- [GG07] GUENNEBAUD G., GROSS M.: Algebraic point set surfaces. *ACM Trans. Graph.* 26, 3 (2007). 3
- [GGG08] GUENNEBAUD G., GERMANN M., GROSS M.: Dynamic sampling and rendering of algebraic point set surfaces. *Computer Graphics Forum* 27, 2 (2008). 3, 10
- [HFG*06] HUANG Q.-X., FLÖRY S., GELFAND N., HOFER M., POTTMANN H.: Reassembling Fractured Objects by Geometric Matching. *ACM Trans. Graph.* 25, 3 (2006), 569–578. 2
- [IT11] ITSKOVICH A., TAL A.: Surface partial matching and application to archaeology. *Computers & Graphics* 35, 2 (2011), 334–341. 2, 7
- [Lev98a] LEVIN D.: The approximation power of moving least-squares. *Math. Comput.* 67, 224 (1998), 1517–1531. 1
- [Lev98b] LEVIN D.: The approximation power of moving least-squares. *Mathematics of Computation* 67 (1998), 1517–1531. 3
- [LG05] LI X., GUSKOV I.: Multi-scale features for approximate alignment of point-based surfaces. In *Proc. of the Eurographics symposium on Geometry processing* (2005). 2
- [Lin94] LINDBERG T.: *Scale-Space Theory in Computer Vision*. Norwell, MA, USA, 1994. 1, 2
- [LZ09] LÉVY B., ZHANG H. R.: Spectral mesh processing. In *ACM SIGGRAPH ASIA 2009 Courses* (2009), pp. 17:1–17:47. 2
- [Mal00] MALLOT H.: *Computational vision: information processing in perception and visual behavior*. 2000. 2
- [Mal08] MALLAT S.: *A Wavelet Tour of Signal Processing, Third Edition: The Sparse Way*, 3rd ed. 2008. 1
- [MH80] MARR D., HILDRETH E.: Theory of edge-detection. In *Proc. of the Royal Society of London* (1980), pp. 187–217. 2
- [MKY01] MOKHTARIAN F., KHALILI N., YUEN P.: Curvature computation on free-form 3-d meshes at multiple scales. *Comput. Vis. Image Underst.* 83 (2001), 118–139. 2
- [MM86] MOKHTARIAN F., MACKWORTH A.: Scale-based description and recognition of planar curves and two-dimensional shapes. *Pattern Analysis and Machine Intelligence, IEEE Transactions on PAMI-8*, 1 (1986), 34–43. 2
- [PKG03] PAULY M., KEISER R., GROSS M. H.: Multi-scale feature extraction on point-sampled surfaces. *Computer Graphics Forum* 22, 3 (2003), 281–290. 2, 6
- [PKG06] PAULY M., KOBELT L. P., GROSS M.: Point-based multiscale surface representation. *ACM Trans. Graph.* 25 (2006), 177–193. 8
- [Pra87] PRATT V.: Direct least-squares fitting of algebraic surfaces. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques* (New York, NY, USA, 1987), pp. 145–152. 3
- [PWHY09] POTTMANN H., WALLNER J., HUANG Q.-X., YANG Y.-L.: Integral invariants for robust geometry processing. *Comput. Aided Geom. Des.* 26 (2009), 37–60. 2
- [Rom09] ROMENY B.: *Front-End Vision and Multi-Scale Image Analysis: Multi-scale Computer Vision Theory and Applications, written in Mathematica*. 2009. 1, 2
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L.: A concise and provably informative multi-scale signature based on heat diffusion. In *Proc. of the Symposium on Geometry Processing* (2009), pp. 1383–1392. 2
- [Wit83] WITKIN A. P.: Scale-space filtering. In *Proc. of the Eighth international joint conference on Artificial intelligence - Volume 2* (1983), pp. 1019–1022. 2
- [WS90] WADA T., SATO M.: Scale-space tree and its hierarchy. In *Pattern Recognition, 1990. Proc., 10th International Conference on* (1990), vol. ii, pp. 103–108 vol.2. 2
- [WSS09] WANG H., SCHEIDEGGER C., SILVA C.: Bandwidth selection and reconstruction quality in point-based surfaces. *IEEE TVCG* 15, 4 (2009), 572–582. 7
- [YLHP06] YANG Y.-L., LAI Y.-K., HU S.-M., POTTMANN H.: Robust principal curvatures on multiple scales. In *Proc. of the Eurographics symposium on Geometry processing* (2006), pp. 223–226. 2
- [ZBVH09] ZAHARESCU A., BOYER E., VARANASI K., HORRAUD R. P.: Surface feature detection and description with applications to mesh matching. In *CVPR* (2009). 2, 6
- [ZH97] ZHANG D., HEBERT M.: Multi-scale classification of 3D objects. In *CVPR* (1997), pp. 864–869. 2
- [ZH99] ZHANG D., HEBERT M.: Harmonic maps and their applications in surface matching. In *CVPR* (1999), pp. 2524–2530. 2

Appendix

Sphere fitting. For ease of implementation we recall the closed-form formula for algebraic sphere fitting [GGG08]:

$$\begin{aligned}
 u_q &= \frac{1}{2} \frac{\sum w_i \mathbf{q}_i^T \mathbf{n}_i - \sum \tilde{w}_i \mathbf{q}_i^T \sum w_i \mathbf{n}_i}{\sum w_i \mathbf{q}_i^T \mathbf{q}_i - \sum \tilde{w}_i \mathbf{q}_i^T \sum w_i \mathbf{q}_i} \\
 \mathbf{u}_\ell &= \sum \tilde{w}_i \mathbf{n}_i - 2u_4 \sum \tilde{w}_i \mathbf{q}_i \\
 u_0 &= -\mathbf{u}_\ell^T \sum \tilde{w}_i \mathbf{q}_i - u_4 \sum \tilde{w}_i \mathbf{q}_i^T \mathbf{q}_i
 \end{aligned} \tag{7}$$

where \tilde{w}_i is the normalized weight of the sample \mathbf{q}_i : $\tilde{w}_i = w_i / \sum_j w_j$. The differentiation of these formulas is straightforward and yield to the differentiation of $w_i(t)$ (Equation 2) over the scale t :

$$\frac{dw_i}{dt}(t) = \frac{4\|\mathbf{p}_i - \mathbf{p}\|^2}{t^3} \left(\frac{\|\mathbf{p}_i - \mathbf{p}\|^2}{t^2} - 1 \right) \tag{8}$$

Reparametrization. After our reparametrization (Equation 4), the scalar field $s_{\hat{\mathbf{a}}}$ can be retrieved by:

$$s_{\hat{\mathbf{a}}}(\mathbf{x}) = s_{\tau, \eta, \kappa}(\mathbf{x}) = \tau + (1 + 2\tau\kappa)^{\frac{1}{2}} \eta \cdot (\mathbf{x} - \mathbf{p}) + \frac{\kappa}{2} (\mathbf{x} - \mathbf{p})^2 \tag{9}$$

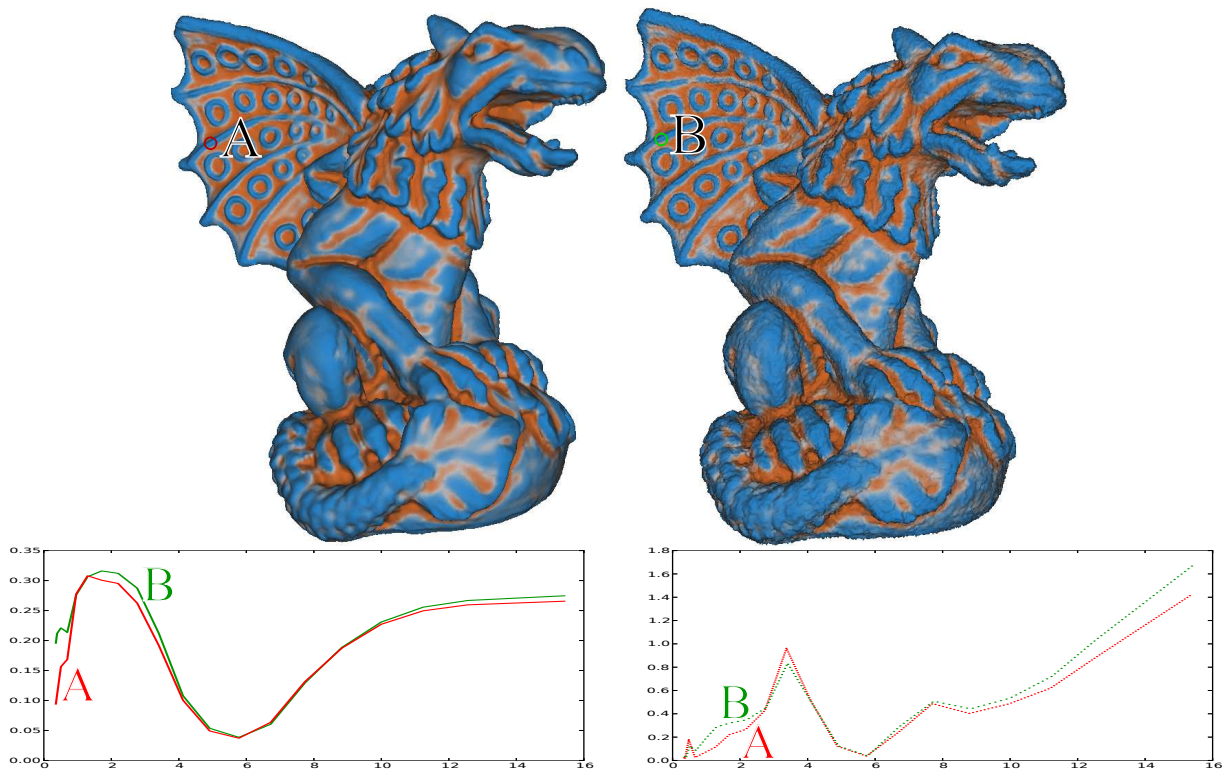


Figure 12: Robustness to noise Top: Gargoyle mesh model with on the left the input version and on the right a corrupted version (uniform noise on the vertices and vertex normals coming from one-ring averages of the resulting face normals). The color code represents curvatures κ at scale $t = 2.2$. Bottom: we show curvature κ (left) and geometric variation ν (right) profiles for the selected point for each version of the model (clean: red-A, noisy: green-B).

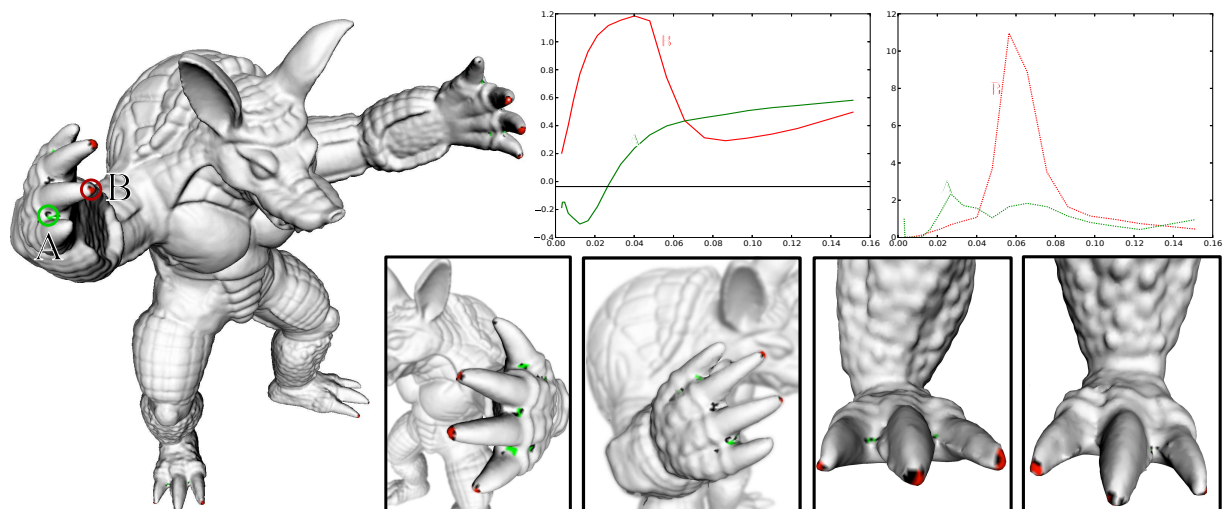


Figure 13: Saddle-like ambiguity Top: given two selected points - a saddle-like region (A-green) and a convexity (B-red) - we show curvature κ (left) and geometric variation ν (right) profiles for each of the two selected points. Bottom: similar points are identified by respective colors. Even though our descriptor is isotropic and cannot explicitly characterize saddle-like shapes, using multi-scale profile comparison permits to disambiguate them from other shapes. Indeed, only saddle-like points are considered similar to point A; whereas not a single saddle-like point is considered similar to point B.