

Scale-Space Representation of Scalar Functions on 2D Manifolds

Andrei Zaharescu[†], Edmond Boyer^{*} and Radu Horaud^{*}

[†]Aimetis Corporation, 500 Weber Street North, Waterloo, ON, Canada

^{*}INRIA Grenoble, 655 Avenue de l'Europe, Montbonnot, 38 334 Saint Ismier Cedex, France

andrei.zaharescu@aimetis.com edmond.boyer@inria.fr radu.horaud@inria.fr

Abstract

In this paper we develop a novel approach for the scale-space representations of scalar functions defined over Riemannian manifolds. One of the main interest in such representations stems from the task of 3D modelling where 2D surfaces, endowed with various physical properties, are recovered from images. Multi-scale analysis allows to structure the information with respect to its intrinsic scale, hence enabling a wide range of low-level computations, similar to what is usually used for representing images. In contrast to the Euclidean image domain, where scale spaces can be easily obtained through convolutions with Gaussian kernels, surfaces require a more general approach that must handle non-Euclidean spaces. Such a generalized scale-space framework is the main contribution of this paper, which builds on the spectral decomposition available with the heat-diffusion framework to derive a computational approach for representing scalar functions on 2D Riemannian manifolds using an intrinsic scale parameter. In addition, we propose a feature detector and a region descriptor, based on these representations, extending the widely used DOG detector and HOG descriptor to manifolds. Experiments on real datasets with various physical properties, i.e., scalar functions, demonstrate the validity and the interest of this approach.

1. Introduction

Recent computer vision methodologies allowed the development of tools for modelling shapes of real objects, e.g., multiple-view stereo. The resulting surfaces are usually associated with numerous object-related features. These features may well be seen as functions defined over the surface domain. They describe either intrinsic properties of the surface, such as various curvature information, or extrinsic properties, such as color, motion information, etc. In the process of analyzing these functions, the multi-scale representation framework appears to be a fundamental tool that




Figure 1. Fine-to-coarse (from up to down) representations of scalar functions defined over 2D Riemannian manifolds.

can build rich low-level descriptors. This is particularly true for the analysis of signals that correspond to physical measurements, as mentioned by Lindeberg [12]. In that case, signals usually carry information on different physical phenomena, at various scales, and the extraction of meaningful information relies on the ability to identify the appropriate scales. Applications of this principle are numerous in image interpretation, including smoothing, feature detection and segmentation, to mention but a few. In order to extend this type of analysis to functions defined on 2D surfaces of 3D shapes, a more general framework is required that generalizes scale-space representation to functions defined over Riemannian manifolds.

Based on the seminal work of Witkin [20], scale-space

representations of images are traditionally built by convolving images with Gaussian kernels at different scales, the scale parameter being associated with the variance of the Gaussian function. This yields a continuous family of convolved images, indexed by the scale, where spectral or Fourier components of the intensity signals are progressively removed, from high to low frequencies, as the scale increases. The extension of this scheme to a function defined over a Riemannian manifold requires the generalization of Gaussian kernels to such non-metric spaces. While such generalizations do not necessarily present closed form expressions, yet they result from the fact, as pointed out by Koenderink [9], that scale-space representations of images can also be seen as the solutions of a heat-diffusion process. This process, governed by the heat-diffusion equation, progressively diffuses image-intensity information over the image domain, starting with the original image-intensity values seen as the initial state. This principle extends to functions defined over Riemannian manifolds, yielding a scale-space family of functions that are solutions of the heat-diffusion equation on non-metric spaces [2]. The time parameter in the heat-diffusion process plays the role of scale in the resulting representations. As time elapses, diffusion blurs the original function, thus progressively eliminating higher frequency components. This is the approach followed in this paper: The solution of the heat-diffusion equation on Riemannian manifolds is used to derive a very general framework for representing any function at multiples scales.

The original contribution of this paper is two-fold. First, a computational approach is proposed to build scale-space representations of any function defined on discrete manifolds, e.g., meshes and graphs. The approach is grounded on the formal framework of heat diffusion on non-metric spaces. It guarantees, in conjunction with the associated Laplace-Beltrami operator, isotropic and homogeneous diffusion of the initial function values over the manifold and hence it yields geometrically consistent scale-space representations. Second, we propose feature detection and description algorithms that illustrate the practical usefulness of our novel scale-space framework. These algorithms are generalizations of the well known difference-of-Gaussian (DOG) image detector and of the histogram of gradient (HOG) region descriptor to any function defined over a manifold. Similar to the image-DOG, the manifold-DOG is able to capture the scale of a feature, hence providing useful information for feature-based descriptions

The paper is organized as follows. Section 2 reviews previous works. Section 3 introduces the heat diffusion framework and the computational approach. Section 4 presents the feature detector and region descriptor. Experiments are shown in Section 5, before concluding in Section 6.

2. Related Work

Fine-to-coarse image descriptions have been widely used in image analysis since the early image pyramid representations [5, 6]. The continuous scale-space representations introduced by Witkin [20] formalized this concept, hence providing the theoretical foundation for multi-scale image analysis. As mentioned earlier, scale-space images can be obtained by convolutions with Gaussian filters, but they are also solutions of the heat equation [9, 8]. This idea was used to build image pyramids [8] as well as to perform anisotropic diffusion in order to preserve edges when filtering images [15, 22]. In particular, in [22], scale-space images are estimated through spectral decompositions of the original images, hence removing the need for iterative filtering operations on the images. In this paper, we extend this line of work to functions defined on Riemannian manifolds, e.g., 2D surfaces of 3D shapes. Since there is no closed-form filter formulation, we build on heat diffusion over non-metric spaces and over the associated spectral decompositions, in order to extract scale-space representations.

Closely related to the material presented and developed in our paper, geometric diffusion methods applied to shapes have recently received growing interest. They allow for spectral representations of shapes using the eigenfunctions of various discretizations of the Laplace-Beltrami operator [11, 17]. These representations exhibit invariance to isometric deformations as well as robustness to noise. This appears to be of particular interest when looking for invariant shape descriptions to estimate shape signatures [16, 18] or shape distances [4]. Laplace-Beltrami eigenfunctions play a fundamental role in heat diffusion over the surface of a shape, since they are also eigenfunctions of the heat-kernel, thus providing an eigenbasis for the heat-diffusion solutions. We use them exactly for that purpose, as an orthogonal basis function over which spectral representations at different scales of any function can then be determined.

While the first category of the aforementioned approaches considers the external intensity function defined over the image domain \mathbb{R}^2 , the second category considers the intrinsic geometric information of a 2D surface. Nevertheless spectral analysis allows to handle both extrinsic and intrinsic data in an elegant way. A recent attempt in that direction is proposed in [10], where geometric and photometric information is merged by considering surfaces as 2D manifolds embedded in a higher dimensional space with both attributes. Although the purpose of [10] is to define shape descriptors, our approach shares similarities that will be discussed in Section 3.3. Another work of interest is that of Luo *et al.* [14], which aims to estimate the gradient of a function defined over a manifold. They propose to perform this estimation in the eigenspace spanned by the Laplacian

eigenfunctions. The interest lies in the robustness to noise in both the manifold and the input function, but at the cost of having many intermediate approximations. We discuss gradient estimation in Section 3.4 and we propose an alternative method.

The recently works of [21] and [7] are probably the closest, in spirit, to our approach. In these works both a scale-space representation and SIFT-like detectors [13] are proposed. The scale-space descriptions of functions defined over 2D discrete surfaces (meshes) are obtained using local Gaussian filters as approximations of the heat kernel. In this paper we take a radically different strategy, grounding the scale-space descriptions on the theoretically well-founded formalism of the heat-diffusion solution on Riemannian manifolds. We derive a computationally efficient formulation, illustrate the method with numerous examples and we thoroughly compare our results with [21].

3. Scale-Space Representation

In this section we explain how to build a scale-space representation of any scalar function defined over a manifold. This representation is based on the heat-kernel that generalizes the Gaussian kernel to manifolds. It applies to any square integrable function¹ defined over a Riemannian manifold, however, in the context of this paper, we limit ourselves to scalar functions defined over 2D manifolds, e.g., surfaces.

3.1. The Heat Kernel

Let \mathcal{M} be a closed (compact and without boundaries) 2D manifold embedded in \mathbb{R}^3 and $f : \mathcal{M} \rightarrow \mathbb{R}$ a square integrable function on \mathcal{M} , colour or curvature for example. The heat operator \mathcal{H}_t is defined as the convolution with a kernel K :

$$\mathcal{H}_t f(x) = \int_{\mathcal{M}} K(x, y; t) f(y) dy, \quad (1)$$

where the heat kernel $K(x, y; t)$ satisfies the partial differential heat equation:

$$\frac{\partial \mathcal{H}_t f}{\partial t} = \Delta_{\mathcal{M}} \mathcal{H}_t f \text{ with } \lim_{t \rightarrow 0} \mathcal{H}_t f = f, \quad (2)$$

and $\Delta_{\mathcal{M}} = \text{div grad}$ denotes the Laplace-Beltrami operator that generalizes the Laplace operator to Riemannian manifolds. Intuitively, $\mathcal{H}_t f$ measures the heat distribution at time t , starting with an initial heat distribution f on \mathcal{M} . The kernel K has no closed form expression in general but is known to be a Gaussian kernel when \mathcal{M}

is a d -dimensional Euclidean space \mathbb{R}^d , e.g. the image plane. Consequently, the heat kernel is often referred as the extension of the Gaussian Kernel to general manifolds.

Although the closed form expression of K is not known, we can use instead its spectral decomposition that can be determined on \mathcal{M} . The operator $\Delta_{\mathcal{M}}$ is self-adjoint negative-semidefinite, assuming that $\{\phi_i\}$'s are its eigenfunctions then:

$$\Delta_{\mathcal{M}} \phi_i = -\lambda_i \phi_i, \quad \lambda_0 = 0 < \dots < \lambda_{\infty}.$$

From the above heat equation (2) we have:

$$\mathcal{H}_t f = e^{t \Delta_{\mathcal{M}}} f,$$

hence:

$$\mathcal{H}_t \phi_i = e^{-t \lambda_i} \phi_i, \forall i.$$

Thus eigenfunctions of the Laplace-Beltrami operator are also eigenfunctions of the heat operator on \mathcal{M} and it can be shown that the heat kernel has the following analytical expression:

$$K(x, y; t) = \sum_{i=0}^{\infty} e^{-t \lambda_i} \phi_i(x) \phi_i(y).$$

Since the $\{\phi_i\}$'s form an orthonormal basis for square integrable functions $\mathcal{L}^2(\mathcal{M})$ on the manifold \mathcal{M} and assuming that $f \in \mathcal{L}^2(\mathcal{M})$:

$$f(x) = \sum_{i=0}^{\infty} \beta_i \phi_i(x), \quad (3)$$

where $\beta_i = \langle f, \phi_i \rangle$ are the spectral -Fourier- coefficients of f . Therefore:

$$\mathcal{H}_t f(x) = \sum_{i=0}^{\infty} \beta_i e^{-t \lambda_i} \phi_i(x), \quad (4)$$

and the scale-space representation $f(\cdot; t)$ of f is then defined as the family of functions:

$$f(x; t) = \mathcal{H}_t f(x) = \sum_{i=0}^{\infty} \beta_i e^{-t \lambda_i} \phi_i(x), \quad (5)$$

that represents convolutions of the function f on \mathcal{M} performed with the heat kernel at scale t . In Euclidean spaces, this is the analog to a scale-space representation based on a Gaussian kernel with variance $\sigma^2 = 2t$.

¹i.e., for which the integral of the square of the absolute value is finite.

3.2. Estimation on meshes

In order to estimate the scale-space representation of f , we need to determine the eigenfunctions and eigenvalues $\{\phi_i\}$ and $\{\lambda_i\}$ of $\Delta_{\mathcal{M}}$ as well as the spectral coefficients $\{\beta_i\}$ of f .

In the case of a discrete representation M of \mathcal{M} , i.e. a mesh or a graph with n vertices (or nodes) v , the $n \times n$ *unnormalized* graph Laplacian matrix L_M approximates the negative of the Laplace-Beltrami operator $\Delta_{\mathcal{M}}$:

$$L_M(i, j) = \begin{cases} \deg(v_i) & \text{if } i = j, \\ -w_{ij} & \text{if } i \neq j \text{ and } v_i \text{ is adjacent to } v_j, \\ 0 & \text{otherwise,} \end{cases}$$

with w_{ij} the edge weights and $\deg(v_i) = \sum_{j \neq i} w_{ij}$. The matrix L_M can be purely combinatorial, i.e. $w_{ij} \in \{0, 1\}$, or involve weights $w_{ij} \geq 0$, e.g. cotangent weights, often used in computer graphics [19]. Weights are discussed below in the practical context of this work. Since L_M is symmetric positive-semidefinite, its eigen or spectral decomposition writes:

$$L_M = \Phi \Lambda \Phi^\top,$$

where $\Lambda = \text{diag}(\lambda_0, \dots, \lambda_{n-1})$ is the diagonal matrix of eigenvalues $\lambda_0 = 0 < \dots < \lambda_{n-1}$ and Φ is the matrix of eigenvectors $\{\phi_0, \dots, \phi_{n-1}\}$ that form an orthonormal basis of \mathbb{R}^n .

The discrete heat operator H_t is then the $n \times n$ matrix such that:

$$\frac{\partial H_t}{\partial t} = -L_M H_t,$$

thus:

$$H_t = e^{-tL_M} = \Phi e^{-t\Lambda} \Phi^\top.$$

Let $F = (f(v_0) \dots f(v_j) \dots f(v_{n-1}))^\top$ be the n vector of discrete values of the function f evaluated at vertices v_j , i.e. $F(j) = f(v_j)$. The scale-space representation of F is then the set of vectors F_t parameterized by t :

$$F_t = H_t F = \Phi e^{-t\Lambda} \Phi^\top F.$$

F can be represented in the orthonormal basis formed by the eigenvectors; Ignoring the null eigenvalue and associated constant eigenvector we obtain:

$$F = \sum_{i=1}^{n-1} \beta_i \phi_i, \quad (6)$$

where the *spectral* coefficients of F are:

$$\beta_i = \langle F, \phi_i \rangle = F^\top \phi_i.$$

Hence:

$$F_t = \sum_{i=1}^{n-1} (F^\top \phi_i) e^{-t\lambda_i} \phi_i. \quad (7)$$

The above expression allows us to build a scale-space representation of any function $f : M \rightarrow \mathbb{R}$. At vertex v_j , the value $f(v_j; t)$ is simply the j th element $F_t(j)$ of vector F_t .

3.3. Computational Strategy

The decomposition (6) of the function F requires in principle the computation of all eigenvectors ϕ_i . Such computation is however prohibitive when n becomes large, as is often the case with meshed surfaces. As a consequence, only the smallest non-zero eigenvalue-eigenvector pairs are usually considered in order to reduce the complexity. In this paragraph, we discuss the resulting approximation for F and therefore its scale-space representation F_t .

Considering only the k smallest non-zero eigenvalue-eigenvector pairs, i.e., $\{\lambda_i, \phi_i\}_{i=1}^k$, $k \ll n$, is equivalent to removing high frequency components in F . The rationale being that high frequency components do not significantly contribute to spectral decompositions. Such an example can be observed in Figure 2. However, depending on the weights w_{ij} , spectral decompositions without high frequency components are more or less accurate. In particular, weights that encode only combinatorial and geometric information about the mesh may lead to eigenvectors that poorly represent the function F in (6), unless many of them are considered.

In order to compensate for such a behavior, weights can be chosen such that they reflect the properties of F as well, e.g.:

$$w_{ij} = -G(|v_i - v_j|, \sigma) - G(|f_i - f_j|, \tau),$$

where G denotes the Gaussian kernel, $|\cdot|$ is a metric, the Euclidean distance in our case, and where standard deviations σ and τ can be learned from the input data. Gaussian kernel weights have been advocated by Belkin *et al.* [1] as providing good properties such as convergence to the continuous Laplace-Beltrami operator. Interestingly, the above weights correspond to a spectral decomposition where the original manifold M is considered in a higher dimensional space \mathbb{R}^{3+l} where l is the dimension of the co-domain of F , e.g. 1 in case of a scalar function. In this space, vertices on the discrete manifold M correspond to points with $3+l$ coordinates: (v_j, f_j) . This is very similar, in principle, to the recent work of Knovnatsky *et al.* [10] on the fusion of geometric and photometric information for shape analysis. Nevertheless, our purpose is different, since we want to build scale-space representations of scalar functions, instead of estimating shape signatures.




Figure 2. Reconstructions of the original scalar function f using k eigenvalue / eigenvector pairs.

Embedding M in a higher dimensional space allows for better spectral decompositions of the function F with fewer eigenvectors than in \mathbb{R}^3 . However, it should be noticed that the eigenvalues and eigenvectors must be recomputed when the values f_j are changing on M . Hence, the computational cost benefit may lower when scale-space representations of several functions, or functions with different values, are to be estimated on the same manifold.

3.4. Gradients

In order to build local descriptors of the function F , e.g. SIFT, spatial gradients on M are often required. Using the scale-space representation $f(\cdot; t)$ instead of f for that purpose allows gradients to be computed at a given scale t , as provided for instance by a feature detection approach, e.g. DOG, hence enabling scale (sampling) invariant descriptions.

Denoting $f(v; t)$ and $\phi_i(v)$ the scalar values of $f(\cdot; t)$ and ϕ_i at vertex v respectively and using expression (7), we have:

$$\nabla_M f(v; t) = \sum_{i=1}^{n-1} \beta_i e^{-t\lambda_i} \nabla_M \phi_i(v), \quad (8)$$

where ∇_M is the gradient operator on M . Thus the gradient of $f(\cdot; t)$ simply relates to the gradients of the eigenvectors ϕ_i at v . In order to estimate the discrete gradient of any scalar function l at any vertex $v \in M$, directional derivatives are usually considered. By definition, the directional derivative on M in the direction $(w - v)$ around v writes:

$$\frac{(w - v)^\top}{\|w - v\|} \nabla_M l(v) = \frac{l(w) - l(v)}{\|w - v\|}.$$

Hence $\nabla_M l(v)$ can be approximated by the vector g that best fits the directional derivatives at v in the least square sense:

$$\nabla_M l(v) = \min_g \sum_{w \in \mathcal{N}(v)} \|g^\top (w - v) - l(w) + l(v)\|^2,$$

where $\mathcal{N}(v)$ is usually the first ring of vertices around v . Note that g can advantageously be constrained to belong to the tangent plane at v when this plane is known. The above least square estimation can be approximated:

1. On the eigenfunctions ϕ_i at v , which then yield $\nabla_M f(v; t)$ through expression (8), or
2. Directly on $f(v; t)$.

While the second approach presents a lower computational cost, the former better accounts for frequency components individually making it more robust to noise in f . In addition the computational cost in that case is balanced by the fact that the estimation is done once for all scales; it can also be reduced by considering less frequency components in the decompositions.

In [14] gradients are estimated in the eigenspace spanned by the ϕ_i s, where a different metric applies, thus requiring a pullback metric to transfer the estimated gradients in the original space. We take a simpler strategy and estimate gradients directly in the original space, yet still keeping the advantages of spectral decompositions, but without the need for transfers between spaces.

4. Features

Using the framework developed in the previous section, we propose a new feature detector and a new region descriptor, that is similar in spirit to [13, 21].

4.1. Feature Detection (Spectral-DoH)

Using the scale-space representation based on the spectral decomposition and on the heat diffusion kernel, we adopt a similar approach to [13] in the image domain and search for the extrema of the function's Laplacian across scales. Scale-space representations of images are traditionally built by convolving images with Gaussian kernels at different scales, where the scale parameter being associated with the variance σ of the Gaussian function. Similarly, the case of the spectral decomposition, the scale is associated with time parameter t of the heat kernel. For Euclidean spaces, the relationship is given by $\sigma^2 = 2t$.

The scale-space representation of the scalar function f is constructed using the spectral decomposition with $k = 100$ eigenvalue/eigenvectors pairs. The initial time diffusion parameter t_0 from the spectral domain corresponds to a mesh distance measure d_0 , related by $t_0 = d_0^2/2$. d_0 is chosen

based on a global surface measure r , corresponding to 1% of the diagonal of the bounding box encompassing the surface. The value is further normalized by the average edge size e_{avg} , in order to account for different mesh samplings. Thus, $d_0 = r/e_{avg}$.

The scale-space representations considers $C = 3$ octaves (i.e. doubling of the t), each octave being divided into $S = 6$ scales, separated by the same constant factor (see [13] for more details). The formula for computing the t at octave c and scale s is given by $t = t_0 2^{c+s/S}$. The main advantage is that, instead of performing iterative convolutions as in [21], where higher octaves take longer to compute, the function representation at any scale can now be computed directly using the spectral decomposition and the appropriate t .

The Difference of Heat (DoH) operator is constructed by subtracting adjacent functions in the scale-space. The feature points are selected as the maxima of the scale-space across the DoH scales, followed by non-maximum-suppression, using a two ring neighbourhood in the current and in the adjacent scales. In addition, the feature point gradient has to be at least $x = 1$ standard deviations away from the gradient mean for that scale.

The top $\gamma = 5\%$ of the maximum number of vertices are being considered, sorted descending by magnitude, and chosen proportionally across scales, according to the number of detections in each scale. Choosing such a schema, versus opting for either a value cut-off or a fixed number of detections, makes the thresholding flexible for different scalar functions on meshes with variable number of vertices.

4.2. Region Descriptor (Spectral-HoG)

The descriptor is very similar in spirit to MeshHoG [21]. The region descriptor for a given feature point is computed using a geodesic support region of radius d_0 . At first, a local coordinate system is chosen, in order to make the descriptor invariant to rotation. Then, a two-level histogram of gradients is computed, both spatially, at a coarse level, in order to maintain a certain high-level spatial ordering, and using orientations, at a finer level. For more details, please see [21]. The major improvements to [21] are the more principled way of computing the gradient and the fact that gradients for each descriptor are computed at the scale at which the feature point was detected.

5. Experiments

The scale-space framework we propose is a generic method to analyze any scalar function defined over a 2D manifold. However, in order to give quantitative evidence

of its benefits, we perform an extensive comparison with already existing discrete local approximate approach [21].

We have created our own dataset, as supposed to use an existent one, such as [3], in order to be able to test the repeatability of the detector and the robustness of the descriptor under *both* photometric and geometric deformations.

5.1. Dataset

The dataset consists of 3 base shapes (also called *null* shapes), with simulated transformations applied to them. The photometric transformations are noise and shot noise. The geometric transformations are noise, shot noise, rotation, scale, local scale, sampling, holes, micro-holes, topology and isometry. Each transformation has 5 levels of noise applies to it. Generally, the level corresponds to the noise amplitude. Therefore, for one base shape, a total of $12 \times 5 = 60$ shapes are obtained.

In the case of photometric transformations (i.e. *color noise* and *color shot noise*), the noise modifies the scalar function (colour), whereas in the case of geometric transformations (i.e. *geometry noise* and *geometry shot noise*), the noise modifies 3-D positions of the mesh vertices. In the case of *topology*, the original mesh was sliced with thin planes. For *isometry*, we used non-rigid semi-elastic transformations obtained either automatically or manually in a modelling software, such as Blender.

In the cases of *noise*, *shot noise*, *rotation*, *holes*, *micro-holes*, *local scale*, *topology*, *isometry* the level corresponds to the amplitude of the noise. For *isometries*, the noise level indicates how far apart in time the 3-D models are. For *topology*, the noise level corresponds to the number of evenly spaced planes that split the mesh. For *scale*, the noise levels correspond to scale factors $\{0.5, 0.83, 1.25, 1.62, 2.0\}$. For *sampling*, the noise level x corresponds to a target mesh with the average edge size is computed as $e_{avg} * (1.0 + x)$, obtained using edge split, edge collapse and edge swap operations.

5.2. Evaluation Methodology

The performance is measured by comparing the features and descriptors obtained for the null shape with the ones obtained for the different transformations.

Feature Detection. The criteria employed for quantifying the quality of feature detection is *repeatability*. Given that the ground-truth (one-to-one correspondence) is known for each transformed shape B of the null shape A , the repeatability is calculated as the percentage of detected feature points in B that are within a geodesic ball of radius $r = 1\%$ of the surface area, from one of the detected interest points in A .




Figure 3. Invariance to manifold sampling: each row shows a mesh with the top 150 interest points and the scale-space representation.

Feature Description. The *robustness* of the feature description is measured as the average normalized L_2 norm between descriptors corresponding to matched feature points.

5.3. Results

Two different scalar fields are considered: photometric information (colour intensity) and geometric information (mean curvature).

The results for Spectral-DoH and Spectral-HoG, the newly proposed method, are shown in Table 1. The results obtained with Mesh DoG and Mesh HoG, one of the state-of-the-art methods proposed by Zaharescu *et al.* [21], are presented in Table 2. As it can be observed, the newly proposed spectral method outperforms [21] for both feature detection and region description, exhibiting a higher resilience to noise, thanks to its more global approach.

6. Conclusion

We have proposed a new scale-space representation and gradient computation method for scalar functions defined over 2D manifolds using a spectral decomposition with the heat diffusion kernel. A great advantage of the proposed method is that, once such a decomposition is computed, all scales responses can be computed with very little extra computational cost. In addition, we have proposed novel feature detection and region description methods based on this representation and compared favourably it with another state-of-the-art implementation.

References

- [1] M. Belkin, J. Sun, and Y. Wang. Constructing laplace operator from point clouds in rd. In *ACM -SIAM Symposium on Discrete Algorithms*, pages 1031–1040, 2009. 4
- [2] P. Bérard, G. Besson, and S. Gallot. Embedding Riemannian manifolds by their heat kernel. *Geometric and Functional Analysis*, 4(4):373–398, 1994. 2
- [3] A. M. Bronstein, M. M. Bronstein, B. Bustos, U. Castellani, M. Crisani, B. Falcidieno, L. J. Guibas, I. Kokkinos, V. Murino, M. Ovsjanikov, G. Patané, I. Sipiran, M. Spagnuolo, and J. Sun. Shrec 2010: robust feature detection and description benchmark. In *Proc. EUROGRAPHICS Workshop on 3D Object Retrieval (3DOR)*, 2010. 6
- [4] M. Bronstein and A. Bronstein. Shape Recognition with Spectral Distances. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2011. To appear. 2
- [5] P. Burt and T. Adelson. The Laplacian Pyramid as a Compact Image Code. *IEEE Trans. Communications*, 9(4):532–540, 1983. 2
- [6] J. Crowley and A. Parker. A Representation for Shape Based on Peaks and Ridges in the Difference of Low Pass Transform. *IEEE Transactions on PAMI*, 6(2):156–170, 1984. 2
- [7] T. Hou and H. Qin. Efficient computation of scale-space features for deformable shape correspondences. In *Proceedings of the European Conference on Computer Vision*, pages 384–397, 2010. 3
- [8] R. Hummel. *The scale-space formulation of pyramid data structures*, pages 107–123. Academic Press, 1987. 2
- [9] J. Koenderink. The structure of images. *Biological Cybernetics*, 50, 1984. 2
- [10] A. Kovnatsky, M. M. Bronstein, A. M. Bronstein, and R. Kimmel. Diffusion framework for geometric and pho-

Strength	Colour Intensity										Mean Curvature									
	Repeatability					Robustness					Repeatability					Robustness				
	1	<2	<3	<4	<5	1	<2	<3	<4	<5	1	<2	<3	<4	<5	1	<2	<3	<4	<5
Noise(F)	0.99	0.99	0.98	0.97	0.95	0.07	0.10	0.13	0.16	0.19	0.99	0.99	0.99	0.99	0.99	0.01	0.01	0.01	0.01	0.01
Shot Noise(F)	0.98	0.94	0.94	0.92	0.90	0.10	0.19	0.25	0.29	0.32	1.00	1.00	1.00	1.00	1.00	0.00	0.00	0.00	0.00	0.00
Noise(G)	1.00	1.00	1.00	1.00	1.00	0.17	0.22	0.25	0.28	0.29	0.94	0.93	0.93	0.94	0.92	0.24	0.28	0.30	0.32	0.34
Shot Noise(G)	1.00	1.00	1.00	1.00	1.00	0.04	0.07	0.11	0.16	0.26	0.99	0.96	0.96	0.95	0.92	0.08	0.14	0.23	0.33	0.45
Rotation	1.00	1.00	1.00	1.00	1.00	0.01	0.02	0.02	0.02	0.02	1.00	1.00	1.00	0.99	1.00	0.01	0.02	0.02	0.02	0.01
Scale	1.00	1.00	1.00	1.00	1.00	0.01	0.01	0.01	0.01	0.00	1.00	1.00	1.00	1.00	1.00	0.01	0.01	0.01	0.01	0.01
Local Scale	1.00	1.00	1.00	1.00	1.00	0.12	0.16	0.18	0.20	0.22	0.97	0.95	0.94	0.92	0.93	0.19	0.24	0.27	0.29	0.30
Sampling	0.98	0.95	0.97	0.92	0.90	0.27	0.33	0.35	0.37	0.36	0.92	0.91	0.92	0.90	0.94	0.29	0.33	0.36	0.37	0.37
Holes	0.98	0.97	0.96	0.96	0.95	0.04	0.07	0.11	0.11	0.12	0.98	0.97	0.97	0.94	0.94	0.05	0.07	0.12	0.11	0.13
Micro-Holes	0.99	0.99	0.99	0.99	0.98	0.04	0.04	0.09	0.08	0.09	0.99	0.98	0.98	0.97	0.97	0.04	0.05	0.09	0.11	0.12
Topology	0.96	0.96	0.92	0.89	0.94	0.16	0.21	0.24	0.29	0.28	0.95	0.96	0.92	0.89	0.95	0.17	0.18	0.21	0.28	0.28
Isometry	0.94	0.96	0.94	0.94	0.95	0.20	0.20	0.21	0.22	0.21	0.94	0.94	0.93	0.93	0.94	0.24	0.24	0.24	0.26	0.24
Average	0.98	0.98	0.97	0.97	0.96	0.10	0.13	0.16	0.18	0.20	0.97	0.97	0.96	0.95	0.96	0.11	0.13	0.15	0.18	0.19

Table 1. Performance evaluation of the Spectral-DoH detector (repeatability) and of the Spectral-HOG descriptor (robustness).

Strength	Colour Intensity										Mean Curvature									
	Repeatability					Robustness					Repeatability					Robustness				
	1	<2	<3	<4	<5	1	<2	<3	<4	<5	1	<2	<3	<4	<5	1	<2	<3	<4	<5
Noise(F)	0.94	0.93	0.93	0.92	0.85	0.05	0.06	0.08	0.11	0.16	0.97	0.97	0.97	0.97	0.97	0.03	0.03	0.03	0.02	0.03
Shot Noise(F)	0.98	0.94	0.88	0.83	0.73	0.12	0.28	0.32	0.36	0.37	1.00	1.00	1.00	1.00	1.00	0.00	0.00	0.00	0.00	0.00
Noise(G)	0.98	0.98	0.98	0.97	0.96	0.21	0.26	0.29	0.32	0.34	0.93	0.90	0.88	0.87	0.85	0.27	0.31	0.34	0.36	0.38
Shot Noise(G)	0.99	0.99	0.97	0.97	0.94	0.07	0.13	0.22	0.31	0.39	0.98	0.96	0.94	0.92	0.91	0.09	0.16	0.25	0.34	0.43
Rotation	1.00	1.00	1.00	1.00	1.00	0.01	0.01	0.01	0.01	0.01	1.00	1.00	1.00	1.00	1.00	0.01	0.01	0.01	0.05	0.01
Scale	1.00	1.00	1.00	1.00	1.00	0.01	0.01	0.01	0.01	0.01	1.00	1.00	1.00	1.00	1.00	0.01	0.01	0.01	0.01	0.01
Local Scale	0.99	0.98	0.97	0.97	0.97	0.16	0.20	0.23	0.25	0.27	0.97	0.95	0.93	0.92	0.91	0.24	0.30	0.34	0.36	0.37
Sampling	0.94	0.93	0.89	0.89	0.92	0.31	0.36	0.36	0.38	0.39	0.88	0.88	0.89	0.83	0.89	0.33	0.34	0.34	0.34	0.34
Holes	0.99	1.00	0.97	0.99	0.95	0.02	0.05	0.10	0.08	0.14	0.99	0.99	0.98	0.97	0.97	0.01	0.04	0.10	0.08	0.10
Micro-Holes	1.00	1.00	0.98	0.99	0.98	0.03	0.03	0.09	0.10	0.11	0.99	0.99	0.98	0.98	0.97	0.02	0.02	0.10	0.11	0.12
Topology	0.97	0.89	0.90	0.79	0.82	0.21	0.29	0.30	0.32	0.35	0.96	0.95	0.85	0.90	0.91	0.17	0.19	0.26	0.25	0.28
Isometry	0.89	0.92	0.91	0.91	0.90	0.22	0.23	0.23	0.25	0.24	0.91	0.92	0.90	0.90	0.90	0.26	0.27	0.25	0.28	0.27
Average	0.97	0.96	0.95	0.93	0.92	0.12	0.16	0.19	0.21	0.23	0.96	0.96	0.94	0.94	0.94	0.12	0.14	0.17	0.18	0.19

Table 2. Performance evaluation of the Mesh-DoG detector (repeatability) and of the Mesh-HoG descriptor (robustness).

- tometric data fusion in non-rigid shape analysis. *CoRR*, abs/1101.4301, 2011. [2](#), [4](#)
- [11] B. Levy. Laplace-Beltrami Eigenfunctions Towards an algorithm that “understands” geometry. In *Shape Modeling and Applications*, 2006. [2](#)
- [12] T. Lindeberg. Scale-space Theory: A basic tool for analysing structures at different scales. *Journal of Applied Statistics*, 21(2):225–270, 1994. [1](#)
- [13] D. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. *International Journal of Computer Vision*, 60(2):91–110, 2004. [3](#), [5](#), [6](#)
- [14] C. Luo, I. Safa, and Y. Wang. Approximating Gradients for Meshes and Point Clouds via Diffusion Metric. In *Eurographics Symposium on Geometry Processing*, 2009. [2](#), [5](#)
- [15] P. Perona and J. Malik. Scale-Space and Edge Detection Using Anisotropic Diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.*, 12:629–639, 1990. [2](#)
- [16] M. Reuter, F. Wolter, and N. Peinecke. Laplace-spectra as fingerprints for shape matching. In *ACM Symposium on Solid and Physical Modeling*, pages 101–106, 2005. [2](#)
- [17] O. Sorkine. Differential representations for mesh processing. *Computer Graphics Forum*, 25(4):789–807, 2006. [2](#)
- [18] J. Sun, M. Ovsjanikov, and L. Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Symposium on Geometric Processing*, 2009. [2](#)
- [19] M. Wardetzky, S. Mathur, F. Kälberer, and E. Grinspun. Discrete Laplace operators: No free lunch. In *Eurographics Symposium on Geometry Processing*, pages 33–37, 2007. [4](#)
- [20] A. Witkin. Scale Space Filtering. In *8th IJCAI*, pages 1019–1022, 1983. [1](#), [2](#)
- [21] A. Zaharescu, E. Boyer, K. Varanasi, and R. Horaud. Surface Feature Detection and Description with Applications to Mesh Matching. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2009. [3](#), [5](#), [6](#), [7](#)
- [22] F. Zhang and E. R. Hancock. Image Scale-Space from the Heat Kernel. In *Progress In Pattern Recognition, Image Analysis And Applications*, pages 181–192, 2005. [2](#)