

Anisotropic Laplace-Beltrami Operators for Shape Analysis

Mathieu Andreux¹, Emanuele Rodolà², Mathieu Aubry³,
and Daniel Cremers²

¹ École polytechnique, Palaiseau, France
andreux@in.tum.de

² Technische Universität München, Munich, Germany

³ Université Paris Est LIGM - École des Ponts Paristech, Champs-sur-Marne, France

Abstract. This paper introduces an *anisotropic* Laplace-Beltrami operator for shape analysis. While keeping useful properties of the standard Laplace-Beltrami operator, it introduces variability in the directions of principal curvature, giving rise to a more intuitive and semantically meaningful diffusion process. Although the benefits of anisotropic diffusion have already been noted in the area of mesh processing (*e.g.* surface regularization), focusing on the Laplacian itself, rather than on the diffusion process it induces, opens the possibility to effectively replace the omnipresent Laplace-Beltrami operator in many shape analysis methods. After providing a mathematical formulation and analysis of this new operator, we derive a practical implementation on discrete meshes. Further, we demonstrate the effectiveness of our new operator when employed in conjunction with different methods for shape segmentation and matching.

Keywords: Shape analysis · Anisotropic diffusion · Curvature · Non-rigid matching · Segmentation · Laplace-Beltrami operator

1 Introduction

Over the last decade, tools derived from harmonic analysis have been successfully used in three-dimensional computer vision for numerous tasks such as shape segmentation, classification, or matching.

The Laplace-Beltrami (LB) operator, or Laplacian, is the natural operator to introduce when studying diffusion processes on shapes. The study of heat diffusion led, for instance, to the definition of the Heat Kernel Signature (HKS) [24], one of the most relevant shape signatures to date. Since the theoretical guarantees of this signature come from the operator itself and its eigen-decomposition, subsequent studies have tried to use the same analytical objects to model different physical phenomena. These include the definition of new signatures such as the Wave Kernel Signature (WKS) [1], learning combinations of LB eigenfunctions to deal with more general classes of deformations [13], or directly employing the eigenfunctions themselves for segmentation tasks [21]. However, all these works rely on the notion of an *isotropic* LB operator.




Fig. 1. Example of stable regions detected via consensus segmentation [21] using the standard Laplacian (first row) and the proposed anisotropic operator (second row). Notice how the local changes in curvature directly affect the detected regions in the second case, giving rise to a more stable and semantically meaningful segmentation. Corresponding regions are colored consistently for visualization purposes.

On the other hand, anisotropic diffusion processes have provided interesting results for shape regularization tasks by taking into account *local* details, but to the best of our knowledge, an anisotropic LB operator by itself has not been used in three-dimensional shape analysis.

In this paper, we propose to study an anisotropic LB operator in the context of shape segmentation and matching. More precisely, we change the diffusion speed along the directions of principal curvature on the surface. Intuitively, this anisotropy can lead to more *semantically consistent* shape segmentations, and can improve the accuracy of point signatures.

1.1 Related Work

3D Shape Analysis: Sometimes called the “swiss army knife” of 3D shape analysis, the Laplacian is used for matching, segmentation and classification. Computing the LB operator directly on meshes with approaches such as [15, 18, 20], numerous works exploit the information contained in its eigen-decomposition to build a new representation of a shape, either by explicit formulas [1, 23, 24] or by learning [13, 22]. Based on these descriptors, several methods have been developed for segmentation purposes, for instance [14, 21]. Recently, Kovnatsky *et al.* [12] modified the metric tensor of the manifold to take into account photometric information: they build another isotropic LB operator upon it and subsequently use the same techniques as presented above.

Anisotropic Denoising and Regularization: All the previously mentioned applications are based on the isotropic LB operator. On the other hand, anisotropic phenomena, such as diffusion, have been analyzed in the fields of image restoring and geometry processing. From the seminal work of Perona *et al.* [16] to the ameliorations of Black *et al.* on 2D images [2], there have been multiple extensions to three-dimensional smoothing and fairing, see *e.g.* [6, 8, 25]. In all

cases, these methods do not compute an anisotropic Laplacian but rather solve the associated anisotropic heat equation.

Medical Image Analysis: Several approaches have also used anisotropic diffusion either on textures or directly on the surface in order to analyze MRI data [3,9]. However, as for denoising and regularization, these methods focus on computing the solutions of the associated anisotropic heat diffusion, without computing the anisotropic LB operator *per se*.

High Dimensional Data Analysis: Spectral clustering based on the LB operator is also a tool used for the analysis of high dimensional data, for instance sets of images. Recently, Kim *et al.* [11] have extended these methods by approximating an anisotropic LB operator in point clouds of high dimension using graph reweighting.

1.2 Contributions

In this work, we propose to transfer the use of anisotropic LB operators to the analysis of 3D shapes. The operator we introduce has a simple mathematical formulation, it can be straightforwardly discretized and provides better results in our experiments than existing alternatives. The key contributions of our work can be summarized as follows:

- We introduce an anisotropic LB (ALB) operator for shape analysis: it preserves the key properties of the isotropic Laplacian, while taking into account extrinsic geometric information;
- We derive a discrete version of this operator for 3D meshes: it can be implemented efficiently and boils down to a simple matrix multiplication;
- We perform an experimental evaluation of the impact of this operator for several standard shape analysis methods.

2 Anisotropic Diffusion on Surfaces

In this section, we introduce the general framework of anisotropic LB operators based on curvature for smooth surfaces, analyze their mathematical properties and derive a discrete version of our selected ALB operator for a 3D mesh.

2.1 Continuous Formulation

We consider a smooth, compact surface $\mathcal{S} \subset \mathbb{R}^3$. After briefly recalling some notions of Riemannian geometry, we introduce an anisotropic LB operator on the surface which takes into account its extrinsic curvature.

Mathematically, we model \mathcal{S} as a Riemannian manifold of dimension 2. This implies that locally around each point $p \in \mathcal{S}$, the surface is very close to the tangent plane at p , which we denote by $T_p\mathcal{S}$. The way in which \mathcal{S} differs by bending from $T_p\mathcal{S}$ is encoded in the second fundamental form, represented by




Fig. 2. Representation of the maximal principal curvature κ_M at each point of two shapes from the TOSCA dataset. Blue parts correspond to regions of low maximal curvature.

at 2×2 matrix. The eigenvalues of this matrix are called *principal curvatures*, noted κ_m and κ_M where $\kappa_m \leq \kappa_M$. The corresponding normalized eigenvectors v_m, v_M are called directions of principal curvature, and form a basis of $T_p\mathcal{S}$. Intuitively, in the direction v_M the surface \mathcal{S} bends the most around p , and conversely for v_m . See Figure 2 for a depiction of the maximal curvature κ_M on two shapes. The quantities κ_m and κ_M are called *extrinsic* since they encode the relationship between \mathcal{S} and its embedding space \mathbb{R}^3 . Moreover, since \mathcal{S} looks locally like a plane, we are able to import most of the calculus known in \mathbb{R}^2 onto \mathcal{S} ; in particular, the gradient (∇) and divergence (div) operators carry the same properties as in the Euclidean case.

Several important results and tools in the field of shape analysis are based on the Laplace-Beltrami operator Δ , defined as:

$$\Delta f = \text{div}(\nabla f). \tag{1}$$

The Laplace-Beltrami operator is an *intrinsic* quantity of the surface, *i.e.* it is invariant under isometric deformations of the manifold. It is also *isotropic*: the corresponding diffusion process does not depend on the direction.

More generally, in this paper we are interested in *anisotropic* LB operators of the form

$$\Delta_D f = \text{div}(D(\nabla f)), \tag{2}$$

where D is a 2×2 matrix acting on tangent vectors. We call D the *anisotropic tensor*, since its deviation from the identity encodes deviation from the isotropic case. In the case of a diffusion process, D controls both the *direction* and the *magnitude* of the diffusion on the surface \mathcal{S} , as shown in Figure 3.

Inspired by successful applications of curvature-aware anisotropic diffusion to denoising of images and surfaces [2, 6, 8, 16, 25], and by recent advances in the analysis of high-dimensional data [11], we attempt to incorporate *extrinsic* information into the anisotropic tensor D for shape analysis purposes. We do so by introducing the following generic linear operator, defined in the orthonormal basis (v_m, v_M) :

$$D_\alpha = \begin{pmatrix} \Psi_\alpha^m(\kappa_m, \kappa_M) & 0 \\ 0 & \Psi_\alpha^M(\kappa_M, \kappa_m) \end{pmatrix}. \tag{3}$$




Fig. 3. Solutions of the heat equation $\Delta u = \frac{\partial}{\partial t} u$ computed at time $t = 3 \cdot 10^{-3}$, with a highly peaked Gaussian on the back of the horse as initial condition. We show the isotropic case (middle), anisotropic case favoring high curvature (left), and deflecting high curvature (right).

By choosing appropriate functions Ψ_α^m and Ψ_α^M , it is possible to favor directions of high curvature or directions of low curvature. In our experiments, we set $\Psi_\alpha^m(\kappa_m, \kappa_M) = \psi_\alpha(\kappa_M)$ and $\Psi_\alpha^M(\kappa_m, \kappa_M) = \psi_\alpha(\kappa_m)$, where ψ_α is defined as:

$$\psi_\alpha(x) = \frac{1}{1 + \alpha|x|}. \tag{4}$$

In the equation above, $\alpha > 0$ directly controls deviation from isotropy, whereas $\alpha \rightarrow 0$ leads to the common LB operator.

In general, operator D_α modifies both the direction and the norm of a given input vector. This differentiates our approach from the anisotropic LB operator $\tilde{\Delta}_{\tilde{D}}$ introduced in [11], where the anisotropic factor \tilde{D} only modifies the norm of a given input vector, but not its direction (see Eq. (5) below).

2.2 Mathematical Analysis

In this section we state the general properties of the ALB operator Δ_{D_α} . Essentially, Δ_{D_α} conserves the strong analytical properties of the conventional Laplace-Beltrami operator Δ . This is in contrast with the previous proposal of [11]; in the following, we will also provide a brief comparison of the two variants.

Our operator has all the key properties of the isotropic case: $-\Delta_{D_\alpha}$ is linear, symmetric, positive semi-definite and therefore has a discrete spectrum $0 = E_0 \leq E_1 \leq \dots \leq E_k \leq \dots$, with associated eigenfunctions ϕ_0, ϕ_1, \dots , and $E_k \rightarrow +\infty$. These properties simply stem from the divergence theorem and the symmetry of D_α (3). However, our operator is *not isometry-invariant*. Indeed, although the product of the principal curvatures, or Gaussian curvature, is preserved under such deformations, it is not the case for each of the principal curvatures. Let us imagine for instance a surface initially planar, and then folded across a line in an isometric fashion: a non-zero principal curvature will appear along the folding line, although the Gaussian curvature will remain equal to zero there. Since our operator directly depends on these extrinsic quantities, it is not in general



Fig. 4. The tenth eigenfunction of Δ (shapes on the left) and Δ_{D_α} (on the right, here $\alpha = 10$). We observe that while the LB operator is invariant under near-isometries, this property is lost for the ALB operator: the eigenfunction is preserved more accurately in the first case.

invariant under isometries. This phenomenon is reflected on the eigenfunctions of Δ_{D_α} , as showed in Figure 4.

It should be noted that the operator $\tilde{\Delta}_{\tilde{D}}$ as defined in [11] is also not invariant to isometries, furthermore it loses other key properties of Δ . In particular, the anisotropic factor \tilde{D} has the form:

$$\tilde{D}(v) = \left\| Q \left(\frac{v}{\|v\|} \right) \right\| v, \tag{5}$$

where $Q : T_p\mathcal{S} \rightarrow T_p\mathcal{S}$ is a linear application acting on tangent vectors, also depending on the second fundamental form in a similar manner as D_α .

We remark that if D_α itself is plugged into (5) as the matrix Q , it follows that D_α and \tilde{D} roughly have the same qualitative properties, apart from the fact that \tilde{D} does not modify the direction of input vectors. Nevertheless, \tilde{D} is not linear, hence the loss of linearity of $\tilde{\Delta}_{\tilde{D}}$. Moreover, in this case Green’s theorem leads to the non-symmetric expression:

$$\langle -\tilde{\Delta}_{\tilde{D}}f, g \rangle = \int_S \left\| Q \left(\frac{\nabla f}{\|\nabla f\|} \right) \right\| \nabla f \cdot \nabla g \neq \langle f, -\tilde{\Delta}_{\tilde{D}}g \rangle. \tag{6}$$

The difference between the mathematical properties of Δ_{D_α} and $\tilde{\Delta}_{\tilde{D}}$ also becomes evident in our evaluation: in most of our experiments, Δ_{D_α} performs better than $\tilde{\Delta}_{\tilde{D}}$, although the effect of curvature on the diffusivity factors is the same.

2.3 Numerical Implementation

In this part, we use a finite-element method to derive a numerical scheme for Δ_D , where in all generality D is a symmetric matrix, allowing us to compute this operator on triangulated meshes. While more sophisticated approaches using

discrete exterior calculus may be employed for this step [10], finite elements are a common tool to discretize differential operators, and have been applied with success to anisotropic operators [6]. In the following, we first recall the basics of finite elements, then provide the formula for our discretized operator in this framework, and finally prove it.

Throughout this discretization, we consider a triangulated mesh consisting of nodes $(z_i)_{1 \leq i \leq N}$, seen as a sample of a surface \mathcal{S} . Any scalar function f on \mathcal{S} can be represented by a vector $\underline{f} = (f_i)$, where $f_i = f(z_i)$. Our goal is to express $\underline{\Delta_D f}$ in terms of geometric values related to the mesh and the vector f . In other words, by linearity, we need to find a matrix L_D acting on \underline{f} and representing the anisotropic operator Δ_D by the relation $\underline{\Delta_D f} = L_D \underline{f}$.

For this purpose, we use finite elements of the first order : the approach consists in approximating f by a piecewise affine version of it taking the same values at the vertices z_i . Formally, we define "hat" functions Φ_i satisfying: $\Phi_i(z_j) = \delta_{ij}$, where Φ_i is affine on each triangle of the mesh, and write $f \approx \sum_i f_i \Phi_i$. Further, we assume that D is constant on each triangle. The weak formulation of the operator Δ_D boils down by linearity to $C \underline{f} = A (\underline{\Delta_D f})$ where the symmetric *mass matrix* A and *stiffness matrix* C are defined as:

$$A_{ij} = \langle \Phi_i, \Phi_j \rangle, C_{ij} = -\langle D(\nabla \Phi_i), \nabla \Phi_j \rangle. \quad (7)$$

Inverting A , our discretized operator L_D finally reads:

$$L_D = A^{-1}C. \quad (8)$$

We now state the following results:

$$A_{ij} = \mathcal{A}_i^{\text{Vor}} \delta_{ij} \quad (9)$$

$$C_{ij} = \begin{cases} \frac{1}{2} \left(\Gamma_{ij} \frac{\cos(\gamma_{ij})}{\sin(\alpha_{ij})} + \Xi_{ij} \frac{\cos(\xi_{ij})}{\sin(\beta_{ij})} \right) & i \neq j \\ -\sum_{T \ni z_j} I(T)_j & i = j \end{cases}, \quad (10)$$

where $\mathcal{A}_i^{\text{Vor}}$ denotes the Voronoi area around z_i . Symbols Γ_{ij} and γ_{ij} are defined as in Figure 5, and Ξ_{ij} , ξ_{ij} are their counterparts in the triangle (z_i, z_j, z_p) . For all triangles $T = (z_k, z_j, z_i)$, we write $I(T)_j = \frac{1}{2} \frac{D^\perp e_j}{\|e_j\|} \cdot \frac{e_j}{\|e_j\|} (\cot(\alpha_{ij}) + \cot(\theta_{jk}))$ with D^\perp defined in (12) and α_{ij} , θ_{ij} as in Figure 5.

Note that whenever $D \equiv Id$, this scheme boils down to the popular cotangent scheme [8, 15, 18]. Moreover, the matrices involved are sparse, which makes their computation simple and efficient.

In order to express numerically the anisotropic tensor D , we computed at each vertex of the mesh the principal curvatures and corresponding directions using the method described in [7] as implemented in [17], and then averaged over each triangle to obtain a constant operator.

In the remainder of this section, we derive the results stated in (9), (10). We fix j and compute C_{ij} , A_{ij} for all i . Clearly, we have C_{ij} , $A_{ij} = 0$ if z_i, z_j are not neighbors. We now assume that this is the case, with $i \neq j$, and decompose the integral over the surface on the different triangles containing the edge (z_i, z_j) .




Fig. 5. Left: Adjacent edges and angles. Right: Voronoi area around z_i , *i.e.* the area of the set of points such that their closest point among the vertices of the mesh is z_i . $\Gamma_{ij} = \frac{\|D^\perp e_i\|}{\|e_i\|}$ and $\gamma_{ij} = -\arccos\left(\frac{D^\perp e_i \cdot e_j}{\|D^\perp e_i\| \cdot \|e_j\|}\right)$.

Let us be given a triangle $T = (z_k, z_j, z_i)$, and let $e_i = \overrightarrow{z_k z_j}$, $e_j = \overrightarrow{z_i z_k}$, $e_k = \overrightarrow{z_j z_i}$. We refer to the heights of the triangle to the points z_p as h_p . Finally, let $\alpha_{ij} = \widehat{z_j z_k z_i}$. It is straightforward that $\nabla\Phi_i = \frac{1}{h_i \|e_i\|} R_{\frac{\pi}{2}} e_i$ and $\nabla\Phi_j = \frac{1}{h_j \|e_j\|} R_{\frac{\pi}{2}} e_j$ where $R_{\frac{\pi}{2}}$ denotes the rotation of angle $\frac{\pi}{2}$. Therefore:

$$D\nabla\Phi_i \cdot \nabla\Phi_j = \frac{1}{h_i \|e_i\|} DR_{\frac{\pi}{2}} e_i \cdot \frac{1}{h_j \|e_j\|} R_{\frac{\pi}{2}} e_j = \frac{1}{4A_T} D^\perp e_i \cdot e_j, \quad (11)$$

where

$$D^\perp = R_{\frac{\pi}{2}}^T D R_{\frac{\pi}{2}}. \quad (12)$$

Furthermore, if $i \neq j$, we also know that the area A_T of the triangle T satisfies $A_T = \frac{1}{2} \|e_i\| \|e_j\| \sin(\alpha_{ij})$. Consequently, when integrating the constant value $D\nabla\Phi_i \cdot \nabla\Phi_j$ over the whole triangle T , we get:

$$\int_T D\nabla\Phi_i \cdot \nabla\Phi_j dS = \frac{1}{4A_T} D^\perp e_i \cdot e_j = -\frac{1}{2} \Gamma_{ij} \frac{\cos(\gamma_{ij})}{\sin(\alpha_{ij})}, \quad (13)$$

where γ_{ij} is the angle between the $D^\perp e_i$ and $-e_j$ and $\Gamma_{ij} = \frac{\|D^\perp e_i\|}{\|e_i\|}$. Summing up over the two triangles sharing the edge (z_i, z_j) , we obtain

$$C_{ij} = \langle D\nabla\Phi_j, \nabla\Phi_i \rangle = -\frac{1}{2} \left(\Gamma_{ij} \frac{\cos(\gamma_{ij})}{\sin(\alpha_{ij})} + \Xi_{ij} \frac{\cos(\xi_{ij})}{\sin(\beta_{ij})} \right), \quad (14)$$

where Ξ_{ij} and ξ_{ij} are the counterparts of Γ_{ij} and γ_{ij} on the other triangle adjacent to the edge (z_i, z_j) .

For the diagonal coefficient C_{jj} , when considering any triangle $T = (z_k, z_j, z_i)$ containing z_j , the same reasoning as above leads to:

$$I(T)_j := \int_T D\nabla\Phi_j \cdot \nabla\Phi_j dS = \frac{1}{4A_T} D^\perp e_j \cdot e_j = \frac{1}{2} \frac{D^\perp e_j}{\|e_j\|} \cdot \frac{e_j}{\|e_j\|} \frac{\|e_j\|}{h_j}. \quad (15)$$



Fig. 6. Maximally stable components detected on a horse shape equipped with our linear anisotropic LB operator Δ_{D^α} . From left to right, we use $\alpha = 0$ (corresponding to the standard LB operator), $\alpha = 5$ and $\alpha = 10$.

From the relation $\frac{\|e_i\|}{h_j} = \cot(\alpha_{ij}) + \cot(\theta_{jk})$, where α_{ij} and θ_{jk} are the two base angles corresponding to z_j (see Figure 5), we get:

$$I(T)_j = \frac{1}{2} \frac{D^\perp e_j}{\|e_j\|} \cdot \frac{e_j}{\|e_j\|} (\cot(\alpha_{ij}) + \cot(\theta_{jk})) . \quad (16)$$

When $\alpha = 0$, we simply have $D^\perp = Id$, and once again we fall back to the cotangent scheme.

Finally, $\langle D\nabla\Phi_j, \nabla\Phi_j \rangle$ is nothing but the sum of (16) over all triangles T involved in z_j . Combining (7),(14) and (16), the result (10) comes easily. Regarding the expression for the mass matrix A (9), we refer to [15,20]. In practice, we used the method described in [8] to compute the Voronoi areas.

3 Experiments

In this section we investigate the practical benefits of adopting the proposed operator in two problems commonly arising in shape analysis, namely stable region detection and matching of deformable shapes, using the datasets [4,5].

3.1 Segmentation and Region Detection

Recent state-of-the-art approaches in deformable shape matching rely on the ability to detect repeatable regions on the given shapes [19]. In these experiments, we employ the anisotropic Laplacian within two different frameworks for detecting region-based features.

The first approach follows the consensus segmentation technique of [21]. The region detection process operates as follows. First, given a shape \mathcal{S} , several segmentations are produced by a clustering method (k -means) over an intrinsic embedding of the shape points. Different segmentations are obtained by




Fig. 7. Left: Hit ratio for isotropic (red curve) and anisotropic HKS (α varying from blue to green) on nearly-isometrically deformed shapes (*michael* class from the TOSCA dataset). Middle: Repeatability curves of MSER detection [14] using our anisotropic operator, for all deformations over the SHREC’10 dataset. Right: Repeatability averaged over all deformations of the SHREC’10 dataset, for our anisotropic operator (green), Kim *et al.*’s (red), and the standard Laplace-Beltrami operator (blue).

initializing the clustering process with different randomized seeds. The intrinsic embedding is provided by the simple mapping

$$p \mapsto \left(\frac{\phi_1(p)}{\sqrt{E_1}}, \frac{\phi_2(p)}{\sqrt{E_2}}, \frac{\phi_3(p)}{\sqrt{E_3}}, \dots \right) \tag{17}$$

for each $p \in \mathcal{S}$, where ϕ_k are the eigenfunctions and E_k the corresponding eigenvalues of the Laplace-Beltrami operator on \mathcal{S} . The mapping above is commonly referred to as the Global Point Signature embedding of \mathcal{S} [23]. Given this initial collection of segmentations, the corresponding *Fréchet mean* (according to an appropriate notion of distance among segmentations) is then solved for in a robust manner (we refer to [21] for the technical details). Intuitively, this process attempts to achieve an “agreement” between the several segmentations constituting the initial putative set; being based entirely on quantities directly derived from the eigenfunctions of Δ (17), regions produced by the consensus approach tend to reflect the overall trend of the eigenfunctions themselves. As such, apart from yielding a stable segmentation of the shape, this approach can easily provide a visual clue on the general behavior of a given operator.

In Figure 1 we show the regions detected on two nearly-isometric shapes via consensus segmentation, when the two manifolds are equipped with the standard LB operator Δ and our proposed anisotropic variant Δ_{D_α} . Notice how, although the regions resulting from the consensus process are repeatable across the two operators, the outcome of the anisotropic case is more stable (two different poses produce the same regions) and more semantically meaningful (no cross-regions covering semantically different parts).

The second experiment is aimed at providing a quantitative comparison between the segmentations produced when employing the two operators Δ and Δ_{D_α} . For these comparisons we adopted a different region detection technique, namely the Maximally Stable Extremal Regions (MSER) approach of [14]. Note that, differently from the consensus approach, the MSER technique is based on the *diffusion process* induced by the different operators rather than their

simple eigen-decomposition. Thus, this experiment is aimed at evaluating our operator from a diffusion-geometric perspective, when put in comparison to the standard Laplacian within a common segmentation pipeline.

The comparisons were performed on the SHREC'10 benchmark. The dataset consists of three different shape classes (*dog, horse, man*) undergoing 9 different types of deformation (including *e.g.*, topological changes, downsampling, and local changes in scale), each at 5 intensity levels. To measure the stability of a given segmentation, we evaluated its *repeatability* curve across the whole dataset [14]. In particular, let \mathcal{M}_0 be a "null" (*i.e.*, in a canonical pose) and \mathcal{M} be a deformed shape respectively. Given the ground-truth correspondence $g : \mathcal{M}_0 \rightarrow \mathcal{M}$, we can compute the *overlap* between two regions $\mathcal{A}_0 \subset \mathcal{M}_0$ and $\mathcal{B} \subset \mathcal{M}$ as the area ratio:

$$O(\mathcal{B}, \mathcal{A}_0) = \frac{|g^{-1}(\mathcal{B}) \cap \mathcal{A}_0|}{|g^{-1}(\mathcal{B}) \cup \mathcal{A}_0|}, \quad (18)$$

where $|\cdot|$ denotes the surface area. Note that

$$0 \leq O(\mathcal{B}, \mathcal{A}_0) \leq 1, \quad (19)$$

the last inequality being attained if and only if $g^{-1}(\mathcal{B}) = \mathcal{A}_0$. For a fixed overlap value ν , the *repeatability* at ν is defined as the percentage of regions in \mathcal{M} having an overlap greater than ν with regions in \mathcal{M}_0 . For any given segmentation, the repeatability at overlap 0 is 100%, and the more stable the segmentation is, the higher the repeatability remains as ν increases.

In Figure 7 (last two columns) we show the repeatability curves obtained by MSER detection using our linear anisotropic operator Δ_D (green curve), the non-linear variant $\hat{\Delta}_{\bar{D}}$ (red curve), and the standard LB operator Δ_α (blue curve). Both anisotropic operators (here computed with $\alpha = 10$) outperform the standard Laplacian by a large margin, while there is only a minor difference in performance between them. This result directly confirms the observation (Section 2.2) that the two operators carry similar qualitative properties, with our linear proposal demonstrating overall better behaviour on the dataset considered. Finally, in Figure 6 we show some examples of MSER segmentations produced with our operator for different values of α .

3.2 Shape Matching

As observed in Section 2.2, the proposed anisotropic Laplacian is not an isometry invariant; hence, its direct application in the computation of intrinsic descriptors [1, 13] may not lead to an increase of performance in typical non-rigid matching scenarios.

Quantitative comparisons among the matching results obtained with a standard HKS implementation [24] and its anisotropic variant are shown in Figure 7

(left). In this experiment, we first performed a (Euclidean) farthest point sampling of a shape in the standard "null" pose; we computed a descriptor for each of the sampled points, and then looked for its closest matches (in descriptor space) on several nearly-isometric deformations of the shape. We did this for the standard HKS descriptor (red curve), and for the anisotropic HKS with α ranging uniformly between 1 (blue) and 20 (green). Each curve depicts the percentage of points in the null shape having their *exact* correspondence in the first 0 – 1% of all shape points, sorted according to descriptor similarity.

As it can be seen from the plot, the advantage brought by the direct adoption of anisotropic diffusion for shape matching is only minimal. This is to be expected, since as shown in Figure 3, solutions to the heat equation will tend to evolve differently along regions of different curvature. However, as seen in the previous section, when used in conjunction with appropriate stability criteria, this property can become very useful for tasks of segmentation of deformable shapes.

4 Conclusions

In this paper we introduced the use of an *anisotropic* Laplace-Beltrami operator for shape analysis, and we derived a numerical scheme that allows to compute it easily for a triangulated mesh. In extensive evaluations, we showed that the proposed operator can improve either quantitatively or qualitatively several shape analysis methods initially developed with the standard Laplace-Beltrami operator in mind. These include the definition of shape signatures and the detection of stable regions following different approaches. Based on these promising results, we believe that the adoption of anisotropic Laplace-Beltrami operators for three-dimensional shape analysis constitutes a valid future direction of research.

Acknowledgments. This work is supported by an Alexander von Humboldt Fellowship and ANR Project SEMAPOLIS (ANR-13-CORD-0003).

References

1. Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In: ICCV Workshops, pp. 1626–1633 (2011)
2. Black, M.J., Sapiro, G., Marimont, D.H., Heeger, D.: Robust anisotropic diffusion. *Trans. Img. Proc.* **7**(3), 421–432 (1998)
3. Boucher, M., Evans, A., Siddiqi, K.: Anisotropic Diffusion of Tensor Fields for Fold Shape Analysis on Surfaces. In: Székely, G., Hahn, H.K. (eds.) IPMI 2011. LNCS, vol. 6801, pp. 271–282. Springer, Heidelberg (2011)
4. Bronstein, A.M., Bronstein, M.M., Castellani, U., Dubrovina, A., Guibas, L.J., Horaud, R.P., Kimmel, R., Knossow, D., von Lavante, E., Mateus, D., Ovsjanikov, M., Sharma, A.: Shrec 2010: robust correspondence benchmark. In: Proc. EUROGRAPHICS Workshop on 3D Object Retrieval, EG 3DOR 2010 (2010)

5. Bronstein, A., Bronstein, M., Kimmel, R.: Numerical Geometry of Non-Rigid Shapes, 1st edn. Springer Publishing Company, Incorporated (2008)
6. Clarenz, U., Diewald, U., Rumpf, M.: Anisotropic geometric diffusion in surface processing. In: Proc. of the Conference on Visualization 2000, VIS 2000, pp. 397–405 (2000)
7. Cohen-Steiner, D., Morvan, J.M.: Restricted delaunay triangulations and normal cycle. In: Proc. of the Nineteenth Annual Symposium on Computational Geometry, SCG 2003, pp. 312–321 (2003)
8. Desbrun, M., Meyer, M., Schröder, P., Barr, A.H.: Implicit fairing of irregular meshes using diffusion and curvature flow. In: Proc. of the 26th Annual Conference on Computer Graphics and Interactive Techniques. pp. 317–324. SIGGRAPH '99 (1999)
9. Fillard, P., Arsigny, V., Ayache, N., Pennec, X.: A Riemannian Framework for the Processing of Tensor-Valued Images. In: Fogh Olsen, O., Florack, L.M.J., Kuijper, A. (eds.) DSSCV 2005. LNCS, vol. 3753, pp. 112–123. Springer, Heidelberg (2005)
10. de Goes, F., Liu, B., Budninskiy, M., Tong, Y., Desbrun, M.: Discrete 2-tensor fields on triangulations. *Computer Graphics Forum* 33(5) (2014)
11. Kim, K., Tompkin, J., Theobalt, C.: Curvature-aware regularization on Riemannian submanifolds. In: Proc. of the IEEE International Conference on Computer Vision, ICCV 2013, pp. 881–888 (2013)
12. Kovnatsky, A., Raviv, D., Bronstein, M.M., Bronstein, A.M., Kimmel, R.: Geometric and photometric data fusion in non-rigid shape analysis. *Numerical Mathematics: Theory, Methods and Applications (NM-TMA)* 6(1), 199–222 (2013)
13. Litman, R., Bronstein, A.: Learning spectral descriptors for deformable shape correspondence. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 36(1), 171–180 (2014)
14. Litman, R., Bronstein, A.M., Bronstein, M.M.: Diffusion-geometric maximally stable component detection in deformable shapes. *Computers & Graphics* 35(3), 549–560 (2011)
15. Meyer, M., Desbrun, M., Schröder, P., Barr, A.H.: Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization and Mathematics III*, pp. 35–57 (2003)
16. Perona, P., Malik, J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Trans. Pattern Anal. Mach. Intell.* 12(7), 629–639 (1990)
17. Peyré, G.: Toolbox graph - a toolbox to process graph and triangulated meshes. (2008) <https://www.ceremade.dauphine.fr/~peyre/matlab/graph/content.html>
18. Pinkall, U., Polthier, K.: Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2(1), 15–36 (1993)
19. Pokrass, J., Bronstein, A.M., Bronstein, M.M., Sprechmann, P., Sapiro, G.: Sparse modeling of intrinsic correspondences. *Computer Graphics Forum* 32(2pt. 4), 459–468 (2013)
20. Reuter, M., Biasotti, S., Giorgi, D., Patanè, G., Spagnuolo, M.: Discrete laplace-beltrami operators for shape analysis and segmentation. *Computers & Graphics* 33(3), 381–390 (2009)
21. Rodolà, E., Rota Bulò, S., Cremers, D.: Robust region detection via consensus segmentation of deformable shapes. *Computer Graphics Forum* 33(5) (2014)

22. Rodolà, E., Rota Bulò, S., Windheuser, T., Vestner, M., Cremers, D.: Dense non-rigid shape correspondence using random forests. In: IEEE Conference on Computer Vision and Pattern Recognition (CVPR) (2014)
23. Rustamov, R.M.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In: Proc. of the Fifth Eurographics Symposium on Geometry Processing, SGP 2007, pp. 225–233 (2007)
24. 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, SGP 2009, pp. 1383–1392 (2009)
25. Tasdizen, T., Whitaker, R., Burchard, P., Osher, S.: Geometric surface smoothing via anisotropic diffusion of normals. In: Proc. of the Conference on Visualization 2002, VIS 2002, pp. 125–132 (2002)