

Disclaimer: This is the peer reviewed version of the following article: D. Boscaini et al., Anisotropic Diffusion Descriptors, Computer Graphics Forum, Volume 35, Issue 2, Pages 431–441, which has been published in final form at <https://doi.org/10.1111/cgf.12844>. This article may be used for non-commercial purposes in accordance with Wiley Terms and Conditions for Use of Self-Archived Versions. This article may not be enhanced, enriched or otherwise transformed into a derivative work, without express permission from Wiley or by statutory rights under applicable legislation. Copyright notices must not be removed, obscured or modified. The article must be linked to Wiley’s version of record on Wiley Online Library and any embedding, framing or otherwise making available the article or pages thereof by third parties from platforms, services and websites other than Wiley Online Library must be prohibited.

Title: Anisotropic Diffusion Descriptors

Date of Publication: 27 May 2016

Authors: D. Boscaini, J. Masci, E. Rodolà, M. M. Bronstein, D. Cremers

Volume: 35, **Issue:** 2, **Pages:** 431–441

DOI: <https://doi.org/10.1111/cgf.12844>

URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.12844>

Anisotropic Diffusion Descriptors

D. Boscaini[†], J. Masci[†], E. Rodolà[‡], M. M. Bronstein[†] and D. Cremers[‡]

[†] University of Lugano (USI), Switzerland

[‡] Technical University of Munich (TUM), Germany

Abstract

Spectral methods have recently gained popularity in many domains of computer graphics and geometry processing, especially shape processing, computation of shape descriptors, distances, and correspondence. Spectral geometric structures are intrinsic and thus invariant to isometric deformations, are efficiently computed, and can be constructed on shapes in different representations. A notable drawback of these constructions, however, is that they are isotropic, i.e., insensitive to direction. In this paper, we show how to construct direction-sensitive spectral feature descriptors using anisotropic diffusion on meshes and point clouds. The core of our construction are directed local kernels acting similarly to steerable filters, which are learned in a task-specific manner. Remarkably, while being intrinsic, our descriptors allow to disambiguate reflection symmetries. We show the application of anisotropic descriptors for problems of shape correspondence on meshes and point clouds, achieving results significantly better than state-of-the-art methods.

Categories and Subject Descriptors (according to ACM CCS): I.3 [Computer Graphics]: Shape modeling—Shape analysis

1. Introduction

In the past decade, spectral methods have gained popularity in many domains of computer graphics and geometry processing, especially in problems of shape analysis [Lév06, R*13] and synthesis [RCG08, BEKB15], symmetry detection [OSG08], shape retrieval [BBG011], correspondence [O*12], and segmentation [SOCG10]. The main idea of spectral methods is to use the eigen-decomposition of the Laplace-Beltrami operator to represent local and global geometric structures. Examples of such structures include diffusion distances [CL06, LRF10], global point signatures [Rus07], heat- [SOG09, GBAL09] and wave-kernel [ASC11] signatures, heat kernel maps [OMMG10], stable regions [LBB11, RRBC14], functional maps [O*12, R*16], and shape difference operators [R*13].

The popularity of spectral geometric structures is owed to several nice properties. First, they are intrinsic by construction and thus invariant to isometric deformations. Second, they are efficiently computable. Third, they can be constructed on shapes in different representations such as meshes or point clouds, provided one has the right discretization of the Laplace-Beltrami operator. A notable drawback of these constructions, however, is that they are *isotropic*, i.e., insensitive to direction. Furthermore, they are ambiguous under intrinsic symmetries [OSG08, Y*14].

In this paper, we show how to construct direction-sensitive spectral feature descriptors using *anisotropic diffusion* on meshes and point clouds. The core of our construction, referred to in the following as *anisotropic diffusion descriptors* (ADD) are directed local

kernels acting similarly to oriented filter banks, which are learned in a task-specific manner. Remarkably, our descriptors allow to disambiguate reflection symmetries.

Related work. Works on adaptive diffusion in image analysis date back to the classical paper of Perona and Malik [PM90], which introduced inhomogeneous diffusion (erroneously called ‘anisotropic’ by the authors) for edge-preserving image filtering. Follow-up works [KS02] used true anisotropic diffusion, driven e.g., by texture gradients. In the computer graphics and geometry processing community, inhomogeneous diffusion driven by texture discontinuities was used to fuse geometric and photometric information on 3D shapes [KBBK12]. Anisotropic diffusion processes were used for fairing [DMSB99] and smoothing [TWBO02, CRT04] of meshes, as well as for processing of volumetric medical data [BES11]. The construction of anisotropic Laplacians driven by principal curvatures was shown in [ARAC14] on meshes and in [KTT13] for abstract high-dimensional data.

On the other side, our work is related to the recent line of works on applying machine learning methods to geometric data [S*13, R*14, KRCH14, W*14, COC14, LBBC14, W*15, SMKL15]. Litman and Bronstein [LB14] showed that heat and wave kernel signatures can be considered as particular parametric families of transfer functions applied to the Laplace-Beltrami operator eigenvalues and proposed to learn an optimal transfer function. Their work follows the recent trends in the image analysis domain, where hand-crafted descriptors are abandoned in favor of learning approaches [KSH12, CGGS12, SZ14, KF15]. The remarkable success

of deep learning architectures such as convolutional neural networks (CNN) [Fuk80,L*89] in computer vision, has driven recent attempts to bring these methods to non-Euclidean domains. Bruna et al. [BZSL14] used a spectral formulation of convolution to apply CNNs on graphs. Boscaini et al. [B*15] used localized spectral CNNs based on vertex-frequency analysis [SRV15] to learn intrinsic 3D shape descriptors. A similar ShapeNet model [MBBV15] used local geodesic polar coordinates [KBLB12] to define convolutions on meshes.

Contributions. The key contribution of this paper is the construction of shape feature descriptors based on anisotropic oriented learnable diffusion kernels. Unlike [ARAC14], we do not use a single anisotropic diffusion kernel in the direction of the principal curvature, but a multitude of kernels at different directions. Such kernels capture local directional structures similarly to the geodesic polar coordinates used in the intrinsic shape context [KBLB12] and ShapeNet [MBBV15], however, with several notable advantages. First, we do not have an ambiguity in the definition of the angular coordinate. Second, unlike geodesic patches that are not always guaranteed to be valid topological disks, anisotropic kernels are always valid. Third, our construction is not limited to meshes and is advantageous computationally.

Learning optimal anisotropic kernels can be considered as an extension of the optimal spectral descriptors method of [LB14]. One of the main advantages is that anisotropy allows to disambiguate intrinsic symmetry. We show that using a different learning cost, we can learn intrinsic correspondence between shapes.

Finally, we show experimentally that our approach outperforms state-of-the-art intrinsic feature descriptors [SOG09,ASC11, LB14], including the most recent Riemannian CNNs [B*15, MBBV15] that so far have held the record of top performance on standard descriptors benchmarks. In shape correspondence experiments, our approach beats popular non-learnable methods such as blended maps [KLF11] and functional maps [O*12], as well as a recent learning approach based on random forests [R*14].

The rest of the paper is organized as follows. In Section 2, we review basic notions in differential geometry and define anisotropic diffusion on manifolds. Section 3 overviews standard spectral shape descriptors. In Section 4 we show the construction of anisotropic heat kernels as a motivating example leading to more general learnable anisotropic kernels discussed in Section 5. Section 6 discussed the numerical computation of our descriptors. Section 7 shows experimental results, and finally, Section 8 concludes the paper.

2. Background

Manifold. We model a 3D shape as a two-dimensional compact Riemannian manifold (surface) X , for simplicity assumed to have no boundary. Let T_xX denote the *tangent plane* at x , modeling the manifold locally as a Euclidean space, TX denote the *tangent bundle*, and let $\exp_x: T_xX \rightarrow X$ be the *exponential map*, mapping tangent vectors onto the manifold. A *Riemannian metric* is an inner product on the tangent plane $\langle \cdot, \cdot \rangle_{T_xX}: T_xX \times T_xX \rightarrow \mathbb{R}$ depending smoothly on x . The Riemannian metric is represented as a 2×2 ma-

trix referred to as *first fundamental form*. Quantities which are expressible entirely in terms of Riemannian metric, and therefore independent on the way the manifold is embedded, are called *intrinsic*. Such quantities are invariant to isometric (metric-preserving) deformations.

Curvature. Given an embedding of the manifold, the *second fundamental form*, represented as a 2×2 matrix, describes how the manifold locally differs from a plane. The eigenvalues κ_m, κ_M of the second fundamental form are called the *principal curvatures*; the corresponding eigenvectors $\mathbf{v}_m, \mathbf{v}_M$ called the *principal curvature directions* form an orthonormal basis on the tangent plane.

Differential operators on manifolds. Let $f: X \rightarrow \mathbb{R}$ be a smooth scalar field. We define the *intrinsic gradient* as

$$\nabla_X f(x) = \nabla(f \circ \exp_x)(\mathbf{0}),$$

where ∇ denotes the standard Euclidean gradient acting in the tangent plane. The intrinsic gradient can be interpreted as the direction (tangent vector on T_xX) in which f changes the most at point x ; it is thus an operator acting on a scalar field and producing a vector field. First-order Taylor expansion takes the form $(f \circ \exp_x)(\mathbf{v}) \approx f(x) + \langle \nabla_X f(x), \mathbf{v} \rangle_{T_xX}$, where the second term is the *directional derivative* of f in the direction $\mathbf{v} \in T_xX$.

Given a smooth vector field $\mathbf{v}: X \rightarrow TX$, the *intrinsic divergence* is an operator acting on vector fields producing scalar fields, defined as the negative adjoint of the intrinsic gradient operator,

$$\int_X \langle \nabla_X f(x), \mathbf{v}(x) \rangle_{T_xX} dx = - \int_X f(x) \operatorname{div}_X \mathbf{v}(x) dx, \quad (1)$$

where the area element dx is induced by the Riemannian metric.

Combining the two, we can now define the *Laplacian* or *Laplace-Beltrami operator* as

$$\Delta_X f(x) = -\operatorname{div}_X(\nabla_X f(x)). \quad (2)$$

The Laplacian can be interpreted as the difference between the average of f on an infinitesimal sphere around x and $f(x)$.

Spectral analysis on manifolds. The Laplacian is a positive-semidefinite operator, admitting real eigen-decomposition

$$\Delta_X \phi_k = \lambda_k \phi_k, \quad (3)$$

with countable eigenvalues $0 = \lambda_0 \leq \lambda_1 \leq \dots$. The corresponding eigenfunctions ϕ_0, ϕ_1, \dots are orthonormal w.r.t. the standard inner product $\langle \phi_i, \phi_j \rangle_X = \int_X \phi_i(x) \phi_j(x) dx = \delta_{ij}$ and form an orthonormal basis for the functional space $L^2(X) = \{f: X \rightarrow \mathbb{R} : \langle f, f \rangle_X < \infty\}$. Quantities expressed in terms of the Laplacian eigenvalues and eigenfunctions are commonly referred to as *spectral*.

The Laplacian eigenfunctions are a generalization of the classical Fourier basis to non-Euclidean domains: a function $f \in L^2(X)$ can be represented as the *Fourier series*

$$f(x) = \sum_{k \geq 0} \langle f, \phi_k \rangle_X \phi_k(x), \quad (4)$$

where the eigenvalues $\{\lambda_k\}_{k \geq 1}$ play the role of frequencies (the first eigenvalue $\lambda_0 = 0$ corresponds to a constant eigenvector or ‘DC component’).

Heat diffusion on manifolds. The Laplacian operator is ubiquitous in physics, appearing in differential equations describing diffusion, wave, and quantum phenomena. We are interested in heat propagation on manifolds, governed by the *diffusion equation*

$$f_t(x, t) = -\Delta_X f(x, t), \quad (5)$$

which is a mathematical description of a physical principle known as *Newton's cooling law*, stating that the rate of change of the temperature of an object (lhs) is proportional to the difference between its own temperature and the temperature of the surrounding (rhs). Here $f(x, t)$ is the temperature at point x at time t .

Given some initial heat distribution $f_0(x) = f(x, 0)$, the solution of heat equation (5) at time t is obtained by applying the *heat operator* $H^t = e^{-t\Delta_X}$ to f_0 ,

$$f(x, t) = H^t f_0(x) = \int_X f_0(\xi) h_t(x, \xi) d\xi, \quad (6)$$

where $h_t(x, \xi)$ is called the *heat kernel*, and the above equation can be interpreted as a non-shift-invariant version of convolution. In the Euclidean case, the heat kernel has the form $h_t(x - \xi)$ and the solution is given as $f = f_0 * h_t$.

In the spectral domain, the heat kernel is expressed as

$$h_t(x, \xi) = \sum_{k \geq 0} e^{-t\lambda_k} \phi_k(x) \phi_k(\xi); \quad (7)$$

appealing again to the signal processing intuition, $e^{-t\lambda}$ acts as a low-pass filter (larger t , corresponding to longer diffusion, lowers the pass band). We will revisit this intuition in Section 3 in relation to spectral descriptors.

Anisotropic diffusion. The heat diffusion equation (5) assumes that the heat conduction properties of the manifold are constant at every point. A more general diffusion equation has the form

$$f_t(x, t) = \text{div}_X(\mathbf{D}(x)\nabla_X f(x, t)), \quad (8)$$

where $\mathbf{D}(x)$ is the *thermal conductivity tensor* (2×2 matrix) applied to the gradient in the tangent plane. The thermal conductivity tensor allows modeling heat flow that is position- and direction-dependent; the diffusion equation in this case is called *anisotropic*; the special case $\mathbf{D}(x) = c(x)\mathbf{I}$ is called *inhomogeneous*, and $\mathbf{D}(x) = c\mathbf{I}$ is called *homogeneous*.

One of the most celebrated applications of such diffusion equations for edge-preserving image enhancement is the seminal work of Perona and Malik [PM90], where the scalar diffusion coefficient was chosen to be inversely proportional to the edge strength, to avoid diffusing across edges.[†] In a recent work, Andreux et al. [ARAC14] considered anisotropic diffusion driven by the surface curvature. Assuming that at each point x the tangent vectors are

expressed w.r.t. the orthogonal basis $\mathbf{v}_m, \mathbf{v}_M$ of principal curvature directions, they used

$$\mathbf{D}_\alpha(x) = \begin{pmatrix} \Psi_\alpha(\kappa_M(x)) & \\ & \Psi_\alpha(\kappa_m(x)) \end{pmatrix}, \quad (9)$$

with $\Psi_\alpha(x) = \frac{1}{1+\alpha|x|}$. Such a thermal conductivity tensor drives the diffusion in the direction of the maximum curvature, where α controls the degree of anisotropy ($\alpha = 0$ corresponds to the classical isotropic case), and the amount of diffusion depends on the extrinsic principal curvatures κ_m, κ_M . A simpler alternative is

$$\mathbf{D}_\alpha(x) = \begin{pmatrix} \alpha & \\ & 1 \end{pmatrix}, \quad (10)$$

which is independent of the extrinsic curvature.

3. Isotropic spectral descriptors

Several recent works have exploited spectral properties of the Laplace-Beltrami operator to construct local shape descriptors. Many popular spectral shape descriptors are constructed taking the diagonal values of heat-like operators. A generic Q -dimensional descriptor of this kind has the form

$$\mathbf{f}(x) = \sum_{k \geq 0} \tau(\lambda_k) \phi_k^2(x) \approx \sum_{k=0}^K \tau(\lambda_k) \phi_k^2(x), \quad (11)$$

where $\tau(\lambda) = (\tau_1(\lambda), \dots, \tau_Q(\lambda))^T$ is a bank of ‘transfer functions’ acting on the Laplacian eigenvalues. Such descriptors are dense (computed at every point x), intrinsic by construction, and typically can be efficiently computed using a small number K of Laplacian eigenfunctions and eigenvalues.

Heat Kernel Signature (HKS) [SOG09, GBAL09] is a particular case of (11) where $\tau_t(\lambda) = e^{-t\lambda}$ are low-pass filters parametrized over t . In other words, the HKS is the diagonal of the heat kernel (7) evaluated at some time values t_1, \dots, t_Q . As a consequence, it can be interpreted as *autodiffusivity*, i.e., the amount of heat remaining at point x after time t , which is equal (up to constant) to the Gaussian curvature for small t [MP49]. A notable drawback of HKS stemming from the use of low-pass filters is its poor spatial localization.

Wave Kernel Signature (WKS). Aubry et al. [ASC11] proposed to define a shape descriptor by looking at the behavior of a quantum particle on the manifold. The particle is assumed to have a log-normal initial energy distribution with mean energy level ν . The WKS measures the probability of finding a particle of energy ν at point x , and turns out as a particular instance of (11) with band-pass filters of the form $\tau_\nu(\lambda) = \exp\left(-\frac{(\log \nu - \log \lambda)^2}{2\sigma^2}\right)$. Compared to HKS, WKS have better localization, but at the same time tend to produce noisier matches.

Optimal Spectral Descriptor (OSD). Given the great effectiveness of HKS and WKS and the fact that they differ only in the choice of the filter τ , Litman and Bronstein [LB14] proposed to consider a parametric family of filters expressed as

$$\tau_q(\lambda) = \sum_{m=1}^M a_{qm} \beta_{qm}(\lambda) \quad (12)$$

[†] The authors erroneously referred to their equation as ‘non-linear anisotropic diffusion’, where the term non-linear is meant in the sense of non-shift-invariant (i.e., the diffusion cannot be expressed as a linear convolution), and the diffusion is in fact isotropic but inhomogeneous. We stress that the anisotropic Laplacian is still a linear operator.




Figure 1: Examples of anisotropic heat kernels computed on meshes (top) and point clouds (bottom) for different values of α , θ , and τ . Leftmost image shows the isotropic case.

in the B-spline basis $\beta_1(\lambda), \dots, \beta_M(\lambda)$, where a_{qm} ($q = 1, \dots, Q, m = 1, \dots, M$) are the parametrization coefficients. Plugging (12) in (11), one can express the spectral descriptor as

$$f_q(x) = \sum_{m=1}^M a_{qm} \underbrace{\sum_{k \geq 0} \beta_m(\lambda_k) \phi_k^2(x)}_{g_m(x)}, \quad (13)$$

where $\mathbf{g}(x) = (g_1(x), \dots, g_M(x))^T$ is called *geometry vector* and depends only on the intrinsic geometry of the shape. Thus, (11) is parametrized by the $Q \times M$ matrix $\mathbf{A} = (a_{lm})$ and can be written in matrix form as $\mathbf{f}(x) = \mathbf{A}\mathbf{g}(x)$. The main idea of [LB14] is to *learn* the optimal parameters \mathbf{A} by minimizing a task-specific loss which reduces to a Mahalanobis-type metric learning.

Windowed Fourier Transform (WFT). Boscaini *et al.* [B*15] used the generalization of the windowed Fourier transform [SRV15] in order to capture the local context of a signal on the manifold. Given a function $f \in L^2(X)$, its WFT is given by

$$(Sf)_{x,l} = \langle f, g_{x,l} \rangle_X = \sum_{k \geq 0} \hat{g}_k \phi_k(x) \langle f, \phi_k \phi_l \rangle_X, \quad (14)$$

where

$$g_{\xi,l}(x) = (M_l T_\xi g)(x) = \phi_l(x) \sum_{k \geq 0} \hat{g}_k \phi_k(\xi) \phi_k(x) \quad (15)$$

are the transform atoms obtained by displacing a window (given in the spectral domain by the coefficients \hat{g}_k) to point ξ by means of the translation operator T_ξ , and modulating it by the l th frequency




Figure 2: Anisotropic HKS descriptors computed on two different poses of a human shape at the three points (marked in red, blue, and green). Solid and dotted lines represent the shape on the left and right, respectively. Each AHKS is computed with $L = 16$ diffusion angles.

by means of the modulation operator M_l . In [B*15], local descriptors were constructed by applying the WFT to geometry vectors; the window function was learned to minimize a task-specific cost.

4. Anisotropic diffusion descriptors

One of the notable drawbacks of spectral descriptors discussed in the previous section is that they ignore directional information. Such information, however, may often carry important cues about the local structure of the surface, which are important for the construction of a good feature descriptor. Furthermore, intrinsic descriptors are *ambiguous under intrinsic symmetries*: given an intrinsic symmetry $\eta: X \rightarrow X$, an intrinsic descriptor would be invariant to it, $\mathbf{f} \circ \eta = \mathbf{f}$.

In this paper, we propose to construct a novel class of descriptors based on anisotropic diffusion. Specifically, we consider diffusion that is driven at each point by an angle θ w.r.t. the principal curvature direction, arising from the thermal conductivity tensor

$$\mathbf{D}_{\alpha\theta}(x) = \mathbf{R}_\theta \mathbf{D}_\alpha(x) \mathbf{R}_\theta^\top, \quad (16)$$

where $\mathbf{D}_\alpha(x)$ is as defined in equation (9) or (10). We refer to the operator $\Delta_{\alpha,\theta} f(x) = -\text{div}_X(\mathbf{R}_\theta \mathbf{D}_\alpha(x) \mathbf{R}_\theta^\top \nabla_X f(x))$ as the *anisotropic Laplacian*, and denote by $\{\phi_{\alpha\theta k}, \lambda_{\alpha\theta k}\}_{k \geq 0}$ its eigenfunctions and eigenvalues.

Note that strictly speaking, our anisotropic Laplacian is not intrinsic: it depends on the principal curvature direction, and additionally, if formula (9) is used, on the principal curvature values κ_m, κ_M . If formula (10) is used and we consider all the possible rotations $\theta \in [0, 2\pi)$, the Laplacian is intrinsic up to the choice of the origin of θ . This ambiguity can be removed in several ways: by fixing the origin using a reference direction (*e.g.*, principal curvature direction) as we do here; by taking the Fourier transform magnitude w.r.t. the θ coordinate as done in [BK10, KBLB12]; or by taking the maximum over all the angles as done in [MBBV15].

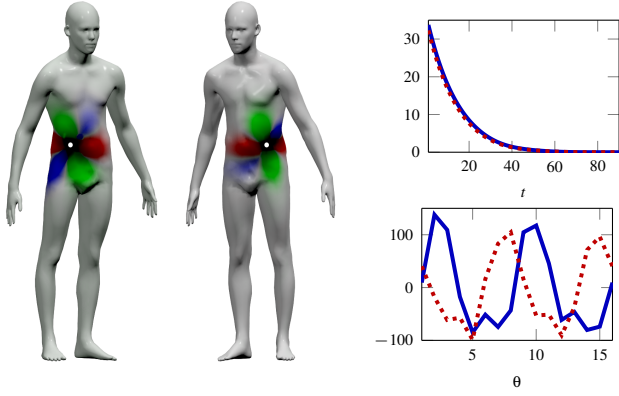


Figure 3: Illustration of the intrinsic symmetry ambiguity. Left: two symmetric points on a shape (white) and the anisotropic heat kernels for different θ computed at these points (similar colors encode similar angles). Right: the values of HKS (top) and anisotropic HKS (bottom) computed at two symmetric points (solid and dotted curves). Note that HKS is fully ambiguous (both curves coincide), while anisotropic HKS allows to distinguish between symmetric points (one curve is the reflection of the other).

4.1. Anisotropic HKS

As a motivating example, we first construct an anisotropic counterpart of the HKS. The anisotropic heat kernel is given by

$$h_{\alpha\theta t}(x, \xi) = \sum_{k \geq 0} e^{-t\lambda_{\alpha\theta k}} \phi_{\alpha\theta k}(x) \phi_{\alpha\theta k}(\xi); \quad (17)$$

similarly to HKS, we define the *anisotropic HKS* by considering the diagonal values $h_{\alpha\theta t}(x, x)$ of the anisotropic heat kernel (17) and sampling t and θ at values t_1, \dots, t_Q and $\theta_1, \dots, \theta_L$, respectively (α is used as a parameter). Examples of anisotropic heat kernels constructed on meshes and point clouds are shown in Figure 1; anisotropic HKS descriptors are shown in Figure 2.

It is important to notice that since we use the principal curvature direction as the origin $\theta = 0$ of the angular coordinate, and since the curvature is an extrinsic property, our descriptor is *not ambiguous under bilateral intrinsic symmetry*. In fact, intrinsic symmetry η reflects the angular coordinate, $\mathbf{f}_\theta \circ \eta = \mathbf{f}_{-\theta}$ (here $\mathbf{f}_\theta(x) = h_{\alpha\theta t}(x, x)$). This phenomenon is illustrated in Figure 3.

4.2. Learnable anisotropic kernels

Anisotropic HKS descriptors show that anisotropic spectral kernels carry rich information about local shape structures. However, the low-pass filters representing the heat kernels are not necessarily best suited for some applications, and one could try to find the optimal application-specific kernel by learning. Even more generally, we propose learning non-linear anisotropic descriptors using deep neural networks with a tailored architecture. The network is built of several layers combined one after another; it is called *deep* if many layers are used. The output of the network is a vector containing the descriptor at a point. A simple example of a rather shallow network is shown in Figure 4.




Figure 4: Example of a simple single-layer neural network architecture implementing the anisotropic descriptors proposed in this paper. The inputs are spectral decompositions of anisotropic Laplacians with anisotropy α at angle $\theta_1, \dots, \theta_L$ w.r.t. the principal curvature direction. In the spectral stage, a learnable spectral bank of M filters (shared for all the L directions) is used to create ML directed kernels. ξ denotes the ReLU nonlinearity. The dimensionality reduction stage then reduces the descriptor to Q output dimensions. Deeper architectures may contain additional layers in both stages.

As the input to the network, we provide the eigenvalues and eigenvectors of the anisotropic Laplacians; we use K eigenvectors and L equally-spaced rotation angles $\theta_1, \dots, \theta_L$. Optionally, we may also use multiple degrees of anisotropy, $\alpha_1, \dots, \alpha_A$.

Anisotropic Spectral Filter (ASF_M) is the first layer, applying a bank of M filters to produce for each of the LA inputs a kernel of the form

$$f_{ijm}(x) = \sum_{k=1}^K \tau_m(\lambda_{\alpha_i\theta_j k}) \phi_{\alpha_i\theta_j k}^2(x), \quad (18)$$

for $m = 1, \dots, M$, $i = 1, \dots, A$, and $j = 1, \dots, L$, where filter transfer functions $\tau_m(\lambda)$ are parametrized as in formula (12) and are learned. The filters are shared across all rotations and anisotropies, resembling the shared connectivity of convolutional networks and allowing to reduce the number of degrees of freedom in the model, thus reducing the chance of overfitting. Particular choices of low-pass and band-pass filters in the ASF layer would produce HKS and WKS at the output, respectively.

Stack Channels (SC) is applied to input channels ordered by angle and anisotropy (e.g., output of the ASF layer) and ‘flattens’ them to be used in the dimensionality reduction. From this layer onward, the model loses any geometrical interpretation and acts as a non-linear extraction of features from spectral information.

ReLU is a non-parametric layer which preserves only the non-negative part of the input vector, $\xi(t) = \max(0, t)$. This helps in learning deep models as it does not suffer from saturation problems typical of sigmoid-type functions, which result in difficult gradient back-propagation.

Fully connected (FC_Q) layer, given a P -dimensional input $f_1^{\text{in}}(x), \dots, f_P^{\text{in}}(x)$ produces a Q -dimensional output $f_1^{\text{out}}(x), \dots, f_Q^{\text{out}}(x)$ as linear combination of the input channels with a set of learnable parameters w ,

$$f_q^{\text{out}}(x) = \sum_{p=1}^P w_{qp} f_p^{\text{in}}(x), \quad q = 1, \dots, Q. \quad (19)$$

Softmax layer applies the softmax function to the P -dimensional input,

$$\sigma(f_p^{\text{in}}) = \frac{e^{f_p^{\text{in}}}}{\sum_{p=1}^P e^{f_p^{\text{in}}}}, \quad p = 1, \dots, P. \quad (20)$$

It is used together with the negative log-likelihood loss function in classification tasks, as described in the following section.

The architecture of our ADD networks can be divided into two stages: spectral filtering and dimensionality reduction. The former applies one or more ASF layers to produce a non-linear spectral filtering (as opposed to previous approaches where this stage was kept purely linear). The latter stage reduces the dimensionality to obtain compact descriptors which require little storage and that are fast to compute.

5. Applications

The neural network architecture described in the previous section allows defining descriptors at each point x of a shape, given by a parametric map $\mathbf{f}_\Theta(x)$, where Θ denotes the learnable parameters of the network. The choice of the parameters is done by an optimization process, minimizing a task-specific cost. Here, we consider two tasks: design of invariant descriptors and computation of shape correspondence.

Invariant descriptors. Given a set of shapes, our goal is to compute a local feature descriptor that is insensitive to the variability of shapes across the collection. We assume to be provided with examples of points x, x^+ from different shapes that are known to be similar (referred to as *positives* and denoted by \mathcal{T}^+), as well as pairs of points x, x^- known to be dissimilar (*negatives* \mathcal{T}^-). The most straightforward way to obtain such training sets is from known correspondence between some of the shapes in the collection. Given T shapes with n known corresponding points, the number of positive and negative pairs is $O(nT^2)$ and $O(n^2T^2)$, respectively. Therefore, in practice just a few shapes with known correspondence are sufficient to produce training sets containing millions of pairs.

We construct a descriptor that is as similar as possible on the positive pairs and as dissimilar as possible on the negatives. For this purpose, we use a *siamese network* configuration [B*94, HCL06, S*14], composed of two identical copies of the same network model sharing the same parametrization and fed by pairs of knowingly similar or dissimilar samples. We minimize the *siamese loss*

$$\begin{aligned} \ell_{\text{siam}}(\Theta) = & (1 - \gamma) \sum_{(x, x^+) \in \mathcal{T}^+} \|\mathbf{f}_\Theta(x) - \mathbf{f}_\Theta(x^+)\|^2 \\ & + \gamma \sum_{(x, x^-) \in \mathcal{T}^-} (\mu - \|\mathbf{f}_\Theta(x) - \mathbf{f}_\Theta(x^-)\|)_+^2, \end{aligned} \quad (21)$$

where $\gamma \in [0, 1]$ is a parameter trading off between the positive and negative losses, μ is a margin, and $(t)_+ = \max\{0, t\}$.

Shape correspondence. Finding the correspondence in a collection of shapes can be posed as a labelling problem, where one tries to label each vertex of a given *query* shape X with the index of a corresponding point on some common *reference* shape Y [R*14]. Let us denote by $y^*(x)$ the ground-truth correspondence of x on the




Figure 5: Discretization of the anisotropic Laplace-Beltrami operators on a triangular mesh. The orthogonal basis vectors $\hat{\mathbf{u}}_M, \hat{\mathbf{u}}_m$, as well as their rotated counterparts (in red), lie on the tangent plane of the respective triangle.

reference shape. We assume to be provided with examples of points from shapes across the collection and their ground-truth correspondence, $\mathcal{T} = \{(x, y^*(x))\}$.

For a point x on a query shape, we compute a descriptor $\mathbf{f}_\Theta(x)$ encoding the probability distribution on Y , which acts as a ‘soft correspondence’. The optimal parameters of the descriptor are found by minimizing the *multinomial regression loss*

$$\ell_{\text{reg}}(\Theta) = - \sum_{(x, y^*(x)) \in \mathcal{T}} \langle \delta_{y^*(x)}, \log f_\Theta(x) \rangle_Y, \quad (22)$$

where $\delta_{y^*(x)}$ is a delta-function on Y at point $y^*(x)$.

6. Discretization

Meshes. In the discrete setting, the manifold X is sampled at n points $V = \{x_1, \dots, x_n\}$. The points are connected by edges E and faces F , forming a manifold triangular mesh (V, E, F) . To each triangle $ijk \in F$, we attach an orthonormal reference frame $\mathbf{U}_{ijk} = (\hat{\mathbf{u}}_M, \hat{\mathbf{u}}_m, \hat{\mathbf{n}})$, where $\hat{\mathbf{n}}$ is the unit normal vector to the triangle and $\hat{\mathbf{u}}_M, \hat{\mathbf{u}}_m \in \mathbb{R}^3$ are the directions of principal curvature, computed using the method of [CSM03]. The thermal conductivity tensor for the triangle ijk operating on tangent vectors is expressed w.r.t. \mathbf{U}_{ijk} as a 3×3 matrix $\mathbf{D} = \begin{pmatrix} \mathbf{D}^\alpha & \\ & \mathbf{1} \end{pmatrix}$, with the 2×2 block \mathbf{D}^α set according to equation (9).

Now let $\hat{\mathbf{e}}_{ab} \in \mathbb{R}^3$ denote the oriented edge pointing from vertex a to vertex b , normalized to unit length, and consider the triangle ijk as in Figure 5. We define the \mathbf{H} -weighted inner product between edges $\hat{\mathbf{e}}_{kj}$ and $\hat{\mathbf{e}}_{ki}$ as

$$\langle \hat{\mathbf{e}}_{kj}, \hat{\mathbf{e}}_{ki} \rangle_{\mathbf{H}} = (\mathbf{U}_{ijk}^\top \hat{\mathbf{e}}_{kj})^\top \mathbf{D} (\mathbf{U}_{ijk}^\top \hat{\mathbf{e}}_{ki}) = \hat{\mathbf{e}}_{kj}^\top \underbrace{\mathbf{U}_{ijk} \mathbf{D} \mathbf{U}_{ijk}^\top}_{\mathbf{H}} \hat{\mathbf{e}}_{ki}, \quad (23)$$

where the *shear matrix* \mathbf{H} encodes the anisotropic scaling, operated via \mathbf{D} , up to an orthogonal basis change. Note that in the case $\mathbf{D} = \mathbf{I}$ we also have $\mathbf{H} = \mathbf{I}$, such that the \mathbf{H} -weighted inner product simplifies to the standard inner product $\langle \hat{\mathbf{e}}_{kj}, \hat{\mathbf{e}}_{ki} \rangle_{\mathbf{H}} = \cos \alpha_{ij}$.

The discretization of the anisotropic Laplacian takes the form of an $n \times n$ sparse matrix $\mathbf{L} = -\mathbf{S}^{-1} \mathbf{W}$. The *mass matrix* \mathbf{S} is a diagonal matrix of area elements $s_i = \frac{1}{3} \sum_{jk:ijk \in F} A_{ijk}$, where A_{ijk}




Figure 6: Qualitative evaluation of our mADD3 descriptors on meshes. Shown is the normalized Euclidean distance between the descriptor at a reference point on the shoulder (white point) and the descriptors computed at the rest of the points for different transformations. Cold and hot colors represent small and large distances, respectively; distances are saturated at the median value. Ideal descriptors would produce a distance map with a sharp minimum at the corresponding point and no spurious local minima at other locations.

denotes the area of triangle ijk . The *stiffness matrix* \mathbf{W} is composed of weights

$$w_{ij} = \begin{cases} \frac{1}{2} \left(\frac{\langle \hat{\mathbf{e}}_{kj}, \hat{\mathbf{e}}_{ki} \rangle_{\mathbf{H}}}{\sin \alpha_{ij}} + \frac{\langle \hat{\mathbf{e}}_{hj}, \hat{\mathbf{e}}_{hi} \rangle_{\mathbf{H}}}{\sin \beta_{ij}} \right) & ij \in E; \\ -\sum_{k \neq i} w_{ik} & i = j; \\ 0 & \text{else,} \end{cases} \quad (24)$$

where the notation is according to Figure 5. A complete derivation of (24) is omitted due to space reasons; we refer to [ARAC14] for a way to derive a similar expression using the standard FEM.

The isotropic Laplace-Beltrami operator is obtained by setting $\mathbf{D} = \mathbf{I}$, in which case we have $\frac{\langle \hat{\mathbf{e}}_{kj}, \hat{\mathbf{e}}_{ki} \rangle_{\mathbf{H}}}{\sin \alpha_{ij}} = \frac{\cos \alpha_{ij}}{\sin \alpha_{ij}} = \cot \alpha_{ij}$, thus reducing equation (24) to the classical cotangent formula [Mac49, Duf59, PP93, MDSB03].

Finally, note that the discretization given above only encodes anisotropy along the directions of principal curvature $\hat{\mathbf{u}}_M, \hat{\mathbf{u}}_m$. In order to allow arbitrary directions used in our anisotropic descriptors, it is sufficient to rotate the basis vectors \mathbf{U}_{ijk} on each triangle around the respective normal $\hat{\mathbf{n}}$ by the angle θ , equal for all triangles (see Figure 5, red). Denoting by \mathbf{R}_θ the corresponding 3×3 rotation matrix, this is equivalent to modifying the \mathbf{H} -weighted inner product with the directed shear matrix $\mathbf{H}_\theta = \mathbf{R}_\theta \mathbf{H} \mathbf{R}_\theta^T$. The resulting weights w_{ij} in equation (24) are thus obtained by using the inner products $\langle \hat{\mathbf{e}}_{kj}, \hat{\mathbf{e}}_{ki} \rangle_{\mathbf{H}_\theta} = \hat{\mathbf{e}}_{kj}^T \mathbf{H}_\theta \hat{\mathbf{e}}_{ki}$.

Point clouds. In order to construct the anisotropic Laplace-Beltrami operator on point clouds, we follow a procedure similar to [CRT04]. For each point \mathbf{x} in the cloud, we define its tangent space by estimating the normal vector at \mathbf{x} via total least squares [MN03] using points within an ε -ball $B_\varepsilon(\mathbf{x})$. The points in $B_\varepsilon(\mathbf{x})$ are then projected onto the tangent plane, where a local mesh construction is carried out by Delaunay triangulation. From the triangulation, only the triangle fan containing \mathbf{x} is retained and projected back onto the original coordinates in \mathbb{R}^3 . Finally, the edges in this local patch receive weights according to equation (24). The directions of principal curvature are computed as eigenvectors of the covariance matrix of neighboring normals, which provides an approximation to the shape operator [BC94].

7. Results

In this section, we show how our anisotropic diffusion descriptors (ADD) framework can be used to learn task-specific features.

Datasets. We used FAUST [BRLB14] and SCAPE [A*05] datasets, containing scanned human shapes in different poses. For all the datasets, vertex-wise correspondence between shapes was given. The meshes in SCAPE were resampled to 12.5K vertices; FAUST shapes contained 6.9K points. All shapes were scaled to unit geodesic diameter. Point clouds were produced by subsampling the meshes regularly to 4K points using the farthest point method, in order to break the local connectivity structure of the mesh.

Settings. Isotropic Laplacians were computed using the cotangent formula; anisotropic Laplacians were computed according to (24). We used $K = 300$ and 900 eigenfunctions for meshes and point clouds, respectively. $L = 16$ angles were used for anisotropic kernels. For single anisotropy case, we used $\alpha = 50$; for multiple anisotropies case, we used $\alpha = 0, 25$, and 50. In the shape descriptors experiment, we used the following three architectures, in increasing complexity order: single layer (SC + FC16 referred to as ADD1); 2-layer (ASF256 + ReLU + SC + FC16 referred to as ADD2), and 4-layer (ASF256 + ReLU + ASF256 + ReLU + SC + FC512 + ReLU + FC16, referred to as ADD3). Architectures with multiple anisotropies are denoted as mADD1–mADD3. For the correspondence experiment, we used a 5-layer architecture with multiple anisotropies, ASF512 + ReLU + ASF512 + ReLU + SC + FC2048 + ReLU + FC2048 + ReLU + FC6890 + Softmax (mADD4).

Implementation. Neural networks were implemented in Theano [B*10]. In all experiments, the ADADELTA [Zei12] stochastic optimization algorithm was used with initial learning rate of 1.0 and $\rho = 0.95$. Convergence was usually reached in ≈ 50 training epochs. For the shape descriptors experiments, we used mini-batches of 100 pairs of points, whereas for the correspondence task the batch size was set to the number of vertices in the shapes. Code to reproduce all the experiments in this paper, and the full framework will be released upon publication.

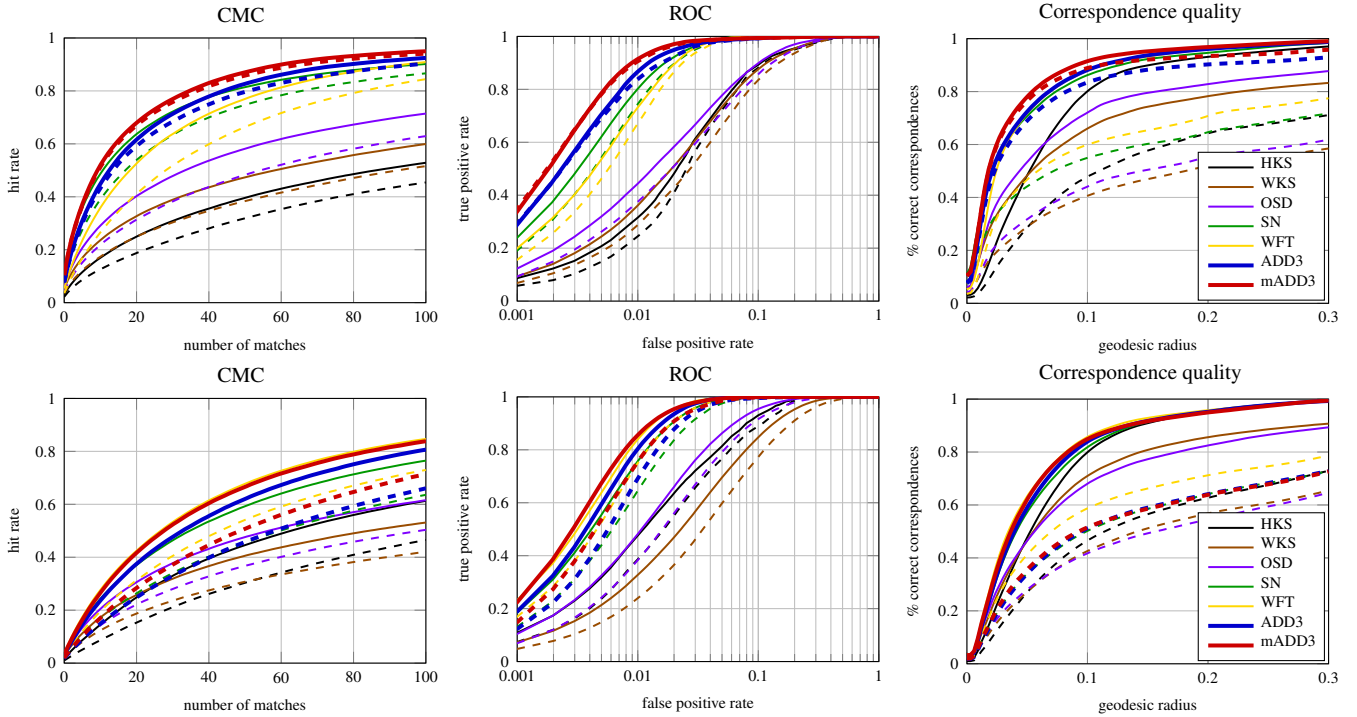


Figure 7: Performance of different descriptors measured using the CMC (left), ROC (center) and Princeton protocol for nearest-neighbor correspondence (right); higher curves correspond to better performance. Shown are symmetric (solid) and asymmetric (dashed) settings. Learnable descriptors were trained and tested on disjoint sets of the FAUST (top) and SCAPE (bottom) datasets, respectively. All descriptors in these plots are 16-dimensional.

Complexity. The pre-computation step includes the construction of L Laplacians, one per diffusion direction. The computation of each Laplacian ($n \times n$ sparse matrix, where n is the number of vertices) and its eigendecomposition requires $O(n^2)$ operations. The computational complexity of the descriptor (forward propagation in the neural network at each vertex) is $O(cn)$, where c is determined by the complexity of the network (number of layers, etc.) independent of n .

Timing. The following are typical timings for FAUST shapes with $n = 6.9K$ vertices. Laplacian computation and eigendecomposition took 1 sec and 4 seconds per angle respectively, on a i7 workstation with 64GB of RAM. For the three architectures ADD1–ADD3, training efficiency were 43, 29, and 17 minibatches of size 100 per second, respectively. The computation of the descriptor at all the points (forward propagation) took approximately 0.5 sec per shape using a NVIDIA Titan Black GPU.

Evaluation. Quantitative evaluation was done using three criteria: *cumulative match characteristic* (CMC), *receiver operator characteristic* (ROC), and the *Princeton protocol* [KLF11]. The CMC evaluates the probability of a correct correspondence among the k nearest neighbors in the descriptor space. The ROC measures the percentage of positives and negatives pairs falling below various thresholds of their distance in the descriptor space (*true positive* and *negative rates*, respectively). The Princeton protocol counts

the percentage of nearest-neighbor matches that are at most r -geodesically distant from the ground-truth correspondence. Furthermore, all the above criteria were evaluated in two settings: *symmetric* (considering symmetric points as correct matches), and *asymmetric* (considering symmetric points as incorrect matches).

7.1. Shape descriptors

In the first experiment, we computed dense intrinsic pose- and subject-invariant anisotropic descriptors for human shapes using the proposed approach. The experimental setup followed verbatim [LB14, MBBV15, B*15]. For reference, we show the results of HKS [SOG09], WKS [ASC11], OSD [LB14], ShapeNet [MBBV15] and WFT [B*15] using the code and settings provided by the respective authors. All the descriptors were $Q = 16$ -dimensional as in [LB14, MBBV15, B*15].

Training was done using the loss (21) with positive and negative sets of vertex pairs generated on the fly from known correspondence between training shapes. Symmetric points were considered as negatives. On the FAUST dataset, we used subjects 1–8 for training, and subject 9–10 for testing. On SCAPE, we used shapes 10–70 for training, and the rest for testing.

Figure 6 shows a qualitative evaluation of the anisotropic spectral descriptors constructed with mADD3 architecture on meshes by visualizing the distance between the descriptor at a selected




Figure 8: Performance of different architectures of our anisotropic descriptor measured using CMC (left), ROC (center) and Princeton protocol for nearest-neighbor correspondence (right); higher curves correspond to better performance. Shown are symmetric (solid) and asymmetric (dashed) settings. Training and testing were done on disjoint sets of the FAUST dataset.

point and the rest of the points on the same shape as well as its transformations. Our descriptor manifests good localization and specificity, and is also not ambiguous to symmetry.

Figure 7 shows the evaluation of our descriptors on the FAUST and SCAPE datasets. Our approach outperforms other descriptors, including the learnable ShapeNet and WFT. Further, note that when switching from the harder asymmetric evaluation setting to the easier symmetric one (dashed and solid curves, respectively), our descriptor manifests significantly smaller drop compared to other methods. Figure 8 compares the performance of different architectures, showing that the multiple anisotropy deep model (mADD3) achieves the best performance.

7.2. Shape correspondence

In the second experiment, we reproduced the settings of Rodolà *et al.* [R*14] on the FAUST dataset, replacing their random forest with our ADD neural network model. Zeroth FAUST shape containing $n = 6890$ vertices was used as reference; for each point on the query shape, the output of mADD4 network representing the soft correspondence as a 6890-dimensional vector was converted into a point correspondence by taking the maximum. Training was done by minimizing the loss (22); training and test sets were as in the previous experiment.

Figures 9–10 (left) quantify the quality of the correspondence learned with our method on FAUST meshes and point clouds. For comparison, we show the performance of blended intrinsic maps (BIM) [KLF11], functional maps (FM) [O*12], and random forest (RF) [R*14]. Note that blended maps use orientation information and thus can distinguish bilaterally symmetric points (therefore, the dashed and solid black curves in Figure 9 coincide). On the other hand, random forests in [R*14] were learned on WKS input, which is ambiguous to symmetry; this explains the significant drop when passing from the symmetric to the asymmetric evaluation setting.

Figures 9–10 (right) depict the quality of the obtained correspondence. To visualize the colors transferred from the reference shape to the query shapes, we use the raw point-wise correspondence produced by our method as an input to the functional maps algorithm.

© 2016 The Author(s)

Computer Graphics Forum © 2016 The Eurographics Association and John Wiley & Sons Ltd.

8. Conclusions

We presented a novel class of shape feature descriptors based on anisotropic diffusion. The main properties of our descriptors are that they are learnable to carry out a specific task, allow to capture directional structures, are not ambiguous to symmetries, and can be efficiently computed on different shape representations such as meshes or point clouds. We showed that our descriptors can be successfully used to address problems such as intrinsic shape correspondence, achieving significantly better performance than previous methods, including the most recent intrinsic convolutional neural network models. In a broad sense, our approach follows the recent line of works on formulating intrinsic versions of successful machine learning schemes used in image analysis and computer vision. We hope that our work brings more attention to deep learning in geometric problems.

Limitations. Note that in this paper we have only used our deep learning architecture in the *supervised* setting. The main limitation of supervised learning stems from the need to provide a sufficiently large and rich training set, which in our case requires a set of shapes with known correspondence. In some scenarios, obtaining such correspondences may be difficult, especially if the shapes have no clear ground-truth correspondence. In practice, however, only a modest number of labelled shapes are actually required, as the number of positive and negative pairs grows quadratically with the training set size.

Acknowledgement

D. B., J. M., and M. B. are supported by the ERC Starting grant No. 307047. E.R. is supported by an Alexander von Humboldt Fellowship. The authors thankfully acknowledge Zorah Löhner and Giuseppe Patanè for their insightful comments.

References

- [A*05] ANGUELOV D., ET AL.: SCAPE: Shape completion and animation of people. *TOG* 24, 3 (2005), 408–416. 7
- [ARAC14] ANDREUX M., RODOLÀ E., AUBRY M., CREMERS D.:




Figure 9: Correspondence on FAUST meshes. Left: evaluation of the correspondence using the symmetric (solid) and asymmetric (dashed) Princeton protocol. Right: example of correspondence obtained using our method (similar colors encode corresponding points, where the leftmost shape is the reference).



Figure 10: Correspondence on FAUST point clouds. Left: evaluation of the correspondence using the symmetric (solid) and asymmetric (dashed) Princeton protocol. Right: example of correspondence obtained using our method (similar colors encode corresponding points, where the leftmost shape is the reference). Note how the correspondence is still accurate even when the underlying manifold has ambiguous topology (second shape from the left).

Anisotropic Laplace-Beltrami operators for shape analysis. In *Proc. NORDIA* (2014). 1, 2, 3, 7

[ASC11] AUBRY M., SCHLICKWEI U., CREMERS D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proc. ICCV* (2011). 1, 2, 3, 8

[B*94] BROMLEY J., ET AL.: Signature verification using a “Siamese” time delay neural network. In *Proc. NIPS* (1994). 6

[B*10] BERGSTRÄ J., ET AL.: Theano: a CPU and GPU math expression compiler. In *Proc. SciPy* (2010). 7

[B*15] BOSCAINI D., ET AL.: Learning class-specific descriptors for deformable shapes using localized spectral convolutional networks. *CGF* 34, 5 (2015), 13–23. 2, 4, 8

[BBGO11] BRONSTEIN A. M., BRONSTEIN M. M., GUIBAS L. J., OVSJANIKOV M.: Shape google: Geometric words and expressions for invariant shape retrieval. *TOG* 30, 1 (2011), 1–20. 1

[BC94] BERKMANN J., CAELLI T.: Computation of surface geometry and segmentation using covariance techniques. *TPAMI* 16, 11 (1994), 1114–1116. 7

[BEKB15] BOSCAINI D., EYNARD D., KOUROUNIS D., BRONSTEIN M. M.: Shape-from-operator: Recovering shapes from intrinsic operators. *CGF* 34, 2 (2015), 265–274. 1

[BES11] BOUCHER M., EVANS A., SIDDIQI K.: Anisotropic diffusion of tensor fields for fold shape analysis on surfaces. In *Information Processing in Medical Imaging* (2011), pp. 271–282. 1

[BK10] BRONSTEIN M. M., KOKKINOS I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In *Proc. CVPR* (2010). 4

[BRLB14] BOGO F., ROMERO J., LOPER M., BLACK M. J.: FAUST: Dataset and evaluation for 3D mesh registration. In *Proc. CVPR* (2014). 7

[BZSL14] BRUNA J., ZAREMBA W., SZLAM A., LECUN Y.: Spectral networks and locally connected networks on graphs. In *Proc. ICLR* (2014). 2

[CGGS12] CIRESAN D. C., GIUSTI A., GAMBARDELLA L. M., SCHMIDHUBER J.: Deep neural networks segment neuronal membranes in electron microscopy images. In *Proc. NIPS* (2012). 1

[CL06] COIFMAN R. R., LAFON S.: Diffusion maps. *Applied and Computational Harmonic Analysis* 21, 1 (2006), 5–30. 1

[COC14] CORMAN É., OVSJANIKOV M., CHAMBOLLE A.: Supervised descriptor learning for non-rigid shape matching. In *Proc. NORDIA* (2014). 1

[CRT04] CLARENZ U., RUMPF M., TELEA A.: Finite elements on point based surfaces. In *Proc. SPBG* (2004), pp. 201–211. 1, 7

- [CSM03] COHEN-STEINER D., MORVAN J.-M.: Restricted delaunay triangulations and normal cycle. In *Proc. SCG* (2003), pp. 312–321. 6
- [DMSB99] DESBRUN M., MEYER M., SCHRÖDER P., BARR A. H.: Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proc. Computer Graphics and Interactive Techniques* (1999). 1
- [Duf59] DUFFIN R. J.: Distributed and lumped networks. *Journal of Mathematics and Mechanics* 8, 5 (1959), 793–826. 7
- [Fuk80] FUKUSHIMA K.: Neocognitron: A self-organizing neural network model for a mechanism of pattern recognition unaffected by shift in position. *Biological Cybernetics* 36, 4 (1980), 193–202. 2
- [GBAL09] GEBAL K., BÆRENTZEN J. A., ANÆS H., LARSEN R.: Shape analysis using the auto diffusion function. *CGF* 28, 5 (2009), 1405–1413. 1, 3
- [HCL06] HADSELL R., CHOPRA S., LECUN Y.: Dimensionality reduction by learning an invariant mapping. In *Proc. CVPR* (2006). 6
- [KBBK12] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., KIMMEL R.: Photometric heat kernel signatures. In *Proc. SSVM* (2012). 1
- [KBLB12] KOKKINOS I., BRONSTEIN M. M., LITMAN R., BRONSTEIN A. M.: Intrinsic shape context descriptors for deformable shapes. In *Proc. CVPR* (2012). 2, 4
- [KF15] KARPATY A., FEI-FEI L.: Deep visual-semantic alignments for generating image descriptions. *Proc. CVPR* (2015). 1
- [KLF11] KIM V. G., LIPMAN Y., FUNKHOUSER T.: Blended intrinsic maps. *TOG* 30, 4 (2011), 1–12. 2, 8, 9
- [KRCH14] KANEZAKI A., RODOLÀ E., CREMERS D., HARADA T.: Learning similarities for rigid and non-rigid object detection. In *Proc. BMVC* (2014). 1
- [KS02] KIMMEL R., SOCHEN N.: Orientation diffusion or how to comb a porcupine. *Journal of Visual Communication and Image Representation* 13, 1 (2002), 238–248. 1
- [KSH12] KRIZHEVSKY A., SUTSKEVER I., HINTON G. E.: ImageNet classification with deep convolutional neural networks. In *Proc. NIPS* (2012). 1
- [KTT13] KIM K. I., TOMPKIN J., THEOBALT C.: Curvature-aware regularization on Riemannian submanifolds. In *Proc. ICCV* (2013). 1
- [L*89] LECUN Y., ET AL.: Backpropagation applied to handwritten zip code recognition. *Neural Comp.* 1, 4 (1989), 541–551. 2
- [LB14] LITMAN R., BRONSTEIN A. M.: Learning spectral descriptors for deformable shape correspondence. *TPAMI* 36, 1 (2014), 170–180. 1, 2, 3, 4, 8
- [LBB11] LITMAN R., BRONSTEIN A. M., BRONSTEIN M. M.: Diffusion-geometric maximally stable component detection in deformable shapes. *Computers and Graphics* 35, 3 (2011), 549–560. 1
- [LBBC14] LITMAN R., BRONSTEIN A. M., BRONSTEIN M. M., CASTELLANI U.: Supervised learning of bag-of-features shape descriptors using sparse coding. *CGF* 33, 5 (2014), 127–136. 1
- [Lév06] LÉVY B.: Laplace-Beltrami eigenfunctions towards an algorithm that “understands” geometry. In *Proc. SMI* (2006). 1
- [LRF10] LIPMAN Y., RUSTAMOV R. M., FUNKHOUSER T. A.: Biharmonic distance. *TOG* 29, 3 (2010), 27. 1
- [Mac49] MACNEAL R. H.: *The solution of partial differential equations by means of electrical networks*. PhD thesis, California Institute of Technology, 1949. 7
- [MBBV15] MASCI J., BOSCAINI D., BRONSTEIN M. M., VANDERGHEYNST P.: Geodesic convolutional neural networks on riemannian manifolds. In *Proc. 3dRR* (2015). 2, 4, 8
- [MDSB03] MEYER M., DESBRUN M., SCHRÖDER P., BARR A. H.: Discrete differential-geometry operators for triangulated 2-manifolds. *Visualization&Mathematics* (2003), 35–57. 7
- [MN03] MITRA N. J., NGUYEN A.: Estimating surface normals in noisy point cloud data. In *Proc. SCG* (2003), pp. 322–328. 7
- [MP49] MINAKSHISUNDARAM S., PLEIJEL A.: Some properties of the eigenfunctions of the Laplace-operator on Riemannian manifolds. *Canadian J. Math* 1, 242–256 (1949), 9. 3
- [O*12] OVSJANIKOV M., ET AL.: Functional maps: a flexible representation of maps between shapes. *TOG* 31, 4 (2012), 1–11. 1, 2, 9
- [OMMG10] OVSJANIKOV M., MÉRIGOT Q., MÉMOLI F., GUIBAS L.: One point isometric matching with the heat kernel. *CGF* 29, 5 (2010), 1555–1564. 1
- [OSG08] OVSJANIKOV M., SUN J., GUIBAS L.: Global intrinsic symmetries of shapes. *CGF* 27, 5 (2008), 1341–1348. 1
- [PM90] PERONA P., MALIK J.: Scale-space and edge detection using anisotropic diffusion. *TPAMI*, 12 (1990), 629–639. 1, 3
- [PP93] PINKALL U., POLTHIER K.: Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics* 2, 1 (1993), 15–36. 7
- [R*13] RUSTAMOV R. M., ET AL.: Map-based exploration of intrinsic shape differences and variability. *TOG* 32, 4 (2013), 72. 1
- [R*14] RODOLÀ E., ET AL.: Dense non-rigid shape correspondence using random forests. In *Proc. CVPR* (2014). 1, 2, 6, 9
- [R*16] RODOLÀ E., ET AL.: Partial functional correspondence. *CGF* 36 (2016), 222–236. 1
- [RCG08] RONG G., CAO Y., GUO X.: Spectral mesh deformation. *The Visual Computer* 24, 7–9 (2008), 787–796. 1
- [RRBC14] RODOLÀ E., ROTA BULÒ S., CREMERS D.: Robust region detection via consensus segmentation of deformable shapes. *CGF* 33, 5 (2014). 1
- [Rus07] RUSTAMOV R. M.: Laplace-Beltrami eigenfunctions for deformation invariant shape representation. In *Proc. SGP* (2007). 1
- [S*13] SHOTTON J., ET AL.: Real-time human pose recognition in parts from single depth images. *Comm. ACM* 56, 1 (2013), 116–124. 1
- [S*14] SIMO-SERRA E., ET AL.: Fracking deep convolutional image descriptors. *arXiv:1412.6537* (2014). 6
- [SMKL15] SU H., MAJI S., KALOGERAKIS E., LEARNER-MILLER E.: Multi-view convolutional neural networks for 3D shape recognition. In *Proc. ICCV* (2015). 1
- [SOCG10] SKRABA P., OVSJANIKOV M., CHAZAL F., GUIBAS L.: Persistence-based segmentation of deformable shapes. In *Proc. NORDIA* (2010). 1
- [SOG09] SUN J., OVSJANIKOV M., GUIBAS L. J.: A concise and provably informative multi-scale signature based on heat diffusion. *CGF* 28, 5 (2009), 1383–1392. 1, 2, 3, 8
- [SRV15] SHUMAN D. I., RICAUD B., VANDERGHEYNST P.: Vertex-frequency analysis on graphs. *Applied and Computational Harmonic Analysis* (2015). 2, 4
- [SZ14] SIMONYAN K., ZISSERMAN A.: Very deep convolutional networks for large-scale image recognition. *arXiv:1409.1556* (2014). 1
- [TWBO02] TASHDIZEN T., WHITAKER R., BURCHARD P., OSHER S.: Geometric surface smoothing via anisotropic diffusion of normals. In *Proc. Visualization* (2002). 1
- [W*14] WINDHEUSER T., ET AL.: Optimal intrinsic descriptors for non-rigid shape analysis. In *Proc. BMVC* (2014). 1
- [W*15] WU Z., ET AL.: 3D ShapeNets: A deep representation for volumetric shape modeling. In *Proc. CVPR* (2015). 1
- [Y*14] YOSHIYASU Y., ET AL.: Symmetry-aware nonrigid matching of incomplete 3D surfaces. In *Proc. CVPR* (2014). 1
- [Zei12] ZEILER M. D.: ADADELTA: An adaptive learning rate method. *arXiv:1212.5701* (2012). 7