

Finsler-Laplace-Beltrami Operators with Application to Shape Analysis

Simon Weber^{1,2*}

sim.weber@tum.de

Thomas Dages^{3*}

thomas.dages@cs.technion.ac.il

Maolin Gao^{1,2}

maolin.gao@tum.de

Daniel Cremers^{1,2}

cremers@tum.de

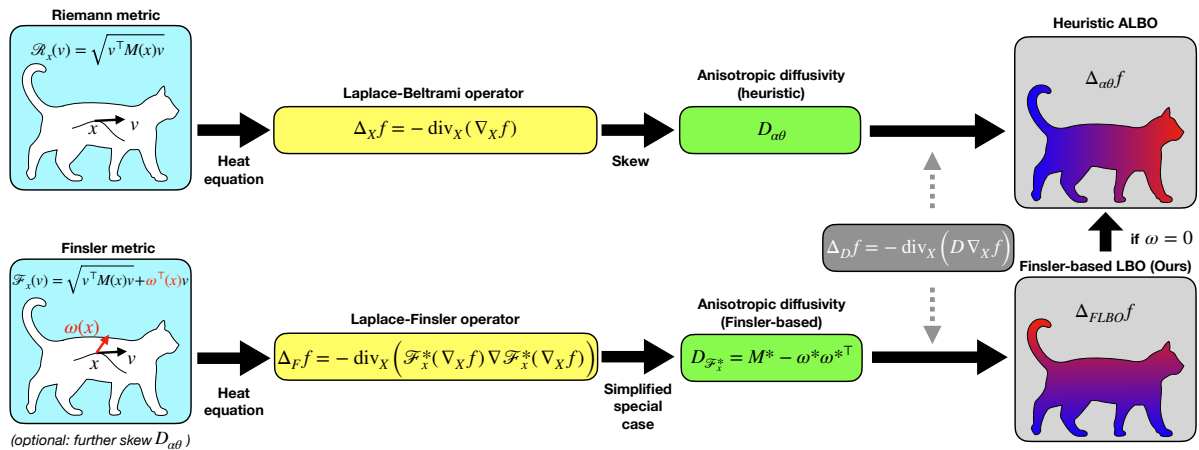


Figure 1. Simplified overview of the derivation of the Finsler-Laplace-Beltrami operator (FLBO), built from a Finsler metric. It generalizes the traditional heuristic anisotropic Laplace-Beltrami operators (ALBO) [6, 29] (top). The FLBO can directly replace the ALBOs in surface processing tasks like shape correspondence, notably by constructing shape-dependent anisotropic convolutions in their spectral domain.

Abstract

The Laplace-Beltrami operator (LBO) emerges from studying manifolds equipped with a Riemannian metric. It is often called the swiss army knife of geometry processing as it allows to capture intrinsic shape information and gives rise to heat diffusion, geodesic distances, and a multitude of shape descriptors. It also plays a central role in geometric deep learning. In this work, we explore Finsler manifolds as a generalization of Riemannian manifolds. We revisit the Finsler heat equation and derive a Finsler heat kernel and a Finsler-Laplace-Beltrami Operator (FLBO): a novel theoretically justified anisotropic Laplace-Beltrami operator (ALBO). In experimental evaluations we demonstrate that the proposed FLBO is a valuable alternative to the traditional Riemannian-based LBO and ALBOs for spatial filtering and shape correspondence estimation. We hope that the proposed Finsler heat kernel and the FLBO will

inspire further exploration of Finsler geometry in the computer vision community.

1. Introduction

1.1. Finsler versus Riemann

In his PhD thesis of 1918, Paul Finsler introduced the notion of a differentiable manifold equipped with a Minkowski norm on each tangent space [22]. It constitutes a generalization of Riemannian manifolds [39] that Elie Cartan later referred to as *Finsler manifolds* [12]. While Riemannian manifolds are omni-present in computer vision and machine learning, Finsler manifolds have remained largely unexplored.

Essentially, a Finsler metric is the same as a Riemannian one without the quadratic assumption. The consequence is that the metric not only depends on the location, but also on the tangential direction of motion. This implies that the length of a curve between two points A and B may change whether we are traversing it from A to B or the other way round, thus geodesic curves from A to B may differ from

* equal contribution.

¹Technical University of Munich

²Munich Center for Machine Learning

³Technion - Israel Institute of Technology

those from B to A . This means that on a Finsler manifold,




Figure 2. Due to non-uniform wind currents, the geodesic curve from A to B differs from the one from B to A . While not possible in Riemannian geometry, such asymmetries are characteristic of Finsler manifolds.

there is a built-in asymmetric anisotropy, which contrasts with the symmetric Riemannian case. In the Riemannian world, moving from an isotropic to a non isotropic metric allows more flexibility to capture interesting features on the manifolds. For instance, on a human hand shape, we may want a point on a finger to be close in some sense to all the points along the ring of the finger more than those along the length of the finger. Shifting from a Riemannian world to a Finsler world allows to increase the flexibility of anisotropy by also allowing non-symmetric distances. For many physical systems the Finsler metric is actually the natural choice of distance. For instance, a boat on a river or a plane in the wind might find that a path along the maximum current provides the shortest travel time whereas the return voyage might require moving away from the maximum current – see Fig. 2.

1.2. Contribution

The aim of this paper is to explore Finsler manifolds, study the Finsler heat diffusion, introduce a Finsler-Laplace-Beltrami Operator (FLBO), and demonstrate its applicability to shape analysis and geometric deep learning via the shape matching problem. In particular:

- We present a theoretical exploration of Finsler manifolds in a concise and self-contained form that is accessible to the computer vision and machine learning community.
- We revisit the heat diffusion on Finsler manifolds and derive a tractable equation.
- We generalize the traditional (anisotropic) Laplace-Beltrami operator to a Finsler-Laplace-Beltrami operator (FLBO).
- By leveraging the relationship between the diffusivity coefficient and the Riemannian metric, we propose an easy-to-implement discretization of the FLBO.
- We perform extensive evaluation of the proposed approach for shape correspondence based on the FLBO. We

highlight the benefits of this new operator in terms of accuracy on both full and partial shape datasets.

- We release our code as open source to facilitate further research: <https://github.com/tum-vision/flbo>.

2. Related work

The research domains most relevant to our work are geometric deep learning for shape correspondence and the use of the Finsler metric in computer vision.

Geometric deep learning for shape correspondence

Finding correspondences between shapes is a fundamental task in computer vision and a key component of many geometry processing applications, such as object recognition, shape reconstruction, texture transfer, or shape morphing. Initially solved with handcrafted approaches [16, 38, 44], the shape correspondence problem has recently benefited from the success of neural networks [23, 25, 27, 28, 42, 43] to learn features achieving tremendous results.

Initially seen in the Euclidean space, shapes are nowadays mostly considered as embedded manifolds. Intrinsic information, such as geodesic distances, are products of the heat diffusion paradigm, which has mostly been studied on manifolds equipped with a Riemannian metric. For such manifolds, heat diffusion is governed by the Poisson equation, and thus naturally the Laplace-Beltrami operator (LBO) plays a central role [10]. The LBO is traditionally isotropic on the manifold, but recent approaches have heuristically injected anisotropy into it, producing anisotropic LBOs (ALBO) [1]. Coupled with spatial filtering [29] ALBOs have shown impressive accuracy in shape correspondence.

We here focus on methods centred around anisotropic ALBOs. Nevertheless, many other types of approaches exist for full and partial shape correspondence. In particular, we point out the vast literature based on the functional map framework [3, 8, 9, 11, 20, 24, 30, 35, 36, 41].

In [31], the well-known heat kernel signature and wave kernel signature are generalized by leveraging the spectral properties of deformable shapes and learning task-specific spectral filters. This work is extended in [7] to direction-sensitive feature descriptors by using the eigendecomposition of anisotropic Laplace-Beltrami Operators (ALBO) [1]. Complementary to such spectral approaches, a line of works builds on spatial filtering methods. The idea is to construct patch operators directly on manifolds or graphs in the spatial domain. In [33], a convolution operator acting on vector fields without constraints on the filters themselves is defined. In [6], a local intrinsic representation of a function defined on a manifold is extracted with anisotropic heat kernels as spatial weighting functions. This idea is extended

in [29] by introducing an anisotropic convolution operator in the spectral domain. In this work, the eigendecomposition of multiple ALBOs proposed in [1] is filtered via Chebychev polynomials with learnable coefficients. This method achieves impressive results for shape correspondence in terms of accuracy.

Finsler computer vision

Computer vision on Finsler manifolds, a natural generalization of Riemannian manifolds, is a largely uncharted research area. In [37], robotics applications are pushed beyond Riemannian geometry and the Euler-Lagrange equation is reviewed from a Finsler point of view. A few works in image segmentation use a Finsler metric, either to add the orientation of an image as an extra spatial dimension [14], or to reformulate geodesic evolutions for region-based active contours [13, 15]. The heat equation on Finsler manifolds has been defined in mathematics [34]. In [46], a geodesic distance based on Finsler metrics is applied to contour detection. Nevertheless the chosen heat equation is theoretically unmotivated and is solely addressed from a heuristic point of view.

3. Background

We traditionally model a 3D shape as a two-dimensional compact Riemannian manifold (X, M) , with X a smooth manifold equipped with an inner product $\langle u, v \rangle_M = u^\top M(x)v$ for vectors u, v on the tangent space $T_x X$ at each point $x \in X$. We alternatively model the shape as a Finsler manifold (X, \mathcal{F}_x) equipped with a Randers metric $\mathcal{F}_x : T_x X \rightarrow \mathbb{R}$ at point x . This manifold can be viewed as a generalization of the Riemannian model, as the metric does not necessarily induce an inner product. We point out that we recover a Riemannian manifold if \mathcal{F}_x is symmetric. In this section we briefly introduce the Riemannian heat diffusion, the Randers metric and the Finsler heat diffusion.

3.1. Riemannian heat diffusion

Laplace-Beltrami operator. The Laplace-Beltrami Operator (LBO) on the Riemannian shape X is defined as:

$$\Delta_X f(x) = -\operatorname{div}_X(\nabla_X f(x)), \quad (1)$$

with $\nabla_X f$ and $\operatorname{div}_X f$ being the intrinsic gradient and divergence of $f(x) \in L^2(X)$. The LBO is a positive semi-definite operator with a real eigen-decomposition (λ_k, ϕ_k) :

$$\Delta_X \phi_k(x) = \lambda_k \phi_k(x), \quad (2)$$

with countable eigenvalues $0 = \lambda_0 \leq \lambda_1 \leq \dots$ growing to infinity. The corresponding eigenfunctions $\phi_0, \phi_1 \dots$ are orthonormal for the standard inner product on $L^2(X)$:

$$\langle \phi_i, \phi_j \rangle = \delta_{ij}, \quad (3)$$

and form an orthonormal basis for $L^2(X)$. It follows that any function $f \in L^2(x)$ can be expressed as:

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

and $\hat{f}(\lambda_k) = \langle \phi_k, f \rangle$ are the coefficients of the manifold Fourier transform of f . The Convolution Theorem on manifolds is given by:

$$(f * g)(x) = \sum_{k \geq 0} \hat{f}(\lambda_k) \hat{g}(\lambda_k) \phi_k(x), \quad (5)$$

where $f * g$ is the convolution of signal f with filter kernel g and $\hat{f}(\lambda_k) \hat{g}(\lambda_k)$ is the spectral filtering of f by g .

Heat diffusion. The LBO governs the diffusion equation:

$$\frac{\partial f(x, t)}{\partial t} = -\Delta_X f(x, t), \quad (6)$$

where $f(x, t)$ is the temperature at point x at time t . Given some initial heat distribution $f_0(x) = f(x, 0)$ the solution of Eq. (6) at time t is given by the convolution:

$$f(x, t) = \int_X f_0(\xi) h_t(x, \xi) d\xi, \quad (7)$$

with the heat kernel $h_t(x, \xi)$. In the spectral domain it is expressed as:

$$h_t(x, \xi) = \sum_{k \geq 0} e^{-t\lambda_k} \phi_k(x) \phi_k(\xi). \quad (8)$$

Anisotropic LBO. Extending the traditional LBO presented above, an anisotropic LBO (ALBO) is created by changing the diffusion speed along chosen directions on the surface, typically the principal directions of curvature [1]:

$$\Delta_\alpha f(x) = -\operatorname{div}_X(D_\alpha(x) \nabla_X f(x)), \quad (9)$$

where $D_\alpha(x)$ is a thermal conductivity tensor acting on the intrinsic gradient direction in the tangent plane and α is a scalar expressing the level of anisotropy. The ALBO can be extended to multiple anisotropic angles θ [6]:

$$\Delta_{\alpha\theta} f(x) = -\operatorname{div}_X(D_{\alpha\theta}(x) \nabla_X f(x)), \quad (10)$$

by defining the thermal conductivity tensor:

$$D_{\alpha\theta}(x) = R_\theta D_\alpha(x) R_\theta^\top, \quad (11)$$

where R_θ is a rotation in the tangent plane (around surface normals) with angle θ .

Anisotropic convolution operator Similarly, Eq. (5) is extended by introducing the anisotropic convolution operator [29]:

$$(f * g)_{\alpha\theta}(x) = \int_{y \in X} f(y)g_{\alpha\theta,x}(y)dy, \quad (12)$$

where $g_{\alpha\theta,x}(y)$ is an oriented kernel function centered at point x from the spectral domain

$$g_{\alpha\theta,x}(y) = \sum_{k \geq 0} \hat{g}(\lambda_{\alpha\theta,k})\phi_{\alpha\theta,k}(x)\phi_{\alpha\theta,k}(y). \quad (13)$$

A straightforward derivation shows that:

$$(f * g)_{\alpha\theta}(x) = \sum_{k \geq 0} \hat{g}(\lambda_{\alpha\theta,k})\hat{f}(\lambda_{\alpha\theta,k})\phi_{\alpha\theta,k}(x), \quad (14)$$

where $\hat{f}(\lambda_{\alpha\theta,k}) = \langle \phi_{\alpha\theta,k}, f \rangle$ are the coefficients of the anisotropic Fourier transform of f . To extract features along all directions, the anisotropic convolution operators are summed and then:

$$(f * g)_{\alpha}(x) = \int_0^{2\pi} (f * g)_{\alpha\theta}(x)d\theta. \quad (15)$$

The design of $\hat{g}(\lambda)$ is crucial when defining the anisotropic convolution operators. Following [18], Chebychev polynomials are adopted to define polynomial filters $\hat{g}(\lambda)$ in [29].

3.2. Randers metric

Riemannian manifolds are equipped with a quadratic metric $\mathcal{R}_x(v) = \|v\|_{M(x)}$ that naturally induces a positive-definite inner product on the tangent space $T_x M$ at each point x . Finsler geometry is often said to be the same as Riemannian geometry without the quadratic constraint. A Finsler metric thus no longer induces an inner-product and is not required to be symmetric. Formally, a Finsler metric is given by the following properties.

Definition 1 (Finsler metric). *A Finsler metric $\mathcal{F}_x : T_x X \rightarrow \mathbb{R}_+$ at point x is a smooth metric satisfying the triangle inequality, positive homogeneity, and positive definiteness axioms:*

$$\begin{cases} \mathcal{F}_x(v + v') \leq \mathcal{F}_x(v) + \mathcal{F}_x(v'), \quad \forall v, v' \in T_x X & (16) \\ \mathcal{F}_x(\gamma v) = \gamma \mathcal{F}_x(v), \quad \forall (\gamma, v) \in \mathbb{R}_+ \times T_x X & (17) \\ \mathcal{F}_x(v) = 0 \iff v = 0. & (18) \end{cases}$$

We consider Randers metrics, one of the simplest type of Finsler metrics generalizing the Riemannian one.

Definition 2 (Randers metric). *Let $M : X \rightarrow S_d^{++}$ be a tensor field of symmetric positive definite matrices and a vector field $\omega : X \rightarrow \mathbb{R}^d$ such that $\|\omega(x)\|_{M(x)^{-1}} < 1$ for all $x \in X$. The Randers metric associated to (M, ω) is given by*

$$\mathcal{F}_x(v) = \|v\|_{M(x)} + \langle \omega(x), v \rangle, \quad \forall x \in X, v \in T_x X. \quad (19)$$

The Randers metric \mathcal{F} is the sum of the Riemannian metric associated to M and an asymmetric linear part governed by ω . Notably when $\omega \equiv 0$, \mathcal{F} is a traditional Riemannian metric. The condition $\|\omega(x)\|_{M(x)^{-1}} < 1$ ensures the positivity of the metric (see supplementary material).

From any metric we can construct a new metric associated to it, its dual, which plays a central role in differential geometry.

Definition 3 (Dual Finsler metric). *The dual metric of a Finsler metric \mathcal{F}_x is the metric $\mathcal{F}_x^* : T_x X \rightarrow \mathbb{R}_+$ defined by*

$$\mathcal{F}_x^*(v) = \max\{\langle v, b \rangle; b \in \mathbb{R}^d, \mathcal{F}_x(b) \leq 1\}. \quad (20)$$

When considering only Randers metrics, a subtype of Finsler metrics, the dual metric becomes explicit.

Lemma 1 (Dual Randers metric). *Let \mathcal{F}_x be a Randers metric given by (M, ω) . The dual metric of \mathcal{F}_x is also a Randers metric \mathcal{F}_x^* associated to (M^*, ω^*) that satisfies*

$$\begin{pmatrix} \alpha M^* & \omega^* \\ \omega^{*\top} & \frac{1}{\alpha} \end{pmatrix} = \begin{pmatrix} M & \omega \\ \omega^\top & 1 \end{pmatrix}^{-1}, \quad (21)$$

where $\alpha = 1 - \langle \omega, M^{-1}\omega \rangle > 0$.

A proof is provided in the supplementary material.

From Eq. (21) we can derive:

$$M^* = \frac{1}{\alpha^2} \left(\alpha M^{-1} + (M^{-1}\omega)(M^{-1}\omega)^\top \right), \quad (22)$$

and

$$\omega^* = -\frac{1}{\alpha} M^{-1}\omega, \quad (23)$$

It is simple to show that $\|\omega^*\|_{M^*(x)^{-1}} < 1$ (see supplementary material). The dual Randers metric \mathcal{F}_x^* is the Randers metric on the manifold X associated to the new Riemannian and linear components (M^*, ω^*) :

$$\mathcal{F}_x^*(v) = \|v\|_{M^*(x)} + \langle \omega^*(x), v \rangle. \quad (24)$$

3.3. Finsler heat equation

The classical heat equation in the Riemannian case is the gradient flow of the Dirichlet energy, and by analogy we can derive a Finsler heat equation as the gradient descent of the dual energy $\frac{1}{2}\mathcal{F}_x^*(v)^2$, in the following way [5, 34].

Definition 4 (Finsler heat equation). *Given a manifold X equipped with a Finsler metric \mathcal{F}_x , the heat equation on this Finsler manifold for smooth vector fields $u(x, t)$ with $(x, t) \in X \times \mathbb{R}_+$ is given by the gradient descent of the dual energy*

$$\partial_t u = \operatorname{div}(\mathcal{F}_x^*(\nabla u)\nabla \mathcal{F}_x^*(\nabla u)). \quad (25)$$

4. Finsler heat kernel

We now analyse the Finsler heat equation to derive a novel Laplacian operator for shape matching. Our plan is the following:

- Analyse the heat equation on a Finsler manifold equipped with a Randers metric,
- From a specific case where the solution is explicit, we exhibit a simplified equation behaving like a regular heat equation with an external heat source,
- The solution to this new problem is explicit and governed by the homogeneous term, which is an anisotropic heat diffusion with diffusivity involving the Randers metric,
- The homogeneous solution is given by a convolution with an anisotropic kernel depending on the Randers metric,
- This heat kernel allows the construction of a new Finsler-Laplace-Beltrami operator (FLBO) that benefits from all the properties of ALBOs.

4.1. Finsler heat diffusion

The heat equation is well-known to be related to distances and geodesics in short time, and thus it is often analyzed when $t \rightarrow 0$. We look at the following specific case in short time. To provide a tractable solution, we also here assume that we slightly depart from Riemannian geometry by taking small $\|\omega\|$.

Proposition 1 (Simplified Randers heat equation - distance trick). *If the initial solution $u_0 = u(x, 0)$ of the heat equation satisfies $\mathcal{F}_x^*(\nabla u_0) = 1$, then in short time the Finsler heat equation on a Randers metric associated to (M, ω) with small $\|\omega\|$ simplifies to*

$$\partial_t u = \operatorname{div} \left((M^* - \omega^* \omega^{*\top}) \nabla u \right) + \operatorname{div}(\omega^*). \quad (26)$$

Proof. We provide the proof in the supplementary material. \square

The initial condition $\mathcal{F}_x^*(\nabla u_0) = 1$ holds when $-u_0$ is a distance function, as the Finsler Eikonal equation is given by $\mathcal{F}_x^*(-\nabla u) = 1$ [32]. In practice, the solution to Eq. (26) is not our primary concern. We only use this particular initialisation to exhibit a tractable solution that allows us to define a family of Finsler-based LBOs. We name the coefficient $D_{\mathcal{F}_x^*} = M^* - \omega^* \omega^{*\top}$ as the Finsler diffusivity. We can interpret Eq. (26) as a homogeneous heat equation:

$$\partial_t u(x, t) = \operatorname{div}(D_{\mathcal{F}_x^*} \nabla u(x, t)), \quad (27)$$

to which we add an external heat source $\operatorname{div}(\omega^*(x))$. Given some initial heat condition $u_0(x) = u(x, 0)$, the solution of Eq. (26) is the sum of the solution of Eq. (27),

$$u(x, t) = \int_X u_0(\xi) h_t(x, \xi) d\xi, \quad (28)$$

where h_t is the heat kernel for Eq. (27), and of any particular solution, such as

$$u^d(x, t) = \int_X \operatorname{div}(\omega^*(\xi)) \hat{h}_t(x, \xi) d\xi, \quad (29)$$

where \hat{h} is the average over time of the heat kernel h_t :

$$\hat{h}_t(x, \xi) = \int_{s=0}^t h_{(t-s)}(x, \xi) ds. \quad (30)$$

4.2. Finsler-Laplace-Beltrami operator

In Riemannian geometry, the heat kernel is by definition given by the convolution kernel of the solution of the homogeneous equation. An external source introduces a convolution with the time average of the heat kernel. By analogy, from the homogeneous Eq. (27) we define the Finsler-Laplace-Beltrami operator (FLBO) as:

$$\Delta_{FLBO} u(x) = -\operatorname{div}_X(D_{\mathcal{F}_x^*} \nabla_X u(x)). \quad (31)$$

The FLBO is a special case of an ALBO on a Riemannian manifold, that was derived after analyzing the heat equation when the manifold is equipped with a Finsler metric instead. Unlike previous ALBOs, it is theoretically justified rather than empirically designed based solely on heuristics. Since the FLBO in an ALBO, it inherits all of its properties, such as the spectral decomposition of a compact positive semi-definite operator

$$\Delta_{FLBO} \phi_k^{FLBO}(x) = \lambda_k^{FLBO} \phi_k^{FLBO}(x). \quad (32)$$

The FLBO also differs from the Laplace-Finsler operator (LF) deriving from Eq. (25):

$$\Delta_F u(x) = -\operatorname{div}(\mathcal{F}_x^*(\nabla u) \nabla \mathcal{F}_x^*(\nabla u)), \quad (33)$$

In particular, the LF operator acts on a Finsler manifold and not a Riemannian one like the FLBO. Thus, it cannot be used as a direct replacement of LBOs in most applications.

In the previous section we have shown that the solution of the simplified Randers heat diffusion equation exhibits the Finsler heat kernel $h_t(x, \xi)$ associated with the homogeneous Finsler heat diffusion equation.

Proposition 2. *In the spectral domain, the Finsler heat kernel can be expressed as*

$$h_t^{FLBO}(x, \xi) = \sum_{k \geq 0} e^{-t \lambda_k^{FLBO}} \phi_k^{FLBO}(x) \phi_k^{FLBO}(\xi). \quad (34)$$

We have unified the Finsler heat diffusion equation and the anisotropic Laplace-Beltrami operator formulation. The link comes from the Finsler heat kernel that was exhibited in the explicit solution of the simplified equation Eq. (26) in a special case. We point out that the asymmetric component ω of the Randers metric influences the diffusivity coefficient and then the eigendecomposition of the Finsler heat kernel.

5. Experiments

5.1. Discretization

We present a possible discretization of our FLBO given a manifold discretized by sampling n vertices. As the FLBO is a positive semi-definite operator, and by analogy with the LBO it can be approximated by an $n \times n$ sparse matrix $L_{FLBO} = -S_{FLBO}^{-1}W_{FLBO}$ with the mass matrix S_{FLBO} and the stiffness matrix W_{FLBO} . Our FLBO is not restricted to a particular discrete surface representation. Following [6, 29], we focus on triangular meshes to discretize the Finsler metrics, but representations for other discretizations can be easily deduced. As in [6, 29], we introduce various angles θ , each giving a Finsler metric $(M_\theta, \omega_\theta)$, to produce convolution operators from FLBOs sensitive to θ , and then sum them up to get a final Finsler-based convolution operator incorporating features from all directions. We use the same notations as in [6, 29].

Our manifold is discretized as a triangular mesh with vertices V , edges E , and faces F . For each triangle $ijk \in F$, we define $U_{ijk} = (\hat{u}_M, \hat{u}_m, \hat{n})$ as an orthonormal reference frame associated to the face unit normal \hat{n} and the directions of principal curvature $\hat{u}_M, \hat{u}_m \in \mathbb{R}^3$. We denote $\hat{e}_{ij} \in \mathbb{R}^3$ the unit vector along the embedding of the edge $(i, j) \in E$ pointing from i to j , and α_{ij} and β_{ij} the angles at k and h of the adjacent triangles $ijk \in F$ and $ijh \in F$ respectively. We write $R_\theta \in \mathbb{R}^3$ the rotation matrix of angle θ around the third basis vector (to be identified as \hat{n}). The notations match those of Fig. 3.




Figure 3. Notations for unit vectors and angles on the triangular mesh. The figure mimics the illustration in [6].

In all our experiments, the mass matrix S_{FLBO} is computed as a diagonal matrix with diagonal element i equal to one third of the areas of faces having i as a vertex [19].

Our chosen Riemannian metric is also anisotropic using ideas from the ALBO community. When $\omega \equiv 0$ we obtain the traditional ALBO [29]. The diffusivity $D_\alpha(x)$ of ALBOs is usually created by introducing a scaling factor $\alpha > 0$ into one of its eigenvalues by various heuristic formulae to break the isotropy along eigendirections [6, 29]. We firstly build on the strategy of [29] to create the uniform anisotropic Riemannian diffusivity

$$D_\alpha = \begin{pmatrix} \frac{1}{1+\alpha} & \\ & 1 \end{pmatrix}, \quad (35)$$

with $\alpha = 10$ in our experiments as in [29]. For each angle θ , we rotate this diffusivity to create the oriented diffusivity matrix $H_{\alpha\theta}$, also called shear matrix,

$$H_{\alpha\theta} = R_\theta U_{ijk} D_\alpha U_{ijk}^\top R_\theta^\top. \quad (36)$$

We can now define our anisotropic Riemannian metric at angle θ . As we want our construction to generalise the Riemannian case $\omega \equiv 0$, and since $D_{\mathcal{F}_x^*}$ equals $M^{-1}(x)$ when $\omega \equiv 0$, we choose

$$M_{\alpha\theta} = H_{\alpha\theta}^{-1}. \quad (37)$$

Note that our scheme recovers the isotropic Riemannian metric when $\alpha = 0$ as the shear matrix is then the identity. For each angle θ , we introduce the Randers anisotropic component rotated by θ from the direction of maximal curvature

$$\omega_\theta = \tau R_\theta \hat{u}_M, \quad (38)$$

where τ is a hyperparameter to ensure $\|\omega\|_{M_{\alpha\theta}^{-1}} < 1$. We have thus constructed the Randers metric $(M_{\alpha\theta}, \omega_\theta)$, from which we can easily compute its dual $(M_{\alpha\theta}^*, \omega_\theta^*)$ and its Finsler diffusivity $D_{\mathcal{F}_x^*}^{\alpha\theta}$. We then use the generalisation of the classical cotangent weight scheme for anisotropic diffusivities [6] to compute the weights of $W_{FLBO}^{\alpha\theta}$:

$$w_{ij} = \begin{cases} \frac{1}{2} \left(\frac{\langle \hat{e}_{kj}, \hat{e}_{ki} \rangle_{D_{\mathcal{F}_x^*}^{\alpha\theta}}}{\sin \alpha_{ij}} + \frac{\langle \hat{e}_{hj}, \hat{e}_{hi} \rangle_{D_{\mathcal{F}_x^*}^{\alpha\theta}}}{\sin \beta_{ij}} \right) & (i, j) \in E, \\ -\sum_{k \neq i} w_{ik} & i = j, \\ 0 & \text{else.} \end{cases} \quad (39)$$

The FLBO at angle θ is finally given by $\Delta_{FLBO}^{\alpha\theta} = -S_{FLBO}^{-1}W_{FLBO}^{\alpha\theta}$. Note that if we had taken $\omega_\theta \equiv 0$, then we would have recovered the ALBO from [29]. See Fig. 1 for an illustrative summary of how we derive the FLBO.

Finsler-based anisotropic convolution operator Since our FLBOs are ALBOs, we can use the method of [29] to create θ -oriented convolution kernels $g_{\alpha\theta, x}$ according to Eq. (13) by working in the spectral domain of $\Delta_{FLBO}^{\alpha\theta}$ and taking $\phi_{\alpha\theta, 0}, \phi_{\alpha\theta, 1} \dots$ to be the eigenvectors of $\Delta_{FLBO}^{\alpha\theta}$. Convolution $(f * g)_{\alpha\theta}$ of any function f with this kernel is then computed in the spectral domain (Eq. (14)) and then oriented convolutions are summed to get the final convolution operator $(f * g)_\alpha$ (Eq. (15)). In our experiments, we use eight uniformly spaced samples of $\theta \in [0, \pi)$ as in [29]. Following [29], we encode the spectral coefficients $\hat{g}(\lambda_{\alpha\theta, k})$ of $g_{\alpha\theta, x}$ using Chebychev polynomials

$$\hat{g}(\lambda_{\alpha\theta, k}) = \sum_{s=0}^{S-1} c_{\alpha\theta, s} T_s(\lambda_{\alpha\theta, k}), \quad (40)$$

where $S = 16$ in our tests as in [29], T_s is the Chebychev polynomial of the first kind of order s , i.e. $T_s(\cos \gamma) =$

$\cos(s\gamma)$ for any γ , and $c_{\alpha\theta,s} \in \mathbb{R}$ are coefficients to be learned.

For visualisation purposes, we provide in the supplementary a plot of convolution filters $g_{\alpha\theta,x}$ for various θ when choosing the spectral coefficients of the kernel to be equal to one of the Chebychev polynomials, similarly to [29].

5.2. An Application: Shape matching

As a concrete example, we use our FLBO for shape matching, which is a common shape analysis application where LBOs play a central role. We compare the results to those based on traditional Riemannian (anisotropic) LBOs, namely FieldConv [33] and ACSCNN [29]. Note that ACSCNN coincides with our method when $\tau = 0$.

Datasets and evaluation. We work on four publicly available datasets: FAUST Remeshed [4, 38], SCAPE Remeshed [2, 38], and SHREC’16 Partial Cuts [17] and Holes [17]. We follow the Princeton benchmark protocol [26]. We provide the percentage of matches that are at most r -geodesically distant from the groundtruth correspondence on the reference shape, with $r \in [0, 1]$.

Setting. Following [29], our chosen architecture is SplineCNN [21]. We train each dataset for 100 epochs with Adam optimizer and we set the learning rate to 0.001. We take as input SHOT descriptors [40]. In particular, we use the open-source implementation of ACSCNN¹. The preprocessing step to get the Finsler-based filters is done with our custom implementation.

5.3. Full mesh correspondence.

FAUST (Remeshed) dataset contains 100 human shapes corresponding to 10 different human beings remeshed with the LRVD [45] to generate more challenging shapes with different mesh connectivity [4, 38]. It does not have one-to-one correspondence between shapes and they do not share the same number of vertices. As suggested in [20], we use the first 80 shapes as the training set and we test on the remaining 20 ones.

SCAPE (Remeshed) dataset contains 71 human shapes corresponding to the same human being in different positions. Each shape is remeshed as in FAUST Remeshed [2, 38]. The first 51 shapes form the training set and we test on the other 20 shapes.

5.4. Partial correspondence.

SHREC’16 Partial Cuts and **SHREC’16 Partial Holes** datasets contain respectively 135 and 90 shapes of human

Method		Accuracy ($r = 0$)	Accuracy ($r = 0.01$)
FieldConv [33]		53.39	75.01
ACSCNN [29]		61.95	82.46
ours	$\tau = 0.1$	62.61	82.90
	$\tau = 0.2$	62.58	82.90
	$\tau = 0.3$	62.67	82.94
	$\tau = 0.4$	62.27	82.66
	$\tau = 0.5$	62.38	82.68

Table 1. Accuracy on FAUST Remeshed.

Method		Accuracy ($r = 0$)	Accuracy ($r = 0.01$)
FieldConv [33]		37.41	61.37
ACSCNN [29]		46.13	71.70
ours	$\tau = 0.1$	46.94	72.44
	$\tau = 0.2$	46.27	71.91
	$\tau = 0.3$	45.91	71.92
	$\tau = 0.4$	45.82	71.44
	$\tau = 0.5$	45.26	71.27

Table 2. Accuracy on SCAPE Remeshed.

and animals divided in 8 categories. The shapes are deformed following two kinds of partiality, respectively cuts and holes. For each category in Cuts, we use the first 10 shapes as a training dataset and we test on the last 5 shapes. For each category in Holes, we use the first 6 shapes as a training dataset and we test on the last 4 shapes.

5.5. Results

We present quantitative results in Tabs. 1 to 4 for FAUST Remeshed, SCAPE Remeshed, SHREC’16 Partial Cuts, and SHREC’16 Partial Holes respectively. For the first two datasets we provide the accuracy within $r \in \{0, 0.01\}$ geodesic radius of the groundtruth, whereas we only display the performance for $r = 0$ in the last two datasets due to space constraints as each of these datasets are decomposed into 8 subdatasets. Clearly, FieldConv struggles compared to the more advanced ACSCNN. Our method systematically outperforms, or is on par with, ACSCNN for a default choice of $\tau = 0.1$. This demonstrates that the FLBO can be used in complement to advanced approaches to boost performance, requiring only a small additional preprocessing overhead for computing the metric and the FLBO, which is negligible compared to neural network training times. We test our method with other values of τ and find that the results remain consistent, proving that our approach is not highly sensitive to finely tuning this hyperparameter. In the literature, providing the results as a plot with respect to all $r \in [0, 1]$ is sometimes performed. However, since our results are close to those of ACSCNN, the curves tend to superimpose and do not provide any additional insight, so we do not provide them here.

¹<https://github.com/GCVGroup/ACSCNN>

Method	Cat	Centaur	David	Dog	Horse	Michael	Victoria	Wolf	
FieldConv [33]	0.054	1.18	0	0.094	1.93	0.33	0.010	45.81	
ACSCNN [29]	38.10	74.99	23.78	59.87	44.37	12.82	31.23	93.75	
ours	$\tau = 0.1$	40.13	75.41	22.34	59.31	42.72	16.04	27.95	94.02
	$\tau = 0.2$	37.72	72.83	22.48	60.02	43.17	14.30	29.72	94.14
	$\tau = 0.3$	40.26	73.40	21.59	57.61	43.49	13.15	31.10	94.17
	$\tau = 0.4$	39.63	72.29	21.51	57.12	44.83	13.55	29.35	93.33
	$\tau = 0.5$	40.94	73.16	20.49	60.17	43.85	15.04	28.46	94.06

Table 3. Accuracy on SHREC’16 Partial Cuts.

Method	Cat	Centaur	David	Dog	Horse	Michael	Victoria	Wolf	
FieldConv [33]	N.A.	0.038	N.A.	0.11	N.A.	0.11	0.049	0.39	
ACSCNN [29]	19.98	32.80	16.09	41.70	46.85	12.29	11.32	84.53	
ours	$\tau = 0.1$	19.95	35.85	17.96	42.68	47.87	12.00	11.70	84.64
	$\tau = 0.2$	17.99	33.10	16.40	42.84	49.77	13.90	10.02	85.23
	$\tau = 0.3$	21.27	35.48	18.43	43.29	45.07	11.31	11.52	85.74
	$\tau = 0.4$	20.08	32.54	16.41	42.98	45.71	13.16	11.41	85.72
	$\tau = 0.5$	18.62	35.04	16.28	42.68	45.38	11.89	10.19	84.23

Table 4. Accuracy on SHREC16’ Partial Holes. N.A. indicates a failure to process the dataset.

We plot some qualitative correspondence results in Fig. 4. We find that our method is able to provide a convincing mapping between shapes undergoing various highly non-rigid deformations. It is also able to handle missing parts such as cuts and holes, although performance naturally deteriorates as the challenge is significantly harder.

6. Conclusion

We revisited Finsler manifolds, a generalization of Riemannian manifolds that naturally allow for anisotropies. By exploring their heat diffusion equation we derived a novel Finsler heat kernel and a Finsler-Laplace-Beltrami operator (FLBO). This FLBO generalizes traditional heuristic anisotropic LBOs while being theoretically motivated. It can be used in place of LBOs or ALBOs and is compatible with various advanced techniques such as spatial filtering or geometric deep learning. We tested the pertinence of the FLBO in shape correspondence experiments and found that our approach can easily complement state-of-the-art geometric deep learning methods using ALBOs. We hope that the concise and self-contained review of Finsler geometry and the derivation of a tractable Finsler heat equation and FLBO will encourage further exploration of the Finsler geometry in the field of computer vision.

Limitations and future works. We have used some strong assumptions to derive the Finsler heat equation. While we acknowledge this limitation, it is necessary, in our work, to exhibit some tractable operators. However, we believe that a further theoretical focus on this stage is an exciting and promising research direction. Our experiments are also built on a straightforward discretization. Although it is not a major drawback, as our goal is to show that traditional LBOs can be generalized to Finsler-based LBOs, an extension of our example to other choices of (M, ω) would also be crucial. Finally, beyond our simple application to shape



Figure 4. Visual shape correspondence results on the FAUST Remeshed (row 1), SCAPE remeshed (row 2), and SHREC’16 Partial Cuts (columns 1 and 2 in rows 3 to 5), and SHREC’16 Partial Holes (columns 1 and 3 in rows 3 to 5) datasets. The source shape is on the left, on which we perform dense correspondence estimation on the shapes to the right. See the supplementary material for more visual results on other shapes.

matching, we do believe that the importance of Finsler manifolds in computer vision is largely uncharted and requires further stimulating works.

Acknowledgement. This work was supported by the ERC Advanced Grant SIMULACRON.

References

- [1] Mathieu Andreux, Emanuele Rodola, Mathieu Aubry, and Daniel Cremers. Anisotropic Laplace-Beltrami operators for shape analysis. In *Sixth Workshop on Non-Rigid Shape Analysis and Deformable Image Alignment (NORDIA)*, 2014. 2, 3
- [2] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. In *ACM Transactions on Graphics (TOG)*, pages 408–416. ACM, 2005. 7
- [3] Souhaib Attaiki, Gautam Pai, and Maks Ovsjanikov. Dpfm: Deep partial functional maps. In *2021 International Conference on 3D Vision (3DV)*, pages 175–185. IEEE, 2021. 2
- [4] Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. Faust: Dataset and evaluation for 3d mesh registration. In *CVPR*, pages 3794–3801, 2014. 7
- [5] J Frederic Bonnans, Guillaume Bonnet, and Jean-Marie Mirebeau. A linear finite-difference scheme for approximating Randers distances on cartesian grids. *ESAIM: Control, Optimisation and Calculus of Variations*, 28:45, 2022. 4
- [6] David Boscaini, Jonathan Masci, Emanuele Rodola, and Michael Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in Neural Information Processing Systems*, pages 3189–3197, 2016. 1, 2, 3, 6
- [7] David Boscaini, Jonathan Masci, Emanuele Rodola, Michael Bronstein, and Daniel Cremers. Anisotropic diffusion descriptors. In *Computer Graphics Forum*, pages 431–441, 2016. 2
- [8] Amit Bracha, Oshri Halim, and Ron Kimmel. Shape Correspondence by Aligning Scale-invariant LBO Eigenfunctions. In *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2020. 2
- [9] Amit Bracha, Thomas Dagès, and Ron Kimmel. On partial shape correspondence and functional maps. *arXiv preprint arXiv:2310.14692*, 2023. 2
- [10] Alexander M Bronstein, Michael M Bronstein, and Ron Kimmel. *Numerical geometry of non-rigid shapes*. Springer Science & Business Media, 2008. 2
- [11] Dongliang Cao, Paul Roetzer, and Florian Bernard. Unsupervised learning of robust spectral shape matching. *arXiv preprint arXiv:2304.14419*, 2023. 2
- [12] Élie Cartan. *Les espaces métriques fondés sur la notion d'aire*. Hermann, 1933. 1
- [13] Da Chen, Jean-Marie Mirebeau, and Laurent D. Cohen. Finsler geodesics evolution model for region based active contours. In *Proc. of the British Machine Vision Conference (BMVC)*, 2016. 3
- [14] Da Chen, Jean-Marie Mirebeau, and Laurent D. Cohen. A new Finsler minimal path model with curvature penalization for image segmentation and closed contour detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 355–363, 2016. 3
- [15] Da Chen, Jean-Marie Mirebeau, and Laurent D. Cohen. Global minimum for a Finsler elastica minimal path approach. pages 458–483, 2017. 3
- [16] Qifeng Chen and Vladlen Koltun. Robust nonrigid registration by convex optimization. In *Proceedings of the IEEE International Conference on Computer Vision*, pages 2039–2047, 2015. 2
- [17] Luca Cosmo, Emanuele Rodola, Michael M. Bronstein, Andrea Torsello, Daniel Cremers, Y Sahillioglu, et al. Shrec’16: Partial matching of deformable shapes. *Proc. 3DOR*, 2(9): 12, 2016. 7
- [18] Michaël Defferrard, Xavier Bresson, and Pierre Vandergheynst. Convolutional neural networks on graphs with fast localized spectral filtering. In *Advances in Neural Information Processing Systems (NeurIPS)*, pages 3844–3852, 2016. 4
- [19] Mathieu Desbrun, Mark Meyer, Peter Schröder, and Alan H. Barr. Implicit fairing of irregular meshes using diffusion and curvature flow. In *Proceedings of the ACM Siggraph’99*, pages 317–24, 1999. 6
- [20] Nicolas Donati, Abhishek Sharma, and Maks Ovsjanikov. Deep geometric functional maps: Robust feature learning for shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8592–8601, 2020. 2, 7
- [21] Matthias Fey, Jan Eric Lenssen, Frank Weichert, and Heinrich Müller. Splinecnn: Fast geometric deep learning with continuous B-spline kernels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 869–877, 2018. 7
- [22] Paul Finsler. *Über Kurven und Flächen in allgemeinen Räumen*. Philosophische Fakultät, Georg-August-Univ., 1918. 1
- [23] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep learning*. MIT press, 2016. 2
- [24] Oshri Halimi, Or Litany, Emanuele Rodola, Alex M. Bronstein, and Ron Kimmel. Unsupervised learning of dense shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4370–4379, 2019. 2
- [25] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 770–778, 2016. 2
- [26] Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. *ACM transactions on graphics (TOG)*, 30(4):1–12, 2011. 7
- [27] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25, 2012. 2
- [28] Yann LeCun, Léon Bottou, Yoshua Bengio, and Patrick Haffner. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998. 2
- [29] Qinsong Li, Shengjun Liu, Ling Hu, and Xinru Liu. Shape correspondence using anisotropic Chebyshev spectral cnns. In *CVPR*, pages 14658–14667, 2020. 1, 2, 3, 4, 6, 7, 8
- [30] Or Litany, Tal Remez, Emanuele Rodola, Alex M. Bronstein, and Michael M. Bronstein. Deep functional maps: Struc-

- tered prediction for dense shape correspondence. In *Proceedings of the IEEE international conference on computer vision*, pages 5659–5667, 2017. 2
- [31] Rooe Litman and Alexander M Bronstein. Learning spectral descriptors for deformable shape correspondence. *IEEE transactions on pattern analysis and machine intelligence*, 36(1):171–180, 2013. 2
- [32] Jean-Marie Mirebeau. Efficient fast marching with Finsler metric. pages 515–557, 2014. 5
- [33] Thomas W. Mitchel, Vladimir G. Kim, and Michael Kazhdan. Field convolutions for surface cnns. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 10001–10011, 2021. 2, 7, 8
- [34] Shin-ishi Otha and Karl-Theodor Sturm. Heat flow on Finsler manifolds. *Commun. Pure Appl. Math.*, 62(10):1386–1433, 2009. 3, 4
- [35] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (ToG)*, 31(4):1–11, 2012. 2
- [36] Maks Ovsjanikov, Etienne Corman, Michael Bronstein, Emanuele Rodolà, Mirela Ben-Chen, Leonidas Guibas, Frederic Chazal, and Alex Bronstein. Computing and processing correspondences with functional maps. In *SIGGRAPH ASIA 2016 Courses*, pages 1–60. 2016. 2
- [37] Nathan D. Ratliff, Karl Van Wyk, Mandy Xie, Anqi Li, and Muhammad Asif Rana. Generalized nonlinear and Finsler geometry for robotics. In *IEEE International Conference on Robotics and Automation (ICRA)*, pages 10206–10212, 2021. 3
- [38] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (ToG)*, 37(6):1–16, 2018. 2, 7
- [39] Bernhard Riemann. Über die Hypothesen, welche der Geometrie zu Grunde liegen. *Abhandlungen der königlichen Gesellschaft der Wissenschaften zu Göttingen*, 13:304–319, 1854. 1
- [40] Samuele Salti, Federico Tombari, and Luigi Di Stefano. Shot: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding*, 125:251–264, 2014. 7
- [41] Abhishek Sharma and Maks Ovsjanikov. Weakly supervised deep functional maps for shape matching. *Advances in Neural Information Processing Systems*, 33:19264–19275, 2020. 2
- [42] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014. 2
- [43] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the inception architecture for computer vision. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826, 2016. 2
- [44] Matthias Vestner, Zorah Löhner, Amit Boyarski, Or Litany, Ron Slossberg, Tal Remez, Emanuele Rodola, Alex Bronstein, Michael Bronstein, Ron Kimmel, and Daniel Cremers. Efficient deformable shape correspondence via kernel matching. In *2017 International Conference on 3D vision (3DV)*, pages 517–526. IEEE, 2017. 2
- [45] Dong-Ming Yan, Guanbo Bao, Xiaopeng Zhang, and Peter Wonka. Low-resolution remeshing using the localized restricted voronoi diagram. *IEEE Transactions on Visualization and Computer Graphics (TVCG)*, 2014. 7
- [46] Fang Yang, Li Chai, Da Chen, and Laurent D. Cohen. Geodesic via asymmetric heat diffusion based on Finsler metric. In *Asian Conference on Computer Vision (ACCV)*, Springer, pages 371–386, 2018. 3