Functional Maps Representation On Product Manifolds

E. Rodolà[1] , Z. Lähner[2], A. M. Bronstein[3,6], M. M. Bronstein[4,5,6] and J. Solomon[7]

[1]Sapienza University of Rome, Italy
rodola@di.uniroma1.it
[2]TU Munich, Germany
laehner@in.tum.de
[3]Technion, Haifa, Israel
bron@cs.technion.ac.il
[4]USI Lugano, Switzerland
bronstein@imperial.ac.uk
[5]Imperial College London, London, UK
[6]Intel, Haifa, Israel
[7]MIT, Cambridge, Massachusetts
jsolomon@mit.edu

Abstract

We consider the tasks of representing, analysing and manipulating maps between shapes. We model maps as densities over the product manifold of the input shapes; these densities can be treated as scalar functions and therefore are manipulable using the language of signal processing on manifolds. Being a manifold itself, the product space endows the set of maps with a geometry of its own, which we exploit to define map operations in the spectral domain; we also derive relationships with other existing representations (soft maps and functional maps). To apply these ideas in practice, we discretize product manifolds and their Laplace–Beltrami operators, and we introduce localized spectral analysis of the product manifold as a novel tool for map processing. Our framework applies to maps defined between and across 2D and 3D shapes without requiring special adjustment, and it can be implemented efficiently with simple operations on sparse matrices.

Keywords: shape matching, functional maps, product manifolds

ACM CCS: Computational Geometry and Object Modelling–Shape Analysis; 3D Shape Matching; Geometric Modelling

1. Introduction

Three-dimensional acquisition continues to reach new levels of sophistication and is rapidly being incorporated into commercial products ranging from the Microsoft Kinect for gaming to LIDAR for autonomous cars and MRI for medical imaging. An essential building block for application design in many of these domains is fast and reliable recovery of 3D shape correspondences. This problem arises in applications as diverse as character animation, 3D avatars, pose and style transfer, or texture mapping, to mention a few.

A modern theme in shape correspondence involves the *representation* of a map from one shape to another. Although the most obvious representation maintains pairs of source and target points, this is by no means the only option. Our paper is mainly related to two frameworks developed for establishing correspondence between shapes: optimization on *product manifolds* and *functional maps*.

The first class of methods represents the correspondence on the Cartesian product of the two shapes. First methods of this type were formulated using graph matching [ZWW*10]. Windheuser *et al.* optimize in a product space [WSSC11], preserving important differential geometric properties. A similar approach was applied in [LRS*16] for 2D-to-3D matching. In [VLR*17], correspondence is formulated as kernel density estimation on the product manifold, interpreted as an alternating diffusion-sharpening process in [VLB*17]. A product between more than two shapes is considered in [CRA*17], but the resulting optimization problem is restricted to yield only sparse correspondences.

Soft maps [SNB*12] represent correspondence between shapes as a distribution on the product manifold with prescribed marginals reflecting area preservation. Non-convex objectives can be used to incorporate metric information into optimization for soft maps [Mém11, SPKS16], while other objectives on soft maps can be understood as probabilistic relaxations of classical distortion measures from differential geometry [SGB13, MCSK*17]. These methods suffer from high complexity, usually quadratic in the number of shape vertices.

Functional maps [OCB*17] abandon pointwise correspondence, instead modelling correspondences as linear operators between spaces of functions. An approximation of such operators in a pair of truncated orthogonal bases. One of the key innovations of this framework is allowing bringing a new set of algebraic methods into the domain of shape correspondence. Several follow-up works tried to improve the framework by employing sparsity-based priors [PBB*13], manifold optimization [KBB*13, KGB16], non-orthogonal [KBBV15] or localized [CSBK17, MRCB18] bases, coupled optimization over the forward and inverse maps [ERGB16, EBC17, HO17], combination of functional maps with metric-based approaches [ADK16, SK17] and kernelization [WGBS18]. Recent works of [NO17, NMR*18] considered functional algebra (function point-wise multiplications together with addition). Generalizations addressing the settings of multiple shapes [HWG14, KGB16], partial correspondence [RCB*17, LRB*16] and cluttered correspondence [CRM*16] have been proposed as well. Most recently, functional maps have also been used in conjunction with intrinsic deep learning methods [LRR*17]. For a comprehensive survey of functional maps and related techniques, we refer the reader to [OCB*17].

1.1. Motivation and contribution

In this paper, we advocate posing correspondence—and understanding relationships between the existing representations above—in terms of functions on the product manifold of the source and target. A motivating observation is that functional maps approximate a distribution representing the correspondence in the product space as a linear combination of *separable* tensor-product basis functions. This distribution, however, is supported on a manifold with a dimension *lower* than that of the product space: For a pair of two-dimensional shapes, the distribution is supported on a two-dimensional manifold embedded in a four-dimensional space. Consequently, most of the support of the basis functions is wasted on 'empty' regions of the product space. Localized bases on the individual domains improve this situation, but still most of their support is wasted.

We show how point-to-point maps, functional maps and soft maps all can be understood as (signed) measures on the product and how these representations might be converted to one another. More importantly, this viewpoint suggests new techniques to *represent* and *approximate* mappings directly on the product, e.g. by building a basis from Eigenfunctions of the product Laplace–Beltrami operator potentially after filtering undesirable matches.

Our theoretical contributions have practical bearing on the design of correspondence techniques. After discretizing product manifolds and their Laplace–Beltrami operators, we consider map design and processing problems among 2D and 3D shapes. Reasoning about the product manifold leads to compact, understandable bases for map design that focus resolution in the part of the product most relevant to a correspondence task. One of such means is the construction of *inseparable* bases. To this end, we propose to compute localized harmonics on the product manifold, and discuss a numerical scheme that keeps the complexity of such a computation feasible and, in particular cases, comparable to that of the construction of a separable localized basis. We finally showcase our framework applied to the task of map refinement.

2. Background

2.1. Manifolds

We model shapes as Riemannian d-manifolds $(\mathcal{M}, g_{\mathcal{M}})$ (possibly with boundary $\partial \mathcal{M}$) equipped with area elements dx induced by the standard metric $g_{\mathcal{M}}$; we do not restrict our focus to surfaces but rather allow $\mathcal{M}$ and $\mathcal{N}$ to have different intrinsic dimensions. We denote by $T_x\mathcal{M}$ the tangent plane at $x \in \mathcal{M}$, modelling the manifold locally as a Euclidean space. Given two scalar functions $f, g : \mathcal{M} \to \mathbb{R}$ belonging to an appropriate functional space $\mathcal{F}(\mathcal{M})$, we use the standard manifold inner product $\langle f, g \rangle_{\mathcal{M}} = \int_{\mathcal{M}} f(x)g(x)\,dx$.

In analogy to the Laplace operator in flat spaces, the positive semidefinite Laplace–Beltrami (LB) operator $\Delta_{\mathcal{M}}$ equips us with the tools to extend Fourier analysis to manifolds. The manifold Laplacian admits an Eigendecomposition $\Delta_{\mathcal{M}}\phi_i = \lambda_i \phi_i$ for $i \geq 1$, with real eigenvalues $0 = \lambda_1 \leq \lambda_2 \leq \ldots$ and Eigenfunctions $\{\phi_i\}_{i \geq 1}$ forming an orthonormal basis of $L^2(\mathcal{M}) = \{f : \mathcal{M} \to \mathbb{R} \mid \langle f, f \rangle_{\mathcal{M}} < \infty\}$. Any function $f \in L^2(\mathcal{M})$ can thus be represented via the Fourier-like series expansion

$$f(x) = \sum_{i \geq 1} \langle f, \phi_i \rangle_{\mathcal{M}} \phi_i(x). \tag{1}$$

2.2. Product manifolds

Given two Riemannian manifolds $(\mathcal{M}, g_{\mathcal{M}})$, $(\mathcal{N}, g_{\mathcal{N}})$ of dimension $d_{\mathcal{M}}$ and $d_{\mathcal{N}}$ with metric tensors $g_{\mathcal{M}}, g_{\mathcal{N}}$, respectively, their product $(\mathcal{M} \times \mathcal{N}, g_{\mathcal{M}} \oplus g_{\mathcal{N}})$ is a manifold of dimension $d_{\mathcal{M}} + d_{\mathcal{N}}$, where $g_{\mathcal{M}} \oplus g_{\mathcal{N}} = \begin{pmatrix} g_{\mathcal{M}} & 0 \\ 0 & g_{\mathcal{N}} \end{pmatrix}$ is the direct sum of the individual metric tensors [GP10], inducing the area element $da = dx\,dy$. By this definition of product, to each point $(x, y) \in \mathcal{M} \times \mathcal{N}$ is attached a tangent space derived by the canonical isomorphism $T_{(x,y)}\mathcal{M} \times \mathcal{N} = T_x\mathcal{M} \times T_y\mathcal{N}$ (see [Tu11], ex. 8.7]). For tangent vectors $\xi, \eta \in T_x\mathcal{M}$ and $\zeta, \mu \in T_y\mathcal{N}$, the inner product $\langle \cdot, \cdot \rangle$ of $(\xi, \zeta), (\eta, \mu) \in T_{(x,y)}\mathcal{M} \times \mathcal{N}$ is given by

$$\langle (\xi, \zeta), (\eta, \mu) \rangle_{T_{(x,y)}\mathcal{M} \times \mathcal{N}} = \langle \xi, \eta \rangle_{T_x\mathcal{M}} + \langle \zeta, \mu \rangle_{T_y\mathcal{N}}. \tag{2}$$

Now let $f \in \mathcal{F}(\mathcal{M})$, $g \in \mathcal{F}(\mathcal{N})$ for some functional space $\mathcal{F}$, and denote by $f \wedge g$ the outer product of f and g defined by the mapping

$$f \wedge g : (x, y) \mapsto f(x)g(y). \tag{3}$$

The LB operator $\Delta_{\mathcal{M} \times \mathcal{N}}$ obeys the (outer) product rule identity [Cha84]:

$$\Delta_{\mathcal{M} \times \mathcal{N}}(f \wedge g) = (\Delta_{\mathcal{M}} f) \wedge g + f \wedge (\Delta_{\mathcal{N}} g). \qquad (4)$$

Given Eigenvectors (ϕ, ψ) with corresponding Eigenvalues (α, β) satisfying $\Delta_{\mathcal{M}} \phi = \alpha \phi$ and $\Delta_{\mathcal{N}} \psi = \beta \psi$, application of the product rule yields

$$\Delta_{\mathcal{M} \times \mathcal{N}}(\phi \wedge \psi) = (\Delta_{\mathcal{M}} \phi) \wedge \psi + \phi \wedge (\Delta_{\mathcal{N}} \psi)$$
$$= (\alpha + \beta)(\phi \wedge \psi). \qquad (5)$$

This observation leads to a characterization of LB Eigenvalues for product manifolds:

Theorem 1 [BGM71, Proposition A.II.3]. *Let ξ be an Eigenfunction of the product LB operator $\Delta_{\mathcal{M} \times \mathcal{N}}$ with the corresponding Eigenvalue γ. Then, there exist some Eigenfunctions ϕ of $\Delta_{\mathcal{M}}$ and ψ of $\Delta_{\mathcal{N}}$ with the Eigenvalues α and β, respectively, such that $\xi = \phi \wedge \psi$ and $\gamma = \alpha + \beta$.*

It is also easy to check that the set of Eigenfunctions $\{\phi_i \wedge \psi_j\}_{i,j}$ is orthogonal, since:

$$\int_{\mathcal{M} \times \mathcal{N}} (\phi_i \wedge \psi_j)(\phi_k \wedge \psi_\ell) \, da = \int_{\mathcal{M} \times \mathcal{N}} \phi_i(x)\psi_j(y)\phi_k(x)\psi_\ell(y) \, da$$

$$= \int_{\mathcal{M}} \phi_i \phi_k \, dx \int_{\mathcal{N}} \psi_j \psi_\ell \, dy \qquad (6)$$

$$= \delta_{ik}\delta_{j\ell} = \begin{cases} 1 & (i = k) \quad \text{and} \quad (j = \ell); \\ 0 & \text{otherwise,} \end{cases} \qquad (7)$$

where δ_{ij} is the Kronecker delta.

2.3. Soft maps

A soft map $\tilde{\mu} : \mathcal{M} \to \text{Prob}(\mathcal{N})$ is a function assigning a probability measure over $\mathcal{N}$ to each point in $\mathcal{M}$ [SNB*12]. Soft maps can be equivalently represented by their densities, i.e. non-negative scalar functions $\mu : \mathcal{M} \times \mathcal{N} \to [0, 1]$ defined on the product manifold $\mathcal{M} \times \mathcal{N}$ satisfying $\tilde{\mu}(x)(B) = \int_{B \subseteq \mathcal{N}} \mu(x, y) \, dy$ for all $x \in \mathcal{M}$ and all measurable subsets $B \subseteq \mathcal{N}$.

As a particular case, a bijection $\tilde{\pi} : \mathcal{M} \to \mathcal{N}$ induces a soft map $\tilde{\mu}$ by requiring, for all $x \in \mathcal{M}$, that $\tilde{\mu}(x)(B) = 1$ if and only if $\tilde{\pi}(x) \in B \subseteq \mathcal{N}$, i.e. the image $\tilde{\mu}(x)$ is a unit Dirac mass $\delta_{\tilde{\pi}(x)}$ centred at $\tilde{\pi}(x)$.

2.4. Functional maps

A functional map T associated to a map $\tilde{\pi} : \mathcal{M} \to \mathcal{N}$ is a linear mapping $T : \mathcal{F}(\mathcal{N}) \to \mathcal{F}(\mathcal{M})$ defined as [OBCS*12]:

$$T(g) = g \circ \tilde{\pi}. \qquad (8)$$

Note how this construction allows to move from identifying a map between manifolds to identifying a linear operator between Hilbert spaces. The functional map T admits a matrix representation with


Figure 1: *Discretization of the Laplace–Beltrami operator on a cycle graph (a) and on a triangle mesh (b) for interior (green) and boundary edges (red). We also show the hat basis function in (a).*

respect to orthogonal bases $\{\phi_i\}_{i \geq 1}$, $\{\psi_j\}_{j \geq 1}$ on $\mathcal{F}(\mathcal{M})$ and $\mathcal{F}(\mathcal{N})$, respectively, with coefficients $\mathbf{C} = (c_{ij})$ determined as follows:

$$T(g) = \sum_{ij \geq 1} \langle \psi_j, g \rangle_{\mathcal{N}} \underbrace{\langle \phi_i, T(\psi_j) \rangle_{\mathcal{M}}}_{c_{ij}} \phi_i. \qquad (9)$$

3. Discretization

We show how to discretize the main quantities involved in our framework on 1D and 2D manifolds, as well as their products.

3.1. 1D shapes (curves)

We model 1D manifolds as closed contours with circular topology (no boundary), discretized as 2-regular cycle graphs $\mathcal{G} = (\mathcal{N}, \mathcal{E})$ with $n \geq 3$ nodes $\mathcal{N}$ and as many edges $\mathcal{E}$. The LB operator Δ is discretized using standard FEM with linear hat functions; in the hat basis, scalar functions on $\mathcal{G}$ are approximated piecewise-linearly on the edges. The Laplacian takes the form of a $n \times n$ sparse matrix $\mathbf{L} = \mathbf{S}^{-1} \mathbf{W}$, where:

$$w_{ij} = \begin{cases} -\frac{1}{\|e_{ij}\|} & e_{ij} \in \mathcal{E} \\ -\sum_{i \neq k} w_{ik} & i = j \\ 0 & \text{otherwise} \end{cases} \qquad (10)$$

$$s_{ij} = \begin{cases} \frac{1}{6} \|e_{ij}\| & e_{ij} \in \mathcal{E} \\ \frac{1}{3} \sum_{k \in \mathcal{N}(i)} \|e_{ik}\| & i = j \\ 0 & \text{otherwise} \end{cases} \qquad (11)$$

and the notation is according to Figure 1, with $\mathcal{N}(i)$ being the set of the neighbours of node i. In our tests, we use non-lumped masses s_{ij}; in applications requiring additional efficiency, lumped mass matrices $\text{diag}(\hat{s}_{ii})$ can be used by setting $\hat{s}_{ii} = \sum_j s_{ij}$.

The product of two boundary-free 1D manifolds $\mathcal{M}, \mathcal{N}$ is a 2D manifold (a surface) $\mathcal{M} \times \mathcal{N}$ with torus topology. For the discretization of the Laplacian on $\mathcal{M} \times \mathcal{N}$, we appeal to the following:

Theorem 2 (Discrete product Laplacian). *Let $\mathcal{M}, \mathcal{N}$ be 1D manifolds with no boundary, discretized as 2-regular cycle graphs, and let $\mathbf{S}_{\mathcal{M}}, \mathbf{W}_{\mathcal{M}}$ and $\mathbf{S}_{\mathcal{N}}, \mathbf{W}_{\mathcal{N}}$ be the mass and stiffness matrices for $\Delta_{\mathcal{M}}$ and $\Delta_{\mathcal{N}}$, respectively, obtained via FEM with respect to piecewise*

linear (hat) basis functions. Then,

$$\mathbf{S}_{\mathcal{M}\times\mathcal{N}} = \mathbf{S}_{\mathcal{M}} \otimes \mathbf{S}_{\mathcal{N}} \tag{12}$$

$$\mathbf{W}_{\mathcal{M}\times\mathcal{N}} = \mathbf{W}_{\mathcal{M}} \otimes \mathbf{S}_{\mathcal{N}} + \mathbf{S}_{\mathcal{M}} \otimes \mathbf{W}_{\mathcal{N}} \tag{13}$$

are the mass and stiffness matrices for the product manifold Laplacian $\Delta_{\mathcal{M}\times\mathcal{N}}$ with respect to piecewise bilinear basis functions, defined on a quad meshing of the toric surface $\mathcal{M} \times \mathcal{N}$. Here, $\otimes$ denotes the Kronecker product.

Proof. See the Appendix.

Corollary 1. *The LB operator $\Delta_{\mathcal{M}\times\mathcal{N}}$ is discretized as:*

$$\mathbf{L}_{\mathcal{M}\times\mathcal{N}} = \mathbf{L}_{\mathcal{M}} \otimes \mathbf{I}_{\mathcal{N}} + \mathbf{I}_{\mathcal{M}} \otimes \mathbf{L}_{\mathcal{N}}, \tag{14}$$

where $\mathbf{I}_{\mathcal{M}}, \mathbf{I}_{\mathcal{N}}$ are $n_{\mathcal{M}} \times n_{\mathcal{M}}$ and $n_{\mathcal{N}} \times n_{\mathcal{N}}$ identity matrices.

Proof. See the Appendix.

The discretization of $\Delta_{\mathcal{M}\times\mathcal{N}}$ does *not* require the explicit construction of a quad mesh embedded in $\mathbb{R}^3$; the toric shapes shown in these pages only serve visualization purposes. Further, the discretization (14) is consistent with the spectral decomposition identities (5); see [Fie73] and [HIK11, Proposition 33.6] for additional discussion.

3.2. 2D shapes (surfaces)

We model 2D surfaces as manifold triangle meshes $(\mathcal{V}, \mathcal{E}, \mathcal{F})$ with n vertices $\mathcal{V}$ connected by edges $\mathcal{E} = \mathcal{E}_i \cup \mathcal{E}_b$ (where $\mathcal{E}_i$ and $\mathcal{E}_b$ are interior and boundary edges, respectively) and triangle faces $\mathcal{F}$. In analogy to the 1D case, the discretization of the LB operator is obtained using FEM with piecewise linear basis functions on triangle elements [Duf59], taking the form of an $n \times n$ sparse matrix $\mathbf{L} = \mathbf{S}^{-1}\mathbf{W}$, where

$$w_{ij} = \begin{cases} (\cot\alpha_{ij} + \cot\beta_{ij})/2 & ij \in \mathcal{E}_i \\ (\cot\alpha_{ij})/2 & ij \in \mathcal{E}_b \\ -\sum_{k \neq i} w_{ik} & i = j \\ 0 & \text{otherwise, and} \end{cases} \tag{15}$$

$$s_{ij} = \begin{cases} (A(T_{hij}) + A(T_{ijk}))/12 & ij \in \mathcal{E}_i \\ A(T_{ijk})/12 & ij \in \mathcal{E}_b \\ \frac{1}{6}\sum_{k \in \mathcal{N}(i)} A(T_k) & i = j \\ 0 & \text{otherwise.} \end{cases} \tag{16}$$

Here, $A(T)$ denotes the area of triangle T and $\mathcal{N}(i)$ is the set of the neighbours of vertex i; see Figure 1 for notation.

Given two 2D manifolds $\mathcal{M}$ and $\mathcal{N}$, their product is a 4D manifold $\mathcal{M} \times \mathcal{N}$. The LB operator on $\mathcal{M} \times \mathcal{N}$ is discretized similarly to the lower dimensional case:

Corollary 2. *Let $\mathcal{M}, \mathcal{N}$ be surfaces discretized as triangle meshes, and let $\mathbf{S}_{\mathcal{M}}, \mathbf{W}_{\mathcal{M}}$ and $\mathbf{S}_{\mathcal{N}}, \mathbf{W}_{\mathcal{N}}$ be the mass and stiffness matrices for*


Figure 2: *The Cartesian product of two edge elements is a quad (a), while taking the product of two triangles yields a 4D geometric structure called a 3-3 (or triangular) duoprism [Cox48] visualized with a Schlegel diagram [Sch83] (b). Note that all these objects are polytopes (i.e. they have faces), not simple graphs.*

$\Delta_{\mathcal{M}}$ and $\Delta_{\mathcal{N}}$. *Then, Equations (12)–(14) provide a valid discretization of the LB operator $\Delta_{\mathcal{M}\times\mathcal{N}}$. This discretization is equivalent to the application of FEM on a 3-3 duoprism tessellation of the 4D product manifold $\mathcal{M} \times \mathcal{N}$ using multi-linear basis functions.*

Proof. See the Appendix.

We emphasize that, as a consequence of the Corollary, the computation of the product Laplacian $\Delta_{\mathcal{M}\times\mathcal{N}}$ does not require constructing a high-dimensional embedding for $\mathcal{M} \times \mathcal{N}$, thus avoiding cumbersome manipulation of duoprismic product elements (see Figure 2 for an illustration of these elements).

Finally, scalar functions on a manifold $\mathcal{M}$ are represented by n-dimensional vectors $\mathbf{f} = (f(x_1), \ldots, f(x_n))^\top$, where $x_1, \ldots, x_n$ denote graph nodes and mesh vertices in the 1D and 2D case, respectively. Inner products $\langle f, g \rangle_{\mathcal{M}}$ are discretized as $\mathbf{f}^\top \mathbf{S} \mathbf{g}$, where $\mathbf{S}$ is the mass matrix. On product manifolds, scalar functions are represented as $n_{\mathcal{M}} \times n_{\mathcal{N}}$ matrices $\mathbf{F}$, usually deriving from an outer product $f \wedge g$ discretized as $\mathbf{fg}^\top$; inner products are computed as $\text{vec}(\mathbf{F})^\top \mathbf{S} \, \text{vec}(\mathbf{G})$ ($\text{vec}(\mathbf{F})$ stacks the columns of $\mathbf{F}$ into a vector).

4. Map Representation on the Product Manifold

4.1. Soft functional maps

It will be instrumental for our purposes to introduce a 'soft' generalization of functional maps. For soft maps $\tilde{\mu} : \mathcal{M} \to \text{Prob}(\mathcal{N})$ with associated density $\mu \in L^1(\mathcal{M} \times \mathcal{N})$, we define a *soft functional map* $T_\mu : \mathcal{F}(\mathcal{N}) \to \mathcal{F}(\mathcal{M})$ as the expectation

$$T_\mu(g)(x) = \int_{\mathcal{N}} g(y)\mu(x, y)\, \mathrm{d}y. \tag{17}$$

It is easy to check that T_μ is linear in g, hence admitting a matrix representation with coefficients defined as in (9); in particular, in the standard basis one obtains a stochastic matrix with each row summing to 1. If the density μ encodes a non-soft map (i.e. whenever $\mu(x, \cdot)$ is concentrated at one point), the definition (17) boils down to the original definition (8), $T(g)(x) = \int_{\mathcal{N}} g(y) \delta_{\tilde{\pi}(x)}(y)\, \mathrm{d}y = (g \circ \tilde{\pi})(x)$, where the last equivalence stems from the sampling property of Dirac deltas.

We begin our discussion by deriving a connection between functional map matrices and expanding soft map measures in the Laplace–Beltrami basis:

Theorem 3 (**Equivalence**). *Let $T_\mu : \mathcal{F}(\mathcal{N}) \to \mathcal{F}(\mathcal{M})$ be a soft functional map* (17) *with underlying density $\mu \in L^1(\mathcal{M} \times \mathcal{N})$. Further, let $c_{ij} = \langle \phi_i, T_\mu(\psi_j) \rangle_\mathcal{M}$ be the matrix coefficients of T_μ in the orthogonal bases $\{\phi_i\}_{i \geq 1}, \{\psi_j\}_{j \geq 1}$, and let $p_{ij} = \langle \phi_i \wedge \psi_j, \mu \rangle_{\mathcal{M} \times \mathcal{N}}$ be the expansion coefficients of μ in the product basis $\{\phi_i \wedge \psi_j\}_{i,j}$, such that $\mu = \sum_{ij} (\phi_i \wedge \psi_j) p_{ij}$. Then, $c_{ij} = p_{ij}$ for all i, j.*

Proof. The functional map matrix coefficients are computed as:

$$c_{ij} = \langle \phi_i, T_\mu(\psi_j) \rangle_\mathcal{M} = \int_\mathcal{M} \phi_i(x) T_\mu(\psi_j)(x) \, \mathrm{d}x \tag{18}$$

$$= \int_\mathcal{M} \phi_i(x) \int_\mathcal{N} \psi_j(y) \mu(x, y) \, \mathrm{d}y \, \mathrm{d}x \tag{19}$$

$$= \int_{\mathcal{M} \times \mathcal{N}} \phi_i(x) \psi_j(y) \mu(x, y) \, \mathrm{d}a, \tag{20}$$

while the expansion coefficients of μ are given by

$$p_{ij} = \langle \phi_i \wedge \psi_j, \mu \rangle_{\mathcal{M} \times \mathcal{N}} = \int_{\mathcal{M} \times \mathcal{N}} \phi_i(x) \psi_j(y) \mu(x, y) \, \mathrm{d}a. \tag{21}$$

Comparing Equations (20) and (21), we see that $c_{ij} = p_{ij}$ for any choice of $i, j \geq 1$. ∎

Note that Theorem 3 applies to any choice of orthogonal bases $\{\phi_i\}_{i \geq 1} \in \mathcal{F}(\mathcal{M})$, $\{\psi_j\}_{j \geq 1} \in \mathcal{F}(\mathcal{N})$.

4.2. Spectral representation

Consider the order-k, band-limited approximation of μ:

$$\mu \approx \sum_{\ell=1}^{k} \xi_\ell p_\ell, \tag{22}$$

where each ξ_ℓ is an Eigenfunction of $\Delta_{\mathcal{M} \times \mathcal{N}}$, which uniquely identifies, via (5), a pair of Eigenfunctions ϕ_i, ψ_j on $\mathcal{M}$ and $\mathcal{N}$, respectively. According to Theorem 3, the expansion coefficients p_ℓ are exactly those appearing in the functional map matrix $\mathbf{C}$, when this is expressed in the Laplacian Eigenbases of $\mathcal{M}$ and $\mathcal{N}$ as originally proposed by Ovsjanikov *et al.* [OBCS*12]. There is, however, a crucial difference in the way the two sets of coefficients are stored. We come to the following observation:

4.3. Truncation

The product Eigenfunctions ξ_ℓ appearing in the summation (22) are associated to the product Eigenvalues $\alpha_i + \beta_j$, which are ordered non-decreasingly. In contrast, in [OBCS*12] it was proposed to truncate the two summations in (9) to $i = 1, \ldots, k_\mathcal{M}$ and $j = 1, \ldots, k_\mathcal{N}$, where indices i and j follow the non-decreasing order of the Eigenvalue sequences α_i and β_j *separately*.

We see that, due to the *different ordering*, the Eigenfunctions ϕ_i, ψ_j involved in the approximation (22) of μ are not


(a)　　　**(b)**　　　**(c)**

Figure 3: *The ground truth map (here the identity) between two shapes approximated according to (a) the standard functional map representation, (b) the (separable) LB Eigenfunctions of the product manifold, ordered according to the product Eigenvalues and (c) the (inseparable) localized harmonics on the product manifold. All three cases use the same amount of coefficients. The black curve in each matrix represents the maximum value for each row. In this example, the product manifold is a flat torus, represented in the parametric domain in (a), (b) and (c).*

necessarily all those involved in the construction of $\mathbf{C}$ (9), assuming $k = k_\mathcal{M} k_\mathcal{N}$. In the former case, we operate with a reduced basis directly on $\mathcal{M} \times \mathcal{N}$, while in the latter case we consider two reduced bases on $\mathcal{M}$ and $\mathcal{N}$ *independently*. This has direct implications on the quality of the approximated maps, as illustrated in Figure 3.

4.4. Relation to finite sections

The functional map representation was originally introduced in [OBCS*12] as a convenient language for solving map inference problems of the type [OCB*17]:

$$\mathbf{CA} = \mathbf{B}, \tag{23}$$

where matrices $\mathbf{B} = (\langle \phi_i, f_j \rangle_\mathcal{M})$, $\mathbf{A} = (\langle \psi_i, g_j \rangle_\mathcal{N})$ contain Fourier coefficients of a given set of corresponding 'probe' functions $f_j, g_j, j = 1, \ldots, q$ on $\mathcal{M}$ and $\mathcal{N}$, respectively (typically, descriptors are used). In the problem above, one is asked to estimate the functional map $\mathbf{C}$.

By truncating the matrix $\mathbf{C}$ to the left upper $k_\mathcal{M} \times k_\mathcal{N}$ submatrix (as in [OBCS*12]), one obtains a finite-dimensional approximation of the infinite linear system (23). This procedure, known as the *finite section method* [GRS10], does not always guarantee convergence, and a series of remedies using rectangular sections ($k_\mathcal{M} \neq k_\mathcal{N}$) have been proposed in the literature (see [GO17] for a discussion pertaining to functional maps).

Recall that, according to Theorem 3, the matrix elements c_{ij} correspond to the expansion coefficients p_{ij} appearing in (22). Thus, due to the different ordering of the p_{ij}'s, the approximation carried out in (22) can be regarded as an 'irregular' finite section (see Figure 4, right); in contrast with purely *algebraic* approaches considering general systems of linear equations such as (23), our approach carries now a *geometric* meaning in that the shape of the section is determined by the geometry of the product manifold.

Figure 4: *Left: The k = 100 frequencies involved in the construction of a 10 × 10 functional map matrix* **C** *correspond to an irregular sampling of the Laplacian spectrum of the product manifold. Right: In turn, only some of the coefficients c_{ij} of matrix* **C** *appear among the first k expansion coefficients p_{ij} of the map in the product Eigenbasis. Here* **C** *is framed in black, while the blue dots identify the first k coefficients p_{ij}.*

5. Spectral Map Processing

In this paper, we consider curves and surfaces as our shapes. Despite their different intrinsic dimensions, our framework applies to both without specific adjustment.

5.1. Localized spectral encoding

Theorem 3 establishes the equivalence between the soft functional map T_μ representation coefficients c_{ij} in the bases $\{\phi_i\}_{i\geq 1} \subseteq \mathcal{F}(\mathcal{M})$ and $\{\psi_j\}_{j\geq 1} \subseteq \mathcal{F}(\mathcal{N})$ and the coefficients p_ℓ of the underlying density μ Fourier series (22) in the Eigenbasis $\{\xi_\ell\}_{\ell\geq 1} \subseteq \mathcal{F}(\mathcal{M} \times \mathcal{N})$ of the product manifold Laplacian $\Delta_{\mathcal{M}\times\mathcal{N}}$. This equivalence directly stems from ξ_ℓ's having the separable form $\phi_i \wedge \psi_j$, by virtue of Theorem 1. It may be advantageous, however, to consider different orthonormal bases on $\mathcal{M} \times \mathcal{N}$ that are not necessarily separable. In particular, we observe that μ tends to be localized on the product manifold $\mathcal{M} \times \mathcal{N}$ (see Figure 3), and thus the standard outer product basis is extremely wasteful as it is supported on the entire $\mathcal{M} \times \mathcal{N}$.

A better alternative is the use of *localized manifold harmonics* [CSBK17, MRCB18]. Assume that we are given a rough indication of the support of μ (for example, coming from a shape matching algorithm) in the form of a step potential function

$$V(x, y) = \begin{cases} \nu & \mu(x, y) \approx 0; \\ 0 & \text{otherwise.} \end{cases} \quad (24)$$

where $\nu \geq 1$. Then, the variational problem

$$\min_{\xi_1,\ldots,\xi_k} \sum_{\ell=1}^{k} \int_{\mathcal{M}\times\mathcal{N}} \left(\|\nabla_{\mathcal{M}\times\mathcal{N}}\xi_\ell\|^2_{g_{\mathcal{M}}\oplus g_{\mathcal{N}}} + V\xi_\ell^2 \right) da$$

$$\text{s.t.} \quad \langle \xi_\ell, \xi_{\ell'} \rangle_{\mathcal{M}\times\mathcal{N}} = \delta_{\ell,\ell'} \quad (25)$$

produces a set of orthonormal functions denoted by $\hat{\xi}_1, \ldots, \hat{\xi}_k$ that, for a sufficiently large value of ν, are also localized in the support of V. Note that this new basis $\{\hat{\xi}_\ell\}_{\ell=1}^{k}$ is *no longer separable*, i.e. the functions $\hat{\xi}$ are not in general expressible as outer products of functions defined on the originating domains. See Figures 5 and 6 for an illustration, and Figures 7 and 8 for practical examples.

The basis $\{\hat{\xi}_\ell\}_{\ell=1}^{k}$ turns out to be the Eigenbasis of the *Hamiltonian operator* [CSBK17] $H = \Delta_{\mathcal{M}\times\mathcal{N}} + V$ and can be computed by the Eigendecomposition of the product Laplacian matrix with the addition of diagonal potential. The size of such problem can be huge (if the shapes are discretized with $n \sim 10^3$ points, the product Laplacian matrix has size $n^2 \times n^2 = 10^6 \times 10^6$; see Theorem 2), and despite its extreme sparsity, computationally expensive.

As an alternative, we consider a *patch* $\mathcal{P} \subset \mathcal{M} \times \mathcal{N}$ of the product manifold with boundary $\partial\mathcal{P}$ corresponding to $\mu(x, y) > 0$, and define the Eigenproblem

$$\begin{array}{ll} \Delta_{\mathcal{P}}\bar{\xi}_\ell(x, y) = \gamma_\ell\bar{\xi}_\ell(x, y) & (x, y) \in \text{int}(\mathcal{P}) \\ \bar{\xi}_\ell(x, y) = 0 & (x, y) \in \partial\mathcal{P} \end{array} \quad (26)$$

of the *product patch Laplacian* $\Delta_{\mathcal{P}}$ with homogeneous Dirichlet boundary conditions. In practice, this is implemented by constructing the stiffness and mass matrices $\mathbf{W}_{\text{int}(\mathcal{P})}, \mathbf{S}_{\text{int}(\mathcal{P})}$ by selecting the rows and columns of $\mathbf{W}_{\mathcal{M}\times\mathcal{N}}, \mathbf{S}_{\mathcal{M}\times\mathcal{N}}$ that correspond to the vertices in $\text{int}(\mathcal{P})$. A generalized Eigenproblem using $\mathbf{W}_{\text{int}(\mathcal{P})}, \mathbf{S}_{\text{int}(\mathcal{P})}$ is solved, yielding Eigenfunctions $\bar{\xi}_{\text{int}(\mathcal{P})}$ defined on $\text{int}(\mathcal{P})$; the final Eigenfunctions $\bar{\xi}$ on the entire patch $\mathcal{P}$ are obtained by setting $\bar{\xi}(x) = \xi_{\text{int}(\mathcal{P})}$ for $x \in \text{int}(\mathcal{P})$ and $\bar{\xi}(x) = 0$ for $x \in \partial\mathcal{P}$.


Figure 5: *Basis functions on the product manifold (here visualized as a torus embedded in $\mathbb{R}^3$) of two 1D shapes. We plot a few standard LB Eigenfunctions (top row) and localized manifold harmonics (bottom row). Here and in the following, we use the present colour scheme (blue denotes small values, red large values, white is zero).*

Figure 6: *Projecting the basis functions on the product manifold of horse and elephant back onto the factor shapes (here only the horse projection is visualized). Top row: Projection of three product LB Eigenfunctions, which correspond exactly to three standard LB Eigenfunctions on the horse shape. Bottom row: Projection of three localized harmonics; these projections do not correspond to any LB Eigenfunction on the horse. Still, note how they capture the geometric features of the underlying shape.*

If the patch is selected in such a way that its size scales as $\mathcal{O}(n)$ rather than $\mathcal{O}(n^2)$ in the size of the shapes (in practice, this can be achieved by taking a fixed-size band around the initial correspondence), the computation of the localized basis $\{\bar{\bar{\xi}}_\ell\}_{\ell=1}^k$ has the same complexity as Eigendecomposition of the individual Laplacians $\Delta_{\mathcal{M}}, \Delta_{\mathcal{N}}$. An example application of this construction is described next.

Despite the computational gains of working with patches $\mathcal{P} \subset \mathcal{M} \times \mathcal{N}$, computing the Eigendecomposition of the full Hamiltonian $\Delta_{\mathcal{M} \times \mathcal{N}} + V$ may still be useful in certain settings. Note, in particular, that one may define a *soft* potential $V(x, y) = 1 - \mu(x, y)$ [MRCB18] by directly reflecting the reliability of the underlying map in terms of its density. Further, it is also possible to define a *patch Hamiltonian* $\Delta_{\mathcal{P}} + V|_{\mathcal{P}}$ with soft potential if desired.

5.2. Example: Map refinement

As an illustrative application of our framework, we propose a simple procedure for map refinement: Given some initial, possibly sparse and noisy correspondence, the task is to produce a dense, denoised map.

We follow an iterative approach. In each iteration k, the map is represented as a density $\mu^{(k)} : \mathcal{M} \times \mathcal{N} \to [0, 1]$. This density is interpreted as a heat distribution throughout the iterations.

At the k-th iteration, a diffusion process is initialized with $u_{t=0}^{(k)} := \mu^{(k)}$ and solved for a given diffusion time $T^{(k)}$ on a patch $\mathcal{P}^{(k)}$. The initial patch can be given or be the entire product manifold. The diffusion process has the effect of spreading correct correspondence information and therefore suppress mismatches, resulting in an effective map denoising approach akin to diffusion-based smoothing from image processing [Wit83, PM90]. The final heat distribution $u_T^{(k)}$ is thresholded to define a patch $\mathcal{P}^{(k)} \subset \mathcal{M} \times \mathcal{N}$, where the correct correspondence is likely to be contained, with likelihood expressed in terms of the diffused density. We then recover a bijective (non-soft) density $\mu^{(k+1)}$ from $u_T^{(k)}$ by solving a linear assignment problem [Ber98] restricted to region $\mathcal{P}^{(k)}$, and use it to initialize the next iteration.

These blur-and-sharpen steps are iterated until convergence while decreasing $T^{(k)}$, resulting in a sequence $\mathcal{P}^{(0)} \supseteq \cdots \supseteq \mathcal{P}^{(k)} \supseteq \mathcal{P}^{(k+1)}$. In practice, we decrease $T^{(k)}$ logarithmically across iterations. At $k = 0$, the density $u^{(0)}$ is the given input, e.g. a mixture of Dirac deltas or a soft map.

The diffusion step in each iteration is realized via the spectral decomposition of the product patch Laplacian $\Delta_{\mathcal{P}}$ with Dirichlet boundary conditions on $\mathcal{P} = \mathcal{P}^{(k)}$ (26); for $p, q \in \mathcal{M} \times \mathcal{N}$:

$$u_T(p) = \int_{\mathcal{M} \times \mathcal{N}} h_T(p, q) u_0(q) \mathrm{d}q \qquad (27)$$

$$h_T(p, q) = \sum_{\ell \geq 0} e^{-T \gamma_\ell} \bar{\bar{\xi}}_\ell(p) \bar{\bar{\xi}}_\ell(q), \qquad (28)$$


Figure 7: *Product space approximation of the correspondence between one-dimensional shapes with $k = 100$ basis functions. Bases constructed on bands of different size (1%, 5%, 25% and 90% of the total product manifold area) around the true correspondence are shown. Separable basis (FM) is shown as a reference. Left: accuracy of the correspondence increases as the product space basis becomes more localized. Right (top row): image of a delta function by the functional maps. Right (bottom row): True correspondence (curve) and its approximation in inseparable product space bases with a varying degree of localization. The product manifold is depicted as a torus.*
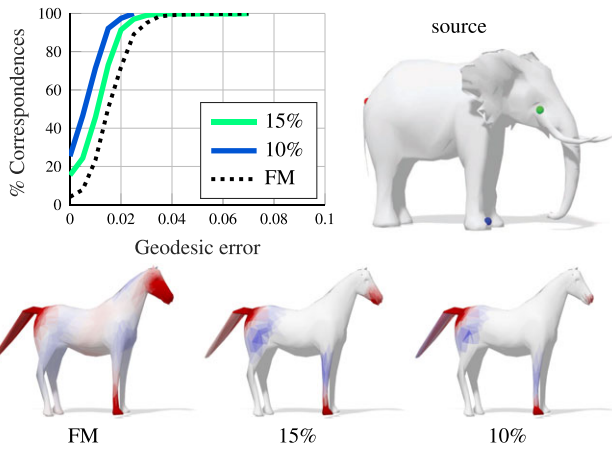
Figure 8: *Map approximation between surfaces with k = 500 basis functions on bands of different size (10% and 15% of the total 4D product manifold area) around the true correspondence. We also show images on the horse of delta functions supported at three points (red, green, blue) on the elephant. Here, the functional map (FM) was calculated using 30 × 30 = 900 basis functions.*

Figure 9: *Example of map refinement. We show the input sparse correspondence above and the recovered dense map below. The heatmap on the bottom right encodes geodesic error of the recovered correspondence.*

where h_T is the heat kernel at time T on the product manifold $\mathcal{M} \times \mathcal{N}$. Throughout the iterations we keep the number of Eigenfunctions for the approximation (28) constant.

The refinement process described above simultaneously improves the correspondence and reduces the support of the density around the most likely bijective map. This is similar in spirit to the kernel matching approaches of [VLR*17, VLB*17], however, with the additional step of 'carving out' the relevant portion $\mathcal{P} \subset \mathcal{M} \times \mathcal{N}$ throughout the iterations.

Illustrative results are reported qualitatively in Figure 9 and quantitatively in Figure 10.

6. Discussion and Conclusions

We introduced a novel perspective on map representation and processing, where pointwise, functional and soft maps can be understood as densities on the product of the input shapes. We showed how to discretize the Laplace-Beltrami operator on the product manifold and proposed the adoption of (inseparable) localized harmonics for compactly encoding correspondences while ensuring minimal energy dispersion, i.e. the resulting harmonics are not 'wasted' on portions of the product manifold that carry no information on the map to be encoded. Our theoretical and applied contributions suggest a new perspective on properties of the correspondence manifold as well as new ways to pose algorithmic design for map inference and processing.

6.1. Limitations

Perhaps the main limitation of our framework lies in the scalability of our current numerical scheme. While we showed that one can reduce the computational complexity to $\mathcal{O}(n)$ by appropriately

selecting a localization region, in practical applications involving very noisy maps where the localization region tends to be spread out across the entire product manifold, the advantage might be less evident. For this reason, considering as a possible extension higher dimensional products to encode cycle-consistent maps in shape collections may soon become prohibitive. With the current approach we trade off scalability for accuracy: Maps are encoded much more precisely in the localized basis, but this requires the explicit computation of inseparable basis functions that do not admit an efficient representation in terms of outer products. As a possible remedy, an efficient solution to the Eigenproblem might be sought via approximation methods similar to [NBH18]. A second limitation is in our map refinement scheme, which has limited resilience to particularly noisy input. We presented our algorithm as an illustrative tool for map denoising, but more effective schemes operating on the product manifold are likely possible.

6.2. Future work

From an investigative standpoint, it might be worth considering a notion of optimal transport *between maps* as a means of exploring the space of maps between given shapes, a natural choice given our modelling of maps as measures on a manifold. Related constructions could extend distortion measures like the Dirichlet energy [Bre03, SGB13, Lav17] to the functional regime.

Another particularly interesting direction will be to consider general *graphs* (as opposed to manifolds) and their products in the

Figure 10: *Sensitivity of map refinement to heat diffusion times and noisy input. The legend reports diffusion time ranges $(t_{\max}, t_{\min})$; within each range, time is decreased logarithmically over iterations. Left: The input is a sparse correspondence of 10% of correct matches. We see that high diffusion times are detrimental due to the excessive spread of correspondence information. Right: The input sparse correspondence is further corrupted with 30% random mismatches.*

context of network analysis, machine learning and applications. While many of our results may be directly translated to graphs, the lack of differentiable structure poses new theoretical challenges and at the same time provides a richer spectrum of possibilities; for example, several different notions of product exist between graphs [HIK11].

Finally, a promising direction is the introduction of product spaces within geometric deep learning [BBL*17] pipelines, where the data is in the form of signals defined on top of a manifold. Our proposed discretization of the (product) Laplace–Beltrami operator, as well as its spectral decomposition, can be directly employed in such pipelines, enabling new forms of structured prediction in a range of challenging problems in vision and graphics.

Acknowledgements

Appendix A: Proofs

We provide proofs for the main propositions of the paper.

Proof of Theorem 2. Following standard FEM, we discretize the Poisson equation $\Delta_{\mathcal{M} \times \mathcal{N}} f = g$ via the weak formulation

$$\langle \Delta_{\mathcal{M} \times \mathcal{N}} f, H_j \rangle = \langle g, H_j \rangle, \tag{A.1}$$

where functions are expressed in the hat basis $\{H_j : \mathcal{M} \times \mathcal{N} \to \mathbb{R}\}$, and are thus approximated piecewise-linearly via the expansion $f(x) \approx \sum_{i=1}^n f(v_i) h_i(x)$. The left-hand side of (A.1) can be written as

$$\langle \Delta f, H_j \rangle = -\langle \nabla f, \nabla H_j \rangle = -\sum_i f(v_i) \underbrace{\langle \nabla H_i, \nabla H_j \rangle}_{w_{ij}}, \tag{A.2}$$

where w_{ij} are elements of the *stiffness matrix* $\mathbf{W}$. The right-hand side of (A.1) can be written as

$$\langle g, H_j \rangle = \left\langle \sum_i g(v_i) H_i(x), H_j \right\rangle = \sum_i g(v_i) \underbrace{\langle H_i, H_j \rangle}_{s_{ij}}, \tag{A.3}$$

where s_{ij} are elements of the *mass matrix* $\mathbf{S}$.

The Cartesian product of the two graphs discretizing $\mathcal{M}$ and $\mathcal{N}$ has grid topology, as illustrated in Figure A1, and the resulting bilinear hat basis functions are expressed via the outer product $H_e = h_j \wedge h_q$. We can then compute the mass values (refer to the Figure for the color notation):

$$
\begin{aligned}
s_{ee} = \langle H_e, H_e \rangle &= \langle h_j \wedge h_q, h_j \wedge h_q \rangle \\
&= \int_{Q_{abde} \cup Q_{bcef} \cup Q_{degh} \cup Q_{efhi}} h_j(x) h_q(y) h_j(x) h_q(y) dx dy \\
&= \int_{E_{ijk}} h_j(x) h_j(x) dx \int_{E_{pqr}} h_q(y) h_q(y) dy \\
&= s_{jj} s_{qq}
\end{aligned}
\tag{A.4}
$$

Figure A1: *Left: The product of two closed contours discretized as cycle graphs (in blue and red) is a quad mesh with toric topology (in grey). Uniform edge lengths are used for illustration purposes. Right: Two overlapping bilinear hats H_e and H_f. On the quad element Q_{efhi} (marked in red), there is non-zero overlap, hence it contributes to the computation of mass and stiffness values.*

$$s_{ae} = \langle H_a, H_e \rangle = \langle h_i \wedge h_r, h_j \wedge h_q \rangle$$
$$= \int_{Q_{abde}} h_i(x) h_r(y) h_j(x) h_q(y) dx dy$$
$$= \int_{E_{ij}} h_i(x) h_j(x) dx \int_{E_{qr}} h_r(y) h_q(y) dy$$
$$= s_{ij} s_{qr} \tag{A.5}$$

$$s_{de} = \langle H_d, H_e \rangle = \langle h_i \wedge h_q, h_j \wedge h_q \rangle$$
$$= \int_{Q_{abde} \cup Q_{degh}} h_i(x) h_q(y) h_j(x) h_q(y) dx dy$$
$$= \int_{E_{ij}} h_i(x) h_j(x) dx \int_{E_{pqr}} h_q(y) h_q(y) dy.$$
$$= s_{ij} s_{qq} \tag{A.6}$$

Similarly, the stiffness integrals read:

$$w_{ee} = \langle \nabla H_e, \nabla H_e \rangle = \langle \nabla h_j \wedge h_q, \nabla h_j \wedge h_q \rangle$$
$$= \langle \nabla h_j h_q, \nabla h_j h_q \rangle + 2 \langle h_j \nabla h_q, \nabla h_j h_q \rangle + \langle h_j \nabla h_q, h_j \nabla h_q \rangle$$
$$= \int_{Q_{abde} \cup Q_{bcef} \cup Q_{degh} \cup Q_{efhi}} \langle \nabla h_j(x) h_q(y), \nabla h_j(x) h_q(y) \rangle dx dy$$
$$+ \cdots = \int_{Q_{abde} \cup Q_{bcef} \cup Q_{degh} \cup Q_{efhi}} h_q(y) h_q(y) \langle \nabla h_j(x), \nabla h_j(x) \rangle$$
$$\times dx dy + \cdots = \int_{E_{ijk}} \langle \nabla h_j(x), \nabla h_j(x) \rangle dx \int_{E_{pqr}} h_q(y)$$
$$\times h_q(y) dy + \cdots + \cdots = w_{jj} s_{qq} + s_{jj} w_{qq} \tag{A.7}$$

$$w_{ae} = \langle \nabla H_a, \nabla H_e \rangle = \langle \nabla h_i \wedge h_r, \nabla h_j \wedge h_q \rangle$$
$$= \langle \nabla h_i h_r, \nabla h_j h_q \rangle + \langle h_i \nabla h_r, h_j \nabla h_q \rangle$$
$$= w_{ij} s_{qr} + s_{ij} w_{qr} \tag{A.8}$$

$$w_{de} = \langle \nabla H_d, \nabla H_e \rangle = \langle \nabla h_i \wedge h_q, \nabla h_j \wedge h_q \rangle$$
$$= \langle \nabla h_i h_q, \nabla h_j h_q \rangle + \langle h_i \nabla h_q, h_j \nabla h_q \rangle$$
$$= w_{ij} s_{qq} + s_{ij} w_{qq}, \tag{A.9}$$

where we applied the outer product rule for the gradient operator, and used the fact that $\langle \nabla f, \nabla g \rangle = 0$ for any pair of functions on the two cycle graphs. Note the integrals s_{ae} and w_{ae} are non-zero even if nodes a and e are not connected in the product graph.

In matrix notation, formulas (A.4)–(A.9) can be succinctly written as:

$$\mathbf{S} = \mathbf{S} \otimes \mathbf{S}$$

$$\mathbf{W} = \mathbf{W} \otimes \mathbf{S} + \mathbf{S} \otimes \mathbf{W},$$

completing the proof. ∎

Proof of Corollary 1. The proof is straightforward and follows from substituting the expressions (12), (13) into the general formula $\mathbf{L} = \mathbf{S}^{-1} \mathbf{W}$:

$$\mathbf{L}_{\mathcal{M} \times \mathcal{N}} = \mathbf{S}_{\mathcal{M} \times \mathcal{N}}^{-1} \mathbf{W}_{\mathcal{M} \times \mathcal{N}}$$
$$= (\mathbf{S}_{\mathcal{M}} \otimes \mathbf{S}_{\mathcal{N}})^{-1} (\mathbf{W}_{\mathcal{M}} \otimes \mathbf{S}_{\mathcal{N}} + \mathbf{S}_{\mathcal{M}} \otimes \mathbf{W}_{\mathcal{N}})$$
$$= (\mathbf{S}_{\mathcal{M}}^{-1} \otimes \mathbf{S}_{\mathcal{N}}^{-1}) (\mathbf{W}_{\mathcal{M}} \otimes \mathbf{S}_{\mathcal{N}}) + (\mathbf{S}_{\mathcal{M}}^{-1} \otimes \mathbf{S}_{\mathcal{N}}^{-1}) (\mathbf{S}_{\mathcal{M}} \otimes \mathbf{W}_{\mathcal{N}})$$
$$= (\mathbf{S}_{\mathcal{M}}^{-1} \mathbf{W}_{\mathcal{M}}) \otimes (\mathbf{S}_{\mathcal{N}}^{-1} \mathbf{S}_{\mathcal{N}}) + (\mathbf{S}_{\mathcal{M}}^{-1} \mathbf{S}_{\mathcal{M}}) \otimes (\mathbf{S}_{\mathcal{N}}^{-1} \mathbf{W}_{\mathcal{N}})$$
$$= \mathbf{L}_{\mathcal{M}} \otimes \mathbf{I}_{\mathcal{N}} + \mathbf{I}_{\mathcal{M}} \otimes \mathbf{L}_{\mathcal{N}}. \quad ∎$$

Proof of Corollary 2. Since triangular (3-3) duoprisms are, by definition, the Cartesian product of two triangles, we can define a multi-linear basis function on the product complex as the outer product of two standard hats defined on triangle meshes. We are now in the same setting as the lower dimensional case, and in particular Equations (A.4)–(A.9) remain valid. ∎

References

[ADK16] AFLALO Y., DUBROVINA A., KIMMEL R.: Spectral generalized multi-dimensional scaling. *IJCV 118*, 3 (2016), 380–392.

[BBL*17] BRONSTEIN M. M., BRUNA J., LECUN Y., SZLAM A., VANDERGHEYNST P.: Geometric deep learning: Going beyond Euclidean data. *IEEE Signal Processing Magazine 34*, 4 (July 2017), 18–42.

[Ber98] BERTSEKAS D. P.: *Network Optimization: Continuous and Discrete Models*. Athena Scientific, 1998.

[BGM71] BERGER M., GAUDUCHON P., MAZET E.: *Le spectre d'une variété Riemannienne*. Lecture Notes in Mathematics. Springer-Verlag, Berlin, Germany, 1971.

[Bre03] BRENIER Y.: Extended monge-kantorovich theory. In *Optimal Transportation and Applications*. L. A. Caffarelli and S. Salsa (Eds.). Springer, Berlin, Germany, 2003, pp. 91–121.

[Cha84] CHAVEL I.: *Eigenvalues in Riemannian Geometry* (2nd edition). Academic Press, Cambridge, MA, 1984.

[Cox48] COXETER H. S. M.: *Regular Polytopes*. Dover Publications, New York, NY, 1948.

[CRA*17] COSMO L., RODOLÀ E., ALBARELLI A., MÉMOLI F., CREMERS D.: Consistent partial matching of shape collections via sparse modeling. *Computer Graphics Forum 36*, 1 (2017), 209–221.

[CRM*16] COSMO L., RODOLÀ E., MASCI J., TORSELLO A., BRONSTEIN M. M.: Matching deformable objects in clutter. In *Proceedings - 2016 4th International Conference on 3D Vision, 3DV 2016* (Stanford, 2016), IEEE, pp. 1–10.

[CSBK17] CHOUKROUN Y., SHTERN A., BRONSTEIN A., KIMMEL R.: Hamiltonian operator for spectral shape analysis. *arXiv: 1611.01990v2* (2017).

[Duf59] DUFFIN R. J.: Distributed and lumped networks. *Journal of Mathematics and Mechanics 8*, 5 (1959), 793–826.

[EBC17] EZUZ D., BEN-CHEN M.: Deblurring and denoising of maps between shapes. *Computer Graphics Forum 36*, 5 (2017), 165–174.

[ERGB16] EYNARD D., RODOLÀ E., GLASHOFF K., BRONSTEIN M. M.: Coupled functional maps. In *Proceedings - 2016 4th International Conference on 3D Vision, 3DV 2016* (Stanford, 2016), IEEE, pp. 399–407.

[Fie73] FIEDLER M.: Algebraic connectivity of graphs. *Czechoslovak Mathematical Journal 23*, 98 (1973), 298–305.

[GO17] GLASHOFF K., ORTLIEB C. P.: Composition operators, matrix representation, and the finite section method: A theoretical framework for maps between shapes. *arXiv:1705.00325* (2017).

[GP10] GUILLEMIN V., POLLACK A.: *Differential Topology*. AMS Chelsea Publishing Series. American Mathematical Society, Providence, RI, 2010.

[GRS10] GRÖCHENIG K., RZESZOTNIK Z., STROHMER T.: Convergence analysis of the finite section method and Banach algebras of matrices. *Integral Equations and Operator Theory 67*, 2 (2010), 183–202.

[HIK11] HAMMACK R., IMRICH W., KLAVZAR S.: *Handbook of Product Graphs* (2nd edition). CRC Press, Boca Raton, FL, 2011.

[HO17] HUANG R., OVSJANIKOV M.: Adjoint map representation for shape analysis and matching. *Computer Graphics Forum 36*, 5 (2017), 151–163.

[HWG14] HUANG Q., WANG F., GUIBAS L.: Functional map networks for analyzing and exploring large shape collections. *ACM Transactions on Graphics (TOG) 33*, 4 (2014), 36.

[KBB*13] KOVNATSKY A., BRONSTEIN M. M., BRONSTEIN A. M., GLASHOFF K., KIMMEL R.: Coupled quasi-harmonic bases. *Computer Graphics Forum 32*, 2 (2013), 439–448.

[KBBV15] KOVNATSKY A., BRONSTEIN M. M., BRESSON X., VANDERGHEYNST P.: Functional correspondence by matrix completion. In *Proceedings/CVPR* (Boston, 2015), IEEE, pp. 905–914.

[KGB16] KOVNATSKY A., GLASHOFF K., BRONSTEIN M. M.: MAD MM: a generic algorithm for non-smooth optimization on manifolds. In *Proceedings/ECCV* (Amsterdam, 2016), Springer.

[Lav17] LAVENANT H.: Harmonic mappings valued in the Wasserstein space. *arXiv:1712.07528* (2017).

[LRB*16] LITANY O., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M., CREMERS D.: Non-rigid puzzles. *Computer Graphics Forum 35*, 5 (2016), 135–143.

[LRR*17] LITANY O., REMEZ T., RODOLÀ E., BRONSTEIN A. M., BRONSTEIN M. M.: Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE International Conference on Computer Vision* (Venice, 2017), vol. *2*, IEEE, pp. 5660–5668.

[LRS*16] LÄHNER Z., RODOLÀ E., SCHMIDT F. R., BRONSTEIN M. M., CREMERS D.: Efficient globally optimal 2D-to-3D deformable shape matching. In *Proceedings of the IEEE Computer Society Conference on Computer Vision and Pattern Recognition* (Las Vegas, 2016), IEEE, pp. 2185–2193.

[MCSK*17] MANDAD M., COHEN-STEINER D., KOBBELT L., ALLIEZ P., DESBRUN M.: Variance-minimizing transport plans for inter-surface mapping. *ACM Transactions on Graphics 36*, 4 (2017), 39:1–39:14.

[Mém11] MÉMOLI F.: Gromov–Wasserstein Distances and the Metric Approach to Object Matching. *Foundations of Computational Mathematics 11*, 4 (2011), 417–487.

[MRCB18] MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M.: Localized manifold harmonics for spectral shape analysis. *Computer Graphics Forum 37*, 6 (2018), 20–34.

[NBH18] NASIKUN A., BRANDT C., HILDEBRANDT K.: Fast approximation of laplace–beltrami eigenproblems. *Computer Graphics Forum 37*, 5 (2018), 121–134.

[NMR*18] NOGNENG D., MELZI S., RODOLÀ E., CASTELLANI U., BRONSTEIN M., OVSJANIKOV M.: Improved functional mappings via product preservation. *Computer Graphics Forum 37*, 2 (2018), 179–190.

[NO17] NOGNENG D., OVSJANIKOV M.: Informative descriptor preservation via commutativity for shape matching. *Computer Graphics Forum 36*, 2 (2017), 259–267.

[OBCS*12] OVSJANIKOV M., BEN-CHEN M., SOLOMON J., BUTSCHER A., GUIBAS L.: Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphphics 31*, 4 (July2012), 30:1–30:11.

[OCB*17] OVSJANIKOV M., CORMAN E., BRONSTEIN M., RODOLÀ E., BEN-CHEN M., GUIBAS L., CHAZAL F., BRONSTEIN A.: Computing and processing correspondences with functional maps. In *ACM SIGGRAPH 2017 Courses* (2017), pp. 5:1–5:62.

[PBB*13] POKRASS J., BRONSTEIN A. M., BRONSTEIN M. M., SPRECHMANN P., SAPIRO G.: Sparse modeling of intrinsic correspondences. *Computer Graphics Forum 32*, 2 (2013), 459–468.

[PM90] PERONA P., MALIK J.: Scale-space and edge detection using anisotropic diffusion. *IEEE Transactions on Pattern Analysis and Machine Intelligence 12*, 7 (July 1990), 629–639.

[RCB*17] RODOLÀ E., COSMO L., BRONSTEIN M. M., TORSELLO A., CREMERS D.: Partial functional correspondence. *Computer Graphics Forum 36*, 1 (2017), 222–236.

[Sch83] SCHLEGEL V.: *Theorie der homogen zusammengesetzten Raumgebilde*. Nova Acta, Ksl. Leop.-Carol. Deutsche Akademie der Naturforscher, Band XLIV, Nr. 4, Druck von E. Blochmann und Sohn, Dresden, 1883.

[SGB13] SOLOMON J., GUIBAS L., BUTSCHER A.: Dirichlet energy for analysis and synthesis of soft maps. *Computer Graphics Forum 32*, 5 (2013), 197–206.

[SK17] SHAMAI G., KIMMEL R.: Geodesic distance descriptors. In *Proceedings/CVPR* (Honolulu, 2017), IEEE.

[SNB*12] SOLOMON J., NGUYEN A., BUTSCHER A., BEN-CHEN M., GUIBAS L.: Soft maps between surfaces. *Computer Graphics Forum 31*, 5 (2012), 1617–1626.

[SPKS16] SOLOMON J., PEYRÉ G., KIM V. G., SRA S.: Entropic metric alignment for correspondence problems. *ACM Transactions on Graphics (TOG) 35*, 4 (2016), 72:1–72:13.

[Tu11] TU L. W.: *An Introduction to Manifolds* (2nd edition). Springer-Verlag New York, 2011.

[VLB*17] VESTNER M., LÄHNER Z., BOYARSKI A., LITANY O., SLOSSBERG R., REMEZ T., RODOLÀ E., BRONSTEIN A., BRONSTEIN M., KIMMEL R., CREMERS D.: Efficient deformable shape correspondence via kernel matching. In *Proceedings - 2017 International Conference on 3D Vision, 3DV 2017* (Qingdao, 2017), IEEE, pp. 517–526.

[VLR*17] VESTNER M., LITMAN R., RODOLÀ E., BRONSTEIN A., CREMERS D.: Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *Proceedings - 30th IEEE Conference on Computer Vision and Pattern Recognition, CVPR 2017* (Honolulu, 2017), IEEE, pp. 6681–6690.

[WGBS18] WANG L., GEHRE A., BRONSTEIN M. M., SOLOMON J.: Kernel functional maps. *Computer Graphics Forum 37*, 5 (2018), 27–36.

[Wit83] WITKIN A. P.: Scale-space filtering. In *Proceedings of the Eighth International Joint Conference on Artificial Intelligence - Volume 2* (San Francisco, CA, USA, 1983), IJCAI'83, Morgan Kaufmann Publishers Inc., pp. 1019–1022.

[WSSC11] WINDHEUSER T., SCHLICKEWEI U., SCHMIDT F. R., CREMERS D.: Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In *Proceedings/ICCV* (Barcelona, 2011), IEEE.

[ZWW*10] ZENG Y., WANG C., WANG Y., GU X., SAMARAS D., PARAGIOS N.: Dense non-rigid surface registration using high-order graph matching. In *Proceedings/CVPR* (San Francisco, 2010), IEEE.