

Deep orientation-aware functional maps: Tackling symmetry issues in Shape Matching

Nicolas Donati
LIX, École Polytechnique

nicolas.donati@polytechnique.edu

Etienne Corman
CNRS, Inria

etienne.corman@cnrs.fr

Maks Ovsjanikov
LIX, École Polytechnique

maks@lix.polytechnique.fr

Abstract

State-of-the-art fully intrinsic network for non-rigid shape matching are unable to disambiguate between shape inner symmetries. Meanwhile, recent advances in the functional map framework allow to enforce orientation preservation using a functional representation for tangent vector field transfer, through so-called complex functional maps. Using this representation, we propose a new deep learning approach to learn orientation-aware features in a fully unsupervised setting. Our architecture is built on DiffusionNet, which makes our method robust to discretization changes, while adding a vector-field-based loss, which promotes orientation preservation without using (often unstable) extrinsic descriptors. Our source code is available at: <https://github.com/nicolasdonati/DUO-FM>.

1. Introduction

Learning for non-rigid shape correspondence is a key problem in 3D shape analysis with applications ranging from statistical shape analysis [6, 36] to deformation or texture transfer [4]. Early approaches have focused either on learning informative features so that corresponding points have similar feature descriptors, e.g., [26], or modeling shape correspondence as a semantic segmentation problem. Approaches in the latter category, e.g., [28, 31, 37, 51] aim to predict, for every point on the surface, the corresponding vertex id on some ground truth template shapes. Unfortunately both approaches impose very little consistency between individual point correspondence predictions, while possibly being sensitive to the underlying shape discretization [43].

More recently, techniques have focused on both predicting and imposing a training loss on the entire *map* between each pair of shapes. This has been greatly facilitated by spectral approaches and especially the functional map representation [33], which encodes a map as a small matrix using the spectral (Laplacian) eigen-basis. A wide range of




Figure 1. Our method aims at producing *orientation preserving* maps for non-rigid 3D shape matching in a fully unsupervised setting through the estimation of descriptors whose gradients also align on source and target shape.

approaches based on both supervised [9, 13, 23] and unsupervised losses [22, 40] have been proposed using the functional map representation. Key to all of these methods is learning feature functions that are then used to predict the functional map *as a whole*. As was shown across multiple recent works, this reduces the amount of necessary training data [13], provides strong regularization promoting smooth maps, makes the learned features robust to changes in discretization [43], and alleviates the requirement of the existence of a fixed template shape.

While using the compact functional map representation introduces a strong bias towards smooth approximately isometric correspondences, it nevertheless leaves room for both orientation-preserving and orientation-reversing correspondences. Although this orientation-agnostic property of functional maps can be useful, e.g., in symmetry detection tasks, in most practical scenarios, the underlying sought correspondence is expected to preserve orientation. Unfortunately, restricting to only orientation-preserving maps is not straightforward while using the functional map representation, and the maps obtained using this framework can easily introduce local and global symmetry flipping (i.e., left/right ambiguity present in many organic shapes). As a result, existing state-of-the-art learning networks require either a supervised loss [13, 23, 43] or a rigid pre-alignment [42] in order to disambiguate symmetries.

In this paper, we demonstrate that these limitations can be overcome by using the recently-proposed *complex functional map* representation [12] that is based on alignment of tangent vector fields (represented as complex functions) rather than real-valued functions.

To achieve this, we propose the first architecture that uses complex functional maps and learns specific features that align tangent bundles on surfaces. The use of the complex structure makes our approach fully orientation-aware, and helps to restrict the space of allowed correspondences to only globally orientation-preserving maps, while helping to regularize learning. We introduce losses adapted to complex functional maps and demonstrate that our network can be trained in a fully unsupervised manner without relying on rigid pre-alignments or ground truth correspondences. More broadly, the vector-valued features learned by our approach provide a novel and informative signal for non-rigid shape analysis tasks.

2. Related works

Non-rigid shape matching is a very rich and well-established research area. Below we review works that are most closely related to ours, focusing on learning-based, and especially unsupervised techniques. We also refer the interested readers to surveys including [5, 41] for a more in-depth overview.

Functional Maps Our method build upon the functional map representation, which was originally introduced in [33] and then extended in a very wide range of follow-up works, e.g., [8, 17, 19, 32, 38, 46, 49, 50] among others. The key advantage of this framework is that it allows to represent and optimize for maps as small-sized matrices, enables strong linear-algebraic regularization, and can even be adapted to the partial setting [25, 39].

An essential step in works using this representation, are the “descriptor” (also known as “probe” [34]) functions that

are used to estimate the underlying functional maps and that must be provided *a priori*. Early methods have exploited axiomatic descriptors such as heat or wave kernel signatures [2, 47], with several attempts aiming to optimize the weights of such descriptors through optimization techniques [9].

Learning-based methods Learning shape correspondence has also been done by treating it as a dense semantic segmentation problem, e.g., [7, 18, 28, 31, 37, 51], among many others, or *via* template alignment [21]. However, such works tend to require significant amount of training data, establish a map to a template, and can fail to generalize under connectivity changes [43].

More closely related to our approach are methods that use learning together with the functional map representation, thus evaluating the map as a whole and allowing to directly train and test on arbitrary shape *pairs*. This was first introduced in FMNet [24], which has demonstrated that it is possible to *refine* given descriptor functions such as SHOT [48] with a deep neural network, to minimize a supervised loss, given some ground truth correspondences. This was later extended to use [13] to extract descriptor functions for functional map estimation directly from the shape’s geometry using point-based feature extractors, with a regularized the map estimation layer.

Unsupervised spectral learning Even closer to ours are spectral approaches that use *unsupervised learning* while exploiting the functional map representation. This was first done on top of FMNet by replacing the supervised loss with either geodesic distance preservation [22] or desirable structural properties in the spectral domain [40]. Other properties such as cycle consistency [20] or unsupervised alignment of heat kernels [3] have also been used to improve efficiency and accuracy.

These methods are attractive since they are fully intrinsic and thus tend to generalize well across pose changes. This, however, is an issue in the presence of intrinsic symmetries such as those present in human shapes. To address this, previous functional maps methods, typically only *refine* given descriptors such as SHOT which carry some extrinsic information [3, 20, 22, 40] but can unfortunately be highly unstable under connectivity changes. More recently, “weak supervision” was advocated in the form of extrinsic pre-alignment [15, 42] to resolve symmetry ambiguity. Finally, Deep Shells [16] performs SHOT feature refinement *jointly* with using the 3D coordinates information to guide the correspondence while avoiding symmetry ambiguities.

Unfortunately, despite significant effort, symmetry ambiguity remains a central problem in unsupervised learning for non-rigid shape matching. This is especially problematic since spectral methods tend to generalize much better to new unseen poses compared to extrinsic methods such as [21]. As we argue in this paper, the symmetry ambiguity

problem is inherent in the functional maps approaches, as the losses used are fully intrinsic thus cannot disambiguate orientation-preserving vs. orientation-reversing maps.

Complex Functional Maps Recently a tool was introduced for geometry processing, called *complex functional maps* [12] by focusing on alignment of *tangent vector fields* rather than functions. Crucially, complex functional maps allow to remove orientation reversing maps from the space of allowed correspondences. As a result, as shown in [12], this can help to gain better control on both orientation, and ultimately symmetry in the computed maps. However, the approach in [12] still uses axiomatic descriptors, and an iterative approach in its pipeline. Therefore, it is unclear how to incorporate this representation into a learning framework, while maintaining accuracy and efficiency.

Contributions Our main contributions are as follows:

1. We introduce a new *orientation-aware unsupervised loss*, using the recently proposed complex functional maps representation [12], that exploits the properties of tangent vector fields.
2. We show that computing complex functional maps directly from gradients of learned estimated features and then imposing an additional loss on these maps helps to regularize currently unstable pipelines with respect to symmetry aliasing.
3. By building upon a recent, robust feature extraction backbone [43], we introduce a fully unsupervised approach, without using extrinsic descriptors or coordinate information, while being robust to significant changes in triangulation.

3. Background and Motivation

3.1. Notation, Background & Motivation

Given a pair of non-rigid shapes, M, N , represented as triangle meshes, our main goal is to estimate a map $\varphi : M \rightarrow N$ in an unsupervised manner.

Functional Maps In this work, we use the functional map framework, which has recently given rise to state-of-the-art supervised [13, 23, 43] and unsupervised [16, 20, 22, 40] learning-based non-rigid shape correspondence methods.

The key idea behind the functional maps approaches is that any correspondence can be represented compactly as a small-sized matrix. Specifically, given any map $\varphi : M \rightarrow N$, it can first be encoded as a binary matrix Π_{MN} s.t. $\Pi_{MN}(i, j) = 1$ if and only if $\varphi(i) = j$ where i and j are vertices on M and N respectively. The associated *functional map* \mathbf{C}_{NM} is given as $\mathbf{C}_{MN} = \Phi_M^\dagger \Pi_{MN} \Phi_N$ where Φ_M, Φ_N are matrices storing as columns the first k

eigenfunctions of the Laplace-Beltrami operators of shapes M, N , while \dagger is the Moore-Penrose pseudo-inverse. Note that \mathbf{C}_{MN} is of size $k \times k$ with k typically between 20–100, and is typically orders of magnitude smaller than Π_{MN} .

In addition to allowing to *represent* any map in a reduced basis, the functional map representation also allows to *recover* the underlying map Π_{MN} by exploiting \mathbf{C}_{MN} . Here and throughout the paper we adopt the notation from [34] where objects in the reduced basis are denoted in bold.

The basic pipeline for map recovery, introduced in [33], assumes the presence of some descriptor functions that are expected to be preserved under the unknown mapping. If $\mathbf{A}_N, \mathbf{A}_M$ are the coefficients of descriptors in the basis Φ_M and Φ_N , the optimal functional map \mathbf{C}_{MN} is computed as:

$$\min_{\mathbf{C}_{NM}} \|\mathbf{C}_{NM} \mathbf{A}_N - \mathbf{A}_M\| + \alpha E_{reg}(\mathbf{C}_{NM}). \quad (1)$$

Here the first term promotes preservation of descriptor functions, whereas the second is a regularizer that promotes structural properties; e.g., $E_{reg}(\mathbf{C}_{NM}) = \|\mathbf{C}_{NM} \Delta_N - \Delta_M \mathbf{C}_{NM}\|$, where Δ_M, Δ_N are diagonal $k \times k$ matrices of Laplacian eigenvalues.

The final point-to-point map $\varphi : M \rightarrow N$ can be extracted via nearest neighbor search between the rows of $\Phi_M \mathbf{C}_{NM}$ and those of Φ_N [35]. We refer to [34] for an overview of the functional map representation and its extensions.

Unsupervised learning with functional maps The compactness of the functional map representation \mathbf{C}_{MN} implies that the optimization problem in Eq. (1) reduces to a small scale least squares problem. On the other hand, the quality of the correspondence is intimately tied to the choice of the input descriptor functions. Early approaches have relied on hand-crafted features such as the Wave Kernel Signature [2]. However, more recent methods have focused on *learning* optimal features from data, first in the supervised setting [9, 23] and recently using unsupervised or weakly supervised deep learning, [16, 20, 22, 40, 42].

Our approach is directly inspired by methods in the latter category. The general approach, which is shared by *all* existing unsupervised or weakly supervised methods is to train a neural network \mathcal{F}_Θ that, given a shape M can produce a set of d real-valued functions on M , $\mathcal{F}_\Theta(M) = \{f_1^M, f_2^M, \dots, f_i^M\}$, where $f_i^M : M \rightarrow \mathbb{R}$.

At training time, the network \mathcal{F}_Θ is presented with a set of *pairs* of shapes M, N , and the extracted features $\mathcal{F}_\Theta(M), \mathcal{F}_\Theta(N)$ are used to estimate the functional map \mathbf{C}_{NM} by first projecting the features onto the reduced basis and then solving the optimization in Eq. (1) (typically ignoring the regularization term E_{reg}). The network parameters Θ are then optimized by minimizing a *training loss* which penalizes some structural properties of the estimated functional map \mathbf{C}_{NM} .

The difference between existing methods [16, 20, 22, 40, 42] lies primarily in: a) The choice of feature extractor \mathcal{F}_Θ and b) The training losses used for learning.

Our first observation is that the vast majority of existing unsupervised learning methods have a fundamental limitation in the presence of shapes with intrinsic self-symmetries. We summarize our observation in the following theorem:

Theorem 1. *Given a set of shapes $\{S_i\}$ that all contain an orientation reversing isometric self-symmetry $\{T_i : S_i \rightarrow S_i\}$, s.t. $d_{S_i}(x_j, x_k) = d_{S_i}(T_i(x_j), T_i(x_k))$, then a generic neural network \mathcal{F}_Θ that is trained by any of the losses introduced in [3, 20, 22, 40, 42] has at least two possible solutions that both lead to the global optimum of the loss.*

Proof. See the supplementary materials \square

In this theorem we call a neural network \mathcal{F}_Θ generic if it is capable of producing an *arbitrary* function on the shape. An orientation-reversing self symmetry is a map that is an intrinsic *reflection* such as the left-right symmetry of human shapes, and d_{S_i} is the geodesic distance on S_i .

A direct consequence of this theorem is that regardless of the neural network used, there must be at least two possible global optima, when training the networks using the unsupervised losses in the majority of existing works, in common settings involving symmetric shapes.

Existing methods have primarily tried to overcome this inherent limitation by restricting the power of the neural network, and training it, not from the shape geometry, but from some initial axiomatic features like the SHOT descriptors [48]. Unfortunately, as has been observed in the past, e.g., [37, 43], and as we confirm in our extensive experiments these descriptors are highly sensitive to the triangle mesh structure. Alternatively, some approaches [42] have relied on pre-aligning the shapes in 3D space (and enforcing a consistent forward direction) or used correspondences in 3D space to guide learning [16]. Such an approach, while useful for some categories, can be difficult to enforce for arbitrary non-rigid 3D shapes.

More fundamentally, existing methods do not provide any way to *control* the orientation preservation of the computed correspondences. Indeed, even given pre-aligned shapes, there is no guarantee that the computed features will lead to orientation-preserving maps. Furthermore, in some settings (e.g., symmetry detection) orientation reversing maps might actually be desirable, and existing approaches do not provide direct control to promote such maps.

Finally, several solutions to intrinsic orientation problem have been proposed within the functional map framework. However, most of them are descriptor-based [38, 48] which are unreliable, and form very weak constraints: the map does not have to be orientation preserving but instead encouraged to follow (possibly noisy) descriptors.

3.2. Complex Functional Maps

In this work, we propose to address the limitations mentioned above by exploiting the recently-proposed *complex functional map* representation [12]. The fundamental observation leading to this new tool is that it is challenging to recover orientation, a global information, from a point-to-point mapping giving only local information. Therefore, the construction in [12] relies on global analysis of the *push-forward* $d\varphi : TM \rightarrow TN$ associated to φ whose local properties are related to normal orientation. By definition, a pushforward maps tangent vectors at $p \in M$ to tangent vectors at $\varphi(p) \in N$. $d\varphi$ is also called the differential of φ and is the best linear approximation of the map at point p .

A complex functional map $Q : TM \rightarrow TN$ is a relaxed representation of the pushforward in the sense that it maps any tangent vector field on M to a tangent vector field on N with only one constraint: it must be *complex linear*: $Q(zX) = zQ(X)$, $z \in \mathbb{C}$, $X \in TM$. The adjective complex comes from the fact that tangent vector fields are represented by complex valued functions as in [45].

By definition, the pushforward contains the information of pointwise mapping (which tangent planes on N correspond to tangent planes on M) already carried by φ thus C and Q cannot be independent. Moreover, Q also contains a local orientation information related to the normal orientation. As shown in [12], since Q is complex linear it can only represent *orientation preserving* maps. These fundamental properties have been summarized in Thm. 2 proved in [12] (Sec. 3.5).

Theorem 2. *The complex-linear map Q is a pushforward if and only if there exists an orientation-preserving and conformal diffeomorphism $\varphi : M \rightarrow N$ satisfying:*

$$\langle X, \nabla C(f) \rangle_{TM} = C(\langle Q(X), \nabla f \rangle_{TN}), \quad (2)$$

for all $X \in TM$, $f \in L^2(N)$.

Apart from this complex-linearity, the construction of complex functional maps (which we also call Q -maps) is analogous to standard functional maps described above. They can be written in the spectral basis $\{\Psi^i\}_{i \in (1,k)}$ of the connection Laplacian \mathbf{L} introduced in [45]. In these reduced spaces, Q -maps are small matrices transferring coefficients in the basis Ψ_M to coefficients in Ψ_N . A point-to-point map is extracted from a nearly isometric \mathbf{Q} using a nearest-neighbor search on Dirac functions: $\Pi_{MN} = \text{NNsearch}(\text{div}_M \Psi_M, \text{div}_N \Psi_N \mathbf{Q})$.

Moreover, a pushforward \mathbf{Q} is isometric if and only if it commutes with the connection Laplacian: $\mathbf{Q}\mathbf{L}_M = \mathbf{L}_N\mathbf{Q}$. Finally, as shown in Sec. 3.6 of [12] a necessary condition for \mathbf{Q} to represent a pushforward is that \mathbf{Q} must be orthogonal *i.e.* $\mathbf{Q}^*\mathbf{Q} = \mathbf{I}$ where $*$ denotes the complex transposition. Here, unlike functional maps, orthogonality is not equivalent to area-preservation.




Figure 2. Overview of our unsupervised network. We extract source and target descriptors D_M and D_N using DiffusionNet [43] and then project descriptors onto the Laplace-Beltrami eigenbasis and descriptor gradients onto the connection Laplacian eigenbasis. This leads respectively to the Functional Map block (Section 4.2.1) and the Complex Functional Map block (Section 4.2.2). Losses are imposed on both of these maps (Section 4.3).

A complex functional map can be estimated by minimizing the a simple optimization problem, similar to Eq. (1):

$$\mathbf{Q}_{MN} = \arg \min_{\mathbf{Q}} \|\mathbf{Q}\mathbf{X}_M - \mathbf{X}_N\|_F^2 + E_{\text{reg}}(\mathbf{Q}), \quad (3)$$

where $E_{\text{reg}}(\mathbf{Q}) = w_{\text{Q-ortho}} \|\mathbf{Q}^* \mathbf{Q} - \mathbf{I}\|_F^2 + w_{\text{Q-iso}} \|\mathbf{Q}\mathbf{L}_M - \mathbf{L}_N \mathbf{Q}\|_F^2$

4. Method

In this section, we describe our network in more details. As mentioned in Section 3, a deep functional map pipeline can be decomposed into three different building blocks: the feature extractor (Sec. 4.1), the non-learnable functional map layer (Sec. 4.2) and the loss (Sec. 4.3). We describe our design choices for each of these components, and how they permit orientation-aware unsupervised learning in the subsections below. We also provide a graphic representation of our whole approach in Figure 2.

4.1. Feature extractor

The first major component of our network is the deep feature extractor. Its structure is that of a Siamese network extracting features for both source and target shapes. We use the 3D surface feature extractor introduced in [43] which outputs feature vectors d on n_M, n_N vertices. We denote by $D_M, D_N \in \mathbb{R}^{n_M \times d} \times \mathbb{R}^{n_N \times d}$ the learned source and target features.

The use of DiffusionNet [43] makes our approach highly robust to changes in *shape triangulation* unlike previous unsupervised approaches relying on SHOT descriptors [48]. These methods usually fail if trained and tested triangulations are different, as demonstrated in Section 5. On the

contrary, DiffusionNet is based on robust diffusion, and is consequently largely independent on the choice of shape triangulation. Although learned diffusion is fully intrinsic, the network is aware of the shape orientation because of oriented gradient blocks, as described in [43] (Section 3.4). Therefore, our feature extractor can produce orientation-aware features that we use later to estimate orientation-preserving complex functional maps.

To sum up, we use the network in Figure 2 as a feature extractor to make sure we obtain orientation-preserving maps independently of the triangulation. In the results, we exhibits shapes with anisotropic triangulations hinting a left-right symmetry to the learning pipeline. All methods fall into the bait excepts ours thanks to the orientation-aware unsupervised loss described in Section 4.3.

4.2. The functional map blocks

This block, first introduced in [24] estimates the functional maps in a differentiable way from output source and target features estimated by the feature extractor.

4.2.1 Regularized functional map block

The input features D are projected on spectral Laplace-Beltrami eigenbasis Φ_S to get spectral features $\mathbf{A}_S = \Phi_S^\dagger D_S$ with $S \in \{M, N\}$. The functional map \mathbf{C}_{NM} is then estimated as the solution to the following least-squares problem:

$$\mathbf{C}_{NM} = \arg \min_{\mathbf{C}} \|\mathbf{C}\mathbf{A}_N - \mathbf{A}_M\|_F^2,$$

leading to:

$$\mathbf{C}_{NM} = \mathbf{A}_M \mathbf{A}_N^\dagger. \quad (4)$$

In this work, we use the *regularized* approach, introduced in [13], which incorporates the Laplacian commutativity energy $\|\mathbf{C}_{NM} \mathbf{\Delta}_N - \mathbf{\Delta}_M \mathbf{C}_{NM}\|_F^2$ in a differentiable manner.

4.2.2 Complex functional map block

The complex functional map estimation is analogous to the standard functional map. Input features D are converted to vector fields using the discrete gradient operator G . We visualize these vector field descriptors in Figure 1 (where they are rotated by $\pi/2$ to better see singularities). These vector valued descriptors are then projected in the eigenbasis Ψ of the connection Laplacian. This leads to complex spectral feature vectors $\mathbf{B}_S = \Psi_S^\dagger G_S D_S$ with $S \in \{M, N\}$, and G_S the gradient operator on shape S .

The complex functional map \mathbf{Q}_{MN} is then estimated as the solution to the following least-squares problem:

$$\mathbf{Q}_{MN} = \arg \min_{\mathbf{Q}} \|\mathbf{Q}\mathbf{B}_M - \mathbf{B}_N\|_F^2,$$

whose closed-form solution given simply as:

$$\mathbf{Q}_{MN} = \mathbf{B}_N \mathbf{B}_M^\dagger. \quad (5)$$

As mentioned in Section 3.2, the complex functional map estimated from feature gradients only is a pushforward if the features themselves give rise to an orientation-preserving map. We elaborate on the relation between the two blocks in Section 4.3.1 below. Specifically, we demonstrate that although \mathbf{C} and \mathbf{Q} are estimated independently, they still satisfy the equation of Thm. 2.

4.3. Losses

From estimated \mathbf{C}, \mathbf{Q} we build a loss inspired by SURFMNet [40].

Loss on \mathbf{C} . SURFMNet [40] imposes the estimated functional map \mathbf{C} to be orthogonal, resulting in the first loss L_{ortho} :

$$L_{\text{ortho}}(\mathbf{C}) = \|\mathbf{C}^\top \mathbf{C} - \mathbf{I}\|_F^2 \quad (6)$$

Moreover, it also promotes commutativity between \mathbf{C} and the Laplace-Beltrami operators Δ_M, Δ_N resulting in the second loss L_{iso} :

$$L_{\text{iso}}(\mathbf{C}) = \|\mathbf{C} \Delta_N - \Delta_M \mathbf{C}\|_F^2 \quad (7)$$

We remark that this isometric loss *is not necessary* if we estimate \mathbf{C} with the Laplacian regularizer of [13]. Indeed, the regularizer only produces maps that have a low isometric loss. We therefore only use L_{ortho} in our implementation.

The association of the two losses is generally enough to estimate intrinsically isometric map. The fundamental problem that we propose to remedy here is the fact that these two losses are not in themselves sufficient to rule out intrinsic symmetries. We stress again that many previous works rule out these symmetries based on triangulation only, using the SHOT feature extractor. However these methods are then biased towards training triangulation.

Loss on \mathbf{Q} . As demonstrated in [12] (Section 3.6), a complex functional map will only encode a pointwise map (which will then be orientation preserving) if it is an orthogonal matrix. Hence the complex orthogonal loss $L_{\text{Q-ortho}}$:

$$L_{\text{Q-ortho}}(\mathbf{Q}) = \|\mathbf{Q}^* \mathbf{Q} - \mathbf{I}\|_F^2 \quad (8)$$

Moreover, since we aim for maps which are as isometric as possible, we can also use a complex isometric loss $L_{\text{Q-iso}}$, evaluating the lack of commutativity with the connection Laplacians:

$$L_{\text{Q-iso}}(\mathbf{Q}) = \|\mathbf{Q} \mathbf{L}_M - \mathbf{L}_N \mathbf{Q}\|_F^2 \quad (9)$$

We observe that this isometric loss can also be avoided by computing \mathbf{Q} using a Laplacian regularizer [13]. We therefore only use $L_{\text{Q-ortho}}$ in our implementation.

Loss function. In summary, we use 2 losses that we combine to compute the final loss L_{final} :

$$L_{\text{final}}(\mathbf{C}, \mathbf{Q}) = w_{\text{ortho}} L_{\text{ortho}}(\mathbf{C}) + w_{\text{Q-ortho}} L_{\text{Q-ortho}}(\mathbf{Q})$$

These losses, jointly with the Laplacian regularizers ensure that the learned descriptors result in an isometric map and that this map is orientation-preserving. The whole pipeline remains both light and unsupervised.

4.3.1 Correlation of the two functional map blocks

It has to be noted that we never explicitly use the equation relating \mathbf{C} and \mathbf{Q} unveiled by Theorem 2. In fact, we demonstrate that since both functional maps are nearly isometric and estimated from the same features this relation is always verified.

Theorem 3. *Let M, N be two manifolds, and F_M, F_N surface features such that the functional map C estimated from these features is an isometry. Let Q be the complex functional map computed with the feature gradients as described in Section 4.2.2. Then the maps (C, Q) satisfy Eq. (2), and C is an orientation-preserving isometry.*

Proof. See the supplementary materials □

The isometric assumption is not restrictive in the sense that deep spectral methods already implicitly make this assumption. Moreover, as we demonstrate below, our approach is robust even for shape categories with significant non-isometry.

4.4. Implementation

We implemented our method with Pytorch 1.8 (this version is required to include complex Tensors in the differentiable pipeline) by adapting the open-source implementation of DiffusionNet [43] for the feature extractor and [13] for the functional map block with Laplacian regularizer.

We use WKS descriptors [2] as input signal for the network. We use this descriptor because it is: *a)* Robust to changes in the shape triangulation, and captures the intrinsic geometry of the surface. As shown in the next section that guarantees that the network learns relevant surface information, rather than overfit to the mesh triangulation. *b)* Independent of the embedding of the shape. This makes our approach fully rotationally invariant, and capable of predicting correspondences between arbitrarily rotated shapes. Indeed methods such as [15, 42] depend on pre-aligned datasets to work, which makes them only weakly supervised instead of fully unsupervised.

Our feature extraction network consists of 4 DiffusionNet blocks (a standard DiffusionNet setup [43]), where the 128-dimensional input WKS features are transformed by each block to learned features of same dimension 128, to

Meth / Data	F _r /F _r	F _r /F _a	S _r /S _r	S _r /S _a
SHOT+FMN [24]	5.8	43.	7.0	41.
WKS+GFM [13]	2.0	2.6	2.2	2.3
BCICP [38]	6.1	8.5	11.	14.
ZO [30]	6.1	8.7	7.5	15.
SHOT+UFMN [22]	5.7	42.	9.9	44.
SHOT+DS [16]	1.7	12.	2.5	10.
WKS+DS [16]	8.2	9.5	8.3	20.
WKS+Ours	2.5	3.0	2.6	2.7
Cross Training				
Meth / Data	S _r /F _r	S _r /F _a	F _r /S _r	F _r /S _a
SHOT+FMN [24]	14.	43.	11.	44.
WKS+GFM [16]	9.9	8.4	3.8	3.9
SHOT+UFMN [22]	12.	44.	9.3	43.
SHOT+DS [16]	2.7	15.	5.7	16.
WKS+DS [16]	6.7	12.	9.2	21.
WKS+Ours	2.7	3.1	4.2	4.4

Table 1. Comparative results ($\times 100$) of all main baselines on FAUST and SCAPE re-meshed and anisotropic. Deep Learning methods are displayed with the descriptor input they were fed during training time. The methods shown in the top group are supervised, while the ones below (separated by a double line) are axiomatic or unsupervised. Note that our approach outperforms all unsupervised baselines *without post-processing*, and achieves similar or better performance even to supervised ones.

finally produce 128-dimensional descriptors on source and target shape. As described in Section 4.1, our network is applied in a Siamese way on the two input shapes, using the same weights for feature extraction on source and target.

Parameters In addition to the architecture above, our method has some key hyper-parameters: *a*) The size of both spectral basis: we use $k_C = 50$ for Laplace-Beltrami and $k_Q = 20$ for connection Laplacian *b*) The Laplacian regularizer from [13] in the functional map blocks: we use $\lambda = 10^{-3}$ as recommended in the original paper *c*) The loss hyper-parameters: The loss is focused on map orthogonality since Laplacian-commutativity is enforced previously with the regularizers. We enforce both maps to be “equally” orthogonal, by setting $w_{ortho} = w_{Q-ortho} = 1$.

We train our network with a batch size of 1 for a number of epochs between 5 and 30. We use a learning rate of 10^{-3} with ADAM optimizer [14].

5. Results

In this section, we show that our network can outperform state-of-the-art deep shape matching architectures on standard datasets like FAUST (F_r) [6] and SCAPE (S_r) [1] as-well-as non-isometric datasets like SHREC’19 [29] and SMAL [52]. Following [38], all shapes are remeshed so that

Meth / Data	F _r /Sh _r	S _r /Sh _r	Sh _r /Sh _r
BCICP [38]	15.	15.	15.
ZO [30]	21.	21.	21.
SHOT+DS [16]	27.	24.	24.
WKS+DS [16]	27.	29.	28.
WKS+Ours	6.4	8.4	3.9

Table 2. Comparative results ($\times 100$) on SHREC’19 re-meshed with different train sets, including SHREC’19 re-meshed itself. We compare unsupervised methods on this more challenging dataset, and keep the same notations as in Table 1. We see that our approach gives the best correspondences and their quality is relatively stable with respect to the training set.

they do not share the same connectivity. Moreover, we introduce an anisotropic re-meshing of FAUST (denoted F_a) and SCAPE (denoted S_a), generated with Mmg [10, 11], to demonstrate how some methods overfit to mesh connectivity to disambiguate between intrinsic symmetries. We show our anisotropic remeshings in the supplementary materials.

In all our tables, we denote we denote by F/S a method trained on the dataset F and tested on S.

5.1. Quantitative results

Baselines We compare our method to:

- Axiomatic methods: BCICP [38] and ZoomOut [30] are very efficient for solving close to isometric matching. It should also be noted that these axiomatic methods are slower than a test pass of our method *which does not require post-processing*.
- Supervised methods: FMNet [24] (denoted as FMN) and GeoFmaps [13] (denoted as GFM), using DiffusionNet as feature extractor.
- Unsupervised methods: Unsupervised-FMNet [22] (denoted as UFMN), and the state-of-the-art method Deep Shells [16] (denoted as DS).

For the most relevant baselines, we compare our method, which uses WKS descriptors, to both original networks (trained with SHOT descriptors) and their variant when trained with WKS as input. Consequently, we denote as “SHOT+Net” a Network trained with SHOT and “WKS+Net” its variant with WKS as input.

Anisotropic FAUST and SCAPE For this experiment we train networks on FAUST and SCAPE re-meshed as in [38], and test them on both re-meshed and anisotropic. We report the mean geodesic errors in Table 1.

From the results shown in Table 1, we see that: *a*) Our method is robust to triangulation changes which make SHOT-based methods fail at test time (often because they mistake anisotropy for meaningful geometric information).

Meth / Data	SMAL _r
BCICP [38]	19.
ZO [30]	35.
SHOT+DS [16]	25.
WKS+DS [16]	33.
WKS+Ours	4.8

Table 3. Comparative results ($\times 100$) on SMAL re-meshed dataset. This animal dataset exhibits strong non-isometries, as can be seen on the qualitative result in Figure 3, to which only our method proves to be robust.

Deep Shells [16] is the most robust SHOT-based approach, since its feature extractor uses spectral filters which helps filter out high-frequency overfitting. Still, the quality of the correspondence collapses on anisotropic datasets. *b*) Deep Shells, if presented with an intrinsically symmetric signal (here WKS [2]) as input, fails to learn accurate descriptors, resulting in overall poor correspondence. *c*) Our approach gives the best results among unsupervised networks. Moreover, the quality of the correspondence is often close to that achieved by the best supervised baseline [13].

SHREC’19 re-meshed For this second experiment, we train networks respectively on FAUST, SCAPE and SHREC’19 (denoted as F_r, S_r and Sh_r) re-meshed so that mesh connectivity are different and tested only on SHREC’19. We removed shape 40 from SHREC’19 in this experiment since it is the only partial non-closed shape and therefore outside the scope of this method.

We show in Table 2 that Deep Shells fails to generalize to this test set, even though all meshes have similar number of vertices and well-shaped meshes. In comparison, our method gives significantly better results, as it is a purely geometric approach tailored to exploit surface orientation rather than triangulation details and produce well oriented maps.

SMAL re-meshed We test unsupervised methods on SMAL shapes [52] (again re-meshed so that connectivity is different on every mesh as in [38]). The dataset, originally composed of 49 shapes, is split in 32 training shapes and 17 test shapes. This dataset constitutes the hardest of the three experiments, since its shapes are *often strongly non-isometric*. Indeed, they involve animal shapes of different species, which often poses significant challenges for existing (especially spectral) approaches.

The results of this experiment are reported in Table 3. Observe that our method is the *only* one that produces even reasonable correspondence, significantly outperforming the closest competitor. This highlights the significant additional robustness ensured by our combination of an accurate feature extractor with a range of well-founded geometric regu-




Figure 3. Qualitative results for baselines on the SMAL dataset. The areas where baselines gave the most wrong predictions are highlighted in red.

larizers, which allow our approach to accommodate even non-isometric shapes and learn from limited training data.

5.2. Qualitative results

We additionally provide a qualitative result for our third experiment in Figure 3, comparing our method with baselines on SMAL. We see that our method yields a map very close to ground-truth, even on this challenging example with strong non-isometric distortions. Meanwhile, both axiomatic and learning baselines fail to predict accurate correspondences. We provide another qualitative comparison in the supplementary materials.

6. Conclusion, Limitations & Future Work

To conclude, we introduced a new fully unsupervised way to efficiently learn accurate descriptors for shape matching. These descriptors are guaranteed by our complex functional map loss to be orientation aware in the sense that the resulting map is orientation preserving, making the pipeline robust to aliasing due to shape intrinsic symmetries. Besides, the use of DiffusionNet for a feature extractor enables robustness to changes in shape discretization.

Our approach has several limitations: the loss, which aims at as-isometric-as-possible functional maps is dependent on the fact that the input shapes are not too non-isometric. It would therefore be interesting to use spectral bases adapted to non-isometry, as in [27], which we leave for future works. Another current limitation of our method is that it needs manifold meshes as input. However, good Laplacian operators can be built on point clouds [44], and it would be interesting to check how robust our pipeline is with potentially noisy point clouds as input. Finally, we believe it would be interesting to further leverage complex functional maps in other learning applications, while promoting general rotation-invariant and orientation-preserving maps.

Acknowledgements Parts of this work were supported by the ERC Starting Grant No. 758800 (EXPROTEA) and the ANR AI Chair AIGRETTE.

References

- [1] Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. SCAPE: Shape completion and animation of people. *ACM Transactions on Graphics*, 24(3):408–416, July 2005. 7
- [2] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *2011 IEEE international conference on computer vision workshops (ICCV workshops)*, pages 1626–1633. IEEE, 2011. 2, 3, 6, 8
- [3] Mehmet Aygün, Zorah Löhner, and Daniel Cremers. Unsupervised dense shape correspondence using heat kernels. In *2020 International Conference on 3D Vision (3DV)*, pages 573–582. IEEE, 2020. 2, 4
- [4] Ilya Baran, Daniel Vlasic, Eitan Grinspun, and Jovan Popović. Semantic deformation transfer. In *ACM SIGGRAPH 2009 papers*, pages 1–6. 2009. 1
- [5] Silvia Biasotti, Andrea Cerri, A Bronstein, and M Bronstein. Recent trends, applications, and perspectives in 3d shape similarity assessment. In *Computer Graphics Forum*, volume 35, pages 87–119, 2016. 2
- [6] Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. Faust: Dataset and evaluation for 3d mesh registration. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 3794–3801, 2014. 1, 7
- [7] Davide Boscaini, Jonathan Masci, Emanuele Rodolà, and Michael Bronstein. Learning shape correspondence with anisotropic convolutional neural networks. In *Advances in neural information processing systems*, pages 3189–3197, 2016. 2
- [8] Oliver Burghard, Alexander Dieckmann, and Reinhard Klein. Embedding shapes with Green’s functions for global shape matching. *Computers & Graphics*, 68:1–10, 2017. 2
- [9] Etienne Corman, Maks Ovsjanikov, and Antonin Chambolle. Supervised descriptor learning for non-rigid shape matching. In *European Conference on Computer Vision*, pages 283–298. Springer, 2014. 1, 2, 3
- [10] Charles Dapogny, Cécile Dobrzynski, and Pascal Frey. Three-dimensional adaptive domain remeshing, implicit domain meshing, and applications to free and moving boundary problems. *Journal of computational physics*, 262:358–378, 2014. 7
- [11] Cécile Dobrzynski and Pascal Frey. Anisotropic delaunay mesh adaptation for unsteady simulations. In *Proceedings of the 17th international Meshing Roundtable*, pages 177–194. Springer, 2008. 7
- [12] Nicolas Donati, Etienne Corman, Simone Melzi, and Maks Ovsjanikov. Complex functional maps: A conformal link between tangent bundles. *Computer Graphics Forum*, 41(1):317–334, 2022. 2, 3, 4, 6
- [13] 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. 1, 2, 3, 5, 6, 7, 8
- [14] J. Ba D.P. Kingma. Adam: A method for stochastic optimization. In *ICLR*, 2015. 7
- [15] M. Eisenberger, D. Novotny, G. Kerchenbaum, P. Labatut, N. Neverova, D. Cremers, and A. Vedaldi. Neuromorph: Unsupervised shape interpolation and correspondence in one go. In *IEEE International Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021. 2, 6
- [16] Marvin Eisenberger, Aysim Toker, Laura Leal-Taixé, and Daniel Cremers. Deep shells: Unsupervised shape correspondence with optimal transport. *Advances in Neural Information Processing Systems*, 34, 2020. 2, 3, 4, 7, 8
- [17] Davide Eynard, Emanuele Rodola, Klaus Glashoff, and Michael M Bronstein. Coupled functional maps. In *3D Vision (3DV)*, pages 399–407. IEEE, 2016. 2
- [18] 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. 2
- [19] Anne Gehre, M Bronstein, Leif Kobbelt, and Justin Solomon. Interactive curve constrained functional maps. In *Computer Graphics Forum*, volume 37, pages 1–12. Wiley Online Library, 2018. 2
- [20] Dvir Ginzburg and Dan Raviv. Cyclic functional mapping: Self-supervised correspondence between non-isometric deformable shapes. In *European Conference on Computer Vision*, pages 36–52. Springer, 2020. 2, 3, 4
- [21] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. 3d-coded: 3d correspondences by deep deformation. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 230–246, 2018. 2
- [22] 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. 1, 2, 3, 4, 7
- [23] Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE international conference on computer vision*, pages 5659–5667, 2017. 1, 2, 3
- [24] Or Litany, Tal Remez, Emanuele Rodolà, Alexander M. Bronstein, and Michael M. Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 5660–5668, 2017. 2, 5, 7
- [25] Or Litany, Emanuele Rodolà, Alex M Bronstein, and Michael M Bronstein. Fully spectral partial shape matching. In *Computer Graphics Forum*, volume 36, pages 247–258. Wiley Online Library, 2017. 2
- [26] Roei 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. 1
- [27] Riccardo Marin, Marie-Julie Rakotosaona, Simone Melzi, and Maks Ovsjanikov. Correspondence learning via linearly-

- invariant embedding. *Advances in Neural Information Processing Systems*, 2020. 8
- [28] Jonathan Masci, Davide Boscaini, Michael Bronstein, and Pierre Vandergheynst. Geodesic convolutional neural networks on riemannian manifolds. In *Proceedings of the IEEE international conference on computer vision workshops*, pages 37–45, 2015. 1, 2
- [29] Simone Melzi, Riccardo Marin, Emanuele Rodolà, Umberto Castellani, Jing Ren, Adrien Poulernard, Peter Wonka, and Maks Ovsjanikov. SHREC 2019: Matching Humans with Different Connectivity. In *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2019. 7
- [30] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)*, 38(6):155:1–155:14, Nov. 2019. 7, 8
- [31] Federico Monti, Davide Boscaini, Jonathan Masci, Emanuele Rodola, Jan Svoboda, and Michael M Bronstein. Geometric deep learning on graphs and manifolds using mixture model cnns. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 5115–5124, 2017. 1, 2
- [32] Dorian Nogneng and Maks Ovsjanikov. Informative descriptor preservation via commutativity for shape matching. *Computer Graphics Forum*, 36(2):259–267, 2017. 2
- [33] 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. 1, 2, 3
- [34] 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 *ACM SIGGRAPH 2017 Courses*, SIGGRAPH '17, New York, NY, USA, 2017. Association for Computing Machinery. 2, 3
- [35] Gautam Pai, Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Fast sinkhorn filters: Using matrix scaling for non-rigid shape correspondence with functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 384–393, 2021. 3
- [36] Leonid Pishchulin, Stefanie Wuhler, Thomas Helten, Christian Theobalt, and Bernt Schiele. Building statistical shape spaces for 3d human modeling. *Pattern Recognition*, 67:276–286, 2017. 1
- [37] Adrien Poulernard and Maks Ovsjanikov. Multi-directional geodesic neural networks via equivariant convolution. *ACM Transactions on Graphics (TOG)*, 37(6):1–14, 2018. 1, 2, 4
- [38] Jing Ren, Adrien Poulernard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics (TOG)*, 37(6), 2018. 2, 4, 7, 8
- [39] Emanuele Rodolà, Luca Cosmo, Michael M Bronstein, Andrea Torsello, and Daniel Cremers. Partial functional correspondence. In *Computer Graphics Forum*, volume 36, pages 222–236. Wiley Online Library, 2017. 2
- [40] Jean-Michel Roufousse, Abhishek Sharma, and Maks Ovsjanikov. Unsupervised deep learning for structured shape matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 1617–1627, 2019. 1, 2, 3, 4, 6
- [41] Yusuf Sahillioğlu. Recent advances in shape correspondence. *The Visual Computer*, pages 1–17, 2019. 2
- [42] Abhishek Sharma and Maks Ovsjanikov. Weakly supervised deep functional maps for shape matching. In H. Larochelle, M. Ranzato, R. Hadsell, M. F. Balcan, and H. Lin, editors, *Advances in Neural Information Processing Systems*, volume 33, pages 19264–19275. Curran Associates, Inc., 2020. 2, 3, 4, 6
- [43] Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. Diffusionnet: Discretization agnostic learning on surfaces, 2020. 1, 2, 3, 4, 5, 6
- [44] Nicholas Sharp and Keenan Crane. A Laplacian for Non-manifold Triangle Meshes. *Computer Graphics Forum (SGP)*, 39(5), 2020. 8
- [45] Nicholas Sharp, Yousuf Soliman, and Keenan Crane. The vector heat method. *ACM Trans. Graph.*, 38(3), 2019. 4
- [46] Megeed Shoham, Amir Vaxman, and Mirela Ben-Chen. Hierarchical Functional Maps between Subdivision Surfaces. *Computer Graphics Forum*, 2019. 2
- [47] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer graphics forum*, volume 28, pages 1383–1392. Wiley Online Library, 2009. 2
- [48] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique signatures of histograms for local surface description. In *Proc. ECCV*, pages 356–369. Springer, 2010. 2, 4, 5
- [49] Larry Wang, Anne Gehre, Michael M Bronstein, and Justin Solomon. Kernel functional maps. In *Computer Graphics Forum*, volume 37, pages 27–36. Wiley Online Library, 2018. 2
- [50] Yuexuan Wang, B Liu, K Zhou, and Yu Tong. Vector field map representation for near conformal surface correspondence. In *Computer Graphics Forum*, volume 37, pages 72–83. Wiley Online Library, 2018. 2
- [51] Ruben Wiersma, Elmar Eisemann, and Klaus Hildebrandt. Cnns on surfaces using rotation-equivariant features. *ACM Transactions on Graphics (TOG)*, 39(4):92–1, 2020. 1, 2
- [52] Silvia Zuffi, Angjoo Kanazawa, David Jacobs, and Michael J. Black. 3D menagerie: Modeling the 3D shape and pose of animals. In *Proceedings IEEE Conference on Computer Vision and Pattern Recognition (CVPR) 2017*, pages 5524–5532, Piscataway, NJ, USA, July 2017. IEEE. 7, 8