STAG: Smart Tools and Applications in Graphics (2023)
K. Lupinetti, F. Banterle, N. Capece, and U. Erra (Editors)

Adjoint Bijective ZoomOut: Efficient upsampling for learned linearly-invariant embedding

G. Viganò[1] and S. Melzi[1]

[1]University of Milano-Bicocca, Department of Informatics, Systems and Communication (DISCo)

Abstract

In this paper, we present a novel method for refining correspondences between 3D point clouds. Our method is compatible with the functional map framework, so it relies on the spectral representation of the correspondence. Although, differently from other similar approaches, this algorithm is specifically for a particular functional setting, being the only refinement method compatible with a recent data-driven approach, more suitable for point cloud matching. Our algorithm arises from a different way of converting functional operators into point-to-point correspondence, which we prove to promote bijectivity between maps, exploiting a theoretical result. Iterating this procedure and performing spectral upsampling in the same way as other similar methods, ours increases the accuracy of the correspondence, leading to more bijective correspondences. We tested our method over different datasets. It outperforms the previous methods in terms of map accuracy in all the tests considered.

CCS Concepts

• ***Computing methodologies*** → *Computer graphics; Machine learning;*

1. Introduction

3D shape matching is a widely investigated task in computer vision and geometry processing. It has countless technical applications in engineering, biology, medicine, and other fields where shape registration methods are instrumental [BCBB16; SCT*19]. This problem aims to compute correspondences between pairs of 3D shapes.

The correspondence problem is particularly challenging when several parts of the surfaces undergo different non-rigid deformations. These can make it difficult to derive a reliable correspondence between the two surfaces [DYDZ22]. Moreover, if we deal with point clouds, the complexity grows even more since a successful solution must deal with large variability in shape deformations and must be robust to noise and missing information in the input data.

In the seminal work of Functional maps [OBS*12], shape matching is solved by exploiting a low-rank functional representation, which can be optimized with less computational effort than the original correspondence. To improve the representation's accuracy and increase the estimate performance, a widely adopted approach is to find ways to refine a first approximation of the map, promoting additional properties. For instance, a simple and flexible solution is the ZoomOut algorithm [MRR*19], an iterative method to perform spectral upsampling and increase the accuracy of the estimated correspondence, forcing the map to be an isometry between the shapes. Unfortunately, these functional solutions have been tied so far to the Laplace-Beltrami eigenfunctions [Lev06], which give a spectral representation that strongly depends on the geometry of the shapes at hand. This choice is not stable enough to deal with the point clouds class of data, where the discretization of the Laplace Beltrami operator is often an inaccurate approximation.

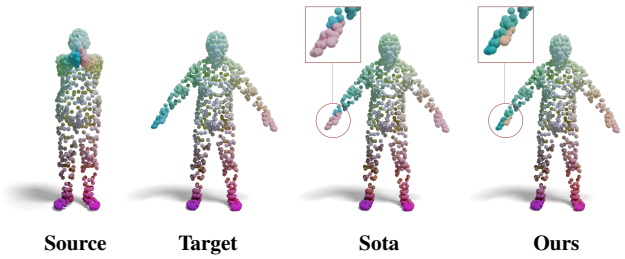

Figure 1: *Performance example of our method. Our method (Ours) helps improve the results of the state-of-the-art (Sota) we consider in our evaluation. In all our visualizations, we depict corresponding points with the same color.*

Linearly-invariant embedding, proposed by [MRMO20], presents a data-driven way of substituting the Laplace Beltrami operator basis with an alternative set of basis learned by a convolutional neural network [QSMG16] to solve the problem of point clouds shape matching, showing state-of-the-art performances. However, this approach suffers from dimensionality limitations, which more standard refinement methods fail to overcome. Indeed,

DOI: 10.2312/stag.20231293

the learned basis is non-orthogonal and optimized for a particular functional representation based on the adjoint operator [HO17]. This suggests that the learned bases are incompatible with the approaches designed for an orthonormal basis in a different functional setting.

In this paper, we target this problem, proposing a *novel* refinement algorithm based on the adjoint operator representation of the correspondence, and so compatible with the data-driven basis. Our refinement deploys a dimensional upsampling to increase the accuracy of the representation, sharing the idea underlying other similar algorithms. At the same time, it exploits a new method to convert the functional map into the actual point-to-point correspondence. We theoretically prove that the proposed method promotes a bijective map, implicitly assuming it as the left inverse of a correspondence between the shapes in the opposite direction. Performing this conversion iteratively makes the map 'more bijective', meaning that the number of points mapped from different preimages is greater, and increases its accuracy. In Figure 1, we visualize the improvement on the state-of-the-art method achieved by our approach for an example pair.

Our method directly promotes a specific feature of the correspondence that is independent of geometrical deformations. For this reason, the proposed algorithm results more stable for point clouds, where geometric information is limited, and performs better for matching scenarios where large deformations occur. Moreover, our method is designed to exploit the adjoint operator representation of the map, so it is compatible with Linearly-invariant embedding. Applying our method to the Linearly-invariant embedding pipeline, we reach the best results in the case of non-isometric point cloud shape matching since we combine an optimal data-driven functional representation with the new refinement properties.

2. Related work

In this section, we briefly review all the approaches to the shape-matching problem most closely related to ours. For an in-depth treatment of the area, we refer to the surveys [VZHC11; Sah20].

Functional maps Our work strongly fits within the functional map framework, originally introduced by [OBS*12]. This work exploits the spectral representation of 3D shapes, converting the point-to-point correspondence between two of them into a map between real-valued functions defined on their surfaces, namely a functional map. Furthermore, by selecting a set of basis functions for these functional spaces, the functional map can be encoded by a more compact matrix. Then, this matrix can be estimated as the linear operator that aligns a given set of probe functions plus other regularization through an optimization process. A general overview of the area is given in [OCB*17]. This seminal work inspired a full range of applications, including [EB17; RCB*17; MMM*20], to name a few.

Different extensions have been developed in recent years. Some methods optimise the conversion between the functional map and the correspondence, such as [RMC15]. A full range of works developed methods to include different properties in the spectral representation, such as isometries [OBS*12], accurate descriptor preservation [NO17], orientation [RPWO18] or bijectivity [ERGB16]. The recent [RMWO21] proposes a new approach to optimize and convert discrete maps to include different properties. Other works combine the spectral representation of the correspondence with different techniques, such as [ELC19], which combines functional maps with extrinsic shape alignment, or [RMOW20], which analyzes the full space of solutions to extract multiple diverse and accurate estimates.

A particularly interesting variation has been made in [HO17], where the adjoint operator has been introduced to give a different functional representation of the correspondence. The role of this operator has been explored in later works as [MRMO20; PRM*21].

Refinement methods The functional map approach represented a revolutionary approach to the shape-matching problem, however, it suffers from diverse limitations. To increase the accuracy of estimated correspondences, different techniques have been developed recently. Our work is strongly related to a class of iterative methods such as the ICP [OBS*12], the ZoomOut [MRR*19], and the Iterative Meta Algorithm [PRM*21]. All these methods operate in the spectral domain performing an iterative upsampling of the spectral dimension of the representation. All these methods aim to give a more precise functional representation without increasing the complexity of the optimization. Various extensions of these works have been proposed recently, such as in the recent [PKO], which proposes a landmark preserving variation of ZoomOut, or in [HRWO20], which exploits dimensional upsampling to promote consistency among a collection of shapes. Moreover, other recent refinement techniques focus on solving scalability [MO23] and smoothness [MRSO22] issues in the functional framework.

Data-driven approaches The majority of spectral approaches rely on axiomatic choices, such as Laplace Beltrami eigenfunctions as spectral basis [Lev06] and dense descriptor fields [ASC11; TSD10] as probe functions. Inspired by the astonishing results that machine learning can reach in other computer vision fields some methods exploit recent machine learning techniques and architecture to solve the correspondence problem [SACO22; GFK*18; TCM*21; LDO22]. Various works tight together the functional map framework with learning approaches. Some methods, including the recent [AO23], inspired by [LRR*17], present data-driven alternatives to the computation of probe functions in a supervised [DSO20], or unsupervised [DCO22] setting. Particularly related to our work is Lineraly-invariant Embedding [MRMO20], which proposes a data-driven alternative to the Laplace Beltrami basis exploiting a point-based architecture [QSMG16] to give a spectral representation of the correspondence for point clouds.

3. Background, notation, and motivations

In this section, we first introduce some necessary elements and fix a notation to understand the following parts of the paper.

3D shapes Theoretically, a 3D shape $\mathcal{X}$ can be modeled as a compact two-dimensional manifold embedded in $\mathbb{R}^3$. Along with the surface, we will consider the space of squared-integrable real-valued functions defined on the shape's surface $\mathcal{L}^2(\mathcal{X})$, defined as

$$\mathcal{L}^2(\mathcal{X}) := \{f : \mathcal{X} \to \mathbb{R}, s.t. \int_{\mathcal{X}} |f(x)|^2 dx < \infty\}$$

This functional space is known to be a Hilbert space, with the inner product

$$\langle f, g \rangle_{\mathcal{L}^2(\mathcal{X})} = \int_{\mathcal{X}} f(x)g(x)dx$$

From a computational point of view, we consider the shape discretized as a point cloud, unordered collections of $n_{\mathcal{X}}$ points, represented by their 3D coordinates. In the discrete setting, each function $f : \mathcal{X} \to \mathbb{R}$ is represented as a vector $\underline{f} = (f_i)_{i=0}^{n_{\mathcal{X}}} = (f(x_i))_{i=0}^{n_{\mathcal{X}}}$

Shape matching If we consider two 3D shapes $\mathcal{X} = \{x_i\}_{i=1}^{n_{\mathcal{X}}}$ and $\mathcal{Y} = \{y_j\}_{j=1}^{n_{\mathcal{Y}}}$, we can assume there exists a meaningful correspondence $T_{\mathcal{XY}}$ which map each point x_i into its corresponding y_j. The goal of shape matching is to estimate this correspondence. $T_{\mathcal{XY}}$ can be represented as an array of vertex index or as a similar permutation matrix $\Pi_{\mathcal{YX}} \in \mathbb{R}^{n_{\mathcal{X}} \times n_{\mathcal{Y}}}$ such that $\Pi_{\mathcal{YX}}(i, j) = 1$, if $T_{\mathcal{XY}}(x_i) = y_j$ and $\Pi_{\mathcal{YX}}(i, j) = 0$ otherwise.

Functional maps Given $T_{\mathcal{XY}} : \mathcal{X} \to \mathcal{Y}$, and considering $\mathcal{L}^2(\mathcal{X})$ and $\mathcal{L}^2(\mathcal{Y})$, $T_{\mathcal{XY}}$ induces via pull-back a linear map $T^F_{\mathcal{YX}} : \mathcal{L}^2(\mathcal{Y}) \to \mathcal{L}^2(\mathcal{X})$, such that

$$\forall g \in \mathcal{L}^2(\mathcal{Y}),\ T^F_{\mathcal{YX}}(g) = g \circ T_{\mathcal{XY}}.$$

If we equip the functional spaces with bases $\phi_{\mathcal{X}} = [\phi^i]_{i=0}^k$ and $\phi_{\mathcal{Y}} = [\phi^j]_{j=0}^k$, we can represent each function as a vector of the coefficients of its representation in basis, thus for each $f \in \mathcal{X}, g \in \mathcal{Y}$ there exist vector of coefficients $\underline{a} = (a_i)$ and $\underline{b} = (b_j)$ such that $f = \sum_l a_i \phi^i_{\mathcal{X}} = \phi_{\mathcal{X}} \underline{a}$ and $g = \sum_l b_j \phi^j_{\mathcal{Y}} = \phi_{\mathcal{Y}} \underline{b}$. Thanks to this representation and the linearity of the functional map we have:

$$T^F_{\mathcal{YX}}(g) = T^F_{\mathcal{YX}}(\sum_j b_j \phi^j_{\mathcal{Y}}) = \sum_j b_j T^F_{\mathcal{YX}}(\phi^j_{\mathcal{Y}})$$

But $T^F_{\mathcal{YX}}(\phi^j_{\mathcal{Y}}) = \sum_i c_{ij} \phi^i_{\mathcal{X}}$ for some $c_{ij} \in \mathbb{R}$, thus

$$T^F_{\mathcal{YX}}(g) = \sum_j b_j T^F_{\mathcal{YX}}(\phi^j_{\mathcal{Y}}) = \sum_j \sum_i b_j c_{ij} \phi^i_{\mathcal{X}}$$

where $T^F_{\mathcal{YX}}(g) = f = \sum_i a_i \phi^i_{\mathcal{X}}$. Therefore, the functional map $T^F_{\mathcal{YX}}$ can be encoded as a compact matrix $C_{\mathcal{YX}} = [c_{ij}]_{i,j=0}^k$ between the coefficient's spaces, denoted as $\xi_{\mathcal{X}}$ and $\xi_{\mathcal{Y}}$.

In this setting, recalling the matrix representation $\Pi_{\mathcal{YX}}$, the functional matrix can be found as

$$C_{\mathcal{YX}} = (\phi_{\mathcal{X}})^{\dagger} \Pi_{\mathcal{YX}} \phi_{\mathcal{Y}} \tag{1}$$

where † denote the Moore-Penrose pseudoinverse.

Note that $C_{\mathcal{YX}} \in \mathbb{R}^{k \times k}$, where k is the dimension of the basis chosen. If we consider the discrete setting, where shapes are composed of n points, we can select just the first $k = n$ bases in the spectral decomposition to have a full representation. However, to decrease the complexity of the original matching problem we have to select $k << n$ bases, introducing an approximation in the problem, but reducing the complexity of the estimation. Indeed the core of this approach is to represent the correspondence as a map of a reduced dimension, optimized as the map that aligns corresponding functions defined on the shape's surfaces.

This representation is the core aspect of the functional map pipeline, which consists of the following steps:

1. Compute a basis of k eigenfunctions $\phi_{\mathcal{X}}$,$\phi_{\mathcal{Y}}$.
2. Compute a set of K corresponding functions $G_{\mathcal{X}} \subset \mathcal{L}^2(\mathcal{X})$ and $G_{\mathcal{Y}} \subset \mathcal{L}^2(\mathcal{Y})$, discretized as matrices of $\mathbb{R}^{n \times K}$.
3. Optimize $C_{\mathcal{YX}}$ via the following optimization problem

$$C_{\mathcal{YX}} = \arg\min_C \|C(\phi^{\dagger}_{\mathcal{Y}} G_{\mathcal{Y}}) - \phi^{\dagger}_{\mathcal{X}} G_{\mathcal{X}}\|_2^2 + E_{reg}(C) \tag{2}$$

   where $E_{reg}(C)$ is a regularization term, that can include different kinds of energies.
4. Finally, convert the optimized map $C_{\mathcal{YX}}$ in the correspondence $\Pi_{\mathcal{XY}}$. The functional operator maps the indicators functions of corresponding points, which are functions centered in each point of the shapes and are represented in basis by the coefficient vector $d_x = (\phi^i_{\mathcal{X}}(x))_i$, so $C_{\mathcal{YX}} d^{\top}_{T_{\mathcal{XY}}(x)} = d^{\top}_x$. So an efficient way to convert the map is by solving the problem:

$$T_{\mathcal{XY}}(x) = \arg\min_y \|\phi_{\mathcal{Y}}(y) C^{\top}_{\mathcal{YX}} - \phi_{\mathcal{X}}(x)\|_2 \quad \forall x \in \mathcal{X} \tag{3}$$

   or equivalently

$$\Pi_{\mathcal{YX}} = \arg\min_{\Pi} \|\Pi \phi_{\mathcal{Y}} C^{\top}_{\mathcal{YX}} - \phi_{\mathcal{X}}\|_2 \tag{4}$$

   This optimization can be solved by a nearest neighbor search between the rows of $\phi_{\mathcal{Y}} C^{\top}_{\mathcal{YX}}$ and $\phi_{\mathcal{X}}$. So that

$$T_{\mathcal{XY}} = NNSearch(\phi_{\mathcal{Y}} C^{\top}_{\mathcal{YX}}, \phi_{\mathcal{X}}) \tag{5}$$

Adjoint operator There is also an alternative way of representing the correspondence with a spectral approach, introduced in [HO17], and extended in [MRMO20; PRM*21]. As the Riesz theorem states, the functional map $T^F_{\mathcal{YX}}$ univocally induces its 'adjoint' operator $T^A_{\mathcal{XY}} : \mathcal{L}^2(\mathcal{X}) \to \mathcal{L}^2(\mathcal{Y})$, as

$$\langle T^A_{\mathcal{XY}} f\ ,\ g \rangle_{\mathcal{L}^2(\mathcal{Y})} = \langle f\ ,\ T^F_{\mathcal{YX}} g \rangle_{\mathcal{L}^2(\mathcal{X})}$$

Given a $x \in \mathcal{X}$, we define δ_x, namely the Dirac delta centered in x as the functional distribution for which $\langle \delta_x, h \rangle = h(x)$ holds for every function $h \in \mathcal{L}^2(\mathcal{X})$. $T^A_{\mathcal{XY}}$ is the linear operator that theoretically maps δ_x to its corresponding delta $\forall x \in \mathcal{X}$ as stated in the following remark.

Remark. *For each $x \in \mathcal{X}$, we denote δ_x as the Dirac delta centered in x. Then, given the adjoint operator, we have that:*

$$T^A_{\mathcal{XY}} \delta_x = \delta_{T_{\mathcal{XY}}(x)} \tag{6}$$

In the discrete setting the equation (6) is translated into

$$A_{\mathcal{XY}} d_x = d_{T_{\mathcal{XY}}(x)},\ \forall x \in \mathcal{X},$$

where $d_{T_{\mathcal{XY}}(x)} = \phi_{\mathcal{Y}}(T_{\mathcal{XY}}(x))^{\top}$ and $d_x = \phi_{\mathcal{X}}(x)^T$ are the full coefficients of the delta functions in the respective basis. Thus, applying the transpose on both sides, we can write:

$$\phi_{\mathcal{X}}(x) A^{\top}_{\mathcal{XY}} = \phi_{\mathcal{Y}}(T_{\mathcal{XY}}(x))\ \ \forall x \in \mathcal{X} \tag{7}$$

or equivalently in matrix notation

$$\phi_{\mathcal{X}} A^{\top}_{\mathcal{XY}} = \Pi_{\mathcal{YX}} \phi_{\mathcal{Y}}$$

This equation leads to two modifications to the functional map pipeline

- The correspondence induces a linear operator

$$A^{\top}_{\mathcal{XY}} = \phi^{\dagger}_{\mathcal{X}} \Pi_{\mathcal{YX}} \phi_{\mathcal{Y}} = C_{\mathcal{YX}} \tag{8}$$

- We can reconstruct the point-to-point map as

$$T_{\mathcal{XY}}(x) = \arg\min_{y} \|\phi_{\mathcal{X}}(x) A^{\top}_{\mathcal{XY}} - \phi_{\mathcal{Y}}(y)\|_2 \ \ \forall x \in \mathcal{X} \tag{9}$$

or equivalently

$$\Pi_{\mathcal{YX}} = \arg\min_{\Pi} \|\phi_{\mathcal{X}} A^{\top}_{\mathcal{XY}} - \Pi \phi_{\mathcal{Y}}\|_2 \tag{10}$$

And this can be solved by

$$T_{\mathcal{XY}} = NNSearch(\phi_{\mathcal{Y}}, \phi_{\mathcal{X}} A^{\top}_{\mathcal{XY}}) \tag{11}$$

On the left side of Figure 2 there is a graphical representation of the two operators in the functional map setting.

Dimensional upsampling The spectral representation of the correspondence helps minimize the complexity of the problem, creating a trade-off between the complexity of the optimization and the accuracy of the solution. The ZoomOut algorithm [MRR*19] is a method to increase the precision of the representation, leading to more accurate correspondences and keeping the problem complexity unchanged.

Starting from a first estimate of the functional map $C^k_{\mathcal{YX}}$, computed considering k eigenfunctions, the algorithm iteratively performs the following steps:

- $C^k_{\mathcal{YX}}$ is converted into $\Pi^k_{\mathcal{YX}}$ via Equation (4) using k eigenfunctions.
- $\Pi^k_{\mathcal{YX}}$ is converted into $C^{k+1}_{\mathcal{YX}}$ via Equation (1) considering $k+1$ eigenfunctions.

This method is proven to have strong performances in a large variety of situations, despite its simple implementation.

A variation of this algorithm is presented by Iterative Meta Algorithm, as called in [PRM*21]. This algorithm considers the adjoint operator and the equation related to this operator, (10) and (8) to convert the maps. This refinement algorithm has some limitations and performs worse than some other more regularized methods such as ZoomOut, which during upsampling promotes the orthonormality of the maps and so more isometric correspondences.

Basis and descriptors choice The quality of the estimations given by these approaches is strongly tied with the corresponding functions G and the basis ϕ chosen. Originally the corresponding functions chosen were axiomatic and we composed of functions describing landmarks points or descriptors of the shapes, such as [ASC11] or [SOG09]. However, recently several approaches have been proposed to learn the probe functions from data [DSO20; HLR*19].

The basis classically chosen is composed of Laplace Beltrami operator eigenfunctions [Lev06]. However, in dealing with point clouds, this choice could be problematic, since its computation can be unstable and unreliable. For this reason, Linearly-invariant Embedding [MRMO20] proposed a data-driven alternative to the Laplace Beltrami operator to learn basis directly from data exploiting PointNet architecture.

Motivations This data-driven approach reaches state-of-the-art results in some challenging point cloud shape-matching scenarios, and most interestingly, it is compatible with the functional map pipeline. However, it suffers from two major limitations:

- The network is able to learn only basis with a limited dimensionality (20), hielding to worse correspondences if the basis dimension increases.
- Although we can refine the correspondences initialized by Linearly-invariant embedding pipeline using ZoomOut algorithm with Laplace Beltrami operator eigenfunctions, it seems impossible to perform upsampling using directly the learned embedding. Indeed the Zoomout algorithm is incompatible since it relies on the orthonormality of the basis and the functional representation, while the learned embedding is learned following the adjoint operator representation. Moreover, other adjoint-based methods such as the Iterative Meta Algorithm aren't enough regularized to give acceptable results.

In this work we target these limitations, proposing a regularized refinement method that performs spectral upsampling, exploiting a different adjoint-based conversion method for the map, and so compatible with Linearly-invariant embedding pipeline.

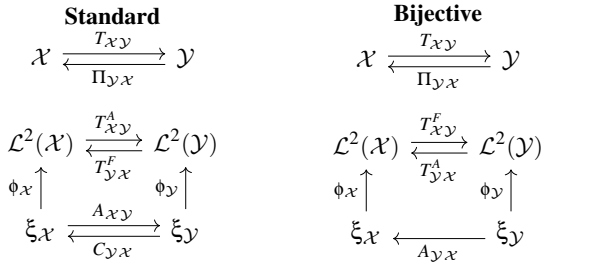

Figure 2: *Functional representation in the two settings. On the left, is the standard functional representation with an adjoint or functional map. On the right, the new representation is considered in which the correspondence induces the functional operators in the opposite direction.*

4. Proposed method

In this section, we show how, by adding a hypothesis on the unknown correspondences and exploiting theoretical results, we can improve conversion and perform a refinement using the bases from Linearly-invariant embedding.

4.1. Adjoint bijective conversion

As stated in [PRM*21], we can extract a pointwise map $T_{\mathcal{XY}}$ from the adjoint $A_{\mathcal{YX}}$, in the opposite direction with respect to what we have seen until now, so that, instead of using (11), we can extract

$$T_{\mathcal{XY}} = NNSearch(\phi_{\mathcal{Y}} A^{\top}_{\mathcal{YX}}, \phi_{\mathcal{X}}) \tag{12}$$

However, this operation is the equivalent of solving the problem

$$T_{\mathcal{XY}}(x) = \arg\min_{y} \|\phi_{\mathcal{Y}}(y) A^{\top}_{\mathcal{YX}} - \phi_{\mathcal{X}}(x)\|_2, \quad \forall x \in \mathcal{X}.$$

This conversion is not justified in the smooth setting since $A_{\mathcal{YX}}$ is not univocally induced by the map $T_{\mathcal{XY}}$, which instead univocally

determines the inverse $A_{\mathcal{XY}}$. For this reason, to adopt the conversion in Equation 12, we need to consider some additional hypotheses on the correspondence. In Figure 2 we can see the differences between this new functional configuration, compared to the standard setting.

Under the assumption that $T_{\mathcal{XY}}$ is bijective, we can prove that the adjoint can also transport the deltas in the opposite direction with respect to equation (6).

Proposition 1. *If $T_{\mathcal{XY}} : \mathcal{X} \to \mathcal{Y}$ is bijective then*

$$T^A_{\mathcal{YX}} \delta_{T_{\mathcal{XY}}(x)} = \delta_x \quad \forall x \in \mathcal{X} \tag{13}$$

Proof. If $T_{\mathcal{XY}}$ is bijective, there exists its inverse $T_{\mathcal{YX}}$ such that

$$T_{\mathcal{YX}}(T_{\mathcal{XY}}(x)) = x \quad \forall x \in \mathcal{X} \tag{14}$$

$T_{\mathcal{YX}}$ induces via pull-back $T^F_{\mathcal{XY}} : \mathcal{L}^2(\mathcal{X}) \to \mathcal{L}^2(\mathcal{Y})$, and its adjoint operator $T^A_{\mathcal{YX}} : \mathcal{L}^2(\mathcal{Y}) \to \mathcal{L}^2(\mathcal{X})$. As we have seen in equation (6), $T^A_{\mathcal{YX}}$ maps the Dirac's δ of $y \in \mathcal{Y}$ into the δ of $x \in \mathcal{X}$, so

$$T^A_{\mathcal{YX}} \delta_y = \delta_{T_{\mathcal{YX}}(y)} = \delta_x$$

, $\forall y \in \mathcal{Y}$ and $x = T_{\mathcal{YX}}(y) \in \mathcal{X}$. So, by definition of the adjoint operator and using condition (14), $\forall f \in \mathcal{L}^2(\mathcal{X})$ and $\forall x \in \mathcal{X}$, we have:

$$\begin{aligned} \langle T^A_{\mathcal{YX}} \delta_{T_{\mathcal{XY}}(x)} \,,\, f \rangle_{\mathcal{L}^2(\mathcal{X})} &= \langle \delta_{T_{\mathcal{XY}}(x)} \,,\, T^F_{\mathcal{XY}} f \rangle_{\mathcal{L}^2(\mathcal{Y})} = \\ &\langle \delta_{T_{\mathcal{XY}}(x)} \,,\, f \circ T_{\mathcal{YX}} \rangle_{\mathcal{L}^2(\mathcal{X})} = \\ &= f(T_{\mathcal{YX}}(T_{\mathcal{XY}}(x))) = f(x) = \langle \delta_x \,,\, f \rangle_{\mathcal{L}^2(\mathcal{X})} \end{aligned}$$

By comparing the first and the last terms in the previous calculation, we have the thesis. □

We should note that we can weaken the bijectivity hypothesis. Indeed the thesis holds if it satisfies the condition (14), which states that there is a left-inverse for $T_{\mathcal{XY}}$. So the proposition holds even for injective correspondences.

In the discrete setting, as we derived Equation (7) from Equation (6), we can derive the following equation from (13):

$$\phi_{\mathcal{Y}}(T_{\mathcal{XY}}(x)) A^{\top}_{\mathcal{YX}} = \phi_{\mathcal{X}}(x), \ \forall x \in \mathcal{X} \tag{15}$$

This means that, if Equation 13 holds, we can extract the point-to-point map $T_{\mathcal{XY}}$ solving the problem

$$T_{\mathcal{XY}}(x) = \arg\min_y \| \phi_{\mathcal{Y}}(y) A^{\top}_{\mathcal{YX}} - \phi_X(x) \|_2, \forall x \in \mathcal{X} \tag{16}$$

or equivalently

$$\Pi_{\mathcal{YX}}(x) = \arg\min_{\Pi} \| \Pi \phi_{\mathcal{Y}} A^{\top}_{\mathcal{YX}} - \phi_X \|_2, \forall x \in \mathcal{X} \tag{17}$$

which can be solved by Equation 12.

We believe that this conversion method helps to find a more constrained correspondence since it implicitly assumes the existence of a left inverse of $T_{\mathcal{XY}}$, by which the adjoint map is induced. This assumption leads to a more injective correspondence, increasing its accuracy. The following proposition formalizes this observation.

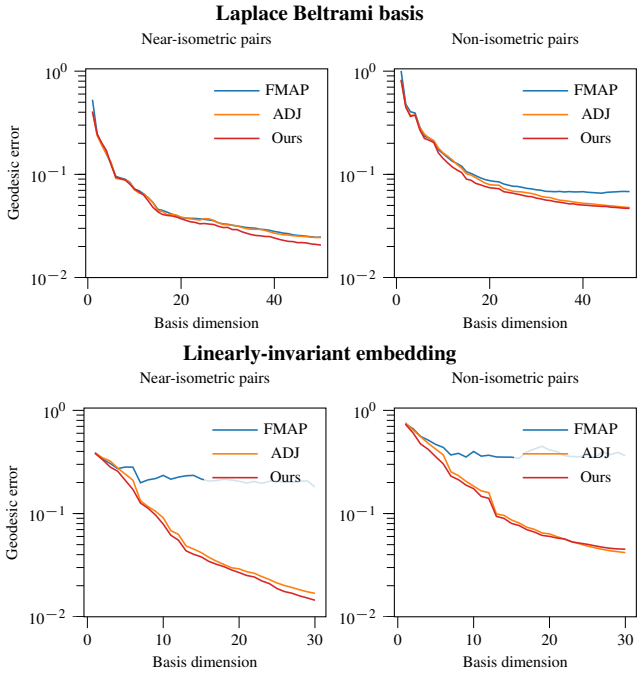

Figure 3: *Recovery error for different basis dimensions in case of near-isometric pairs from the FAUST dataset (left), and non-isometric pairs from the SHREC19 dataset (right), using Laplace Beltrami basis (up) and Linearly-invariant embeddings (down). We compared different ways of converting the maps into the correspondence using Equation (3) (FMAP), (9) (ADJ), (17) (Ours).*

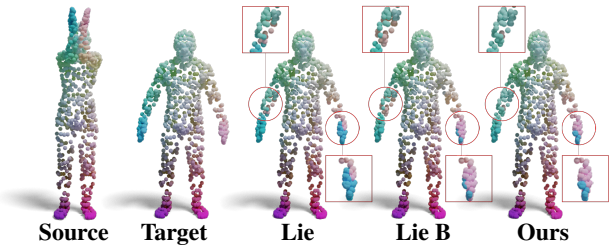

Figure 4: *Qualitative performance example of the new methods. Note that the new approach hields to more precise correspondences, matching more accurately points with unclear positions in the source shape, for example, in the forearm or the hand of these shapes.*

Proposition 2. *Given a surjiective $T_{\mathcal{YX}}$ and $A_{\mathcal{YX}}$ the matrix representation of the adjoint operator $T^A_{\mathcal{YX}}$. The solution to the problem:*

$$T_{\mathcal{XY}}(x) = \arg\min_y \| \phi_{\mathcal{Y}}(y) A^{\top}_{\mathcal{YX}} - \phi_X(x) \|_2 \quad \forall x \in \mathcal{X} \,, \tag{18}$$

satisfies the following:

$$T_{\mathcal{YX}}(T_{\mathcal{XY}}(x)) = x \quad \forall x \in \mathcal{X} \tag{19}$$

Proof. Let $T_{\mathcal{XY}}$ be the exact solution (in the full basis) of the problem 18, then

$$\phi_{\mathcal{Y}}(T_{\mathcal{XY}}(x))A_{\mathcal{YX}}^{\top} - \phi_X(x) = 0$$
$$\implies \phi_X(x) = \phi_{\mathcal{Y}}(T_{\mathcal{XY}}(x))A_{\mathcal{YX}}^{\top}.$$

but $A_{\mathcal{YX}}$ is univocally induced by $T_{\mathcal{YX}}$ by Equation (7). This is the equivalent of saying that for each $y \in \mathcal{Y}$ it holds $\phi_{\mathcal{Y}}(y)A_{\mathcal{YX}}^{\top} = \phi_{\mathcal{X}}(T_{\mathcal{YX}}(y))$. Since this holds for each $y \in \mathcal{Y}$, if $y = T_{\mathcal{XY}}(x)$ for a certain $x \in \mathcal{X}$, then

$$\phi_{\mathcal{X}}(x) = \phi_{\mathcal{Y}}(T_{\mathcal{XY}}(x))A_{\mathcal{YX}}^{\top} = \phi_{\mathcal{X}}(T_{\mathcal{YX}}(T_{\mathcal{XY}}(x)))$$

Finally, thanks to the surjectivity of $T_{\mathcal{YX}}$, this chain of equivalences holds for each $x \in \mathcal{X}$. We have the thesis by comparing the $\phi_{\mathcal{X}}$ argument in the left and right sides of the last equality. □

We remark that the surjectivity hypothesis is necessary to lead the thesis to hold for each point $x \in \mathcal{X}$. In practice, we could have situations where this hypothesis cannot be satisfied. However, the thesis still holds for a subset of points $x \in T_{\mathcal{YX}}(\mathcal{Y}) \subset \mathcal{X}$. This result can be seen as the inverse of Proposition 1. The main consequence is that the correspondence found by our method admits a left inverse, and so it is injective.

Recovery error In Figure 3, we compare this conversion method with the methods arising from the standard setting by analyzing their recovery error. This metric is the accuracy of a correspondence that is converted by a functional map derived directly from ground truth, and it is useful to compare different conversion methods. From the plots, we can state that:

- If we use the functional conversion method based on Equation (3) we cannot get acceptable estimations using the learned embedding.
- The adjoint-based conversions are preferable in the case of non-isometric pairs. In particular, our bijective conversion is slightly better than the standard one in almost all situations.

4.2. Adjoint Bijective ZoomOut

The new relations between $T_{\mathcal{XY}}$ and $A_{\mathcal{YX}}$ highlights how (17) can help find more injective correspondences. So under the bijectivity hypothesis, we can find a new iterative algorithm that promotes bijectivity between the shapes:

- $A_{\mathcal{YX}}^k$ is converted into $\Pi_{\mathcal{YX}}^k$ via equation (17) using k basis.
- $\Pi_{\mathcal{YX}}^k$ is converted into $A_{\mathcal{YX}}^{k+1}$ via

$$A_{\mathcal{YX}}^{k+1} = \left((\Pi_{\mathcal{YX}}^k \phi_{\mathcal{Y}}^{k+1})^{\dagger} \phi_{\mathcal{X}}^{k+1}\right)^{\top} \tag{20}$$

We remark that under the bijectivity hypothesis, Equation (20) is equivalent to the standard way of deriving the adjoint matrix from the correspondence of Equation (8). This can be easily proved since if $T_{\mathcal{XY}}$ is bijective, then $\Pi_{\mathcal{YX}}^{\dagger} = \Pi_{\mathcal{YX}}^{\top} = \Pi_{\mathcal{XY}}$ and so:

$$A_{\mathcal{YX}} = \left((\Pi_{\mathcal{YX}}\phi_{\mathcal{Y}})^{\dagger}\phi_{\mathcal{X}}\right)^{\top} =$$
$$\left(\phi_{\mathcal{Y}}^{\dagger}\Pi_{\mathcal{YX}}^{\dagger}\phi_{\mathcal{X}}\right)^{\top} = \left(\phi_{\mathcal{Y}}^{\dagger}\Pi_{\mathcal{XY}}\phi_{\mathcal{X}}\right)^{\top}$$

This algorithm shares the same "philosophy" of the ZoomOut algorithm with its iterative structure, but, instead of promoting the orthogonality of the functional map and so the isometry of the correspondence, it promotes its bijectivity.

In Figure 4 we give a qualitative example of the performance of the bijective conversion and of the Adjoint Bijective ZoomOut compared to the standard Linearly-invariant embedding pipeline.

5. Results

We evaluate our method on the correspondence problem between non-rigid 3D point clouds, considering the class of human shapes. Although our method is general and can be applied to every shape category, We use this class because of the availability of data and baselines for comparison.

5.1. Experimental setting

Dataset We tested our method on three different datasets. A popular dataset to analyze real identities and poses is FAUST [BRLB14], composed of ten subjects in ten different poses. We considered the version in which 1K points have been sampled for each shape version of it. Additionally, we considered the version in which the points are perturbated by Gaussian noise, introduced in [MRMO20]. Moreover, we challenge our method on SHREC'19 [MMR*19]. Such a dataset is composed of 44 shapes with different discretization, densities, and characteristic features. For each shape in SHREC19, we sampled 500 points uniformly covering the surface via Farthest Point Sampling. Then we add to these 500 points the other 500 points sampled at random to select the final set of 1000 points. Finally, we recovered the ground truth correspondence between these subsampled shapes through the nearest neighbor search in the 3D space. Since the Linearly-invariant embedding approach by itself does not perform well on the majority of these couples, we considered only a portion of pairs (34) in which the Linearly-invariant embedding estimation was reasonable.

Baseline Our main competitor is the Linearly-Invariant Embedding approach [MRMO20] (**LIE**) since we build our refinement explicitly to improve the performance of this method. We considered also a refined version of this approach in which we used the ZoomOut algorithm [MRR*19] with the Laplace Beltrami basis (**LIE+ZO**). We also compared our method with the standard functional map framework (**FMAP**), and its version refined by ZoomOut (**FMAP+ZO**). For the methods that require the Laplace Beltrami basis, we considered the approximation of the operator for point clouds proposed in [SC20]. Our contribution is the application of the new conversion method of (12) applied to the LIE framework (**LIE B**), a version refined via ZoomOut using Laplace Beltrami basis (**LIE B+ZO**) and the application of the new refinement method to the correspondence found by this initialization (**LIE B + Ours**).

Architecture and parameters In the case of Laplace Beltrami operator basis baselines, we computed $\tilde{k} = 30$ eigenvectors and 40 Wave Kernel Signature descriptors [ASC11]. We initialized the map using 20 basis and 40 descriptors, then used the remaining 10 basis functions to perform the refinement.

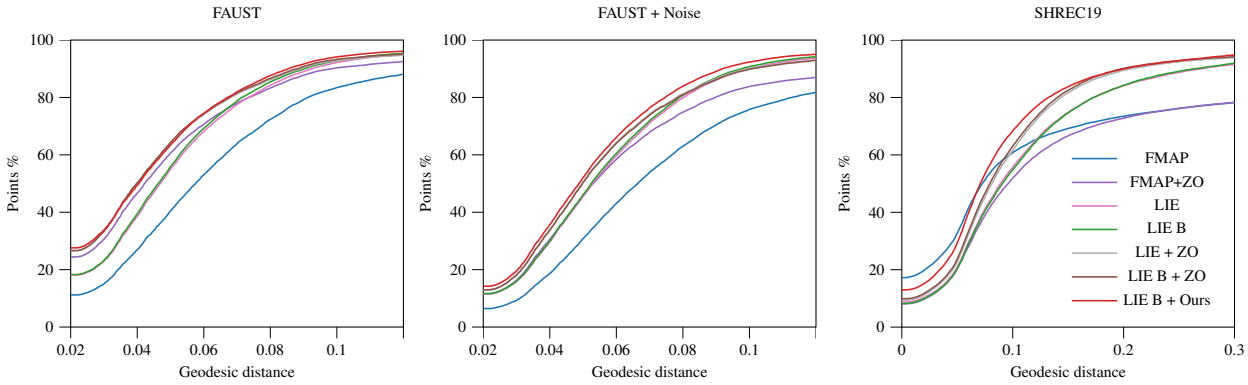

Figure 5: *Error curves for the three datasets considered (from left to right: FAUST, FAUST + Noise, SHREC19). In these plots, the x-axis represents a varying geodesic distance threshold, ρ, and the y-axis shows the average percentage of points p for which* $d(T_{\mathcal{XY}}(p), T^{gt}_{\mathcal{XY}}(p)) <$ ρ. *The higher the curve, the better the performance. The perfect accuracy is represented by a curve that starts from* $(0, 100)$. *We report the legend once, on the right, since the colors are consistent for all the plots.*

In the case of Linearly-invariant embedding, we trained two networks to learn 20 basis and 40 descriptors. Then we trained a different network to have 30 basis atoms, which have been involved in the refinement. All the networks have been trained over 12k Shapes, split in 10k for the training set and 2k for the test set, from the SURREAL dataset [VRM*17] as done in [MRMO20]. We adopt the PointNet architecture [QSMG16] to train every network we compare in our evaluation. A sample code of our implementation can be found online at link `https://github.com/gviga/AB-ZoomOut`.

Metrics To analyze the method's performance we evaluate different properties of the estimate with different metrics. We evaluate the accuracy of $T_{\mathcal{XY}}$ with respect to a known map $T^{gt}_{\mathcal{XY}}$ using the evaluation method presented by [KLF11]. Moreover, we evaluated the bijectivity of the estimated correspondence using the following error:

$$E_{bij}(T_{\mathcal{XY}}) = \|\Pi^{\top}_{\mathcal{YX}} \underline{1} - \underline{1}\|_2 \tag{21}$$

where $\underline{1}$ is a vector of ones, and the orthonormality of the induced map via

$$E_{ortho}(A_{\mathcal{YX}}) = \|A^{\top}_{\mathcal{YX}} A_{\mathcal{YX}} - I\|_2 \tag{22}$$

5.2. Non-rigid point clouds matching

First, we consider the accuracy of the correspondences estimated by the methods considered. In Table 1, we gather the quantitative values of the error metric. In Figure 5, we plot the error curves in the three datasets. Moreover in Figure 6, we give a qualitative comparison of the estimates in two relevant situations. From these results we can conclude that:

- The Adjoint Bijective conversion method leads to more accurate estimates. Indeed, the accuracy of the map found using this conversion method results from 6% to 7% improved.
- The Adjoint Bijective ZoomOut refined maps reach the best accuracy so far, with an increase from 12% up to 18%. In particular, differently from the ZoomOut algorithm, it can increase the accuracy in the FAUST+Noise dataset.

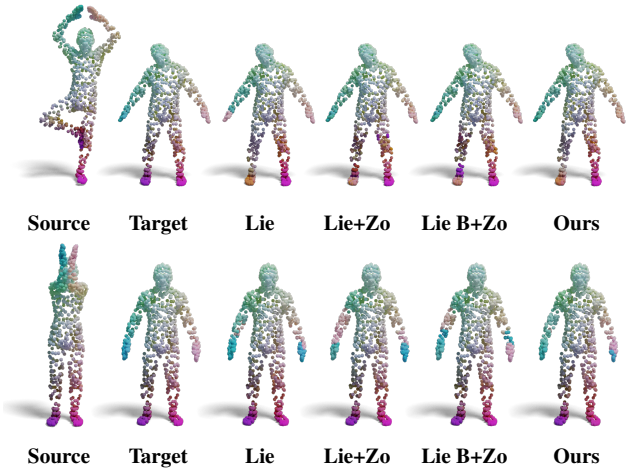

Figure 6: *Two qualitative comparisons. Our method always improves the result of Lie. On the top row, we report an example where our approach does not match correctly some parts of the shapes such as the right foot of the shape. On the bottom row, we show an example where our method improves the accuracy of the correspondence in challenging situations where ZoomOut with Laplace Beltrami basis fails.*

We note that some results for the SHREC'19 dataset are marked with an asterisk. Indeed, in this case, the correspondences computed via the functional map have been recovered directly from the author's [MMR*19] estimations. These estimations have been performed starting from the triangular mesh discretization since they have been proven to be more reliable compared to the one computed directly from the point cloud discretization.

Model	FAUST	FAUST+Noise	SHREC19
FMAP	0.0804	0.1062	0.2260*
FMAP + ZO	0.0581	0.0877	0.2448*
LIE	0.0476	0.0537	0.1495
LIE + ZO	0.0449 (-6%)	0.0568 (+6%)	0.1351 (-10%)
LIE B	0.0440 (-8%)	0.0522 (-3%)	0.1402 (-6%)
LIE B + ZO	0.0418 (-12%)	0.0566 (+5%)	0.1241 (-17%)
LIE B + **Ours**	**0.0364 (-24%)**	**0.0461 (-14%)**	**0.110 (-26%)**

Table 1: *Correspondence accuracy using different methods. The percentages reported are with respect to our main comparison LIE.*

Model	Acc	Bij	Ortho
LIE	0.0476	72.40	10.73
LIE B	0.0440	70.26	14.02
LIE + ZO(lie)	0.354	566.23	17.85
LIE + IMA	0.0482	71.88	6.63
LIE + ZO	0.0449	55.33	1.51
LIE B + ZO	0.0418	55.41	**1.50**
LIE B + **Ours**	**0.0364**	**55.0**	11.81

Table 2: *Values of different metrics on the Faust dataset after different dimensional upsampling methods*

5.3. Refinement analysis

We now focus our analysis on the properties of Adjoint Bijective ZoomOut as a refinement method. In addition to the refinement methods already considered, we also consider the ZoomOut algorithm [MRR*19] performed on the learned embedding (**ZO(lie)**) and the Iterative Meta Algorithm (**IMA**), introduced in [PRM*21]. We report here the results from the experiments on the FAUST dataset. In Figures 7 and 8, we show the value of metrics during the iterations, while in Table 2 we show the quantitative values after the refinement.

First, as anticipated, if we perform the ZoomOut iterations using the learned basis, we get unreliable correspondences, and the accuracy error explodes. We can compare the other refinements only if we rule out this method, which fails. From the right plot of Figure 7 and from the first column of Table 2, we can see that the Adjoint Bijective ZoomOut consistently increases the accuracy of the estimation during the iterations, outperforming the other methods in all the dataset considered.

Finally, we consider the orthonormality of the induced maps. In this case, the ZoomOut algorithm with Laplace Beltrami basis is the only refinement to act on this property. These results show that even if our refinement forces the bijectivity in obtained maps, it does not promote solutions that represent isometries.

Moreover, analyzing the bijectivity error we can see that our method decreases its values compared to relative initialization, and it is the only method that decreases this metric consistently during the iterations, as we can see from the left plot of Figure 8. These results represent quantitative proof of the theoretical motivation that led us to design our novel method.

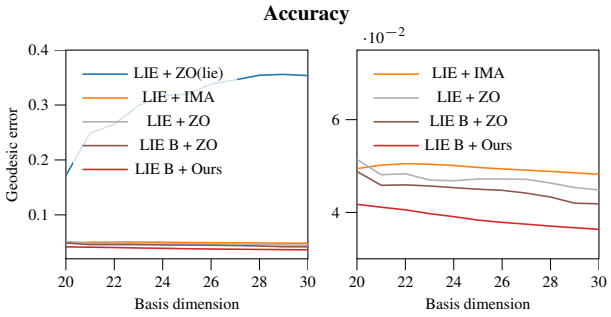

Figure 7: *Accuracy of the estimated correspondence varying the basis dimension. On the left, we include the ZoomOut algorithm performed using Linearly-invariant embedding, on the right we zoom the graph to compare more precisely the other methods.*

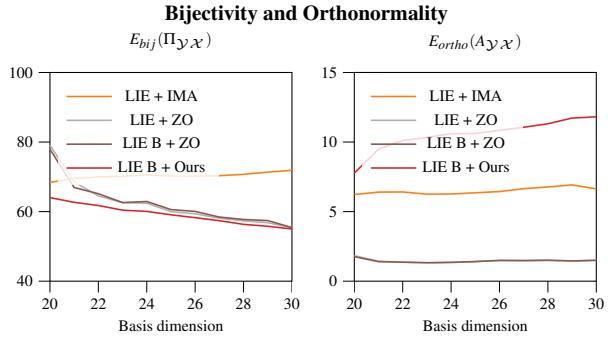

Figure 8: *Bijectivity and orthonormality errors varying the basis dimension.*

6. Conclusions

In the context of data-driven solutions for solving pointwise correspondences among point clouds, we analyzed some limitations of the recently proposed Linearly-invariant embedding pipeline [MRMO20]. To overcome these limitations, we compared different ways of converting a functional operator into a point-to-point correspondence, proving theoretically that our new conversion method promotes bijectivity between shapes. As a consequence, the accuracy of the solution increases using this new method. Moreover, using the bijectivity as a prior constraint, we propose an iterative refinement algorithm to perform spectral upsampling and refine pointwise maps compatible with the Linearly-invariant embedding pipeline.

The results shown by the experiments back up our theoretical analysis. We evaluate the obtained maps through standard metrics, proving the accuracy improvements of the refined maps. Moreover, we select a metric to assess the bijectivity of the estimated correspondence quantitatively, testing the theoretical observations on which we based our work. Our model represents a valid adjoint-based alternative to the ZoomOut algorithm and a preferable choice in the case of Linearly-invariant embedding.

7. Limitations and future work

Our method suffers from some limitations. First, since it is presented as an ad-hoc refinement for the Linearly-invariant embedding pipeline, it cannot perform well in situations where Linearly-invariant embedding does not initialize the map accurately. Moreover, considering the Laplace Beltrami basis, the ZoomOut algorithm is still the preferred choice for spectral upsampling.

However, our work sheds light on some properties of spectral solutions for point cloud correspondence and an application of a new implicit bijectivity constraint, so it has more general value and gives rise to intriguing future directions. In particular, these approaches could be merged with other refinement techniques, such as the work of [HRWO20] in the case of a collection of shapes. Moreover, it could be applied to other frameworks that exploit non-orthonormal bases, such as the recent work of [HSA*23], in which the role of the adjoint operator is analyzed when using this kind of basis.

Acknowledgements

This work was partially supported by the MUR under the grant "Dipartimenti di Eccellenza 2023-2027" of the Department of Informatics, Systems and Communication of the University of Milano-Bicocca, Italy. We gratefully acknowledge the support of NVIDIA Corporation with the RTX A5000 GPUs granted through the Academic Hardware Grant Program to the University of Milano-Bicocca for the project "Learned representations for implicit binary operations on real-world 2D-3D data".

References

[AO23] ATTAIKI, SOUHAIB and OVSJANIKOV, MAKS. *Understanding and Improving Features Learned in Deep Functional Maps*. 2023. arXiv: 2303.16527 [cs.CV] 2.

[ASC11] AUBRY, MATHIEU, SCHLICKEWEI, ULRICH, and CREMERS, DANIEL. "The wave kernel signature: A quantum mechanical approach to shape analysis". *Computer Vision Workshops (ICCV Workshops), 2011 IEEE International Conference on*. IEEE. 2011, 1626–1633 2, 4, 6.

[BCBB16] BIASOTTI, SILVIA, CERRI, ANDREA, BRONSTEIN, ALEXANDER, and BRONSTEIN, MICHAEL. "Recent Trends, Applications, and Perspectives in 3D Shape Similarity Assessment". *Computer Graphics Forum* 36 (Sept. 2016), 87–119. DOI: 10.1111/cgf.12734 1.

[BRLB14] BOGO, FEDERICA, ROMERO, JAVIER, LOPER, MATTHEW, and BLACK, MICHAEL J. "FAUST: Dataset and evaluation for 3D mesh registration". *Proc. CVPR*. Columbus, Ohio: IEEE, 2014, 3794–3801 6.

[DCO22] DONATI, NICOLAS, CORMAN, ETIENNE, and OVSJANIKOV, MAKS. *Deep Orientation-Aware Functional Maps: Tackling Symmetry Issues in Shape Matching*. 2022. arXiv: 2204.13453 [cs.CV] 2.

[DSO20] DONATI, NICOLAS, SHARMA, ABHISHEK, and OVSJANIKOV, MAKS. "Deep Geometric Functional Maps: Robust Feature Learning for Shape Correspondence". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2020, 8592–8601 2, 4.

[DYDZ22] DENG, BAILIN, YAO, YUXIN, DYKE, ROBERTO M., and ZHANG, JUYONG. *A Survey of Non-Rigid 3D Registration*. 2022. arXiv: 2203.07858 [cs.CV] 1.

[EB17] EZUZ, DANIELLE and BEN-CHEN, MIRELA. "Deblurring and Denoising of Maps between Shapes". *Computer Graphics Forum* 36.5 (2017), 165–174 2.

[ELC19] EISENBERGER, MARVIN, LÄHNER, ZORAH, and CREMERS, DANIEL. "Divergence-Free Shape Correspondence by Deformation". *Computer Graphics Forum*. Vol. 38. 5. Wiley Online Library. 2019, 1–12 2.

[ERGB16] EYNARD, D., RODOLÀ, E., GLASHOFF, K., and BRONSTEIN, M. M. "Coupled Functional Maps". *2016 Fourth International Conference on 3D Vision (3DV)*. Oct. 2016, 399–407 2.

[GFK*18] GROUEIX, THIBAULT, FISHER, MATTHEW, KIM, VLADIMIR G., et al. *3D-CODED : 3D Correspondences by Deep Deformation*. 2018. arXiv: 1806.05228 [cs.CV] 2.

[HLR*19] HALIMI, OSHRI, LITANY, OR, RODOLA, EMANUELE, et al. "Unsupervised learning of dense shape correspondence". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 2019, 4370–4379 4.

[HO17] HUANG, RUQI and OVSJANIKOV, MAKS. "Adjoint Map Representation for Shape Analysis and Matching". *Computer Graphics Forum* 36.5 (2017), 151–163 2, 3.

[HRWO20] HUANG, RUQI, REN, JING, WONKA, PETER, and OVSJANIKOV, MAKS. "Consistent ZoomOut: Efficient Spectral Map Synchronization". *Computer Graphics Forum* 39.5 (2020), 265–278 2, 9.

[HSA*23] HARTWIG, FLORINE, SASSEN, JOSUA, AZENCOT, OMRI, et al. "An Elastic Basis for Spectral Shape Correspondence". *ACM SIGGRAPH 2023 Conference Proceedings*. SIGGRAPH '23. Los Angeles, CA, USA: Association for Computing Machinery, 2023. ISBN: 9798400701597. DOI: 10.1145/3588432.3591518. URL: https://doi.org/10.1145/3588432.3591518 9.

[KLF11] KIM, VLADIMIR G, LIPMAN, YARON, and FUNKHOUSER, THOMAS. "Blended intrinsic maps". *ACM Transactions on Graphics (TOG)*. Vol. 30. 4. ACM. 2011, 79 7.

[LDO22] LI, LEI, DONATI, NICOLAS, and OVSJANIKOV, MAKS. *Learning Multi-resolution Functional Maps with Spectral Attention for Robust Shape Matching*. 2022. arXiv: 2210.06373 [cs.CV] 2.

[Lev06] LEVY, B. "Laplace-Beltrami Eigenfunctions Towards an Algorithm That "Understands" Geometry". *IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06)*. 2006, 13–13 1, 2, 4.

[LRR*17] LITANY, OR, REMEZ, TAL, RODOLÀ, EMANUELE, et al. "Deep Functional Maps: Structured Prediction for Dense Shape Correspondence". Oct. 2017, 5660–5668. DOI: 10.1109/ICCV.2017.603 2.

[MMM*20] MELZI, SIMONE, MARIN, RICCARDO, MUSONI, PIETRO, et al. "Intrinsic/extrinsic embedding for functional remeshing of 3D shapes". *Computers & Graphics* 88 (2020), 1–12. ISSN: 0097-8493 2.

[MMR*19] MELZI, SIMONE, MARIN, RICCARDO, RODOLÀ, EMANUELE, et al. "SHREC 2019: Matching Humans with Different Connectivity". *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association, 2019 6, 7.

[MO23] MAGNET, ROBIN and OVSJANIKOV, MAKS. *Scalable and Efficient Functional Map Computations on Dense Meshes*. 2023. arXiv: 2303.05965 [cs.GR] 2.

[MRMO20] MARIN, RICCARDO, RAKOTOSAONA, MARIE-JULIE, MELZI, SIMONE, and OVSJANIKOV, MAKS. *Correspondence Learning via Linearly-invariant Embedding*. 2020. arXiv: 2010.13136 [cs.CV] 1–4, 6–8.

[MRR*19] MELZI, SIMONE, REN, JING, RODOLÀ, EMANUELE, et al. "ZoomOut: Spectral Upsampling for Efficient Shape Correspondence". *ACM Transactions on Graphics (TOG)* 38.6 (Nov. 2019), 155:1–155:14. ISSN: 0730-0301 1, 2, 4, 6, 8.

[MRSO22] MAGNET, ROBIN, REN, JING, SORKINE-HORNUNG, OLGA, and OVSJANIKOV, MAKS. "Smooth Non-Rigid Shape Matching via Effective Dirichlet Energy Optimization". *2022 International Conference on 3D Vision (3DV)*. IEEE, Sept. 2022. DOI: 10.1109/3dv57658.2022.00061. URL: https://doi.org/10.1109%2F3dv57658.2022.00061 2.

[NO17] NOGNENG, DORIAN and OVSJANIKOV, MAKS. "Informative Descriptor Preservation via Commutativity for Shape Matching". *Computer Graphics Forum* 36.2 (2017), 259–267 2.

[OBS*12] OVSJANIKOV, MAKS, BEN-CHEN, MIRELA, SOLOMON, JUSTIN, et al. "Functional maps: a flexible representation of maps between shapes". *ACM Transactions on Graphics (TOG)* 31.4 (2012), 30:1–30:11 1, 2.

[OCB*17] OVSJANIKOV, MAKS, CORMAN, ETIENNE, BRONSTEIN, MICHAEL, et al. "Computing and Processing Correspondences with Functional Maps". *ACM SIGGRAPH 2017 Courses*. 2017, 5:1–5:62 2.

[PKO] PANINE, MIKHAIL, KIRGO, MAXIME, and OVSJANIKOV, MAKS. "Non-Isometric Shape Matching via Functional Maps on Landmark-Adapted Bases". *Computer Graphics Forum* n/a.n/a () 2.

[PRM*21] PAI, G., REN, J., MELZI, S., et al. "Fast Sinkhorn Filters: Using Matrix Scaling for Non-Rigid Shape Correspondence with Functional Maps". *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. Los Alamitos, CA, USA: IEEE Computer Society, June 2021, 384–393 2–4, 8.

[QSMG16] QI, CHARLES RUIZHONGTAI, SU, HAO, MO, KAICHUN, and GUIBAS, LEONIDAS J. "PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation". *CoRR* abs/1612.00593 (2016). arXiv: 1612.00593. URL: http://arxiv.org/abs/1612.00593 1, 2, 7.

[RCB*17] RODOLÀ, EMANUELE, COSMO, LUCA, BRONSTEIN, MICHAEL M, et al. "Partial functional correspondence". *Computer Graphics Forum*. Vol. 36. 1. Wiley Online Library. 2017, 222–236 2.

[RMC15] RODOLÀ, EMANUELE, MOELLER, MICHAEL, and CREMERS, DANIEL. "Point-wise Map Recovery and Refinement from Functional Correspondence". *Proc. Vision, Modeling and Visualization (VMV)*. 2015 2.

[RMOW20] REN, JING, MELZI, SIMONE, OVSJANIKOV, MAKS, and WONKA, PETER. "MapTree: Recovering Multiple Solutions in the Space of Maps". *ACM Trans. Graph.* 39.6 (Nov. 2020) 2.

[RMWO21] REN, JING, MELZI, SIMONE, WONKA, PETER, and OVSJANIKOV, MAKS. "Discrete Optimization for Shape Matching". *Computer Graphics Forum* 40.5 (2021), 81–96 2.

[RPWO18] REN, JING, POULENARD, ADRIEN, WONKA, PETER, and OVSJANIKOV, MAKS. "Continuous and Orientation-preserving Correspondences via Functional Maps". *ACM Transactions on Graphics (TOG)* 37.6 (2018) 2.

[SACO22] SHARP, NICHOLAS, ATTAIKI, SOUHAIB, CRANE, KEENAN, and OVSJANIKOV, MAKS. *DiffusionNet: Discretization Agnostic Learning on Surfaces*. 2022. arXiv: 2012.00888 [cs.CV] 2.

[Sah20] SAHILLIOĞLU, YUSUF. "Recent advances in shape correspondence". *The Visual Computer* 36.8 (2020), 1705–1721 2.

[SC20] SHARP, NICHOLAS and CRANE, KEENAN. "A Laplacian for Nonmanifold Triangle Meshes". *Computer Graphics Forum* 39.5 (2020). Ed. by JACOBSON, ALEC and HUANG, QIXING, 69–80. DOI: 10.1111/cgf.14069 6.

[SCT*19] SÁNCHEZ-BELENGUER, CARLOS, CERIANI, SIMONE, TADDEI, PIERLUIGI, et al. "Global matching of point clouds for scan registration and loop detection". *Robotics and Autonomous Systems* 123 (Oct. 2019). DOI: 10.1016/j.robot.2019.103324 1.

[SOG09] SUN, JIAN, OVSJANIKOV, MAKS, and GUIBAS, LEONIDAS. "A concise and provably informative multi-scale signature based on heat diffusion". *Computer graphics forum* 28.5 (2009), 1383–1392 4.

[TCM*21] TRAPPOLINI, GIOVANNI, COSMO, LUCA, MOSCHELLA, LUCA, et al. *Shape registration in the time of transformers*. 2021. arXiv: 2106.13679 [cs.CV] 2.

[TSD10] TOMBARI, FEDERICO, SALTI, SAMUELE, and DI STEFANO, LUIGI. "Unique signatures of histograms for local surface description". *Proc. ECCV*. Springer. 2010, 356–369 2.

[VRM*17] VAROL, GÜL, ROMERO, JAVIER, MARTIN, XAVIER, et al. "Learning from Synthetic Humans". *CoRR* abs/1701.01370 (2017). arXiv: 1701.01370. URL: http://arxiv.org/abs/1701.01370 7.

[VZHC11] VAN KAICK, OLIVER, ZHANG, HAO, HAMARNEH, GHASSAN, and COHEN-OR, DANIEL. "A survey on shape correspondence". *Computer Graphics Forum* 30.6 (2011), 1681–1707 2.