

# MDND: Unsupervised Learning Guided by Non-Differentiable Refinement for Shape Correspondence

Qinsong Li<sup>1\*</sup>, Jing Meng<sup>2\*</sup>, Haibo Wang<sup>2</sup>, Shengjun Liu<sup>2†</sup>

<sup>1</sup>Big Data Institute, Central South University, Changsha, Hunan 410083, China

<sup>2</sup>School of Mathematics and Statistics, Central South University, Changsha, Hunan 410083, China  
qinsli.cg@csu.edu.cn, wykqhmj1112@163.com, wang\_haibo2017@163.com, shjliu.cg@csu.edu.cn

## Abstract

Deep functional map frameworks (DFM) for shape correspondence are powerful, yet fundamentally limited by their reliance on end-to-end differentiability. This constraint prevents the integration of highly accurate, non-differentiable refinement techniques, capping their overall performance, especially on challenging non-isometric shapes. To overcome this, we introduce MDND, a novel DFM paradigm built on the principle of merging differentiable and non-differentiable components. Our framework facilitates unsupervised learning guided by an internal, non-differentiable refinement. Specifically, MDND employs a dual-branch architecture: a non-differentiable refinement branch leverages a novel, multi-scale iterative solver to produce highly robust correspondences, acting as a refined target. Concurrently, a fully differentiable branch learns to predict correspondences from features. The entire system is trained end-to-end without supervision by enforcing a consistency loss that compels the differentiable branch to learn from the superior, refined results of the non-differentiable branch. Extensive experiments show that MDND sets a new state-of-the-art, demonstrating remarkable robustness on shapes with non-isometric deformations and topological noise.

**Code** — <https://github.com/AMAWDBAC/MDND>

## Introduction

Establishing correspondences between non-rigid 3D shapes is a fundamental problem in computer vision and graphics, with broad applications in texture transfer (Dinh, Yezzi, and Turk 2005), shape interpolation (Ezuz, Solomon, and Ben-Chen 2019), statistical shape analysis (Bogo et al. 2014), and animation (Sumner and Popović 2004). Despite decades of research, the problem remains highly challenging, especially when shapes undergo severe non-isometric deformations or contain substantial topological noise.

The functional map framework has emerged as a dominant paradigm for non-rigid shape correspondence, revolutionizing the field by recasting the problem from matching points to matching functions between shapes (Ovsjanikov

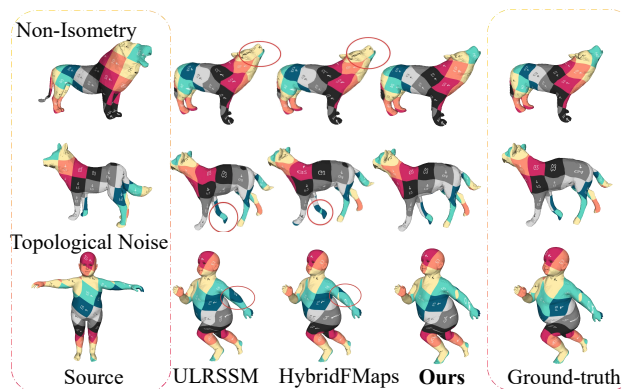


Figure 1: Qualitative comparison on challenging shapes. We visualize correspondences via texture transfer, comparing our method against state-of-the-art approaches like ULRSSM (Cao, Roetzer, and Bernard 2023) and HybridFMaps (Bastian et al. 2024). Our method produces noticeably more accurate and coherent maps in challenging scenarios involving non-isometric deformations (top rows) and significant topological noise (bottom row).

et al. 2012; Pai et al. 2021; Liu et al. 2022). This approach represents the correspondence compactly as a small matrix, the functional map, in a spectral basis typically composed of the eigenfunctions of the Laplace-Beltrami Operator (LBO). The advent of deep learning has given rise to deep functional map (DFM) pipelines (Litany et al. 2017; Donati, Sharma, and Ovsjanikov 2020; Donati, Corman, and Ovsjanikov 2022; Viganò, Ovsjanikov, and Melzi 2025), which learn optimized feature descriptors directly from data to improve matching accuracy. Over time, the field has converged on a powerful, yet highly structured, standard architecture. These pipelines typically employ a Siamese-style network to extract pointwise features, which are then fed into a differentiable layer to solve for the functional map matrix. A key innovation was the introduction of a second, parallel branch that promotes consistency between the functional map and a soft pointwise map derived from feature similarity (Cao, Roetzer, and Bernard 2023; Luo et al. 2025). This two-branch architecture, which enforces consistency between the

\*Equal contribution.

†Corresponding author.

functional and spatial domains, has become the de facto standard for achieving state-of-the-art results in both supervised and unsupervised settings.

This convergence on a standard pipeline, however, exposes a fundamental and shared limitation: an implicit reliance on *end-to-end differentiability*. This constraint is a significant bottleneck, as it precludes the integration of powerful axiomatic refinement methods. These methods (Melzi et al. 2019; Eisenberger, Lahner, and Cremers 2020; Hu et al. 2021; Li et al. 2024), often formulated as iterative optimization procedures, can produce highly accurate correspondences from a reasonable initialization. However, they frequently incorporate complex, non-differentiable operations such as nearest-neighbor searches or discrete optimization steps (Ren et al. 2021), making them incompatible with gradient-based training. Recent efforts to create *differentiable* versions of these axiomatic algorithms have been met with significant trade-offs. Such approaches often require storing and differentiating through large, dense soft-correspondence matrices, leading to quadratic memory complexity that is infeasible for high-resolution meshes (Eisenberger et al. 2020; Li, Donati, and Ovsjanikov 2022; Hu et al. 2023). Furthermore, differentiating through the linear system solvers inherent to many DFM pipelines is known to be numerically unstable (Donati, Sharma, and Ovsjanikov 2020). This suggests a flawed premise in the current research trajectory: by forcing the axiomatic method to become differentiable, its original robustness and elegance are often compromised in favor of a fragile, inefficient, and incomplete approximation that fits the deep learning mold.

In this work, we challenge the necessity of end-to-end differentiability and propose a novel DFM paradigm, termed MDND (**M**erging **D**ifferentiable and **N**on-**D**ifferentiable components). Instead of forcing axiomatic methods into a differentiable framework, we leverage their full, uncompromised power by treating them as non-differentiable supervisory oracles. In our proposed framework, a deep network first predicts an initial correspondence. This map is then refined by a powerful, off-the-shelf axiomatic algorithm, which may contain non-differentiable steps. The resulting high-quality map from this oracle is then used as a pseudo-ground-truth target. A consistency loss between the network’s initial prediction and the oracle’s refined output is backpropagated through the feature-learning network. This process effectively teaches the network to generate features that produce better initializations—ones that the axiomatic method can readily refine to a high-accuracy solution, thereby dramatically improving learning efficiency and final matching precision. A powerful learning framework requires an equally powerful oracle. However, many efficient and popular axiomatic refiners, such as ZoomOut (Melzi et al. 2019) or MWP (Hu et al. 2021), are built upon the LBO eigenbasis. While the LBO basis is intrinsically defined and thus robust to isometries, it fundamentally struggles to characterize the high-frequency, extrinsic details like bending and creasing that define non-isometric deformations. This makes LBO-based refiners inherently ill-suited for the most challenging matching problems. Inspired by recent work showing that a hybrid basis (Bastian et al. 2024)—com-

binning the intrinsic LBO basis with an extrinsic basis derived from an elastic thin-shell energy (ELA) (Hartwig et al. 2023)—is far more expressive for non-isometric shapes, we propose a new iterative refinement by generalizing the principles of state-of-the-art refiners MWP (Hu et al. 2021) to this more powerful hybrid basis. Our refiner is theoretically grounded, efficient, and highly robust to severe non-isometric deformations and topological artifacts. In summary, the main contributions of this work include:

- We propose **MDND**, the first approach to integrate *non-differentiable iterative refinement* into the deep functional map framework, aiming to enhance feature learning and matching accuracy.
- We introduce an effective refinement with theoretical justification, and seamlessly incorporate it as a *supervisory oracle* within the MDND framework, significantly improving robustness in challenging scenarios.
- Extensive experiments across a wide range of challenging conditions demonstrate that our method sets a new state-of-the-art, particularly in cases involving substantial non-isometric deformations and topological noise.

## Related Work

Shape correspondence is a central topic in geometric processing, and we refer the reader to recent surveys for a comprehensive overview (Sahillioğlu 2020; Deng et al. 2022). Our work builds upon three key pillars of research: axiomatic functional maps, deep functional map methods, and the development of expressive spectral bases.

**Axiomatic Functional Maps and Refinement.** The functional map framework, introduced by Ovsjanikov et al. (Ovsjanikov et al. 2012), provides an elegant algebraic representation of correspondences. This foundational work sparked a wave of axiomatic (i.e., non-learning-based) methods aimed at improving map quality. These approaches typically focus on designing sophisticated energy functions to enforce desirable properties, such as orientation preservation (Ren et al. 2018; Donati et al. 2022), consistency with geometric wavelets (Hu et al. 2021; Liu et al. 2022), or multi-shape consistency (Huang et al. 2020; Gao, Zorah, and Bernard 2021). To optimize these energies, powerful iterative refinement strategies have become standard practice. Methods like ZoomOut (Melzi et al. 2019), MWP (Hu et al. 2021), and Smoothshells (Eisenberger, Lahner, and Cremers 2020) iteratively alternate between solving for a functional map and updating a pointwise correspondence. While these axiomatic techniques can achieve high accuracy, their performance is fundamentally limited by their reliance on hand-crafted feature descriptors (Aubry, Schlickewei, and Cremers 2011; Sun, Ovsjanikov, and Guibas 2009; Salti, Tombari, and Di Stefano 2014; Liu et al. 2024a) and spectral bases of LBO, which often fail in the presence of strong non-isometric deformations.

**Deep Functional Maps.** To mitigate the reliance on hand-crafted features, the field has shifted towards Deep Functional Maps (DFM). FMNet (Litany et al. 2017) was the first to learn feature descriptors for functional maps

in a supervised manner. Unsupervised learning was subsequently introduced, using losses based on geodesic distances (Halimi et al. 2019) or structural properties of the functional map (Roufousse, Sharma, and Ovsjanikov 2019). The DFM pipeline has progressively matured with architectural innovations. GeomFmaps (Donati, Sharma, and Ovsjanikov 2020) introduced a differentiable regularized map solver, DUOFMNet (Donati, Corman, and Ovsjanikov 2022), which learned orientation-aware features using complex functional maps (Donati et al. 2022). AttentiveFMaps (Li, Donati, and Ovsjanikov 2022) employed spectral attention to handle varying resolutions. A significant breakthrough came with the introduction of dual-branch architectures (Cao, Roetzer, and Bernard 2023, 2024; Sun et al. 2023; Luo et al. 2025), which enforce consistency between the functional map domain and the pointwise spatial domain. However, a common thread unites these advanced pipelines: their complete reliance on end-to-end differentiability prevents them from incorporating the powerful, non-differentiable solvers developed in the axiomatic literature. Compared to the above works, which only optimize descriptors, recent research has attempted to use generative models to directly learn functional mappings simultaneously (Zhuravlev, Löhner, and Golyanik 2025; Emery et al. 2025). However, this work requires training on large labeled datasets and is less effective than the former.

**Intrinsic-Extrinsic Bases in Functional Maps.** The functional map framework is built upon spectral bases. The eigenvectors of LBO are the conventional choice due to their intrinsic nature, which provides robustness to isometric deformations. However, this very property is a limitation in non-isometric settings, where crucial extrinsic information (e.g., bending and creases) is lost. To address this, recent work has explored more expressive bases. Hartwig et al. (Hartwig et al. 2023) introduced an extrinsic basis derived from the Hessian of a thin-shell elastic energy (the ELA-basis), which is highly sensitive to such fine-grained details. Building on this, HybridFMaps (Bastian et al. 2024) demonstrated that a hybrid spectral space combining the intrinsic LBO basis and the extrinsic ELA basis is significantly more expressive for non-isometric shapes. Our work is motivated by these advancements, and we leverage a hybrid basis to construct a refinement oracle that is robust to the challenging deformations where purely intrinsic methods fail.

## Background

In this section, we provide an overview of the background knowledge related to several key modules integral to our approach.

### Functional Map

Let  $T : \mathcal{M} \rightarrow \mathcal{N}$  be a pointwise map from shape  $\mathcal{M}$  to shape  $\mathcal{N}$ . The induced functional map  $T_F : \mathcal{L}^2(\mathcal{N}) \rightarrow \mathcal{L}^2(\mathcal{M})$  transforms square-integrable real-valued functions from  $\mathcal{N}$  to  $\mathcal{M}$ . Specifically, for any function  $f_{\mathcal{N}} \in \mathcal{L}^2(\mathcal{N})$ , the corresponding function  $f_{\mathcal{M}} \in \mathcal{L}^2(\mathcal{M})$  is defined by the composition  $f_{\mathcal{M}} = T_F(f_{\mathcal{N}}) = f_{\mathcal{N}} \circ T$ . Assuming that  $\mathcal{L}^2(\mathcal{M})$  and  $\mathcal{L}^2(\mathcal{N})$  are equipped with basis functions

$\{\phi_i^{\mathcal{M}}\}_{i \geq 1}$  and  $\{\phi_j^{\mathcal{N}}\}_{j \geq 1}$ , respectively, the functional map can be represented as a matrix  $\mathbf{C} = (c_{ij})$ , where  $c_{ij} = \langle T_F(\phi_j^{\mathcal{N}}), \phi_i^{\mathcal{M}} \rangle$ . Each element of  $\mathbf{C}$  captures the relationship between the two sets of basis functions.

In the discrete setting, shapes  $\mathcal{M}$  and  $\mathcal{N}$  are typically represented as triangular meshes with  $m$  and  $n$  vertices, respectively. The pointwise map  $T$  is denoted by  $\Pi_{\mathcal{M}\mathcal{N}} \in \mathbb{R}^{m \times n}$ , where  $\Pi_{\mathcal{M}\mathcal{N}}(i, j) = 1$  if  $T(i) = j$ , and 0 otherwise. Here,  $i$  and  $j$  represent vertex indices on shape  $\mathcal{M}$  and  $\mathcal{N}$ , respectively. And we regard it as a proper pointwise map. Let  $\Phi_{\mathcal{M}}^{\text{LBO}} \in \mathbb{R}^{m \times k}$  and  $\Phi_{\mathcal{N}}^{\text{LBO}} \in \mathbb{R}^{n \times k}$  denote the matrices containing the first  $k$  discretized Laplacian eigenfunctions for each shape. The functional map  $\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}}$  is given by the projection of  $\Pi_{\mathcal{M}\mathcal{N}}$  onto the corresponding functional basis:

$$\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}} = (\Phi_{\mathcal{M}}^{\text{LBO}})^\dagger \Pi_{\mathcal{M}\mathcal{N}} \Phi_{\mathcal{N}}^{\text{LBO}}, \quad (1)$$

where  $\dagger$  denotes the Moore-Penrose pseudo-inverse. Since the functional map  $\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}}$  in Eq.(1) arises from a pointwise correspondence, we call it *proper functional map* (Ren et al. 2021).

When the pointwise map (and, by extension, the functional map) is unknown, the functional map can be computed by solving the following optimization problem:

$$\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}} = \arg \min_{\mathbf{C}} E_{\text{desc}}(\mathbf{C}) + \alpha E_{\text{reg}}(\mathbf{C}), \quad (2)$$

where  $E_{\text{desc}}(\mathbf{C}) = \|\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}} (\Phi_{\mathcal{N}}^{\text{LBO}})^\dagger \mathbf{F}_{\mathcal{N}} - (\Phi_{\mathcal{M}}^{\text{LBO}})^\dagger \mathbf{F}_{\mathcal{M}}\|_{\text{F}}^2$  enforces descriptor preservation. Here,  $\mathbf{F}_{\mathcal{M}} \in \mathbb{R}^{m \times d}$  and  $\mathbf{F}_{\mathcal{N}} \in \mathbb{R}^{n \times d}$  are  $d$ -dimensional feature matrices for  $\mathcal{M}$  and  $\mathcal{N}$ , respectively. The term  $E_{\text{reg}}(\mathbf{C})$  represents a regularization function that promotes structural consistency in  $\mathbf{C}$ , and  $\alpha$  is a regularization parameter. Once  $\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}}$  is obtained, the pointwise map is derived through a nearest neighbor search in the spectral embedding spaces ( $\Phi_{\mathcal{N}}^{\text{LBO}}$  and  $\Phi_{\mathcal{M}}^{\text{LBO}} \mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}}$ ). However, the quality of the resulting map is often suboptimal. To improve this, various refinement techniques have been proposed (Melzi et al. 2019; Magnet et al. 2022; Donati et al. 2022; Ren et al. 2018; Hu et al. 2021), which iteratively alternate between optimizing the functional and pointwise maps to enhance accuracy.

### Deep Functional Map

Deep functional map (DFM) methods have become state-of-the-art for non-rigid shape correspondence. The standard DFM pipeline, illustrated in Figure 2, consists of three main stages: feature extraction, differentiable map estimation, and unsupervised loss computation.

**Feature Extraction.** Given two input shapes,  $\mathcal{M}$  and  $\mathcal{N}$ , a trainable Siamese network  $\mathcal{F}_\theta$  is employed to extract pointwise feature descriptors,  $\mathbf{F}_{\mathcal{M}}$  and  $\mathbf{F}_{\mathcal{N}}$ , respectively. Here,  $\theta$  represents the learnable network parameters. DiffusionNet (Sharp et al. 2022) has emerged as the de facto standard backbone for this task, as it excels at learning robust features that are invariant to discretization and aware of orientation.

**Differentiable Map Estimation.** The learned features are then passed to a differentiable solver, which computes a functional map,  $\mathbf{C}$ , by minimizing an energy function such as the one specified in Equation (2). The differentiability of

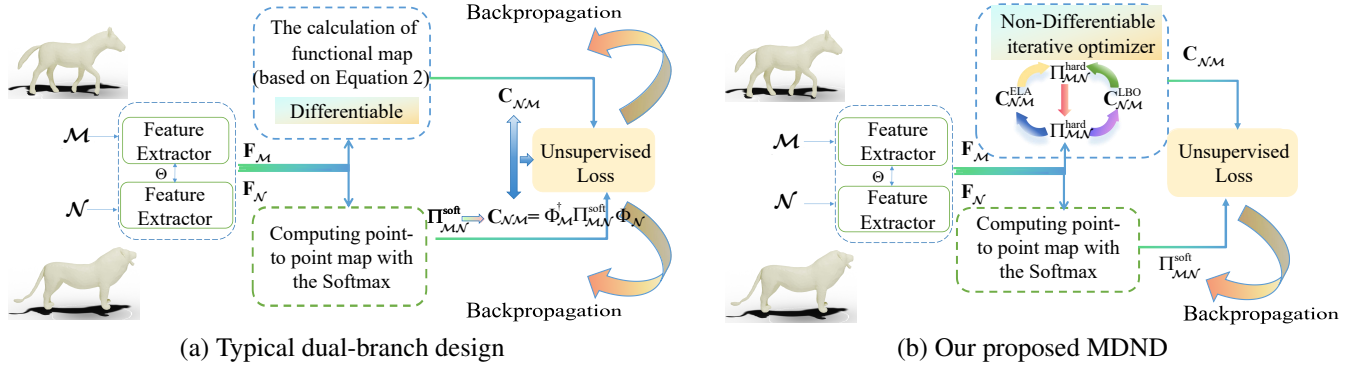


Figure 2: A comparison of deep functional map architectures. (a) The conventional dual-branch framework which is fully differentiable and thus limited to differentiable solvers. (b) Our proposed MDND framework, which introduces a paradigm shift by incorporating a non-differentiable iterative refinement oracle. This oracle generates a high-quality target map, while gradients for training the feature extractor propagate exclusively through the parallel, differentiable branch. This design decouples refinement from learning and avoids potential gradient conflicts.

this step is crucial as it allows gradients to flow back to the feature extraction network during training.

**Unsupervised Losses.** To train the feature extractor  $\mathcal{F}_\theta$  without ground-truth correspondences, several structural regularizers are imposed on the computed functional maps. Foundational losses include an orthogonality regularizer, which encourages the map to be area-preserving, and a bijectivity regularizer, which enforces cycle consistency (i.e., mapping from  $\mathcal{M}$  to  $\mathcal{N}$  and back should approximate the identity). These are formulated as:

$$\mathcal{L}_{\text{orth}} = \|\mathbf{C}_{\mathcal{N}\mathcal{M}}^T \mathbf{C}_{\mathcal{N}\mathcal{M}} - \mathbf{I}\|_F^2 + \|\mathbf{C}_{\mathcal{M}\mathcal{N}}^T \mathbf{C}_{\mathcal{M}\mathcal{N}} - \mathbf{I}\|_F^2, \quad (3)$$

$$\mathcal{L}_{\text{bij}} = \|\mathbf{C}_{\mathcal{N}\mathcal{M}} \mathbf{C}_{\mathcal{M}\mathcal{N}} - \mathbf{I}\|_F^2 + \|\mathbf{C}_{\mathcal{M}\mathcal{N}} \mathbf{C}_{\mathcal{N}\mathcal{M}} - \mathbf{I}\|_F^2. \quad (4)$$

More recently, (Cao, Roetzer, and Bernard 2023) introduced a coupling loss used to enforce consistency between the map computed by the differentiable solver and the one derived from the soft pointwise map  $\Pi_{\mathcal{M}\mathcal{N}}^{\text{soft}}$ , which is derived from the feature similarity using a Softmax operation.:

$$\mathcal{L}_{\text{couple}} = \|\mathbf{C}_{\mathcal{N}\mathcal{M}} - \Phi_{\mathcal{M}}^T \Pi_{\mathcal{M}\mathcal{N}}^{\text{soft}} \Phi_{\mathcal{N}}\|_F^2 \quad (5)$$

This encourages the learned functional map to correspond to a valid pointwise map, significantly improving matching accuracy.

Our MDND framework fundamentally departs from this standard pipeline. Instead of relying on a differentiable solver to compute an initial map, *we leverage a powerful, non-differentiable iterative refinement algorithm*. This algorithm directly optimizes the functional map to a much higher quality, and its output is then used as a supervisory signal to guide the learning of the feature network. This key difference allows us to break free from the constraints of end-to-end differentiability and integrate the strengths of axiomatic refinement into the deep learning process.

## Method

Traditional iterative optimization methods are crucial in functional map computations; however, their direct integra-

tion into deep learning frameworks has been hindered by their non-differentiable nature. To overcome this limitation, we introduce a novel paradigm called MDND (**M**erging **D**ifferentiable and **N**on-**D**ifferentiable **B**ranches), which facilitates information exchange between differentiable and non-differentiable components within deep functional maps. This approach challenges the prevailing assumption that only differentiable operations are suitable for such frameworks. Figure 2 offers a concise overview of our method. Specifically, the first branch generates hard correspondence derived from learned features and serves as the input to produce a well-structured functional map using a novel iterative optimization solver (see Algorithm 1). This map subsequently serves as a supervisory signal to guide the backpropagation of the second Differentiable branch, which operates on soft correspondences. A detailed explanation of our method is provided in the following.

### Feature Extractor

The first core component of our network is the Deep Feature Module, implemented as a Siamese network with shared weights. This module extracts features from the source and target shapes, which consist of  $m$  and  $n$  vertices, respectively. We employ DiffusionNet (Sharp et al. 2022), a state-of-the-art surface feature extractor that utilizes diffusion across the surface to generate features resilient to discretization variations. Furthermore, DiffusionNet incorporates a spatial gradient operation to effectively address bilateral symmetry. The extracted features for the source and target shapes are denoted by  $\mathbf{F}_{\mathcal{M}} \in \mathbb{R}^{m \times d}$  and  $\mathbf{F}_{\mathcal{N}} \in \mathbb{R}^{n \times d}$ , respectively, where  $d$  signifies the dimensionality of the learned features.

### Non-Differentiable Iterative Refinement

Inspired by the efficiency of spectral filtering techniques like MWP (Hu et al. 2021), we propose a novel and efficient iterative optimization method, which we term **Hybrid Wavelet**

---

**Algorithm 1: HWF for Correspondence**


---

**Input:** Initialize pointwise map  $\Pi_{\mathcal{M}\mathcal{N}}^{\text{hard}}$   
**Output:** Refined  $\Pi_{\mathcal{M}\mathcal{N}}^{\text{hard}}$ ,  $(\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{ELA}})^\wedge$ ,  $(\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}})^\wedge$   
Iterative updates between  $\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}}$ ,  $\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{ELA}}$  and  $\Pi_{\mathcal{M}\mathcal{N}}^{\text{hard}}$   
For  $i = 1$  to  $\text{maxIter}$  do  
 $\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{ELA}} = (\Phi_{\mathcal{M}}^{\text{ELA}})^\dagger \Pi_{\mathcal{M}\mathcal{N}}^{\text{hard}} \Phi_{\mathcal{N}}^{\text{ELA}}$   
 $(\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{ELA}})^\wedge = \sum_{l=1}^L g(s_l \Lambda_{\mathcal{M}}^{\text{ELA}}) \mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{ELA}} g(s_l \Lambda_{\mathcal{N}}^{\text{ELA}})$   
 $\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}} = (\Phi_{\mathcal{M}}^{\text{LBO}})^\dagger \Pi_{\mathcal{M}\mathcal{N}}^{\text{hard}} \Phi_{\mathcal{N}}^{\text{LBO}}$   
 $(\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}})^\wedge = \sum_{l=1}^L g(s_l \Lambda_{\mathcal{M}}^{\text{LBO}}) \mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}} g(s_l \Lambda_{\mathcal{N}}^{\text{LBO}})$   
 $\Pi_{\mathcal{M}\mathcal{N}}^{\text{hard}} = \text{NNsearch} \begin{pmatrix} \Phi_{\mathcal{N}}^{\text{ELA}} ((\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{ELA}})^\wedge)^* & \Phi_{\mathcal{M}}^{\text{ELA}} \\ \Phi_{\mathcal{N}}^{\text{LBO}} ((\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}})^\wedge)^* & \Phi_{\mathcal{M}}^{\text{LBO}} \end{pmatrix}$   
end

---

**Filtering (HWF).** The core idea is that a high-quality correspondence can be recovered through a remarkably simple iterative loop: (1) converting a pointwise map to its functional map representation, (2) refining this functional map via spectral filtering, and (3) converting the refined map back to an updated pointwise map. Starting with an initial correspondence, iterating these steps rapidly converges to an accurate solution at a very low computational cost.

Our primary contribution, which distinguishes HWF from MWP, is the **generalization of this filtering process to hybrid spectral bases**. Instead of operating solely on the standard LBO basis, our method leverages a combined basis of both LBO and ELA eigenfunctions. This is crucial for improving robustness in challenging scenarios involving non-isometric deformations and topological noise, where the LBO basis alone is insufficient. While inspired by MWP, we present a completely different theoretical derivation, which is provided in detail in the appendix.

The complete algorithmic workflow is detailed in Algorithm 1. In the algorithm,  $\Phi_{\mathcal{M}}^{\text{LBO}}$  and  $\Phi_{\mathcal{M}}^{\text{ELA}}$  are matrices composed of the LBO and ELA eigenfunctions, while  $\Lambda^{\text{LBO}}$  and  $\Lambda^{\text{ELA}}$  are the diagonal matrices of their corresponding eigenvalues.  $\{g(s_l \lambda)\}_{l=1}^L$  represents a family of spectral manifold wavelet filters,  $\Pi_{\mathcal{M}\mathcal{N}}^{\text{hard}}$  is a vertex index sequence, and  $(\mathbf{C}_{\mathcal{N}\mathcal{M}}^\wedge)^*$  denotes to the adjoint operator of  $\mathbf{C}_{\mathcal{N}\mathcal{M}}^\wedge$ .

In the following section, we will detail how we embed this powerful, non-differentiable HWF algorithm into our deep functional map framework to serve as a supervisory oracle, guiding the learning of robust feature descriptors.

### The MDND Dual-Branch Architecture

Our framework is built on a dual-branch architecture designed to leverage the strengths of both differentiable learning and non-differentiable optimization. One branch operates in a fully differentiable manner to enable gradient-based training, while the other, non-differentiable branch acts as a powerful refinement oracle to provide high-quality supervision.

**Non-Differentiable Refinement Branch.** The purpose of this branch is to generate a highly accurate target correspondence. It begins by computing an initial hard pointwise map,  $\Pi_{\mathcal{M}\mathcal{N}}^{\text{hard}}$ , via a simple nearest-neighbor search on the learned

features  $\mathbf{F}_{\mathcal{M}}$  and  $\mathbf{F}_{\mathcal{N}}$ :

$$\Pi_{\mathcal{M}\mathcal{N}}^{\text{hard}} = \text{NNsearch}(\mathbf{F}_{\mathcal{N}}, \mathbf{F}_{\mathcal{M}}) \quad (6)$$

This hard map, which can be memory-efficiently represented as a single vector of indices in implementation, serves as the input to our Hybrid Wavelet Filtering (HWF) algorithm (Algorithm 1). The HWF iteratively refines this initial guess, producing high-quality functional maps based on hybrid spectral bases,  $(\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{ELA}})^\wedge$  and  $(\mathbf{C}_{\mathcal{N}\mathcal{M}}^{\text{LBO}})^\wedge$ . Since the entire HWF process is parameter-free and non-differentiable, we do not track gradients through this branch, treating its output solely as a supervisory signal.

**Differentiable Learning Branch.** This branch is responsible for learning the feature extractor  $\mathcal{F}_\theta$ . To maintain a differentiable path for backpropagation, we compute a soft pointwise map,  $\Pi_{\mathcal{M}\mathcal{N}}^{\text{soft}}$ , from the learned features. This is achieved by calculating a feature similarity matrix and applying a temperature-scaled Softmax operator:

$$\Pi_{\mathcal{M}\mathcal{N}}^{\text{soft}} = \text{Softmax}(\mathbf{F}_{\mathcal{M}} \mathbf{F}_{\mathcal{N}}^T / \tau), \quad (7)$$

where  $\tau$  is the temperature parameter that controls the softness of the resulting map. The output of this branch,  $\Pi_{\mathcal{M}\mathcal{N}}^{\text{soft}}$ , is a soft correspondence that can be directly used in our loss function to update the network weights.

### Unsupervised Loss Function

Traditional DFM frameworks often rely on multiple structural regularizers, such as orthogonality and bijectivity losses. Balancing the weights of these competing terms can be challenging and can complicate the optimization landscape. To avoid this, we adopt a single, streamlined unsupervised loss function inspired by recent work (Hu et al. 2023). Our loss directly enforces consistency between the output of our two branches: the soft pointwise map from the differentiable branch and the high-quality functional map from the non-differentiable oracle. The alignment loss constructed upon hybrid spectral bases is defined as:

$$\mathcal{L}_{\text{align}} = \|\Phi_{\mathcal{M}} - \Pi_{\mathcal{M}\mathcal{N}}^{\text{soft}} \Phi_{\mathcal{N}} (\mathbf{C}_{\mathcal{N}\mathcal{M}}^\wedge)^*\|_{\text{F}}^2. \quad (8)$$

During backpropagation,  $\mathbf{C}_{\mathcal{N}\mathcal{M}}^\wedge$  obtained by the HWF refiner is treated as a fixed, detached constant, ensuring that gradients only flow through the differentiable branch to update the feature extractor. This elegant formulation allows the network to learn meaningful correspondences by chasing a high-quality, iteratively refined target, without the need for multiple, hand-weighted loss terms.

## Experiments

In this section, we conduct extensive experiments to evaluate our method. We compare MDND against a diverse set of previous approaches across a broad range of challenging scenarios, from near-isometric matching to settings with significant non-isometric deformations and topological noise.

### Implementation Details

Our method was implemented in PyTorch, and all experiments were run on a single NVIDIA RTX 4090 GPU. Following standard evaluation protocols, we report the mean

Method / Dataset	F <sub>r</sub> /	S <sub>r</sub> /	F <sub>r</sub> /	S <sub>r</sub> /	SMAL	DT4D-H inter	DT4D-H intra	TOPKIDS
	F <sub>r</sub>	S <sub>r</sub>	S <sub>r</sub>	F <sub>r</sub>				
BCICP	6.1	11.0	-	-	28.6	-	-	-
ZoomOut	6.1	7.5	-	-	38.4	29.0	4.0	33.7
SmoothShells	2.5	4.7	-	-	36.1	6.4	1.2	10.8
DiscreteOp	5.6	13.1	-	-	38.1	27.6	3.6	35.5
MWP	3.1	4.1	-	-	20.9	25.4	25.4	5.7
FMNet	11.0	17.0	30.0	33.0	42.0	38.0	9.6	-
GeomFmaps	3.5	4.3	4.8	4.0	8.4	4.2	1.9	-
DeepShells	1.9	4.5	6.8	5.5	28.7	31.1	3.4	13.7
DUOFMNet	2.5	2.6	4.2	2.7	6.7	15.8	2.6	-
AttentiveFMaps	1.9	2.1	<u>2.6</u>	1.9	4.4	11.6	1.7	23.4
ConsistentFMaps	2.3	2.4	<u>2.6</u>	2.5	5.2	6.1	1.2	-
RFMNet	1.6	4.5	5.3	2.1	4.4	5.4	1.8	<u>4.9</u>
ULRSSM	<u>1.6</u>	<b>1.8</b>	6.4	4.5	4.4	<u>4.1</u>	<b>0.9</b>	9.2
MSSFMaps	1.9	2.6	4.6	2.8	4.3	<u>4.1</u>	1.8	-
HybridFMaps	<b>1.5</b>	<b>1.8</b>	8.2	<u>1.8</u>	<u>3.3</u>	<b>3.5</b>	<u>1.0</u>	5.0
MSRFMNet	1.7	2.1	<u>2.6</u>	2.0	4.5	36.2	1.5	33.2
DFAFM	<u>1.6</u>	<u>1.9</u>	2.7	1.9	3.9	4.2	<b>0.9</b>	6.3
Ours	<u>1.6</u>	<u>1.9</u>	<b>2.1</b>	<b>1.6</b>	<b>3.1</b>	4.4	<u>1.0</u>	<b>3.5</b>

Table 1: Quantitative comparison with state-of-the-art methods. We report the mean geodesic error ( $\times 100$ ) across datasets representing near-isometric, non-isometric, and topologically noisy scenarios. Methods are grouped by category. The **best** and second-best results are highlighted in bold and underlined, respectively.

geodesic error, normalized by the square root of the source shape’s area. For fair comparison, we use 128-dimensional HKS descriptors (Sun, Ovsjanikov, and Guibas 2009) as input features for all methods and datasets. We set the LBO basis size to 128 and the ELA basis size to 200 for all datasets, except for SMAL, where it was reduced to 100 due to the nature of the shapes. No post-processing or test-time adaptation was applied to our results. For readability, all reported geodesic errors in our tables are multiplied by 100.

### Comparison with State-of-the-Art

**Baselines.** We compare our method against a comprehensive set of recent and influential works, which can be categorized as follows:

- *Axiomatic methods:* BCICP (Ren et al. 2018), ZoomOut (Melzi et al. 2019), Smooth-Shells (Eisenberger, Lahner, and Cremers 2020), DiscreteOp (Ren et al. 2021), and MWP (Hu et al. 2021).
- *Supervised methods:* FMNet (Litany et al. 2017) and GeomFmap (Donati, Sharma, and Ovsjanikov 2020).
- *Unsupervised methods:* A wide range of recent approaches including DeepShells (Eisenberger et al. 2020), DUOFMNet (Donati, Corman, and Ovsjanikov 2022), AttentiveFMaps (Li, Donati, and Ovsjanikov 2022), ConsistentFMaps (Sun et al. 2023), RFMNet (Hu et al. 2023), ULRSSM (Cao, Roetzer, and Bernard 2023), MSSFMaps (Magnet and Ovsjanikov 2024), HybridFMaps (Bastian et al. 2024), MSRFMNet (Liu et al. 2024b) and DFAFM (Luo et al. 2025).

**Near-isometric Shape Matching.** We first evaluate




Figure 3: Cross-dataset generalization results. Correspondence quality is visualized using texture transfer when training and testing on different datasets.

MDND on two standard near-isometric benchmarks: FAUST (Bogo et al. 2014) and SCAPE (Anguelov et al. 2005). As shown in Table 1, our method achieves performance comparable to the state-of-the-art on the standard remeshed FAUST and SCAPE test sets. Notably, our method excels in the generalization tests ( $F \rightarrow S$  and  $S \rightarrow F$ ), indicating its ability to learn robust features that transfer well across different shape collections.

**Non-isometric Shape Matching.** To assess performance on more challenging non-isometric shapes, we use the SMAL (Zuffi et al. 2017) and DT4D-H (Magnet et al. 2022) datasets. On SMAL, which contains various tetrapod species, Table 1 shows that our method significantly outperforms approaches that rely solely on the LBO ba-

Method/Dataset	SCAPE	SMAL	TOPKIDS
Differentiable Solver	2.0	4.2	9.4
Non-Differentiable			
Iterative Refinement	<b>1.9</b>	<b>3.1</b>	<b>3.5</b>

Table 2: Impact of the non-differentiable refinement branch. To measure the effectiveness of our core contribution, we replaced our non-differentiable oracle with a standard differentiable solver. The results show a clear performance drop for the fully differentiable version, confirming that our architecture is crucial for achieving state-of-the-art accuracy.

sis, underscoring the importance of extrinsic information for non-isometric correspondence. On the DT4D-H dataset, our method demonstrates superiority over most competitors in both intra-class and inter-class matching scenarios.

**Matching with Topological Noise.** Finally, we test the robustness of our method on the TOPKIDS dataset (Löhner et al. 2016), which features near-isometric deformations corrupted by significant topological artifacts. The results in Table 1 are striking: our method achieves the best performance, improving upon the next-best approach by a remarkable 30%. This highlights the exceptional robustness of our refinement oracle and learning framework.

## Ablation Studies

Our primary contributions are twofold: the integration of a non-differentiable iterative refinement method into the deep functional map framework, and the generalization of the MWP algorithm to hybrid bases (HWF). To validate the effectiveness of each component, we conduct the following targeted ablation studies.

**(1) Differentiable Solver vs. Non-Differentiable Iterative Refinement.** To verify the benefit of embedding a non-differentiable optimizer within our learning framework, we compare our full MDND model against a baseline that adheres to a more traditional, fully differentiable pipeline. In this baseline, we replace our non-differentiable iterative refinement branch with a standard differentiable solver that computes the functional map directly from the features via Equation (2). We perform experiments on representative datasets covering near-isometric (SCAPE), non-isometric (SMAL), and topologically noisy (TOPKIDS) shapes.

The results, presented in Table 2, demonstrate that our proposed MDND framework with non-differentiable refinement consistently achieves superior matching accuracy across all three categories of datasets. This confirms that using a powerful, non-differentiable oracle to generate a high-quality supervisory signal is more effective than relying on a purely differentiable solver.

**(2) Analysis of the Refinement Oracle: MWP vs. HWF.** To address the limitations of the LBO basis in non-isometric scenarios, we proposed HWF, which generalizes MWP to a hybrid basis. To isolate and validate the effectiveness of this contribution, we configure the MDND framework with three different refinement oracles:

Method/Dataset	SCAPE	SMAL	TOPKIDS
MWP (LBO only)	2.2	4.9	14.7
MWP (ELA only)	2.3	4.1	5.5
HWF (LBO + ELA)	<b>1.9</b>	<b>3.1</b>	<b>3.5</b>

Table 3: Effectiveness of the Hybrid-Basis Refiner (HWF). To isolate the contribution of our proposed HWF algorithm, we compare its performance against single-basis alternatives. The results confirm that combining both the LBO and ELA bases within our HWF refiner is crucial for achieving the best performance.

- MWP (LBO only): The standard MWP algorithm using only the LBO basis.
- MWP (ELA only): MWP adapted to use only the ELA basis.
- HWF (LBO + ELA): Our proposed method using the hybrid basis.

All other experimental settings remain constant. The results, shown in Table 3, clearly indicate that our HWF (LBO + ELA) achieves the best matching performance across all datasets. This validates our hypothesis that generalizing the refinement to a hybrid basis provides a significant advantage, leading to a more robust oracle capable of handling diverse and challenging geometric settings.

## Conclusions

In this paper, we introduced MDND, a novel deep functional map framework that merges the power of deep learning with the robustness of traditional axiomatic optimization. Our core contribution was a dual-branch architecture that leverages a powerful, non-differentiable iterative refinement oracle to provide high-quality supervision for a feature-learning network. This was enabled by our new HWF algorithm, which operates on expressive hybrid (LBO+ELA) bases. Our approach simplifies the training process and achieves state-of-the-art accuracy, particularly on challenging non-isometric and topologically noisy shapes.

Despite these strong results, we acknowledge the inherent limitations of a purely spectral approach. A promising direction for future work is to integrate our framework with explicit spatial deformation models, potentially bridging the gap between the spectral and spatial domains for even greater robustness.

## Acknowledgments

This work was supported by the Natural Science Foundation of China (No. 62172447, 62302530), the Hunan Provincial Natural Science Foundation of China (No. 2023JJ40769), and the Postgraduate Research and Innovation Project of Hunan Province (No. CX20250157). This work was supported in part by the High Performance Computing Center of Central South University.

## References

- Anguelov, D.; Srinivasan, P.; Koller, D.; Thrun, S.; Rodgers, J.; and Davis, J. 2005. SCAPE: Shape completion and animation of people. *ACM Transactions on Graphics*, 24: 408–416.
- Aubry, M.; Schlickewei, U.; and Cremers, D. 2011. The wave kernel signature: A quantum mechanical approach to shape analysis. In *Proceedings of IEEE/CVF International Conference on Computer Vision Workshops*, 1626–1633.
- Bastian, L.; Xie, Y.; Navab, N.; and Löhner, Z. 2024. Hybrid Functional Maps for Crease-Aware Non-Isometric Shape Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3313–3323.
- Bogo, F.; Romero, J.; Loper, M.; and Black, M. J. 2014. FAUST: Dataset and evaluation for 3D mesh registration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3794–3801.
- Cao, D.; Roetzer, P.; and Bernard, F. 2023. Unsupervised Learning of Robust Spectral Shape Matching. *ACM Transactions on Graphics*, 42: 1–15.
- Cao, D.; Roetzer, P.; and Bernard, F. 2024. Revisiting map relations for unsupervised non-rigid shape matching. In *2024 International Conference on 3D Vision (3DV)*, 1371–1381.
- Deng, B.; Yao, Y.; Dyke, R. M.; and Zhang, J. 2022. A Survey of Non-Rigid 3D Registration. *Computer Graphics Forum*, 41: 559–589.
- Dinh, H. Q.; Yezzi, A.; and Turk, G. 2005. Texture transfer during shape transformation. *ACM Transactions on Graphics*, 24: 289–310.
- Donati, N.; Corman, E.; Melzi, S.; and Ovsjanikov, M. 2022. Complex functional maps: A conformal link between tangent bundles. *Computer Graphics Forum*, 41: 317–334.
- Donati, N.; Corman, E.; and Ovsjanikov, M. 2022. Deep orientation-aware functional maps: Tackling symmetry issues in shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 742–751.
- Donati, N.; Sharma, A.; and Ovsjanikov, M. 2020. Deep geometric functional maps: robust feature learning for shape correspondence. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8592–8601.
- Eisenberger, M.; Lahner, Z.; and Cremers, D. 2020. Smooth shells: Multi-scale shape registration with functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12265–12274.
- Eisenberger, M.; Toker, A.; Leal-Taixé, L.; and Cremers, D. 2020. Deep shells: Unsupervised shape correspondence with optimal transport. In *Proceedings of Advances in Neural Information Processing Systems*, 10491–10502.
- Emery, P.; Lei, L.; Angela, D.; and Maks, O. 2025. DifuMatch: Category-Agnostic Spectral Diffusion Priors for Robust Non-rigid Shape Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1–10.
- Ezuz, D.; Solomon, J.; and Ben-Chen, M. 2019. Reversible harmonic maps between discrete surfaces. *ACM Transactions on Graphics*, 38: 1–12.
- Gao, M.; Zorah, L.; and Bernard, F. 2021. Isometric multi-shape matching. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14183–14193.
- Halimi, O.; Litany, O.; Rodola, E.; Bronstein, A. M.; and Kimmel, R. 2019. Unsupervised learning of dense shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4370–4379.
- Hartwig, F.; Sassen, J.; Azencot, O.; Rumpf, M.; and Ben-Chen, M. 2023. An Elastic Basis for Spectral Shape Correspondence. In *ACM SIGGRAPH Conference Proceedings*, 1–11.
- Hu, L.; Li, Q.; Liu, S.; and Liu, X. 2021. Efficient deformable shape correspondence via multiscale spectral manifold wavelets preservation. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14536–14545.
- Hu, L.; Li, Q.; Liu, S.; Yan, D.-M.; Xu, H.; and Liu, X. 2023. RFMNet: Robust deep functional maps for unsupervised non-rigid shape correspondence. *Graphical Models*, 129: 1–11.
- Huang, R.; Ren, J.; Wonka, P.; and Ovsjanikov, M. 2020. Consistent ZoomOut: Efficient Spectral Map Synchronization. *Computer Graphics Forum*, 39: 265–278.
- Löhner, Z.; Rodolà, E.; Bronstein, M. M.; Cremers, D.; Burghard, O.; Cosmo, L.; Dieckmann, A.; Klein, R.; Sahillioglu, Y.; et al. 2016. SHREC’16: Matching of deformable shapes with topological noise. In *Eurographics Workshop on 3D Object Retrieval, EG 3DOR*, 55–60.
- Li, L.; Donati, N.; and Ovsjanikov, M. 2022. Learning Multi-resolution Functional Maps with Spectral Attention for Robust Shape Matching. In *Proceedings of Advances in Neural Information Processing Systems*, 1–17.
- Li, Q.; Guo, Y.; Liu, X.; Hu, L.; Luo, F.; and Liu, S. 2024. Deformable shape matching with multiple complex spectral filter operator preservation. *The Visual Computer*, 40: 4885–4898.
- Litany, O.; Remez, T.; Rodola, E.; Bronstein, A.; and Bronstein, M. 2017. Deep functional maps: Structured prediction for dense shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 5659–5667.
- Liu, S.; Luo, F.; Li, Q.; Liu, X.; and Hu, L. 2024a. AWEDD: a descriptor simultaneously encoding multiscale extrinsic and intrinsic shape features. *The Visual Computer*, 40: 2537–2554.
- Liu, S.; Meng, J.; Hu, L.; Guo, Y.; Liu, X.; Yang, X.; Wang, H.; and Li, Q. 2024b. Multiscale spectral manifold wavelet regularizer for unsupervised deep functional maps. *Computer Graphics Forum*, 43(7): 1–12.
- Liu, S.; Wang, H.; Hu, L.; Li, Q.; and Liu, X. 2022. Incremental functional maps for accurate and smooth shape correspondence. *The Visual Computer*, 38: 3313–3325.

- Luo, F.; Li, Q.; Hu, L.; Wang, H.; Xu, H.; Liu, X.; Liu, S.; and Chen, H. 2025. Deep Frequency Awareness Functional Maps for Robust Shape Matching. *IEEE Transactions on Visualization and Computer Graphics*, 31: 7781–7794.
- Magnet, R.; and Ovsjanikov, M. 2024. Memory-Scalable and Simplified Functional Map Learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4041–4050.
- Magnet, R.; Ren, J.; Sorkine-Hornung, O.; and Ovsjanikov, M. 2022. Smooth non-rigid shape matching via effective dirichlet energy optimization. In *2022 International Conference on 3D Vision (3DV)*, 495–504.
- Melzi, S.; Ren, J.; Rodolà, E.; Sharma, A.; Wonka, P.; and Ovsjanikov, M. 2019. ZoomOut: Spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics*, 38: 1–20.
- Ovsjanikov, M.; Ben-Chen, M.; Solomon, J.; Butscher, A.; and Guibas, L. 2012. Functional maps: A flexible representation of maps between shapes. *ACM Transactions on Graphics*, 31: 1–11.
- Pai, G.; Ren, J.; Melzi, S.; Wonka, P.; and Ovsjanikov, M. 2021. Fast Sinkhorn filters : Using matrix scaling for non-rigid Shape Correspondence with functional maps. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 384–393.
- Ren, J.; Melzi, S.; Wonka, P.; and Ovsjanikov, M. 2021. Discrete optimization for shape matching. *Computer Graphics Forum*, 40: 81–96.
- Ren, J.; Poulenard, A.; Wonka, P.; and Ovsjanikov, M. 2018. Continuous and orientation-preserving correspondences via functional maps. *ACM Transactions on Graphics*, 37: 1–16.
- Roufousse, J.-M.; Sharma, A.; and Ovsjanikov, M. 2019. Unsupervised deep learning for structured shape matching. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 1617–1627.
- Sahillioğlu, Y. 2020. Recent advances in shape correspondence. *The Visual Computer*, 36: 1705–1721.
- Salti, S.; Tombari, F.; and Di Stefano, L. 2014. SHOT: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding*, 125: 251–264.
- Sharp, N.; Attaiki, S.; Crane, K.; and Ovsjanikov, M. 2022. DiffusionNet: Discretization Agnostic Learning on Surfaces. *ACM Transactions on Graphics*, 41: 1–16.
- Sumner, R. W.; and Popović, J. 2004. Deformation transfer for triangle meshes. *ACM Transactions on Graphics*, 23: 399–405.
- Sun, J.; Ovsjanikov, M.; and Guibas, L. 2009. A concise and provably informative multi-scale signature based on heat diffusion. *Computer Graphics Forum*, 28: 1383–1392.
- Sun, M.; Mao, S.; Jiang, P.; Ovsjanikov, M.; and Huang, R. 2023. Spatially and Spectrally Consistent Deep Functional Maps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 14497–14507.
- Viganò, G.; Ovsjanikov, M.; and Melzi, S. 2025. NAM: Neural Adjoint Maps for refining shape correspondences. *ACM Transactions on Graphics*, 44: 1–15.
- Zhuravlev, A.; Lähner, Z.; and Golyanik, V. 2025. Denoising functional maps: Diffusion models for shape correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 26899–26909.
- Zuffi, S.; Kanazawa, A.; Jacobs, D. W.; and Black, M. J. 2017. 3D menagerie: Modeling the 3D shape and pose of animals. In *Proceedings of IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 6365–6373.