

# DcMatch: Unsupervised Multi-Shape Matching with Dual-Level Consistency

Tianwei Ye, Yong Ma, Xiaoguang Mei\*

Electronic Information School, Wuhan University, Wuhan 430072, China  
 twye2001@gmail.com, mayong@whu.edu.cn, meixiaoguang@whu.edu.cn

## Abstract

Establishing point-to-point correspondences across multiple 3D shapes is a fundamental problem in computer vision and graphics. In this paper, we introduce DcMatch, a novel unsupervised learning framework for non-rigid multi-shape matching. Unlike existing methods that learn a canonical embedding from a single shape, our approach leverages a shape graph attention network to capture the underlying manifold structure of the entire shape collection. This enables the construction of a more expressive and robust shared latent space, leading to more consistent shape-to-universe correspondences via a universe predictor. Simultaneously, we represent these correspondences in both the spatial and spectral domains and enforce their alignment in the shared universe space through a novel cycle consistency loss. This dual-level consistency fosters more accurate and coherent mappings. Extensive experiments on several challenging benchmarks demonstrate that our method consistently outperforms previous state-of-the-art approaches across diverse multi-shape matching scenarios.

**Code** — <https://github.com/YeTianwei/DcMatch>

## Introduction

Shape matching is a fundamental problem in computer vision and graphics (Van Kaick et al. 2011; Sahillioğlu 2020), which aims to establish accurate point-to-point correspondences between 3D shapes. It supports a wide range of applications, including texture transfer in graphics, statistical shape analysis in medical imaging, and 3D reconstruction in computer vision. While most existing work focuses on matching shape pairs, recent advances in 3D scanning technology have made it increasingly common to capture multiple shapes simultaneously—for instance, different scans of the same object under varying conditions. In this context, establishing correspondences across a collection of shapes, known as multi-shape matching, becomes essential.

Although recent pairwise shape matching methods have achieved promising results (Bastian et al. 2024; Zhuravlev, Löhner, and Golyanik 2025), extending them to the multi-shape setting poses significant challenges. First, multi-shape




Figure 1: **Overview of mainstream multi-shape matching paradigms.** (Left) The permutation synchronization paradigm, which consists of two stages: computing pairwise correspondences and enforcing cycle consistency via post-processing. (Right) The universe-based paradigm, which introduces a virtual universe shape and reduces the multi-shape matching problem to a set of pairwise mappings.

matching requires cycle consistency across the shape collection—namely, the composition of maps along any closed cycle should yield the identity map. This global constraint, which is absent in pairwise settings, considerably increases the problem’s complexity. Secondly, the number of shape pairs grows combinatorially with the size of the shape set, leading to substantial computational overhead.

To address the challenges of multi-shape matching, existing methods generally fall into two main paradigms, as illustrated in Fig. 1. The first paradigm imposes cycle consistency constraints directly on pairwise mappings to enforce global consistency. Known as permutation synchronization (Pachauri, Kondor, and Singh 2013; Bernard et al. 2019a), this approach offers flexibility but typically involves a two-stage optimization process and often leads to spatially non-smooth and noisy results. The second paradigm introduces a shared latent domain, commonly referred to as the *universe*, which converts pairwise correspondences into mappings between each shape and the universe (Cosmo et al. 2017; Bernard et al. 2019b). This formulation enforces global consistency by construction. However, most methods within this paradigm typically learn the universe embedding from a single shape, in either the spatial (Cao and Bernard 2022) or spectral (Huang and Guibas 2013) domain, thereby reducing multi-shape matching to a collection of isolated pairwise problems. This often neglects the structural relation-

\*Corresponding author.

ships within the shape collection and leads to suboptimal performance (Eisenberger et al. 2023).

To tackle the above limitations, we propose DcMatch, a novel unsupervised multi-shape matching framework that enforces dual-level cycle consistency. Our pipeline begins by extracting per-vertex features for each shape using a feature extractor. Inspired by the accuracy of pairwise matching in permutation synchronization methods, we directly compute functional maps and point-wise correspondences from these features. In parallel, instead of learning a canonical embedding from a single shape as in prior universe-based approaches, we model the entire shape collection as an undirected graph and employ a graph attention network to extract manifold-aware features that capture the underlying geometric structure. These features are then fed into a universe predictor to estimate shape-to-universe correspondences. By doing so, our model effectively learns a universe embedding enriched with manifold information. Beyond this, by leveraging the inherent cycle consistency of functional maps combined with the global consistency of shape-to-universe matching, we enforce consistency between these two alignment paths, thereby improving robustness and consistency in the shared universe space. Through this dual-level consistency, our method achieves more accurate and coherent multi-shape matching in a fully unsupervised setting. Extensive experiments on diverse datasets demonstrate the competitive performance of our approach.

We summarize our main contributions as follows:

1. We propose a novel unsupervised framework for multi-shape matching that enforces both spectral and spatial cycle consistency, resulting in more accurate and coherent correspondence predictions.
2. We introduce a shape graph attention module that captures the underlying manifold structure of the entire shape collection. This facilitates the construction of a robust and generalizable universe space, in contrast to prior methods that learn from a single reference shape.
3. We conduct extensive experiments under various settings, demonstrating that our method achieves state-of-the-art performance across challenging benchmarks.

## Related Work

We refer interested readers to (Van Kaick et al. 2011; Sahillioğlu 2020) for a comprehensive review. Here, we highlight only the approaches most relevant to our method.

### Functional Maps

Functional maps, first introduced by (Ovsjanikov et al. 2012), have become one of the most widely used frameworks in shape matching. By representing point-wise correspondences as compact matrices in the spectral domain, they offer substantial computational efficiency. Thanks to their concise and flexible formulation, functional maps have been extended in various directions to improve accuracy and robustness (Eynard et al. 2016; Melzi et al. 2019b; Ren et al. 2019), and to handle more challenging scenarios, including partial (Litany et al. 2017b; Rodolà et al. 2017), non-

isometric (Nogneng and Ovsjanikov 2017; Ren et al. 2021, 2018), and non-unique shape matching (Ren et al. 2020).

### Learning Method Based on Functional Maps

Unlike axiomatic functional map methods that depend on hand-crafted descriptors (Sun, Ovsjanikov, and Guibas 2009; Aubry, Schlickewei, and Cremers 2011; Salti, Tombari, and Di Stefano 2014), recent approaches learn feature descriptors directly from data. FMNet (Litany et al. 2017a) introduced a supervised framework that transforms SHOT descriptors into more effective embeddings for functional map estimation. This was extended to the unsupervised setting in (Halimi et al. 2019; Roufousse, Sharma, and Ovsjanikov 2019) by introducing regularization-based losses. More recently, DiffusionNet (Sharp et al. 2022) has enabled robust, resolution-aware feature extraction, inspiring several state-of-the-art methods (Cao and Bernard 2022; Cao, Roetzer, and Bernard 2023; Li, Donati, and Ovsjanikov 2022; Donati, Corman, and Ovsjanikov 2022; Bastian et al. 2024) that perform well across diverse shape datasets.

### Multi-shape Matching

Early multi-shape matching methods enforce cycle consistency via semidefinite programming or convex relaxation, but often yield sparse correspondences and poor scalability. Within the functional map framework, Consistent ZoomOut (Huang et al. 2020) synchronizes maps via a shared basis but depends on good initialization. IsoMuSh (Gao et al. 2021) jointly enforces consistency on pointwise and functional maps but is restricted to near-isometric shapes. CycoMatch (Xia et al. 2025) builds cycle-consistent bases via a graph in a two-stage pipeline. Recently, learning-based approaches have improved scalability. UDMSM (Cao and Bernard 2022) predicts a canonical embedding to promote cycle consistency, while G-MSM (Eisenberger et al. 2023) leverages a heuristic to model the shape collection’s underlying manifold.

## Background

In this section, we first review the deep hybrid functional maps pipeline (Bastian et al. 2024), then revisit cycle consistency with corresponding theoretical insights.

### Deep Hybrid Functional Maps

Given a pair of 3D shapes  $\mathcal{S}_i$  and  $\mathcal{S}_j$ , represented as triangle meshes with  $n_i$  and  $n_j$  vertices respectively, the hybrid functional map framework aims to represent dense correspondences in a compact, linear form. The main pipeline consists of the following steps:

1. Compute two sets of basis functions for each shape: (i) the first  $k_{\text{LB}}$  eigen-functions of the Laplace Beltrami operator (LBO) (Pinkall and Polthier 1993), denoted as  $\Phi_i \in \mathbb{R}^{n_i \times k_{\text{LB}}}$ , and (ii) the first  $k_{\text{Elas}}$  eigen-functions of the elastic thin-shell energy (Hartwig et al. 2023), denoted as  $\Psi_i \in \mathbb{R}^{n_i \times k_{\text{Elas}}}$ .
2. Extract vertex-wise feature descriptors  $\mathcal{F}_i \in \mathbb{R}^{n_i \times d}$  using a scalable network (Sharp et al. 2022), where  $d$  is

the feature dimension. These features are projected onto both sets of basis functions to obtain spectral coefficients  $\mathcal{A}_i^{\text{LB}} = \Phi_i^\dagger \mathcal{F}_i \in \mathbb{R}^{k_{\text{LB}} \times d}$ ,  $\mathcal{A}_i^{\text{Elas}} = \Psi_i^\dagger \mathcal{F}_i \in \mathbb{R}^{k_{\text{Elas}} \times d}$ . The combined hybrid basis is denoted as  $\tilde{\Phi}_i = [\Phi_i; \Psi_i]$ , with corresponding projected features  $\mathcal{A}_i = \tilde{\Phi}_i^\dagger \mathcal{F}_i \in \mathbb{R}^{(k_{\text{LB}} + k_{\text{Elas}}) \times d}$ .

- Solve the block-diagonal map  $C_{ij} = \text{diag}(C_{ij}^{11}, C_{ij}^{22}) \in \mathbb{R}^{(k_{\text{LB}} + k_{\text{Elas}}) \times (k_{\text{LB}} + k_{\text{Elas}})}$  by minimizing regularized objectives per basis:

$$C_{ij}^{11} = \arg \min_C E_{\text{data}}^{\text{LB}}(C) + \lambda_{\text{LB}} E_{\text{reg}}^{\text{LB}}(C),$$

$$E_{\text{data}}^{\text{LB}}(C) = \|C \mathcal{A}_i^{\text{LB}} - \mathcal{A}_j^{\text{LB}}\|_{\text{F}}^2, \quad (1)$$

$$E_{\text{reg}}^{\text{LB}}(C) = \|C \Lambda_i^{\text{LB}} - \Lambda_j^{\text{LB}} C\|_{\text{F}}^2,$$

$$C_{ij}^{22} = \arg \min_C E_{\text{data}}^{\text{Elas}}(C) + \lambda_{\text{Elas}} E_{\text{reg}}^{\text{Elas}}(C),$$

$$E_{\text{data}}^{\text{Elas}}(C) = \|C \mathcal{A}_i^{\text{Elas}} - \mathcal{A}_j^{\text{Elas}}\|_{M_{k_{\text{Elas}}, j}}^2, \quad (2)$$

$$E_{\text{reg}}^{\text{Elas}}(C) = \|C \Lambda_i^{\text{Elas}} - \Lambda_j^{\text{Elas}} C\|_{\text{HS}}^2,$$

where  $M_{k_{\text{Elas}}, j} = \Psi_j^\top M_j \Psi_j$  denotes the reduced mass matrix induced by the elastic basis on shape  $\mathcal{S}_j$ ,  $\|\cdot\|_{\text{HS}}$  represents the Hilbert-Schmidt norm and  $\Lambda_i, \Lambda_j$  are diagonal eigenvalue matrices for each basis.

- Convert the hybrid functional map  $C_{ij}$  into a point-wise correspondence  $\Pi_{ji} \in \{0, 1\}^{n_j \times n_i}$  using nearest-neighbor search or other post-processing techniques (Melzi et al. 2019b; Pai et al. 2021; Vestner et al. 2017; Xia et al. 2024), based on the relation:

$$\tilde{\Phi}_j C_{ij} \approx \Pi_{ji} \tilde{\Phi}_i. \quad (3)$$

### Cycle Consistency Formulation

We begin by defining the notion of cycle consistency. Let  $\mathcal{S} = \{\mathcal{S}_i\}_{i=1}^n$  be a collection of shapes, and let  $\mathcal{P} = \{\Pi_{ij}\}_{i,j=1}^n$  and  $\mathcal{C} = \{C_{ij}\}_{i,j=1}^n$  denote the sets of point-wise and functional maps, respectively. The set  $\mathcal{P}$  is said to be cycle consistent if, for any closed sequence of shapes  $\{i, j, k, \dots, l, i\}$ , the composed map satisfies  $\Pi_{ij} \Pi_{jk} \dots \Pi_{li} = \mathbf{I}$ , where  $\mathbf{I}$  denotes the identity map (Huang and Guibas 2013). A similar definition applies to the set of functional maps  $\mathcal{C}$ .

Rather than explicitly enforcing cycle consistency over all shape pairs, the same principle can be realized implicitly through a *shape-to-universe* formulation (Huang and Guibas 2013; Tron et al. 2017; Gao et al. 2021). In this setting, each shape is associated with a shared canonical domain - referred to as the *universe shape* - via a one-to-one point-wise correspondence. Pairwise correspondences are then constructed by composing the shape-to-universe and universe-to-shape maps, inherently ensuring global consistency (Huang and Guibas 2013). Specifically, let  $\Pi_i$  denote the correspondence between shape  $\mathcal{S}_i$  and the universe. The point-wise correspondence  $\Pi_{ij}$  between  $\mathcal{S}_i$  and  $\mathcal{S}_j$  is given by:

$$\Pi_{ij} = \Pi_i \Pi_j^\top, \quad (4)$$

where first maps  $\mathcal{S}_i$  to the universe and then back to  $\mathcal{S}_j$ .

Likewise, functional maps exhibit natural cycle consistency under the universe-based formulation (Sun et al. 2023), as formalized below.

**Theorem 1.** *Let the total energy over all shape pairs be defined as  $E_{\text{total}}(\mathcal{C}) = \sum_{i,j} E_{\text{data}}(C_{ij}) + \lambda E_{\text{reg}}(C_{ij})$ . If  $E_{\text{total}}(\mathcal{C}) = 0$ , then for any shape  $\mathcal{S}_i$  and any closed cycle  $\{i, j, \dots, l, i\}$ , the composed functional map satisfies  $C_{ii} \mathcal{A}_i = \mathcal{A}_i$ , i.e., it acts as the identity on the subspace spanned by  $\mathcal{A}_i$ .*

In practice,  $\mathcal{A}_i$  is typically of full row rank, as we generally set  $d > k_{\text{LB}} + k_{\text{Elas}}$ . Under this condition, Theorem 1 implies  $C_{ii} = \mathbf{I}$ , from which it follows that the entire set  $\mathcal{C} = \{C_{ij}\}_{i,j=1}^n$  is cycle consistent. In this case,  $\mathcal{A}_i$  can be interpreted as a functional embedding of shape  $\mathcal{S}_i$  into a shared universe. Analogous to Eq. (4), the functional map between any two shapes can be expressed as:

$$C_{ij} = \mathcal{A}_j \mathcal{A}_i^\dagger. \quad (5)$$

### Method

In this section, we detail our proposed method, as illustrated in Fig. 2. The pipeline begins with feature extraction, followed by the functional maps component. We then introduce the shape graph attention module, which captures inter-shape relationships across the shape set. Subsequently, the shape-to-universe correspondences are established through the universe predictor module. Finally, we describe the unsupervised loss functions used to train the network.

#### Feature Extractor

We begin by extracting features from a set of input shapes  $\mathcal{S} = \{\mathcal{S}_i\}_{i=1}^n$ . For each shape, we compute per-vertex features that serve as the foundation for all subsequent components. We employ DiffusionNet (Sharp et al. 2022) as our feature extractor, owing to its robustness to variations in mesh resolution and sampling density. This ensures consistent and discriminative features across diverse shape collections. The extracted features are denoted as  $\mathcal{F} = \{\mathcal{F}_i\}_{i=1}^n$ .

#### Functional Maps Module

For notational simplicity, we consider a shape pair  $\mathcal{S}_i$  (source) and  $\mathcal{S}_j$  (target) from the collection  $\mathcal{S} = \{\mathcal{S}_i\}_{i=1}^n$  when describing operations between two shapes.

The functional maps module is responsible for establishing direct correspondences between shape pairs. Specifically, it computes bidirectional functional maps  $C_{ij}$  and  $C_{ji}$ , along with the corresponding point-wise map  $\Pi_{ij}$  between  $\mathcal{S}_i$  and  $\mathcal{S}_j$ . This module comprises two components: a functional map solver and a point-wise map solver, both of which take as input the features produced by the feature extractor.

**Functional Maps Solver** We employ a regularized functional map formulation to estimate  $C_{ij}$  and  $C_{ji}$ . During training, we incorporate regularization terms promoting bijectivity and orthogonality, following (Ren et al. 2019):

$$\mathcal{L}_{\text{struct}} = \lambda_{\text{bij}} \mathcal{L}_{\text{bij}} + \lambda_{\text{orth}} \mathcal{L}_{\text{orth}}, \quad (6)$$

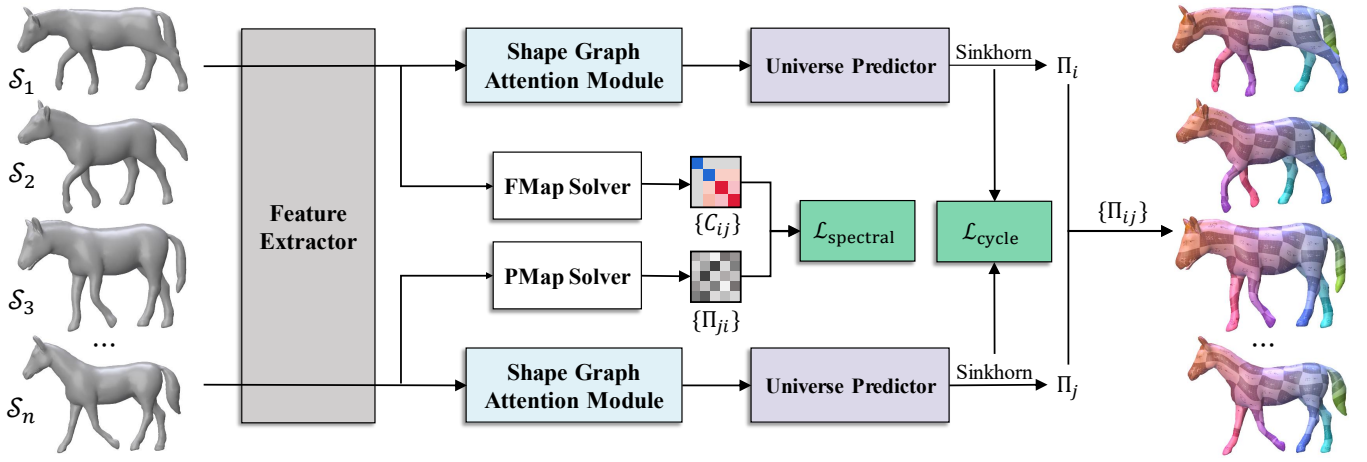


Figure 2: **Method Overview.** Given a collection of shapes  $\mathcal{S} = \{\mathcal{S}_i\}_{i=1}^n$ , we first extract per-vertex features  $\mathcal{F} = \{\mathcal{F}_i\}_{i=1}^n$  using DiffusionNet. These features are then used to compute bidirectional functional maps  $\{C_{ij}\}$  and point-to-point correspondences  $\{\Pi_{ij}\}$ . Meanwhile, the shape graph attention module generates manifold-aware features for each shape, which are passed to the universe predictor to estimate the correspondences between each shape and a shared virtual universe. In addition to the spectral loss, we implement a cycle consistency loss to further align the spatial and spectral consistency in the shared universe space.

where

$$\mathcal{L}_{\text{bij}} = \sum_{i,j}^n \|C_{ij}C_{ji} - \mathbf{I}\|_{\mathbb{F}}^2 + \|C_{ji}C_{ij} - \mathbf{I}\|_{\mathbb{F}}^2, \quad (7)$$

$$\mathcal{L}_{\text{orth}} = \sum_{i,j}^n \|C_{ij}^*C_{ji} - \mathbf{I}\|_{\mathbb{F}}^2 + \|C_{ji}^*C_{ij} - \mathbf{I}\|_{\mathbb{F}}^2. \quad (8)$$

**Point-wise Map Solver** The point-wise map  $\Pi_{ij}$  is theoretically expected to be a (partial) permutation matrix:

$$\left\{ \Pi_{ij} \in \{0, 1\}^{n_i \times n_j} : \Pi_{ij} \mathbf{1}_{n_j} = \mathbf{1}_{n_i}, \mathbf{1}_{n_i}^\top \Pi_{ij} \leq \mathbf{1}_{n_j}^\top \right\}, \quad (9)$$

where  $\Pi_{ij}(s, t)$  indicates the correspondence between the  $s$ -th vertex of  $\mathcal{S}_i$  and the  $t$ -th vertex of  $\mathcal{S}_j$ . Following (Eisenberger et al. 2021), we estimate  $\Pi_{ij}$  by computing the similarity between feature matrices  $\mathcal{F}_i$  and  $\mathcal{F}_j$ :

$$\Pi_{ij} = \text{softmax}(\mathcal{F}_i \mathcal{F}_j^\top / \tau), \quad (10)$$

where  $\tau$  is a temperature parameter that controls the sharpness of the assignment distribution.

### Shape Graph Attention Module

Shape collections typically exhibit structured relationships rather than consisting of independent entities. Some shapes are more similar to each other, and mappings between similar shapes often share intrinsic geometric correlations. To capture these dependencies, we first construct a shape graph and employ a Graph Attention Network (GAT) (Brody, Alon, and Yahav 2022) to refine the per-shape features  $\mathcal{F} = \{\mathcal{F}_i\}_{i=1}^n$ . The edge set  $\mathcal{E}$  of the graph is defined based on the top- $k$  cosine similarities between shape features:

$$\mathcal{E} = \{(i, j) \mid j \in \text{Top-}k(\cos(\mathcal{F}_i, \mathcal{F}_j))\}. \quad (11)$$

This connects each shape to its most similar neighbors, forming a graph that encodes the underlying geometric structure of the shape set.

To extract manifold-aware features, we employ a GAT, which dynamically learns attention weights between connected shapes. Unlike fixed edge weights, the attention mechanism adaptively computes inter-shape affinities, enabling more flexible and expressive feature aggregation. We begin by applying mean pooling to the per-vertex features of each shape, resulting in a shape-level descriptor. Given a pair of shapes  $(i, j)$ , the attention weight  $\alpha_{ij}$  is computed as:

$$\alpha_{ij} = \frac{\exp(\mathbf{a}^\top \text{LeakyReLU}(\mathbf{W} \cdot [\mathcal{F}_i \parallel \mathcal{F}_j]))}{\sum_{j' \in \mathcal{N}_i} \exp(\mathbf{a}^\top \text{LeakyReLU}(\mathbf{W} \cdot [\mathcal{F}_i \parallel \mathcal{F}_{j'}]))}, \quad (12)$$

where  $\mathbf{W}$  is a learnable weight matrix and  $\mathbf{a}$  is a learnable attention vector. The final manifold-aware feature for shape  $\mathcal{S}_i$  is obtained by aggregating features from its neighbors, weighted by attention:

$$\mathcal{F}'_i = \sigma \left( \sum_{j \in \mathcal{N}_i} \alpha_{ij} \cdot \mathbf{W} \mathcal{F}_j \right), \quad (13)$$

where  $\sigma$  denotes a non-linear activation function. We use a two-layer GAT to model the local structure of the shape graph, followed by a LayerNorm and dropout on the final output to improve stability and generalization.

Through message passing on the shape graph, each shape's representation is enriched with contextual information from its neighbors, capturing the underlying manifold structure. To preserve both local geometric detail and global structural context, we concatenate the original features  $\mathcal{F}_i$  with the aggregated features  $\mathcal{F}'_i$  to form the final representation  $\mathcal{G} = \{\mathcal{G}_i\}_{i=1}^n$ , where  $\mathcal{G}_i = [\mathcal{F}_i \parallel \mathcal{F}'_i]$ .

### Universe Predictor Module

The universe predictor takes as input the shape-level features  $\mathcal{G} = \{\mathcal{G}_i\}_{i=1}^n$  produced by the graph attention module, and predicts a correspondence  $\Pi_i$  that maps each shape  $\mathcal{S}_i$  to

Train		FAUST		SCAPE		FAUST+SCAPE		
Test		FAUST	FAUST <sub>a</sub>	SCAPE	SCAPE <sub>a</sub>	FAUST	SCAPE	SHREC'19
<i>Pairwise Matching</i>	ZoomOut	6.1	8.7	7.5	14.0	6.1	7.5	-
	SmoothShells	2.5	5.4	4.7	5.0	2.5	4.7	-
	DiscreteOp	5.6	6.2	13.1	14.6	5.6	13.1	-
	Deep Shells	1.7	12.0	2.5	10.0	1.6	2.4	21.1
	DUO-FMNet	2.5	3.0	2.6	2.7	2.5	4.3	6.4
	AttentiveFMaps	1.9	2.4	2.2	2.3	1.9	2.3	5.8
	ULRSSM	1.6	1.9	1.9	1.9	1.6	2.1	4.8
	HybridFMaps	1.5	1.8	<b>1.8</b>	1.9	1.5	2.0	4.5
	DenoisFMaps	1.7	2.0	2.1	2.2	1.7	2.0	6.3
<i>Multi-Matching</i>	Consistent ZoomOut	2.2	7.3	2.5	12.1	2.2	2.5	-
	IsoMuSh	4.4	-	5.6	-	4.4	5.6	-
	CycoMatch	3.8	9.2	4.7	13.6	3.8	4.6	-
	UDMSM	1.5	15.3	2.0	4.9	1.7	3.2	17.8
	G-MSM	1.5	12.7	<b>1.8</b>	28.1	1.5	2.1	6.8
	<b>Ours</b>	<b>1.4</b>	<b>1.7</b>	<b>1.8</b>	<b>1.8</b>	<b>1.4</b>	<b>1.9</b>	<b>4.2</b>

Table 1: **Quantitative results on near-isometric datasets (FAUST, SCAPE, SHREC'19) and anisotropically remeshed versions (FAUST<sub>a</sub>, SCAPE<sub>a</sub>).** The best results are shown in **bold**.

the shared universe. Following UDMSM (Cao and Bernard 2022), we utilize a DiffusionNet architecture to generate these assignment matrices, where the number of output columns corresponds to the number of universe points.

In the universe-based formulation, each vertex in a shape is assigned to a single universe point, while each universe point maps to at most one shape vertex. Thus,  $\Pi_i$  is a (partial) permutation matrix:

$$\Pi_i \in \left\{ \Pi \in \{0, 1\}^{n_i \times c} : \Pi \mathbf{1}_c = \mathbf{1}_{n_i}, \mathbf{1}_{n_i}^\top \Pi \leq \mathbf{1}_c^\top \right\}, \quad (14)$$

where  $c$  denotes the number of universe points.

To enable end-to-end training, we apply Sinkhorn normalization (Sinkhorn and Knopp 1967; Mena et al. 2018) to the raw network outputs. This iterative normalization projects assignment scores into a doubly stochastic matrix, serving as a smooth, differentiable approximation of a partial permutation. It enables the network to learn discrete constraints in a relaxed yet consistent manner, supporting optimization in a fully unsupervised setting.

## Loss Functions

**Spectral Loss** The spectral loss integrates both structural regularization (Eq. (6)) and a coupling term. The coupling loss  $\mathcal{L}_{\text{couple}}$  encourages the point-wise maps  $\Pi_{ij}$  and  $\Pi_{ji}$  to be consistent via the functional maps:

$$\mathcal{L}_{\text{couple}} = \sum_{i,j}^n \left\| C_{ij} - \Phi_j^\dagger \Pi_{ji} \Phi_i \right\|_F^2 + \left\| C_{ij} - \Psi_i^\dagger \Pi_{ij} \Psi_j \right\|_{\text{HS}}^2. \quad (15)$$

The total spectral loss is given by:

$$\mathcal{L}_{\text{spectral}} = \mathcal{L}_{\text{struct}} + \lambda_{\text{couple}} \mathcal{L}_{\text{couple}}. \quad (16)$$

**Cycle Consistency Loss** In our framework, shape-to-universe alignment can be realized in two compatible forms. The functional map coefficient matrix  $\mathcal{A}_i$  maps the spectral

basis of shape  $\mathcal{S}_i$  to a shared latent universe, resulting in the aligned embedding  $\Phi_i \mathcal{A}_i$ . Alternatively, the point-wise correspondence matrix  $\Pi_i$  projects shape vertices directly onto the universe, yielding a universe-aligned embedding  $\Pi_i^\top \Phi_i$ .

To exploit this structural redundancy, we introduce a novel cycle consistency loss that enforces consistency between these two alignment paths across the shape collection. This constraint leverages the intrinsic cycle consistency of functional maps while regularizing the learning of shape-to-universe matching matrices, enhancing robustness against noise and structural variability.

We consider two variants of this loss. For near-isometric shapes, we use a Frobenius norm formulation:

$$\mathcal{L}_{\text{cycle}} = \sum_{i,j}^n \left\| \Pi_i^\top \tilde{\Phi}_i \mathcal{A}_i - \Pi_j^\top \tilde{\Phi}_j \mathcal{A}_j \right\|_F^2. \quad (17)$$

For more general deformations, we adopt a cosine similarity formulation:

$$\mathcal{L}_{\text{cycle}} = \sum_{i,j}^n \left( 1 - \cos \left( \Pi_i^\top \tilde{\Phi}_i \mathcal{A}_i, \Pi_j^\top \tilde{\Phi}_j \mathcal{A}_j \right) \right). \quad (18)$$

The final loss combines the spectral and cycle terms:

$$\mathcal{L}_{\text{total}} = \mathcal{L}_{\text{spectral}} + \lambda_{\text{cycle}} \mathcal{L}_{\text{cycle}}. \quad (19)$$

## Experiment

In this section, we evaluate our approach on multiple benchmarks and compare it against state-of-the-art methods.

**Baselines** We compare our method with a wide range of baselines, including pairwise methods (ZoomOut (Melzi et al. 2019b), SmoothShells (Eisenberger, Lahner, and Cremers 2020), DiscreteOp (Ren et al. 2021), DeepShells (Eisenberger et al. 2020), DUO-FMNet (Donati, Corman, and Ovsjanikov 2022), AttentiveFMaps (Li, Donati, and Ovsjanikov 2022), ULRSSM (Cao, Roetzer, and




Figure 3: **Proportion of Correct Keypoints (PCK) curves and Area Under Curve (AUC) values** on SHREC'19, SMAL, and DT4D-H inter-class, comparing our method with ULRSSM and HybridFMaps.

Bernard 2023), HybridFMaps (Bastian et al. 2024), DenoisFMaps (Zhuravlev, Löhner, and Golyanik 2025)) and multi-matching approaches (ConsistentZoomOut (Huang et al. 2020), IsoMuSh (Gao et al. 2021), CycoMatch (Xia et al. 2025), UDMSM (Cao and Bernard 2022), G-MSM (Eisenberger et al. 2023)).

**Evaluation Metrics** Following the Princeton benchmark protocol (Kim, Lipman, and Funkhouser 2011), we report the average geodesic error ( $\times 100$ ) as the evaluation metric.

### Near-isometric Shape Matching

**Datasets** We evaluate on three near-isometric datasets: FAUST (Bogo et al. 2014), SCAPE (Anguelov et al. 2005), and SHREC'19 (Melzi et al. 2019a), using their remeshed versions (Ren et al. 2018; Donati, Sharma, and Ovsjanikov 2020). FAUST contains 100 human shapes across 10 subjects and 10 poses, with a 80/20 train-test split. SCAPE includes 71 shapes from a single subject, split 51/20 for training and testing. SHREC'19 consists of 44 human shapes with more variation in body type and pose, used solely for testing, excluding shape 40 due to its non-closed geometry.

**Results** We compare our method with both pairwise shape matching algorithms and multi-shape matching frameworks. Table 1 reports the average geodesic error for all methods. Compared to pairwise approaches, our method consistently achieves superior performance across most settings and datasets. When compared to multi-shape baselines, it demonstrates strong generalization, particularly on the SHREC'19 dataset. We attribute this to the model's ability to capture the underlying manifold structure of the shape collection, leading to more robust correspondences on unseen shapes. The Proportion of Correct Keypoints (PCK) curves on the SHREC'19 dataset are shown in Fig. 3 (left).

### Matching with Anisotropic Meshing

**Datasets** To assess robustness to mesh connectivity variations, we follow DUO-FMNet (Donati, Corman, and Ovsjanikov 2022) and use anisotropically remeshed versions of FAUST and SCAPE (FAUST.a and SCAPE.a), where feature uneven triangulations with small, dense triangles on one

	Geo.error ( $\times 100$ )	SMAL	DT4D-H	
			intra	inter
<i>Pairwise Matching</i>	ZoomOut	38.4	4.0	29.0
	SmoothShells	36.1	1.1	6.3
	DiscreteOp	38.1	3.6	27.6
	Deep Shells	29.3	3.4	31.1
	DUO-FMNet	6.7	2.6	15.8
	AttentiveFMaps	5.4	1.7	11.6
	ULRSSM	3.9	<b>0.9</b>	4.1
	HybridFMaps	3.4	1.0	3.9
	DenoisFMaps	46.1	14.4	22.9
<i>Multi-Matching</i>	Consistent ZoomOut	16.9	8.7	26.4
	CycoMatch	24.7	4.4	24.9
	UDMSM	26.5	2.4	15.8
	G-MSM	43.9	7.8	12.0
	<b>Ours</b>	<b>2.9</b>	1.0	<b>3.8</b>

Table 2: **Quantitative results on non-isometric shape matching.**

side and large, coarse triangles on the other. These distortions challenge methods sensitive to mesh structure.

**Results** Table 1 presents the quantitative results. Most methods exhibit a noticeable drop in performance under anisotropic conditions. In contrast, our method achieves the best performance on both datasets. These results suggest that the combination of pairwise functional maps and manifold-aware universe matching enables our approach to effectively mitigate the challenges posed by irregular mesh structures.

### Non-isometric Shape Matching

**Datasets** We also evaluate on two widely used non-isometric datasets: SMAL (Zuffi et al. 2017) and DT4D-H (Magnet et al. 2022). The SMAL dataset contains 49 animal shapes from 8 species; we use 5 species for training and the remaining for testing, yielding a 29/20 train-test split. The DT4D-H dataset includes 9 human shape categories, with 198 shapes used for training and 95 for testing.

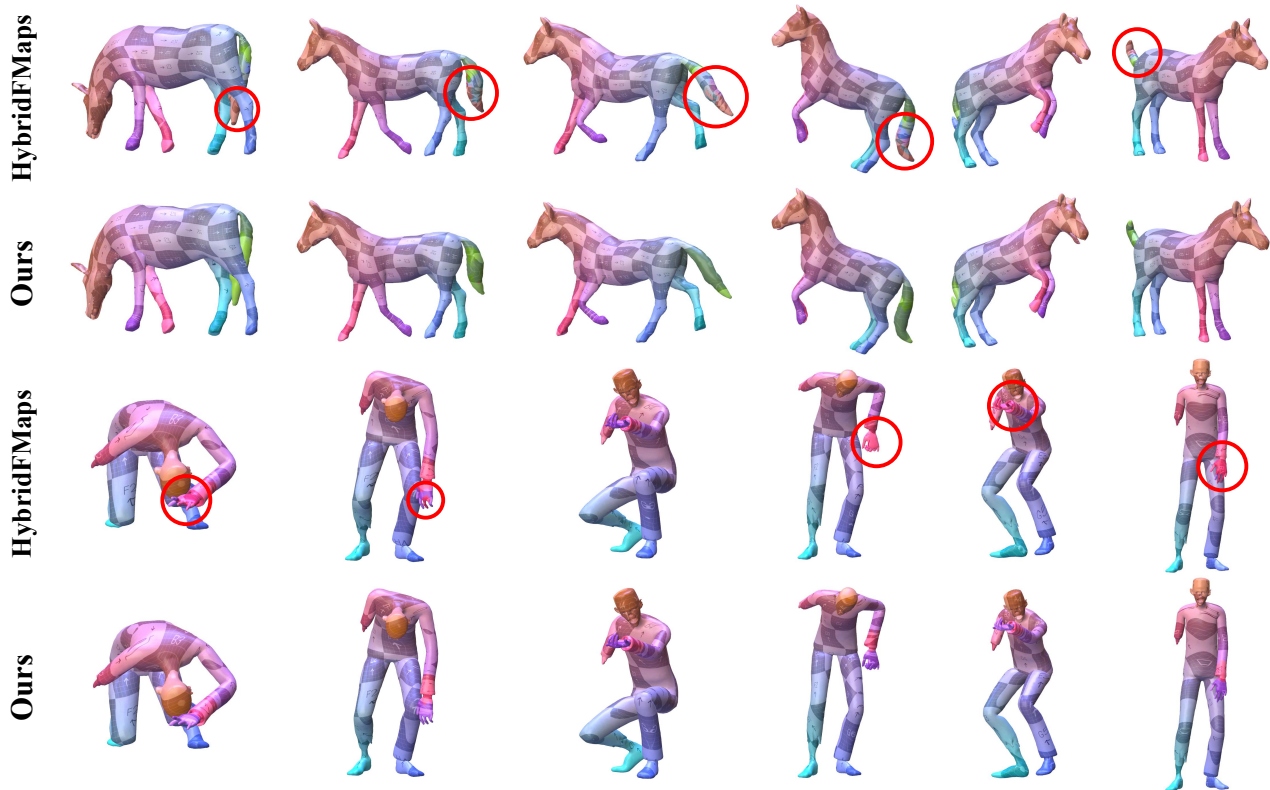


Figure 4: **Qualitative multi-shape matching results** via texture transfer, comparing our method with HybridFMaps on SMAL (top) and DT4D-H inter-class (bottom).

Ablation Setting	Geo.error ( $\times 100$ )
w/o shape graph attention module	3.7
w/o functional maps module	26.5
w/o universe predictor module	3.4
w/o cycle consistency loss	3.8
<b>Ours</b>	<b>2.9</b>

Table 3: **Ablation study results on the SMAL dataset.**

**Results** Table 2 presents results for both intra-class and inter-class settings. The DT4D-inter setup is particularly challenging, as it requires generalization to unseen shape classes. Our method achieves competitive intra-class results and significantly outperforms all baselines in the inter-class case, with notable improvements on the SMAL. These results demonstrate the strong generalization of our approach across diverse, non-isometric categories. PCK curves are shown in Fig. 3, with qualitative comparisons in Fig. 4.

### Ablation Study

We conduct ablation experiments on the non-isometric SMAL dataset to evaluate the contribution of each component in our framework. Specifically, we compare the following configurations: (1) Removing the shape graph attention module; (2) Removing the functional map module; (3) Removing the universe predictor module; (4) Removing the cy-

cle consistency loss. The results in Table 3 show that removing any single component leads to a noticeable performance drop, especially the functional maps module.

### Limitations

Despite its effectiveness, our method has several limitations that merit further exploration. First, although our framework avoids the need to explicitly select a reference shape, it still requires the universe size to be predefined and fixed. Second, processing the entire shape collection as a graph incurs additional computational overhead. While this trade-off leads to improved matching accuracy, future work could focus on designing more efficient and scalable algorithmic alternatives to reduce runtime and memory demands.

### Conclusion

In this paper, we introduce DcMatch, an unsupervised multi-shape matching framework with dual-level cycle consistency. Using a shape graph attention network, our method captures the manifold structure of shape collections, enabling more informed and robust feature aggregation. By enforcing spatial and spectral cycle consistency, DcMatch improves accuracy and coherence in the shared universe space. Extensive experiments across multiple benchmarks and ablation studies validate the superior performance and effectiveness of its components. We believe this work advances multi-shape matching by highlighting the importance of manifold-aware modeling and dual-level consistency.

## Acknowledgments

This work was supported by National Key R&D Program of China (No.2024YFE0202500), the Natural Science Foundation of Guangdong Province, China (2023A1515012834), the Natural Science Foundation of Hubei Province, China (Grant 2025AFB639) and the Natural Science Foundation of China (U23B2050).

## References

- Anguelov, D.; Srinivasan, P.; Koller, D.; Thrun, S.; Rodgers, J.; and Davis, J. 2005. SCAPE: shape completion and animation of people. *ACM ToG*, 24(3): 408–416.
- Aubry, M.; Schlickewei, U.; and Cremers, D. 2011. The wave kernel signature: A quantum mechanical approach to shape analysis. In *ICCV Workshops*, 1626–1633.
- Bastian, L.; Xie, Y.; Navab, N.; and Löhner, Z. 2024. Hybrid functional maps for crease-aware non-isometric shape matching. In *CVPR*, 3313–3323.
- Bernard, F.; Thunberg, J.; Goncalves, J.; and Theobalt, C. 2019a. Synchronisation of partial multi-matchings via non-negative factorisations. *Pattern Recognition*, 92: 146–155.
- Bernard, F.; Thunberg, J.; Swoboda, P.; and Theobalt, C. 2019b. Hippo: Higher-order projected power iterations for scalable multi-matching. In *ICCV*, 10284–10293.
- Bogo, F.; Romero, J.; Loper, M.; and Black, M. J. 2014. FAUST: Dataset and evaluation for 3D mesh registration. In *CVPR*, 3794–3801.
- Brody, S.; Alon, U.; and Yahav, E. 2022. How Attentive are Graph Attention Networks? In *ICLR*.
- Cao, D.; and Bernard, F. 2022. Unsupervised deep multi-shape matching. In *ECCV*, 55–71. Springer.
- Cao, D.; Roetzer, P.; and Bernard, F. 2023. Unsupervised Learning of Robust Spectral Shape Matching. *ACM ToG*, 42(4): 1–15.
- Cosmo, L.; Rodola, E.; Albarelli, A.; Mémoli, F.; and Cremers, D. 2017. Consistent partial matching of shape collections via sparse modeling. In *Computer Graphics Forum*, volume 36, 209–221. Wiley Online Library.
- Donati, N.; Corman, E.; and Ovsjanikov, M. 2022. Deep orientation-aware functional maps: Tackling symmetry issues in shape matching. In *CVPR*, 742–751.
- Donati, N.; Sharma, A.; and Ovsjanikov, M. 2020. Deep geometric functional maps: Robust feature learning for shape correspondence. In *CVPR*, 8592–8601.
- Eisenberger, M.; Lahner, Z.; and Cremers, D. 2020. Smooth shells: Multi-scale shape registration with functional maps. In *CVPR*, 12265–12274.
- Eisenberger, M.; Novotny, D.; Kerchenbaum, G.; Labatut, P.; Neverova, N.; Cremers, D.; and Vedaldi, A. 2021. Neuromorph: Unsupervised shape interpolation and correspondence in one go. In *CVPR*, 7473–7483.
- Eisenberger, M.; Toker, A.; Leal-Taixé, L.; and Cremers, D. 2020. Deep shells: Unsupervised shape correspondence with optimal transport. *NeurIPS*, 33: 10491–10502.
- Eisenberger, M.; Toker, A.; Leal-Taixé, L.; and Cremers, D. 2023. G-msm: Unsupervised multi-shape matching with graph-based affinity priors. In *CVPR*, 22762–22772.
- Eynard, D.; Rodola, E.; Glashoff, K.; and Bronstein, M. M. 2016. Coupled functional maps. In *3DV*, 399–407. IEEE.
- Gao, M.; Lahner, Z.; Thunberg, J.; Cremers, D.; and Bernard, F. 2021. Isometric multi-shape matching. In *CVPR*, 14183–14193.
- Halimi, O.; Litany, O.; Rodola, E.; Bronstein, A. M.; and Kimmel, R. 2019. Unsupervised learning of dense shape correspondence. In *CVPR*, 4370–4379.
- Hartwig, F.; Sassen, J.; Azencot, O.; Rumpf, M.; and Ben-Chen, M. 2023. An Elastic Basis for Spectral Shape Correspondence. In *2023 Special Interest Group on Computer Graphics and Interactive Techniques Conference, SIGGRAPH 2023*, 58. Association for Computing Machinery, Inc.
- Huang, Q.-X.; and Guibas, L. 2013. Consistent shape maps via semidefinite programming. In *Computer Graphics Forum*, volume 32, 177–186. Wiley Online Library.
- Huang, R.; Ren, J.; Wonka, P.; and Ovsjanikov, M. 2020. Consistent zoomout: Efficient spectral map synchronization. In *Computer Graphics Forum*, volume 39, 265–278. Wiley Online Library.
- Kim, V. G.; Lipman, Y.; and Funkhouser, T. 2011. Blended intrinsic maps. *ACM ToG*, 30(4): 1–12.
- Li, L.; Donati, N.; and Ovsjanikov, M. 2022. Learning multi-resolution functional maps with spectral attention for robust shape matching. *NeurIPS*, 35: 29336–29349.
- Litany, O.; Remez, T.; Rodola, E.; Bronstein, A.; and Bronstein, M. 2017a. Deep functional maps: Structured prediction for dense shape correspondence. In *ICCV*, 5659–5667.
- Litany, O.; Rodolà, E.; Bronstein, A. M.; and Bronstein, M. M. 2017b. Fully spectral partial shape matching. In *Computer Graphics Forum*, volume 36, 247–258. Wiley Online Library.
- Magnet, R.; Ren, J.; Sorkine-Hornung, O.; and Ovsjanikov, M. 2022. Smooth non-rigid shape matching via effective dirichlet energy optimization. In *3DV*, 495–504.
- Melzi, S.; Marin, R.; Rodolà, E.; Castellani, U.; Ren, J.; Poulenard, A.; Ovsjanikov, P.; et al. 2019a. SHREC’19: matching humans with different connectivity. In *Proceedings of the Eurographics Workshop on 3D Object Retrieval*, 1–8.
- Melzi, S.; Ren, J.; Rodolà, E.; Wonka, P.; Ovsjanikov, M.; et al. 2019b. ZoomOut: Spectral upsampling for efficient shape correspondence. *ACM ToG*, 38(6).
- Mena, G.; Snoek, J.; Linderman, S.; and Belanger, D. 2018. Learning latent permutations with Gumbel-Sinkhorn networks. In *ICLR 2018 Conference Track*, volume 2018. OpenReview.
- Nogneng, D.; and Ovsjanikov, M. 2017. Informative descriptor preservation via commutativity for shape matching. In *Computer Graphics Forum*, volume 36, 259–267. Wiley Online Library.

- Ovsjanikov, M.; Ben-Chen, M.; Solomon, J.; Butscher, A.; and Guibas, L. 2012. Functional maps: a flexible representation of maps between shapes. *ACM ToG*, 31(4): 1–11.
- Pachauri, D.; Kondor, R.; and Singh, V. 2013. Solving the multi-way matching problem by permutation synchronization. *NeurIPS*, 26.
- 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 *CVPR*, 384–393.
- Pinkall, U.; and Polthier, K. 1993. Computing discrete minimal surfaces and their conjugates. *Experimental Mathematics*, 2(1): 15–36.
- Ren, J.; Melzi, S.; Ovsjanikov, M.; and Wonka, P. 2020. Maptree: Recovering multiple solutions in the space of maps. *ACM ToG*, 39(6): 1–17.
- Ren, J.; Melzi, S.; Wonka, P.; and Ovsjanikov, M. 2021. Discrete optimization for shape matching. In *Computer Graphics Forum*, volume 40, 81–96. Wiley Online Library.
- Ren, J.; Panine, M.; Wonka, P.; and Ovsjanikov, M. 2019. Structured regularization of functional map computations. In *Computer Graphics Forum*, volume 38, 39–53. Wiley Online Library.
- Ren, J.; Poulenard, A.; Wonka, P.; and Ovsjanikov, M. 2018. Continuous and orientation-preserving correspondences via functional maps. *ACM ToG*, 37(6): 1–16.
- Rodolà, E.; Cosmo, L.; Bronstein, M. M.; Torsello, A.; and Cremers, D. 2017. Partial functional correspondence. In *Computer Graphics Forum*, volume 36, 222–236. Wiley Online Library.
- Roufousse, J.-M.; Sharma, A.; and Ovsjanikov, M. 2019. Unsupervised deep learning for structured shape matching. In *ICCV*, 1617–1627.
- Sahillioglu, Y. 2020. Recent advances in shape correspondence. *The Visual Computer*, 36(8): 1705–1721.
- Salti, S.; Tombari, F.; and Di Stefano, L. 2014. SHOT: Unique signatures of histograms for surface and texture description. *CVIU*, 125: 251–264.
- Sharp, N.; Attaiki, S.; Crane, K.; and Ovsjanikov, M. 2022. Diffusionnet: Discretization agnostic learning on surfaces. *ACM ToG*, 41(3): 1–16.
- Sinkhorn, R.; and Knopp, P. 1967. Concerning nonnegative matrices and doubly stochastic matrices. *Pacific Journal of Mathematics*, 21(2): 343–348.
- Sun, J.; Ovsjanikov, M.; and Guibas, L. 2009. A concise and provably informative multi-scale signature based on heat diffusion. In *Computer Graphics Forum*, volume 28, 1383–1392. Wiley Online Library.
- Sun, M.; Mao, S.; Jiang, P.; Ovsjanikov, M.; and Huang, R. 2023. Spatially and spectrally consistent deep functional maps. In *ICCV*, 14497–14507.
- Tron, R.; Zhou, X.; Esteves, C.; and Daniilidis, K. 2017. Fast multi-image matching via density-based clustering. In *ICCV*, 4057–4066.
- Van Kaick, O.; Zhang, H.; Hamarneh, G.; and Cohen-Or, D. 2011. A survey on shape correspondence. In *Computer Graphics Forum*, volume 30, 1681–1707. Wiley Online Library.
- Vestner, M.; Litman, R.; Rodola, E.; Bronstein, A.; and Cremers, D. 2017. Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *CVPR*, 3327–3336.
- Xia, Y.; Lu, Y.; Gao, Y.; and Ma, J. 2024. Locality preserving refinement for shape matching with functional maps. In *AAAI*, volume 38, 6207–6215.
- Xia, Y.; Ye, T.; Zhou, H.; Wang, Z.; and Ma, J. 2025. Multi-Shape Matching with Cycle Consistency Basis via Functional Maps. In *AAAI*, volume 39, 8575–8583.
- Zhuravlev, A.; Löhner, Z.; and Golyanik, V. 2025. Denoising Functional Maps: Diffusion Models for Shape Correspondence. In *CVPR*.
- Zuffi, S.; Kanazawa, A.; Jacobs, D. W.; and Black, M. J. 2017. 3D menagerie: Modeling the 3D shape and pose of animals. In *CVPR*, 6365–6373.