

Unsupervised Contrastive Learning for Efficient and Robust Spectral Shape Matching

Feifan Luo¹, Hongyang Chen^{2*}

¹College of Computer Science and Technology, Zhejiang University, Hangzhou 310027, China

²Research Center for Data Hub and Security, Zhejiang Lab, Hangzhou 311121, China

luoff@zju.edu.cn, dr.h.chen@ieee.org

Abstract

Estimating correspondences between pairs of non-rigid deformable 3D shapes remains a significant challenge in computer vision and graphics. While deep functional map methods have become the go-to solution for addressing this problem, they primarily focus on optimizing pointwise and functional maps either individually or jointly, rather than directly enhancing feature representations in the embedding space, which often results in inadequate feature quality and suboptimal matching performance. Furthermore, these approaches heavily rely on traditional functional map techniques, such as time-consuming functional map solvers, which incur substantial computational costs. In this work, we introduce, for the first time, a novel unsupervised contrastive learning-based approach for efficient and robust 3D shape matching. We begin by presenting an unsupervised contrastive learning framework that promotes feature learning by maximizing consistency within positive similarity pairs and minimizing it within negative similarity pairs, thereby improving both the consistency and discriminability of the learned features. We then design a significantly simplified functional map learning architecture that eliminates the need for computationally expensive functional map solvers and multiple auxiliary functional map losses, greatly enhancing computational efficiency. By integrating these two components into a unified two-branch pipeline, our method achieves state-of-the-art performance in both accuracy and efficiency. Extensive experiments demonstrate that our approach is not only computationally efficient but also outperforms current state-of-the-art methods across various challenging benchmarks, including near-isometric, non-isometric, and topologically inconsistent scenarios—even surpassing supervised techniques.

Code — <https://github.com/LuoFeifan77/ContrastiveFMNet>

Introduction

Non-rigid deformable shape matching is a fundamental problem in shape analysis and related fields, focusing on establishing meaningful correspondences between shapes. This task has a wide range of applications, including deformation transfer (Sumner and Popović 2004), shape interpolation (Eisenberger et al. 2021), and statistical shape analysis (Bogo et al. 2014).

*Corresponding author.

Copyright © 2026, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

With recent advancements in deep learning, numerous learning-based approaches have been proposed for non-rigid deformable shape matching, primarily based on the functional map pipeline (Ovsjanikov et al. 2012). The seminal work FMNet (also known as deep functional maps) (Litany et al. 2017) was the first to leverage learned shape features as optimal descriptors to generate the desired functional maps. Building upon this framework, several deep learning-based methods (Eisenberger et al. 2021; Cao, Roetzer, and Bernard 2023; Bastian et al. 2024; Attaiki, Pai, and Ovsjanikov 2021) have been introduced to address a variety of shape matching scenarios. However, most existing deep functional maps approaches are still concentrating on optimizing pointwise and functional maps individually or simultaneously, e.g., constraining pointwise maps (Halimi et al. 2019; Aygün, Löhner, and Cremers 2020) and functional maps alone (Li, Donati, and Ovsjanikov 2022; Donati, Corman, and Ovsjanikov 2022), or penalizing both pointwise and functional maps to promote properness (Attaiki and Ovsjanikov 2023; Cao, Roetzer, and Bernard 2023). While regularizing functional and pointwise maps is actually to learn informative local feature representations (Cao, Roetzer, and Bernard 2023), these advanced methods overlook the direct supervision of feature representation within the embedding space, resulting in subpar feature representation and inferior matching results. Furthermore, these methods like (Cao, Roetzer, and Bernard 2023; Luo et al. 2025) heavily rely on functional map solvers and multiple auxiliary functional map losses, resulting in complex training paradigms and low computational efficiency. A detailed comparison of methods can be found in Table 1.

To overcome the limitations, we introduce the first unsupervised contrastive learning-based method, offering a novel principle for resolving non-rigid 3D shape matching challenges. We first propose the unsupervised contrastive learning framework designed to promote informative feature learning. Positive and negative sets comprising candidate points with high and low similarity, respectively, are defined, derived from ranking vertex embedding similarities. Using these sets, a hybrid similarity generator constructs corresponding similarity pairs, i.e., positive and negative. By minimizing distance between positive pairs while maximizing distance between negative pairs, two novel unsupervised contrastive losses are introduced: cross-contrastive loss, which improves

	ProF	Unsup	WoS	OI
FMNet	✗	✗	✗	-
GeomFmaps	✗	✗	✗	-
SRFeat	✓	✗	✓	-
DUO-FMNet	✗	✓	✗	✗
AttentiveFMaps	✗	✓	✗	✗
RFMNet	✗	✓	✓	✓
ULRSSM	✗	✓	✗	✗
DiffZO	✗	✓	✓	✗
DeepFAFM	✗	✓	✗	✗
Ours	✓	✓	✓	✓

Table 1: Method comparison. Our method is the first unsupervised contrastive learning approach specifically for non-rigid deformable 3D shape matching, combining a distinctive set of properties that enhance its performance. Where ProF: Promote features in the embedding space. Unsup: Full unsupervised. WoS: Without functional map solver. OI: Only single unsupervised functional map loss.

feature representation consistency between shapes, and self-contrastive loss, which enhances feature self-discriminability. Then, we build the streamlined functional map learning architecture without the computationally intensive functional maps solver or auxiliary functional map losses, significantly reducing reliance on structural assumptions of functional maps (e.g., orthogonality) and minimizing the computational cost. Consequently, building upon the contrastive learning and functional map framework, our method achieves superior performance in both efficiency and accuracy across a wide range of matching scenarios. The main contributions are summarized as follows:

- The first unsupervised contrastive learning framework for non-rigid deformable 3D shape matching.
- A simple yet efficient two-branch architecture without a computationally intensive functional map solver and multiple auxiliary functional losses.
- Extensive experiments demonstrate that our approach achieves state-of-the-art performance across challenging benchmarks.

Related Works

Non-rigid 3D shape matching is a long-standing problem that has been studied extensively over the years. For a comprehensive overview of the field, readers are encouraged to consult surveys such as (Sahillioğlu 2020; Deng et al. 2022). Here, we review approaches most similar to ours.

Axiomatic Functional Map Methods

A landmark approach to non-rigid deformable shape matching is the functional map framework (Ovsjanikov et al. 2012), which has inspired numerous subsequent works (Ren et al. 2018, 2019; Huang et al. 2020; Hu et al. 2021; Ren et al. 2021; Gao et al. 2021; Pai et al. 2021; Fan et al. 2022; Magnet et al. 2022; Donati et al. 2022; Hartwig et al. 2023). These methods incorporate functional or map constraints to

enhance matching accuracy and robustness. However, axiomatic functional map methods rely heavily on the quality of extrinsic (Salti, Tombari, and Di Stefano 2014) and intrinsic (Sun, Ovsjanikov, and Guibas 2009; Aubry, Schlickewei, and Cremers 2011) handcrafted features. Their performance deteriorates under large-scale deformations, often producing unsatisfactory results. In contrast, our approach directly learns distinguishing features from training data, offering improved accuracy and robustness, especially in challenging matching scenarios.

Deep Functional Map Methods

Unlike axiomatic approaches, deep functional map methods aim to eliminate the reliance on handcrafted features by directly extracting shape features from training data. The pioneering FMNet (Litany et al. 2017) improved matching results by optimizing the SHOT descriptor (Salti, Tombari, and Di Stefano 2014) using residual MLP layers. Methods like UnsupFMNet (Halimi et al. 2019) and *Ayguen et al.* (Aygün, Löhner, and Cremers 2020) enhanced pointwise or functional maps within their loss functions, while other works (Roufousse, Sharma, and Ovsjanikov 2019; Sharma and Ovsjanikov 2020) optimized functional maps with properties like bijection, and orthogonality.

Recent two-branch architectures (Cao, Roetzer, and Bernard 2023; Sun et al. 2023) supervise both pointwise and functional maps to enhance properness, inspiring subsequent works (Cao et al. 2024; Cao, Roetzer, and Bernard 2024). Magnet et al. (Magnet and Ovsjanikov 2024) proposed a single-branch network to improve functional maps. However, these methods often neglect feature enhancement in the embedding space, limiting feature learning performance. Additionally, they still depend on computationally expensive solvers, and multiple loss functions, resulting in complex networks and reduced efficiency.

In contrast, we propose a novel pipeline that learns informative feature representations for more accurate and robust correspondences while significantly reducing reliance on functional map techniques. Our approach outperforms state-of-the-art methods (Cao, Roetzer, and Bernard 2023; Magnet and Ovsjanikov 2024) across challenging shape-matching scenarios, including those using supervised techniques (Litany et al. 2017; Donati, Sharma, and Ovsjanikov 2020), while maintaining high computational efficiency.

Contrastive Learning

Contrastive learning has recently emerged as a dominant technique in self-supervised learning, applied across various domains such as computer vision, natural language processing (NLP), and beyond. The core objective is to bring augmented versions of the same sample closer in the embedding space while pushing apart embeddings of different samples. Follow-up works have extended contrastive learning to structured data, such as graphs (You et al. 2020; Chen and Kou 2023; Shen et al. 2023), meshes (Li, Attaiki, and Ovsjanikov 2022), and point clouds (Xie et al. 2020; Cao and Bernard 2023; Jiang, Sun, and Huang 2023), among others. In contrast to Li et al. (Li, Attaiki, and Ovsjanikov 2022), who employed the PointInfoNCE loss (Xie et al. 2020), a common penalty

term in contrastive learning that promotes feature learning in a supervised manner and heavily relies on ground truth, we are the first to systematically introduce a comprehensive *unsupervised contrastive learning framework* specifically for non-rigid 3D shape matching. Our approach not only advances the theoretical foundation of shape matching but also outperforms supervised methods, such as (Li, Attaiki, and Ovsjanikov 2022), in terms of matching performance.

Background

We begin by providing a brief overview of the basic pipeline using deep functional maps and direct interested readers to relevant literature (Ovsjanikov et al. 2016; Donati, Sharma, and Ovsjanikov 2020; Cao, Roetzer, and Bernard 2023) for further details.

Deep Functional Map Pipeline

Given a pair of non-rigid 3D shapes denoted as \mathcal{X} and \mathcal{Y} with $|V_{\mathcal{X}}|$ and $|V_{\mathcal{Y}}|$ vertices, respectively. The goal is to compute a high quality dense correspondence between these shapes in an efficient way. The basic learning pipeline estimates a functional map between \mathcal{X} and \mathcal{Y} using the following five steps.

(1) **Precompute operators.** Compute the first k eigenfunctions $\Phi_{\mathcal{X}} \in \mathbb{R}^{|V_{\mathcal{X}}| \times k}$, $\Phi_{\mathcal{Y}} \in \mathbb{R}^{|V_{\mathcal{Y}}| \times k}$ eigenvalues $\Lambda_{\mathcal{X}}, \Lambda_{\mathcal{Y}} \in \mathbb{R}^{k \times k}$ in matrix notation via generalized eigen-decomposition of the corresponding Laplacian matrices $L_{\mathcal{X}} \in \mathbb{R}^{|V_{\mathcal{X}}| \times |V_{\mathcal{X}}|}$, $L_{\mathcal{Y}} \in \mathbb{R}^{|V_{\mathcal{Y}}| \times |V_{\mathcal{Y}}|}$, respectively. The Moore-Penrose pseudo-inverse of $\Phi_{\mathcal{M}}$ is $\Phi_{\mathcal{M}}^{\dagger} = \Phi_{\mathcal{M}}^{\top} \mathbf{A}_{\mathcal{M}}$, where $\mathbf{A}_{\mathcal{M}}$ is the diagonal matrix of lumped area elements, and \top denote transpose operation.

(2) **Feature extractor.** Compute feature vectors $F_{\mathcal{X}} \in \mathbb{R}^{|V_{\mathcal{X}}| \times d}$, $F_{\mathcal{Y}} \in \mathbb{R}^{|V_{\mathcal{Y}}| \times d}$ defined on each shape via a feature extractor network F_{θ} , where d is the dimension of features.

(3) **Functional map computation.** Compute the functional maps by solving the optimisation problem

$$\mathbf{C}_{\mathcal{Y}\mathcal{X}} = \arg \min_{\mathbf{C}_{\mathcal{Y}\mathcal{X}}} \left\| \mathbf{C}_{\mathcal{Y}\mathcal{X}} \Phi_{\mathcal{Y}}^{\dagger} \mathbf{F}_{\mathcal{Y}} - \Phi_{\mathcal{X}}^{\dagger} \mathbf{F}_{\mathcal{X}} \right\|_{\mathbb{F}}^2 + \lambda E_{reg}(\mathbf{C}_{\mathcal{Y}\mathcal{X}}), \quad (1)$$

The first term represents the descriptor preservation term, while the second term, known as the Laplacian commutativity term, enforces structural consistency in the functional map. The parameter λ acts as a hyperparameter.

(4) **Functional map penalty.** During training stage, structural regularisation (e.g. orthogonality, bijectivity) is imposed on the functional maps, i.e

$$L_{fmap} = \theta_{bi} L_{bi} + \theta_{or} L_{or}, \quad (2)$$

where

$$L_{bi} = \left\| \mathbf{C}_{\mathcal{Y}\mathcal{X}} \mathbf{C}_{\mathcal{X}\mathcal{Y}} - \mathbf{I} \right\|_{\mathbb{F}}^2, L_{or} = \left\| \mathbf{C}_{\mathcal{Y}\mathcal{X}} \mathbf{C}_{\mathcal{Y}\mathcal{X}}^{\top} - \mathbf{I} \right\|_{\mathbb{F}}^2. \quad (3)$$

To promote proper functional maps, Cao et al. (Cao, Roetzer, and Bernard 2023) proposed a coupling loss, i.e.

$$L_{co} = \left\| \mathbf{C}_{\mathcal{Y}\mathcal{X}} - \Phi_{\mathcal{X}}^{\dagger} \Pi_{\mathcal{X}\mathcal{Y}} \Phi_{\mathcal{Y}} \right\|_{\mathbb{F}}^2, \quad (4)$$

where the pointwise correspondence matrix $\Pi_{\mathcal{X}\mathcal{Y}}$ can be formulated either as a differentiable doubly-stochastic matrix enabling probabilistic correspondence.

(5) **Pointwise maps computation.** At inference time, the pointwise map commonly by nearest neighbor search between the aligned spectral embeddings $\Phi_{\mathcal{X}} \mathbf{C}_{\mathcal{X}\mathcal{Y}}^{\top}$ and $\Phi_{\mathcal{Y}}$, namely,

$$\Pi_{\mathcal{X}\mathcal{Y}} = NNsearch(\Phi_{\mathcal{Y}} \mathbf{C}_{\mathcal{Y}\mathcal{X}}^{\top}, \Phi_{\mathcal{X}}). \quad (5)$$

Most existing deep functional map methods focus solely on optimizing pointwise and functional maps either separately or together, rather than enhancing shape features in embedding space, resulting in less than ideal performance for feature learning and matching results. In contrast, our approach enhances feature representation learning through contrastive learning, resulting in consistent and discriminative features and more accurate correspondence performance. Furthermore, our method eliminates complex and time-consuming components of deep functional maps, such as functional map solvers Eq. (1), achieving higher computational efficiency than theirs.

Unsupervised Contrastive Learning Framework

In this section, we comprehensively introduce an unsupervised contrastive learning framework for enhancing the consistency and discriminative power of feature representations in the embedding space.

Positive and Negative Set

Definition 1. Considering a vertex x_i from shape \mathcal{X} and vertex sets $\{y_j\}_{j=1}^{|V_{\mathcal{Y}}|}$ from shape \mathcal{Y} , calculating the feature similarity scores $\mathbf{S}_{i,:} = \{\text{sim}(\mathbf{F}_i, \mathbf{F}_j)\}_{j=1}^{|V_{\mathcal{Y}}|} \in \mathbb{R}^{|V_{\mathcal{Y}}|}$, where \mathbf{F}_i and \mathbf{F}_j denote the feature representation of vertex x_i and y_j , respectively. Then, selecting the top k largest scores as the positive sets. Formally, this is expressed as

$$\mathcal{P}(x_i) = \{y_j \mid y_j \in \text{topk}(\mathbf{S}_{i,:})\} \in \mathbb{R}^k, \quad (6)$$

where sim denotes a function measuring the similarity between features, and topk denotes the top- k sampling function. Contrarily, the negative set can be denoted as

$$\mathcal{N}(x_i) = \{y_j \mid y_j \notin \text{topk}(\mathbf{S}_{i,:})\} \in \mathbb{R}^{|V_{\mathcal{Y}}| - k}. \quad (7)$$

We define the positive set and negative set according to the similarity rank of the feature embeddings. Since the similarity score represents the correspondence probability between shapes, the positive set and negative set actually encode the high and low correspondence probability between the shape vertices, respectively.

Hybrid Similarity Generator

Building on the previously introduced positive and negative sets, we propose a hybrid similarity generator that produces corresponding positive and negative similarity pairs through two main steps: similarity estimation and similarity sampling. The core of the similarity estimation step involves computing the similarity score matrix $\mathbf{S}_{\mathcal{X}\mathcal{Y}} \in \mathbb{R}^{|V_{\mathcal{X}}| \times |V_{\mathcal{Y}}|}$, which measures the similarity between vertex features $\mathbf{F}_{\mathcal{X}}$ and $\mathbf{F}_{\mathcal{Y}}$. The similarity metric can be defined using either a distance-based

or a cosine-based function. In this work, we adopt the cosine similarity function, defined as follows:

$$\mathbf{S}_{\mathcal{X}\mathcal{Y}} = \mathbf{F}_{\mathcal{X}}\mathbf{F}_{\mathcal{Y}}^T \in \mathbb{R}^{|V_{\mathcal{X}}| \times |V_{\mathcal{Y}}|} \quad (8)$$

where the features are L_2 row-wise normalized before computing dot products. s_{ij} represents the feature similarity between vertex x_i and y_j . Clearly, different vertex features result in different similarity score matrices, capturing the relationships between vertex pairs (x_i, y_j) in distinct feature spaces.

After estimating the similarity scores for all vertex pairs, we perform a two-stage sampling process to construct the positive and negative similarity pairs, respectively. The elements of the i -th row of the similarity matrix $\mathbf{S}_{\mathcal{X}\mathcal{Y}}$ represent the similarity scores of vertex x_i and all vertices from shape \mathcal{Y} . Therefore, computing the positive similarity pairs involves selecting the first k largest elements from each row, namely:

$$\mathbf{S}_{\mathcal{X}\mathcal{Y}}^+ = \{s_{ij} | s_{ij} \in \text{topk}(\mathbf{S}_{i,:}), i = 1, 2, \dots, |V_{\mathcal{X}}|\} \in \mathbb{R}^{|V_{\mathcal{X}}| \times k}, \quad (9)$$

$\mathbf{S}_{\mathcal{X}\mathcal{Y}}^+$ encodes the set of top k most similar vertex embeddings from shape \mathcal{Y} .

For the negative similarity sampling stage, we define the remaining similarity pairs as $\mathbf{S}_{\mathcal{X}\mathcal{Y}}^R = \mathbf{S}_{\mathcal{X}\mathcal{Y}} \setminus \mathbf{S}_{\mathcal{X}\mathcal{Y}}^+ \in \mathbb{R}^{|V_{\mathcal{X}}| \times (|V_{\mathcal{Y}}| - k)}$, where \setminus symbol represents the set difference operation. Then we apply a sampling function to select nodes from $\mathbf{S}_{\mathcal{X}\mathcal{Y}}^R$ to construct the negative similarity pairs, defined as follows:

$$\mathbf{S}_{\mathcal{X}\mathcal{Y}}^- = \{s_{ij} | s_{ij} \in \text{Sample}(\mathbf{S}_{\mathcal{X}\mathcal{Y}}^R)\}. \quad (10)$$

Clearly, based on the similarity ranking and sampling function, the hybrid similarity generator effectively divides the feature similarity score $\mathbf{S}_{\mathcal{X}\mathcal{Y}}$ into the positive similarity pair $\mathbf{S}_{\mathcal{X}\mathcal{Y}}^+$ and the negative similarity pair $\mathbf{S}_{\mathcal{X}\mathcal{Y}}^-$ (see Figure 1), which capture high-similarity and low-similarity feature pairs, respectively.

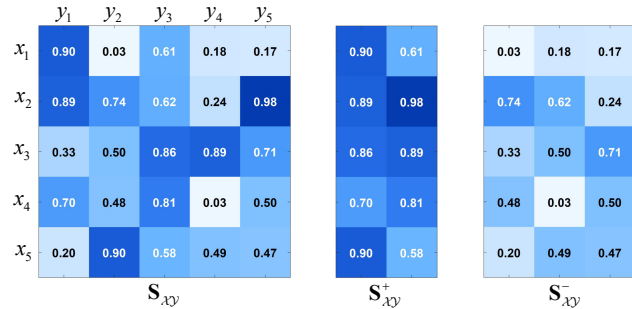


Figure 1: A visual example illustrating the similarity score $\mathbf{S}_{\mathcal{X}\mathcal{Y}}$, the positive similarity pair $\mathbf{S}_{\mathcal{X}\mathcal{Y}}^+$, and the negative similarity pair $\mathbf{S}_{\mathcal{X}\mathcal{Y}}^-$.

Cross and Self Contrastive Learning Loss

The core objective of our contrastive loss is to narrow the distance between the positive similarity pairs while widening the distance between the negative similarity pairs in the feature space. Building on the hybrid similarity generator described

above, we propose cross-contrastive and self-contrastive loss functions, which respectively enhance the consistency and discrimination of the learned feature representations.

Cross-contrastive loss. The cross-contrastive loss is calculated as follows:

$$L_{cross}(\mathbf{F}_{\mathcal{X}}, \mathbf{F}_{\mathcal{Y}}) = \frac{1}{|V_{\mathcal{X}}|} \sum_i -\log \frac{\exp(\hat{s}_{ij}^+/\tau_c)}{\sum_{j=1}^{n_c} \exp(s_{ij}^-/\tau_c)}, \quad (11)$$

where $s_{ij}^+ \in \mathbf{S}_{\mathcal{X}\mathcal{Y}}^+$, $s_{ij}^- \in \mathbf{S}_{\mathcal{X}\mathcal{Y}}^-$, $\hat{s}_{ij}^+ = \frac{1}{p_c} \sum_{j=1}^{p_c} s_{ij}^+$, p_c means top p_c largest scores as the positive pairs, n_c denotes the number of negative pairs for each vertex x_i , and τ_c denotes the temperature parameter that controls the sensitivity of penalties on positive and negative similarity. Eq. (11) enforces the representation of the target features to be close to the central representation of all positive samples and away from all negative samples. Cross-contrastive loss makes the similar features between $\mathbf{F}_{\mathcal{X}}$ and $\mathbf{F}_{\mathcal{Y}}$ more similar and the dissimilar features more dissimilar, which encourages feature consistency between shapes, beneficial for downstream shape matching tasks. Moreover, we can build a bidirectional cross-contrastive loss function by exchanging the positions of $\mathbf{F}_{\mathcal{X}}$ and $\mathbf{F}_{\mathcal{Y}}$ in Eq (8).

Self-contrastive loss. The previously introduced loss L_{cross} enforces similarity consistency between the shape features $\mathbf{F}_{\mathcal{X}}$ and $\mathbf{F}_{\mathcal{Y}}$. In addition, to further enhance the discriminability of the features themselves, we incorporate intra-shape similarity relationships and introduce the self-contrastive loss, which encourages the learning of more discriminative and self-aware feature representations.

Consider shape \mathcal{X} as an illustration, the self-similarity score $\mathbf{S}_{\mathcal{X}\mathcal{X}}$ can be expressed as

$$\mathbf{S}_{\mathcal{X}\mathcal{X}} = \mathbf{F}_{\mathcal{X}}\mathbf{F}_{\mathcal{X}}^T \in \mathbb{R}^{|V_{\mathcal{X}}| \times |V_{\mathcal{X}}|}. \quad (12)$$

For each vertex, we pick the first p_s similar candidates as the positive similarity pairs, and we obtain

$$\mathbf{S}_{\mathcal{X}\mathcal{X}}^+ = \{s_{ij} | s_{ij} \in \text{topk}(\mathbf{S}_{i,:}), i = 1, 2, \dots, |V_{\mathcal{X}}|\} \in \mathbb{R}^{|V_{\mathcal{X}}| \times p_s}. \quad (13)$$

Similarly, the negative similarity pairs can be acquired by downsampling the remaining vertex features, formulated as:

$$\mathbf{S}_{\mathcal{X}\mathcal{X}}^- = \{s_{ij} | s_{ij} \in \text{Sample}(\mathbf{S}_{\mathcal{X}\mathcal{X}}^R)\}, \quad (14)$$

where $\mathbf{S}_{\mathcal{X}\mathcal{X}}^R = \mathbf{S}_{\mathcal{X}\mathcal{X}} \setminus \mathbf{S}_{\mathcal{X}\mathcal{X}}^+ \in \mathbb{R}^{|V_{\mathcal{X}}| \times (|V_{\mathcal{X}}| - p_s)}$.

Unlike cross-contrastive learning, self-contrastive learning leverages the fact that each vertex's most similar feature is itself. Therefore, there is no need to encourage proximity to itself or to the central representation of the positive pairs. Instead, it is sufficient to push the representation away from the negative pairs. Based on this insight, we omit the attraction to positive pairs and define the self-contrastive loss as follows:

$$L_{self}(\mathbf{F}_{\mathcal{X}}) = \frac{1}{|V_{\mathcal{X}}|} \sum_i -\log \frac{1}{\sum_{j \neq i}^{n_s} \exp(s_{ij}^-/\tau_s)}, \quad (15)$$

where $s_{ij}^- \in \mathbf{S}_{\mathcal{X}\mathcal{X}}^-$, n_s denotes negative sample size, and τ_s denotes the temperature parameter. Similarly, we can also apply the self-contrastive loss L_{self} for shape feature $\mathbf{F}_{\mathcal{Y}}$.

Interestingly, the self-contrastive loss can be regarded as a special case of the cross-contrastive loss when considering correspondences between identical shapes; however, they serve distinct roles in feature learning. The cross-contrastive loss encourages consistency in feature representations across different shapes, while the self-contrastive loss promotes the self-discriminative ability of vertex representations. A visual example is shown in Figure 2. Top row: The color of x_1 becomes cooler after adding the self-contrastive loss, indicating an increased feature difference between x_1 and x_0 , which demonstrates that the loss promotes self-discrimination by expanding the distance between dissimilar points in the feature space. Bottom row: After adding the cross-contrastive loss, the color of y_0 becomes hotter, while the color of y_1 becomes cooler, highlighting that the feature difference between y_0 and x_0 has decreased, and between y_1 and x_0 has increased. This illustrates that the cross-contrastive loss promotes feature consistency by reducing the distance between similar points and increasing the distance between dissimilar points. Together, these two losses complement each other, guiding feature learning from different perspectives and leading to more comprehensive and expressive shape representations.




Figure 2: A visual example highlighting the functionality of two contrastive losses. The Euclidean distance between the feature of different points x_1 on shape \mathcal{X} , and y_0, y_1 on shape \mathcal{Y} , to the source point x_0 is computed, where hotter/colder colors mean smaller/larger distances.

Efficient and Robust Unsupervised Contrastive Learning Spectral Shape Matching

In this section, we propose a novel, two-branch unsupervised shape matching approach, which integrates the aforementioned unsupervised contrastive learning framework and the simplified functional map architecture. An overview of the pipeline is illustrated in Figure 3.

Unsupervised Contrastive Learning Branch

The contrastive learning branch follows the previously described unsupervised contrastive learning framework; therefore, we do not reiterate it here but provide guidance on sampling functions and parameter settings. To maintain a simple and efficient shape matching framework, we intentionally omit the sampling functions, although incorporating them could potentially improve our method’s performance. Specifically, we set $\mathbf{S}_{\mathcal{X}\mathcal{Y}}^- = \mathbf{S}_{\mathcal{X}\mathcal{Y}}^R, n_c = |\mathcal{Y}| - p_c$ for cross-contrastive loss Eq.(11), and $\mathbf{S}_{\mathcal{X}\mathcal{X}}^- = \mathbf{S}_{\mathcal{X}\mathcal{X}}^R, n_s = |\mathcal{V}_{\mathcal{X}}| - p_s$

for self-contrastive loss Eq.(15). Additionally, we set $\tau_c = \tau_s$ and $p_c = p_s$ for both losses to further simplify parameter tuning. This approach eliminates the need for designing specialized sampling functions and parameters, demonstrating the robustness of our framework.

Simplified Functional Map Learning Branch

Unlike the functional map learning-based approaches (Li, Donati, and Ovsjanikov 2022; Cao, Roetzer, and Bernard 2023; Luo et al. 2025) that heavily rely on computationally expensive functional map solver and multiple functional map losses to achieve desirable performances. Our functional map learning branch is streamlined, comprising only differentiable soft pointwise and functional map computations, along with a single loss term. For the computation of soft pointwise maps, we forego the higher precision but computationally expensive optimal transport algorithm (Eisenberger et al. 2020; Hu et al. 2023). Instead, we utilize the softmax operator to efficiently generate a soft correspondence matrix (Eisenberger et al. 2021), namely,

$$\Pi_{\mathcal{X}\mathcal{Y}} = \text{Softmax}(\mathbf{F}_{\mathcal{X}}\mathbf{F}_{\mathcal{Y}}^T/\alpha), \quad (16)$$

where the element at position (i, j) represents the probability of correspondence between the i -th point on \mathcal{X} and the j -th point on \mathcal{Y} , and α is the scaling factor to determine the softness of the correspondence matrix.

Differs from regularized optimization approaches Eq. (1), an alternative computational strategy establishes an explicit relationship between the functional maps and pointwise correspondence (Ovsjanikov et al. 2012), namely,

$$\mathbf{C}_{\mathcal{Y}\mathcal{X}} = \Phi_{\mathcal{X}}^\dagger \Pi_{\mathcal{X}\mathcal{Y}} \Phi_{\mathcal{Y}}. \quad (17)$$

The functional map $\mathbf{C}_{\mathcal{Y}\mathcal{X}}$ is calculated by spectral basis projection Eq.(17), rather than the linear system solver Eq.(1), which not only reduces the computational cost but also avoids the instability of the least squares system.

A key component of the deep functional maps framework is the incorporation of loss functions to enforce structural properties of functional maps, such as orthogonality, bijectivity, and so on. Most existing deep functional map methods rely on *multiple* functional map losses to compensate for limitations in their architectures, enabling robust mapping performance across diverse scenarios.

By effectively promoting feature consistency and discriminability through our contrastive learning framework, conventional bijectivity and orthogonality regularizers can be replaced by the proposed contrastive loss. Consequently, our approach operates without relying on explicit structural constraints, requiring only a single, simple loss function (Hu et al. 2023) from to promote consistency between the functional map and soft pointwise map, namely:

$$L_{align} = \left\| \Phi_{\mathcal{X}} - \Pi_{\mathcal{X}\mathcal{Y}} \Phi_{\mathcal{Y}} \mathbf{C}_{\mathcal{Y}\mathcal{X}} \right\|_{\mathbf{F}}^2. \quad (18)$$

Overall, the total unsupervised loss can be expressed as

$$L_{total} = \theta_{cross} L_{cross} + \theta_{self} L_{self} + \theta_{align} L_{align}, \quad (19)$$

where θ denotes the corresponding weight.

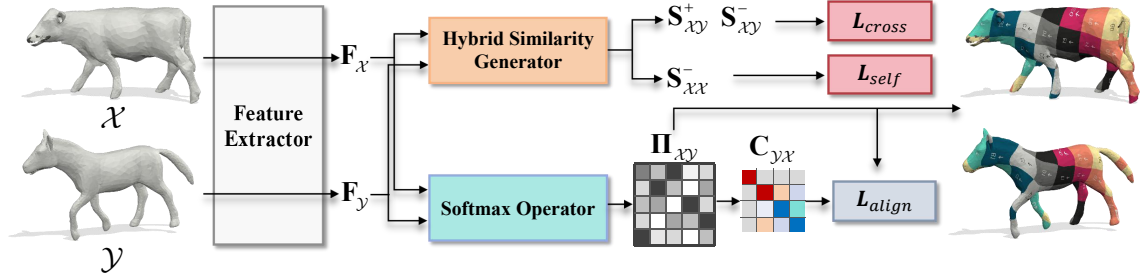


Figure 3: An overview of our method. (1) Feature Extraction: Learned features \mathbf{F}_X and \mathbf{F}_Y are extracted from shapes \mathcal{X} and \mathcal{Y} , respectively. (2) Unsupervised Contrastive Learning Branch: The learned features are used to generate positive and negative similarity pairs \mathbf{S}_{XY}^+ , \mathbf{S}_{XY}^- , \mathbf{S}_{XX}^- via the hybrid similarity generator. Two unsupervised contrastive losses Eq. (11) and Eq. (15) are then applied for feature enhancement. (3) Simplified Functional Map Branch: The differentiable pointwise map Π_{XY} is computed using the softmax operator Eq. (16), and the functional map \mathbf{C}_{YX} is calculated via spectral basis projection Eq. (17). A functional loss Eq. (18) is constructed to supervise both pointwise and functional map learning.

Finally, we adopt the same strategy as ULRSSM (Cao, Roetzer, and Bernard 2023) to recover the pointwise map during inference, ensuring consistency and accuracy in the final matching results.

Experiments and Results

Baselines

We extensively compare our method with existing non-rigid deformable shape matching methods, which we categorize as follows:

- *Axiomatic approaches*, including ZoomOut (Melzi et al. 2019b), Smooth Shells (Eisenberger, Löhner, and Cremers 2020), DiscreteOp (Ren et al. 2021), and MWP (Hu et al. 2021).
- *Supervised approaches*, FMNet (Litany et al. 2017), GeomFmaps (Donati, Sharma, and Ovsjanikov 2020), SRFeat-S (Li, Attaiki, and Ovsjanikov 2022).
- *Unsupervised approaches*, including UnsupFMNet (Halimi et al. 2019), SURFMNet (Roufousse, Sharma, and Ovsjanikov 2019), Deep Shells (Eisenberger et al. 2020), NeuroMorph (Eisenberger et al. 2021), DUO-FMNet (Donati, Corman, and Ovsjanikov 2022), AttentiveFMaps (Li, Donati, and Ovsjanikov 2022), RFMNet (Hu et al. 2023), ULRSSM (Cao, Roetzer, and Bernard 2023), DifZO (Magnet and Ovsjanikov 2024), HybridFMaps (Bastian et al. 2024), and DeepFAFM (Luo et al. 2025). Moreover, ULRSSM introduces a fine-tuning technique called Test-Time Adaptation (TTA), which individually adjusts the network parameters for each test pair during inference. We extend ULRSSM by incorporating TTA, denoted as ULRSSM(+TTA), and perform a comprehensive comparison with this enhanced version. For HybridFMaps (Bastian et al. 2024)—another approach that employs fine-tuning to enhance performance was excluded from comparison. Finally, we will not highlight supervised methods and fine-tuned matching results because the former relies on ground truth and the latter requires training the network on a test set, and the purpose of using these methods as baselines is to further demonstrate that our methods

require neither labeling nor fine-tuning, and still achieve superior performance, even better than them.

Results

Extensive experimental results across multiple datasets, including challenging non-isometric ones, are presented. We use the mean geodesic error (Kim, Lipman, and Funkhouser 2011) to evaluate correspondence accuracy, with all results multiplied by 100 for better readability.

Near-isometric matching. We evaluate our method on the remeshed versions (Ren et al. 2018) of the standard benchmarks FAUST and SCAPE (referred to as F and S, respectively), which are more challenging than the original datasets. FAUST consists of 100 human shapes, representing 10 different people in 10 different poses, and is split into 80 for training and 20 for testing. SCAPE contains 71 human shapes, depicting the same person in various poses, and is divided into 51 for training and 20 for testing.

The results of these benchmarks are provided in Table 2, where our method is compared with current state-of-the-art axiomatic, supervised, and unsupervised learning approaches. The results indicate that our method performs better than the previous state-of-the-art axiomatic, supervised, and unsupervised methods, such as GeomFmaps and DeepFAFM on FAUST, and achieves comparable results on SCAPE.

Cross-dataset generalization. To evaluate the generalization performance of our method, we train networks on remeshed datasets and test them on another remeshed dataset, i.e., training on F and testing on S, and vice versa. Moreover, we use the more challenging SHREC’19 (S19 for short) dataset (Melzi et al. 2019a) exclusively as a test set. This dataset consists of 44 human shapes, providing a rigorous benchmark for our approach.

As shown in Table 2, the quantitative results demonstrate our approach outperforms others in most settings. Nevertheless, the existing cutting edge supervised (Li, Attaiki, and Ovsjanikov 2022) and unsupervised approach (Cao, Roetzer, and Bernard 2023; Bastian et al. 2024) suffers from huge performance drops when testing on the cross-dataset generalisation datasets (e.g., training on F and testing on S19),

Train Test	F			S			F		S		SMAL	DT4D-H	
	F	S	S19	F	S	S19	F _a	S _a	F _a	S _a		intra-class	inter-class
Axiomatic Methods													
ZoomOut	6.1	7.5	-	6.1	7.5	-	8.7	15.0	8.7	15.0	47.7	4.0	29.0
SmoothShells	2.5	4.7	-	2.5	4.7	-	5.4	5.0	5.4	5.0	34.9	1.1	6.3
DiscreteOp	5.6	13.1	-	5.6	13.1	-	6.2	14.6	6.2	14.6	36.1	3.6	27.6
MWP	3.1	4.1	-	3.1	4.1	-	8.2	8.7	8.2	8.7	20.9	1.7	25.4
Supervised Methods													
FMNet	11.1	30.0	-	33.0	17.0	-	42.0	43.0	43.0	41.0	-	9.6	38.0
GeomFmaps	2.6	3.4	9.9	3.0	3.0	12.2	3.2	3.8	8.4	3.1	4.3	2.1	4.1
SRFeat-S	1.1	3.9	13.1	2.5	2.2	8.9	2.3	4.0	3.2	2.1	3.7	2.4	5.0
Unsupervised Methods													
Deep Shells	1.7	5.4	27.4	2.7	2.5	23.4	12.0	16.0	15.0	10.0	21.4	3.4	31.1
DUO-FMNet	2.5	4.2	6.4	2.7	2.6	8.4	3.0	4.4	3.1	2.7	4.8	2.6	15.8
AttentiveFMaps	1.9	2.6	6.4	2.2	2.2	9.9	2.4	<u>2.8</u>	<u>2.5</u>	2.3	4.4	1.7	11.6
RFMNet	1.7	2.3	6.3	1.7	2.1	6.9	3.6	<u>2.6</u>	<u>3.6</u>	3.9	4.4	1.5	13.9
ULRSSM	1.6	6.7	14.5	4.8	1.9	18.5	2.5	8.9	7.0	1.9	4.5	0.9	5.2
ULRSSM(+TTA)	1.6	2.2	5.7	1.6	1.9	6.7	1.9	2.4	2.1	1.9	4.2	0.9	4.1
DiffZO	1.9	<u>2.4</u>	4.2	1.9	2.4	<u>6.9</u>	2.2	3.8	2.7	2.4	4.3	1.8	4.1
HybridFMaps	1.4	4.2	9.5	2.3	1.8	13.0	2.0	4.6	3.4	1.8	3.5	1.0	<u>3.9</u>
DeepFAFM	1.6	2.7	7.0	1.9	<u>1.9</u>	7.9	<u>2.0</u>	2.9	2.6	<u>1.9</u>	3.9	<u>0.9</u>	4.2
Ours	<u>1.6</u>	2.7	<u>5.8</u>	<u>1.8</u>	2.1	5.8	1.9	2.6	2.2	2.1	<u>3.6</u>	0.9	3.9

Table 2: Evaluating the matching results across various benchmarks, including near-isometric shape matching, cross-dataset generalization, anisotropic meshing, and non-isometric shape matching, respectively. The numbers in the table are mean geodesic errors ($\times 100$). **Bold:** Best. Underline: Runner-up.

which substantially demonstrates its inadequate generalization ability compared to existing unsupervised learning-based methods. Additionally, our method achieves superior performance even compared with the fine-tuning matching results of ULRSSM. Clearly, the results also show that the superior performance of ULRSSM heavily relies on the test-time adaptation process, our method is more robust than it.

Matching with anisotropic meshing. To evaluate robustness across different discretizations, we train networks on remeshed datasets and test them on anisotropic remeshed versions (denoted F_a and S_a, respectively), which feature different mesh connectivity compared to the original datasets.

The results presented in Table 2 show that our method achieves the best performance with state-of-the-art in most settings and demonstrates greater resilience to changes in triangulation. For instance, when training on S and testing on F_a, other competitors like HybridFMaps experience significant performance declines, often overfitting to mesh connectivity and producing inaccurate predictions. In contrast, our method maintains strong robustness to varying mesh connectivity and consistently surpasses the current state-of-the-art methods. Furthermore, we achieve performance nearly identical to fine-tuning results, which further underscores the robustness of our approach compared to existing learning-based methods.

Non-isometric shape matching. In the evaluation of non-isometric shape matching, our approach undergoes rigorous testing across two non-isometric datasets: SMAL (Zuffi et al. 2017) and DT4D-H (Magnet et al. 2022), which are significantly more challenging than previous datasets, such as SHREC’19 (Melzi et al. 2019a), due to the large deformations present in the shapes. The SMAL dataset consists of 49 shapes representing four-legged animals from eight species, divided into a training set of 32 instances and a testing set of

17 instances. In contrast, DT4D-H includes nine classes of humanoid shapes, with a training-testing split of 198 and 95 instances, respectively.

As shown in Table 2, our method consistently achieves excellent performance across most settings. On more challenging SMAL, our method surpasses existing supervised approaches and achieves performance comparable to HybridFMaps—which integrates both extrinsic and intrinsic information, while other purely state-of-the-art intrinsic spectral techniques like DiffZO and DeepFAFM are unable to attain such results. For the most challenging DT4D-inter dataset, our approach outperforms all baselines, clearly demonstrating the superiority of our method.

For additional comparison results, including topological noise matching, runtime comparisons, and so on, please refer to the supplementary materials.

Conclusion

We introduce a novel unsupervised contrastive learning-based approach for non-rigid deformable 3D shape matching, comprising the unsupervised contrastive learning framework and the simplified functional map learning module. Extensive experiments demonstrate that our method achieves state-of-the-art performance in both matching results and computational efficiency, establishing a new benchmark for unsupervised 3D shape matching. Additionally, its streamlined design positions it as a foundational framework for future research.

While effective, our method does not yet support partial shape matching or severely non-isometric deformations. A promising research direction involves integrating our framework with explicit spatial deformation methods (Eisenberger et al. 2020, 2021; Cao et al. 2024) to enhance robustness across diverse shape matching scenarios.

Acknowledgments

We extend our gratitude to Dongliang Cao for his generous assistance. This work is supported in part by National Key Research and Development Program of China 2023YFB4502400, in part by National Natural Science Foundation of China under Grant 62271452.

References

- Attaiki, S.; and Ovsjanikov, M. 2023. Understanding and improving features learned in deep functional maps. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1316–1326.
- Attaiki, S.; Pai, G.; and Ovsjanikov, M. 2021. DPFM: Deep Partial Functional Maps. In *International Conference on 3D Vision*, 175–185.
- Aubry, M.; Schlickewei, U.; and Cremers, D. 2011. The wave kernel signature: A quantum mechanical approach to shape analysis. In *IEEE International Conference on Computer Vision Workshops*, 1626–1633.
- Aygün, M.; Lähner, Z.; and Cremers, D. 2020. Unsupervised Dense Shape Correspondence using Heat Kernels. In *International Conference on 3D Vision*, 573–582.
- Bastian, L.; Xie, Y.; Navab, N.; and Lähner, 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 *IEEE Conference on Computer Vision and Pattern Recognition*, 3794–3801.
- Cao, D.; and Bernard, F. 2023. Self-supervised learning for multimodal non-rigid 3d shape matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 17735–17744.
- Cao, D.; Eisenberger, M.; El Amrani, N.; Cremers, D.; and Bernard, F. 2024. Spectral meets spatial: Harmonising 3d shape matching and interpolation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 3658–3668.
- 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*, 1371–1381. IEEE.
- Chen, J.; and Kou, G. 2023. Attribute and structure preserving graph contrastive learning. In *Proceedings of the AAAI conference on artificial intelligence*, 7024–7032.
- Deng, B.; Yao, Y.; Dyke, R. M.; and Zhang, J. 2022. A survey of non-rigid 3D registration. In *Computer Graphics Forum*, volume 41, 559–589.
- 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 732–741.
- Donati, N.; Sharma, A.; and Ovsjanikov, M. 2020. Deep Geometric Functional Maps: Robust Feature Learning for Shape Correspondence. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 8589–8598.
- Eisenberger, M.; Lähner, Z.; and Cremers, D. 2020. Smooth Shells: Multi-Scale Shape Registration With Functional Maps. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 12262–12271.
- 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 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 7473–7483.
- Eisenberger, M.; Toker, A.; Leal-Taixé, L.; and Cremers, D. 2020. Deep shells: Unsupervised shape correspondence with optimal transport. *Advances in Neural information processing systems*, 33: 10491–10502.
- Fan, A.; Ma, J.; Tian, X.; , X.; and Liu, W. 2022. Coherent Point Drift Revisited for Non-Rigid Shape Matching and Registration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 1424–1434.
- Gao, M.; Lahner, Z.; Thunberg, J.; Cremers, D.; and Bernard, F. 2021. Isometric Multi-Shape Matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14183–14193.
- Halimi, O.; Litany, O.; Rodolà, E. R.; Bronstein, A. M.; and Kimmel, R. 2019. Unsupervised Learning of Dense Shape Correspondence. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 4365–4374.
- 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 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 14531–14540.
- 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.
- Jiang, P.; Sun, M.; and Huang, R. 2023. Non-rigid shape registration via deep functional maps prior. *Advances in Neural Information Processing Systems*, 36: 58409–58427.
- Kim, V. G.; Lipman, Y.; and Funkhouser, T. 2011. Blended intrinsic maps. *ACM transactions on graphics*, 30: 1–12.
- Li, L.; Attaiki, S.; and Ovsjanikov, M. 2022. SRFeat: Learning locally accurate and globally consistent non-rigid shape

- correspondence. In *2022 International Conference on 3D Vision*, 144–154.
- Li, L.; Donati, N.; and Ovsjanikov, M. 2022. Learning multi-resolution functional maps with spectral attention for robust shape matching. *Advances in Neural Information Processing Systems*, 35: 29336–29349.
- Litany, O.; Remez, T.; Rodolà, E.; Bronstein, A.; and Bronstein, M. 2017. Deep Functional Maps: Structured Prediction for Dense Shape Correspondence. In *IEEE International Conference on Computer Vision*, 5660–5668.
- 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 *International Conference on 3D Vision*, 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 *Eurographics Workshop on 3D Object Retrieval*, 1–8.
- Melzi, S.; Ren, J.; Rodolà, E.; Sharma, A.; Wonka, P.; and Ovsjanikov, M. 2019b. ZoomOut: Spectral Upsampling for Efficient Shape Correspondence. *ACM Transactions on Graphics*, 38: 1–14.
- 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.
- Ovsjanikov, M.; Corman, E.; Bronstein, M.; Rodolà, E.; Ben-Chen, M.; Guibas, L.; Chazal, F.; and Bronstein, A. 2016. Computing and processing correspondences with functional maps. In *SIGGRAPH ASIA 2016 Courses*, 1–60.
- 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 *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. In *Computer Graphics Forum*, volume 40, 81–96.
- Ren, J.; Panine, M.; Wonka, P.; and Ovsjanikov, M. 2019. Structured regularization of functional map computations. In *Computer Graphics Forum*, volume 38, 39–53.
- 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 *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.
- Sharma, A.; and Ovsjanikov, M. 2020. Weakly supervised deep functional maps for shape matching. *Advances in Neural Information Processing Systems*, 33: 19264–19275.
- Shen, X.; Sun, D.; Pan, S.; Zhou, X.; and Yang, L. T. 2023. Neighbor contrastive learning on learnable graph augmentation. In *Proceedings of the AAAI conference on artificial intelligence*, 9782–9791.
- 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.
- Xie, S.; Gu, J.; Guo, D.; Qi, C. R.; Guibas, L.; and Litany, O. 2020. Pointcontrast: Unsupervised pre-training for 3d point cloud understanding. In *European conference on computer vision*, 574–591.
- You, Y.; Chen, T.; Sui, Y.; Chen, T.; Wang, Z.; and Shen, Y. 2020. Graph contrastive learning with augmentations. *Advances in neural information processing systems*, 33: 5812–5823.
- 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 the IEEE conference on computer vision and pattern recognition*, 6365–6373.