

Probabilistic Deformation Consistency for Unsupervised Shape Matching

Yifan Xia, Tianwei Ye, Jun Huang, Xiaoguang Mei, Jiayi Ma*

Electronic Information School, Wuhan University, Wuhan 430072, China
xiayifan@whu.edu.cn, jyama2010@gmail.com

Abstract

In this paper, we propose a novel unsupervised shape matching framework based on probabilistic deformation consistency in the spectral domain, termed as PDCMatch. Axiomatic optimization methods suffer from expensive geodesic distance calculations and vulnerability to local optima, and learning-based methods typically lack geometric consistency in pointwise correspondences. To overcome both limitations, we develop a non-Euclidean probabilistic deformation model that jointly estimates the underlying deformation and the correspondence probability via a linear Expectation-Maximization procedure. Building on this formulation, we further design a task-specific deformation loss that explicitly encourages geometric smoothness and structural consistency in an unsupervised manner. This tailored loss function plays a central role in improving the matching performance across challenging scenarios. Extensive experiments on public benchmarks involving near-isometric shapes, anisotropic meshing, cross-dataset generalization, topological noise, and non-isometric shapes demonstrate that our method consistently outperforms state-of-the-art methods, highlighting both its effectiveness and generalizability.

Code — <https://github.com/XiaYifan1999/PDCMatch>

Introduction

Shape matching, defined as the automated computation of dense correspondences between non-rigid shapes, is a long-standing fundamental problem in computer vision and graphics. It serves as a foundation for downstream tasks such as shape registration (Bogo et al. 2017), generation (Deng et al. 2022), deformation (Egger et al. 2020), and analysis (Sahillioğlu 2020). This problem has garnered significant attention and has been extensively investigated.

Traditional axiomatic methods (Holzschuh, Löhner, and Cremers 2020; Eisenberger, Löhner, and Cremers 2019) formulate the shape matching task as a quadratic assignment problem, yet this NP-hard problem incurs prohibitively high computational costs and suffers from local optima. Functional maps (Ovsjanikov et al. 2012), by contrast, solve the shape matching problem via linear operators in functional




Figure 1: Visualization of our probabilistic deformation model estimating deformation and predicting probability for given pointwise maps from the source shape to the target.




Figure 2: We propose a nonrigid shape matching method based on deformation consistency. The above four examples of texture transfer from source to target shapes demonstrate that our method is capable of producing accurate pointwise maps across difficult scenarios, including near-isometric cross-dataset testing, anisotropic meshing, topological noise, and non-isometric shape matching. Notably, our method performs well even in the challenging cyan hexagonal regions, such as significant surface discrepancies, intricate fine textures, and salient topological noise.

spaces and enforce isometric consistency via the Laplace-Beltrami (LB) operator. While follow-up works (Eisenberger, Lahner, and Cremers 2020; Magnet et al. 2022; Melzi et al. 2019b; Ren et al. 2021) have extended this framework by imposing structural constraints for more accurate matching, their effectiveness remains constrained by handcrafted descriptor limitations. Although handcrafted descriptors (Bronstein and Kokkinos 2010; Salti, Tombari, and Di Stefano 2014) have been developed for this purpose, they still lack sufficient representational capacity to represent 3D manifolds.

To address the above limitations, learning-based methods directly learn descriptor functions from data within the functional map framework, exemplified by seminal method (Litany et al. 2017a) and follow-up works (Halimi et al. 2019; Sharma and Ovsjanikov 2020; Donati, Sharma,

*Corresponding author.

and Ovsjanikov 2020; Li, Donati, and Ovsjanikov 2022; Cao, Roetzer, and Bernard 2023; Magnet and Ovsjanikov 2024). By integrating map refinement techniques during testing, deep functional maps have achieved state-of-the-art performance on several challenging benchmarks. However, most existing methods operate in the spectral domain and fail to ensure topological consistency, *e.g.*, adjacent vertices on one shape exhibit large deviation on the matched shape, thereby hindering fine-grained matching. Although recent methods (Roetzer and Bernard 2024; Roetzer et al. 2024) enforce neighborhood preservation via geodesic distances, the computational complexity limits their applicability. Thus, a key challenge lies in overcoming the inherent coarseness of low-dimensional functional maps to achieve accurate high-dimensional pointwise correspondences.

To address the above issues, we establish a probabilistic deformation model to evaluate the correspondence set under given pointwise maps. This model jointly estimates the underlying deformation and the matching probability based on topological consistency, which further guides the training of the feature extractors. Specifically, given pointwise maps, the corresponding points on a shape pair are expected to follow an underlying shape deformation (Eisenberger, Lahner, and Cremers 2020). Considering incorrect correspondences in the predicted pointwise maps, we formulate a probabilistic model to evaluate the correctness of the pointwise maps and design an unsupervised loss. Observing that deformation consistency, *i.e.*, neighboring points tend to move coherently during the deformation process, we propose a Gaussian mixture-based probabilistic deformation model and realize its efficient optimization through kernel techniques in the non-Euclidean domain coupled with a linear Expectation-Maximization algorithm. This allows us to jointly estimate a consistent deformation field and the correct probability of each correspondence as shown in Fig. 1. The underlying deformation reflects the smoothness of estimated pointwise maps, while the probability quantifies the matching accuracy. Based on these two predicted components, we design a deformation-consistency-driven unsupervised loss, which guides the feature extractor to learn more accurate and geometry-aware representations, thereby enabling precise dense pointwise correspondences. Extensive experiments demonstrate its matching performance as shown in Fig. 2 across multiple experimental categories. The main contributions of this paper are summarized as follows:

- a) We develop a probabilistic deformation framework that jointly models deformation smoothness and correspondence reliability, enabling robust unsupervised learning of nonrigid shape matching.
- b) We design a deformation-consistency-based unsupervised loss that effectively guides feature learning towards geometry-aware representations, significantly improving the accuracy of dense pointwise correspondence.
- c) Experiments on public datasets involving near-isometric, anisotropic meshing, topological noise, and non-isometric matching demonstrate our competitiveness over state-of-the-art methods.

Related Work

In the following, we review the shape matching works most relevant to our method. A comprehensive review can be found in surveys (Sahillioğlu 2020).

Axiomatic Shape Matching Methods

The simple axiomatic formulation is the linear assignment problem (Munkres 1957), which, however, does not account for the geometric properties of 3D manifolds. Some traditional axiomatic methods attempt to impose geometric constraints (Roetzer et al. 2022; Holzschuh, Löhner, and Cremers 2020), while others formulate the quadratic assignment problem based on non-rigid registration (Bernard, Suri, and Theobalt 2020; Eisenberger, Löhner, and Cremers 2019; Huang et al. 2008; Fan et al. 2022). However, due to the non-convex nature of the optimization problem and its NP-hardness, their solutions heavily rely on reliable initialization and remain difficult to solve for practical shape sizes.

Functional maps (Ovsjanikov et al. 2012) have gained widespread popularity due to their efficiency, achieved by embedding pointwise maps into a small matrix in the functional domain. Numerous follow-up works have extended this framework by incorporating robust mapping constraints (Eynard et al. 2016; Ren et al. 2019; Eisenberger, Lahner, and Cremers 2020; Xia et al. 2024), improving map optimization (Melzi et al. 2019b; Ren et al. 2021), addressing challenging scenarios (*e.g.*, non-isometry (Magnet et al. 2022; Ren et al. 2018) and partiality (Litany et al. 2017b; Rodolà et al. 2017)), and enabling multi-shape matching (Gao et al. 2021). However, axiomatic methods rely on handcrafted features (*e.g.*, HKS (Bronstein and Kokkinos 2010) and SHOT (Salti, Tombari, and Di Stefano 2014)), which fail to capture deep representations of 3D manifolds, significantly limiting their matching accuracy and generalizability.

Deep Functional Map Methods

Deep functional map methods outperform axiomatic approaches by learning more discriminative features from training data. As a pioneering work, FMNet (Litany et al. 2017a) employs a supervised approach to learn nonlinear transformations of SHOT descriptors (Salti, Tombari, and Di Stefano 2014) using multilayer perceptrons. Subsequent methods (Halimi et al. 2019; Roufousse, Sharma, and Ovsjanikov 2019; Sun et al. 2023; Attaiki and Ovsjanikov 2023a) enable unsupervised learning through near-isometric regularization. To enhance representational capacity, several methods (Donati, Sharma, and Ovsjanikov 2020; Sharma and Ovsjanikov 2020) incorporate pointwise networks (Qi et al. 2017; Thomas et al. 2019) for robust feature extraction. More recently, DiffusionNet (Sharp et al. 2022) has achieved state-of-the-art performance by leveraging the diffusion process for spatial communication, and has been applied in various scenarios (Attaiki and Ovsjanikov 2023b, 2022; Attaiki, Pai, and Ovsjanikov 2021; Cao and Bernard 2023). Follow-up works include shape interpolation (Eisenberger et al. 2021), hybridizing operator integration (Bastian et al. 2024), adaptive functional map solver (Cao, Roetzer,

and Bernard 2024), embedding shell solvers (Eisenberger et al. 2020), unsupervised spectral learning (Cao, Roetzer, and Bernard 2023), spectral attention mechanisms (Li, Donati, and Ovsjanikov 2022), and scalable memory architectures (Magnet and Ovsjanikov 2024).

Furthermore, some methods exploit features extracted from deep functional maps as priors, followed by the design of optimization-based shape registration methods (Jiang, Sun, and Huang 2023). Similarly, SpiderMatch (Roetzer and Bernard 2024) and DiscoMatch (Roetzer et al. 2024) devise solvers based on geometric consistency. However, such post-hoc registration approaches typically require computing pairwise geodesic distances for the neighborhood preservation, resulting in high computational overhead during iterative optimization. For instance, SpiderMatch only reports experiments on shapes with 1,000 triangles. In contrast, although our method also adopts EM-based iterative optimization, it is applied solely during training and benefits from linear complexity, thereby avoiding excessive time burden.

Deep Functional Maps

In this section, we briefly introduce the basic learning pipeline known as Deep Functional Maps (Litany et al. 2017a). Given two 3D shapes \mathcal{X} and \mathcal{Y} represented as triangular meshes with $n_{\mathcal{X}}$ and $n_{\mathcal{Y}}$ (*w.l.o.g.* $n_{\mathcal{X}} \leq n_{\mathcal{Y}}$) vertices, the goal is to compute accurate dense pointwise correspondences between them. We briefly outline the main steps of this pipeline, and refer readers to (Litany et al. 2017a; Cao, Roetzer, and Bernard 2023) for full details.

- Compute the first k eigenvector matrices $\Phi_{\mathcal{X}} \in \mathbb{R}^{n_{\mathcal{X}} \times k}$, $\Phi_{\mathcal{Y}} \in \mathbb{R}^{n_{\mathcal{Y}} \times k}$ and diagonal eigenvalue matrix $\Lambda_{\mathcal{X}} \in \mathbb{R}^{k \times k}$, $\Lambda_{\mathcal{Y}} \in \mathbb{R}^{k \times k}$ of LBO on shapes \mathcal{X} and \mathcal{Y} .
- Extract feature descriptors $\mathbf{F}_{\mathcal{X}} \in \mathbb{R}^{n_{\mathcal{X}} \times d}$ and $\mathbf{F}_{\mathcal{Y}} \in \mathbb{R}^{n_{\mathcal{Y}} \times d}$ for each shape using DiffusionNet (Sharp et al. 2022) as feature extraction network \mathcal{F}_{θ} , where d denotes the descriptor dimension.
- Compute pointwise maps Π using deep feature similarities, *i.e.*,

$$\Pi = \text{Softmax}(\mathbf{F}_{\mathcal{X}} \mathbf{F}_{\mathcal{Y}}^{\top}). \quad (1)$$

- Compute the optimal functional map $\mathbf{C} \in \mathbb{R}^{k \times k}$ via spectral embedding by solving a continuous optimization problem:

$$\mathbf{C} = \arg \min_{\mathbf{C}} E_{data}(\mathbf{C}) + \lambda E_{reg}(\mathbf{C}), \quad (2)$$

where $E_{data}(\mathbf{C}) = \left\| \mathbf{C} \Phi_{\mathcal{X}}^{\dagger} \mathbf{F}_{\mathcal{X}} - \Phi_{\mathcal{Y}}^{\dagger} \mathbf{F}_{\mathcal{Y}} \right\|_F^2$ preserves the feature descriptors, and E_{reg} imposes structural properties, *e.g.* commutativity (Ovsjanikov et al. 2012).

- The loss function \mathcal{L} during training is composed of spectral regularization terms about the functional maps, such as orthogonality, bijectivity (Roufousse, Sharma, and Ovsjanikov 2019), and coupling (Cao, Roetzer, and Bernard 2023), *i.e.*,

$$\mathcal{L}_{spectral} = \lambda_{bij} \mathcal{L}_{bij} + \lambda_{orth} \mathcal{L}_{orth} + \lambda_{couple} \mathcal{L}_{couple}. \quad (3)$$

- The pointwise map Π is derived based on $\mathbf{C} = \Phi_{\mathcal{Y}}^{\dagger} \Pi^{\top} \Phi_{\mathcal{X}}$, typically through nearest-neighbor search

and sometimes followed by post-processing steps (Eisenberger, Lahner, and Cremers 2020; Ezuz, Solomon, and Ben-Chen 2019; Melzi et al. 2019b; Vestner et al. 2017).

Method

The framework of our method is illustrated in Fig. 3. Given shapes \mathcal{X} and \mathcal{Y} represented as triangle meshes, deep functional maps learn feature representations for them to estimate pointwise maps. Our method aims to enhance the accuracy of dense pointwise maps based on probabilistic deformation consistency, with a detailed description as follows.

Probabilistic Deformation Model

Given pointwise maps Π as input, our goal is to estimate the underlying deformation $\mathcal{T} : \mathbb{R}^3 \rightarrow \mathbb{R}^3$ between \mathbf{X} and $\Pi \mathbf{Y}$ and evaluate the reliability of correspondences, where matrices $\mathbf{X} = [\mathbf{x}_i]$ and $\mathbf{Y} = [\mathbf{y}_i]$ store the spatial coordinates of vertices on shape \mathcal{X} and \mathcal{Y} , respectively. We define a putative correspondence set by:

$$\mathcal{M} = \{m_i\}_{i=1}^{n_{\mathcal{X}}} = \{\mathbf{x}_i, \mathbf{y}_{i'}\}_{i=1}^{n_{\mathcal{X}}}, \quad (4)$$

$$s.t., i' = \arg \max_j \Pi_{ij},$$

where $\mathbf{y}_{i'} = \mathcal{T}(\mathbf{x}_i)$ when m_i is an ideal correct correspondence and $n_{\mathcal{X}}$ denotes both the number of vertex correspondences and the vertex count of shape \mathcal{X} . We model the source point set \mathbf{X} as samples drawn from a Gaussian Mixture Model (GMM, following (Myronenko and Song 2010)), whose centroids are given by the target coordinates $\Pi \mathbf{Y}$. We assume equal component weights and isotropic covariances with shared variance σ^2 . Outliers are modeled by a uniform distribution with density $\frac{1}{a}$, where a denotes the spatial domain volume. By introducing γ as the weighting factor for outliers and D as coordinate dimensions, we formulate the probability of correspondence m_i being correct as follows:

$$p(m_i) = \frac{1 - \gamma}{(2\pi\sigma^2)^{D/2}} e^{-\frac{\|\mathbf{y}_{i'} - \mathcal{T}(\mathbf{x}_i)\|^2}{2\sigma^2}} + \gamma \frac{1}{a}. \quad (5)$$

Following the maximum likelihood principle, the objective is to maximize the joint probability $p(\mathcal{M}) = \prod_{i=1}^{n_{\mathcal{X}}} p(m_i)$. We adopt the Expectation-Maximization (EM) algorithm, which iteratively alternates between the Expectation (E) step and the Maximization (M) step.

During the E-step, posterior probabilities are computed to assign data points to mixture components:

$$p_i = \frac{(1 - \gamma) e^{-\frac{\|\mathbf{y}_{i'} - \mathcal{T}(\mathbf{x}_i)\|^2}{2\sigma^2}}}{(1 - \gamma) e^{-\frac{\|\mathbf{y}_{i'} - \mathcal{T}(\mathbf{x}_i)\|^2}{2\sigma^2}} + \gamma \frac{(2\pi\sigma^2)^{D/2}}{a}}. \quad (6)$$

In the M-step, model parameters are updated by minimizing the expected negative log-likelihood:

$$\mathcal{J}(\mathcal{T}, \sigma^2) = \sum_{i=1}^{n_{\mathcal{X}}} p_i \frac{\|\mathbf{y}_{i'} - \mathcal{T}(\mathbf{x}_i)\|^2}{2\sigma^2} + \frac{D}{2} \log \sigma^2 \sum_i p_i, \quad (7)$$

which focuses on fitting the underlying deformation \mathcal{T} between \mathbf{X} and $\Pi \mathbf{Y}$, yet overlooks the consistency of local

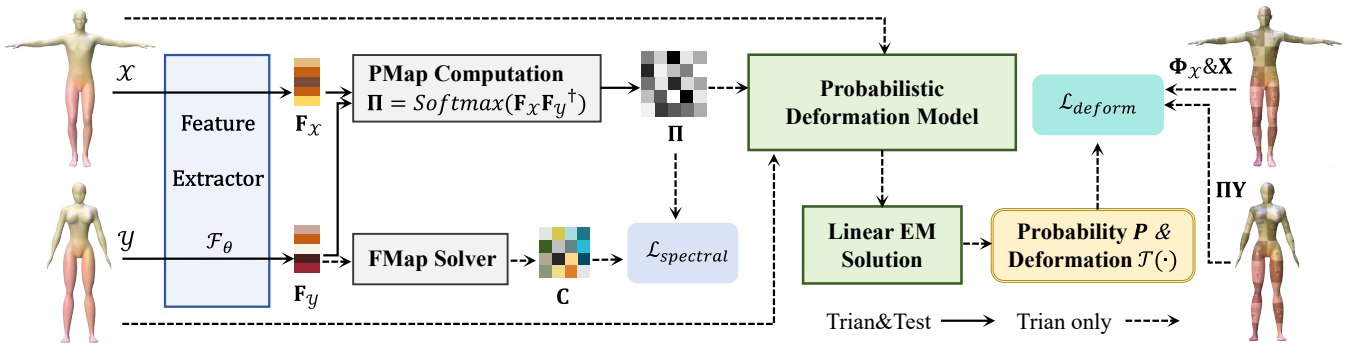


Figure 3: Overview of our unsupervised shape matching method based on probabilistic deformation consistency. Specifically, a shared-weight DiffusionNet serves as a Feature Extractor to obtain features \mathbf{F}_X and \mathbf{F}_Y from shapes \mathcal{X} and \mathcal{Y} , respectively. These features are then employed to predict the pointwise maps Π via PMap Computation and the deformation parameter \mathbf{C} via FMap Solver, with both predictions regularized by the spectral loss $\mathcal{L}_{spectral}$. Subsequently, the predicted Π , combined with the eigensystems and point coordinates of shapes, is input to Probabilistic Deformation Model. Linear EM Solution efficiently computes the probability \mathbf{P} and deformation $\mathcal{T}(\cdot)$. Finally, our designed \mathcal{L}_{deform} guides the optimization of the predicted pointwise maps by enforcing deformation consistency between the mapped shape $\Pi\mathcal{Y}$ and source \mathcal{X} .

deformations among neighboring points during the transformation process. To address this limitation, we introduce a regularization term for deformation consistency on the deformation \mathcal{T} , leading to the final objective as follows:

$$\mathcal{J}'(\mathcal{T}, \sigma^2) = \mathcal{J}(\mathcal{T}, \sigma^2) + \frac{\mu}{2}\phi(\mathcal{T}). \quad (8)$$

Deformation Consistency Regularization

In the Euclidean domain, the mapping between two point sets $\{\mathbf{x}_i, \mathbf{y}_i\}_{i=1}^n$ can be approximated by a smooth function deformation function f , *i.e.*, ideally $\mathbf{y}_i = f(\mathbf{x}_i)$. Within the context of regularization networks (Evgeniou, Pontil, and Poggio 2000), this problem can be formulated as the minimization of an empirical risk function $\mathcal{E}[f]$ with a regularization term:

$$\mathcal{E}[f] = \frac{1}{n} \sum_{i=1}^n c(f(\mathbf{x}_i), \mathbf{y}_i) + \frac{\mu}{2}\phi(f), \quad (9)$$

where $c(\cdot)$ denotes a cost function, μ is a balance hyperparameter, and $\phi(\cdot)$ represents a regularization function. Drawing from the kernel trick in machine learning (Girosi, Jones, and Poggio 1995), the solution of above problem is:

$$\begin{aligned} f(\mathbf{x}) &= \sum_{i=1}^n K(\mathbf{x}_i, \mathbf{x}) \mathbf{w}_i, \\ \phi(f) &= \mathbf{w}^\top \mathbf{K} \mathbf{w}, \end{aligned} \quad (10)$$

where $\mathbf{w}_i \in \mathbb{R}^D$ is the weight vector, $K(\cdot, \cdot)$ represents the kernel function, and the matrix forms of the weights and kernel functions are \mathbf{w} and \mathbf{K} . Based on (Smola, Schölkopf, and Müller 1998), Generalized CPD (Fan et al. 2022) generalizes the regularization framework from Euclidean domains to non-Euclidean domains. For source shape \mathcal{X} embedded in \mathbb{R}^3 , the Laplace-Beltrami operator yields eigenfunctions $\{v_i\}_{i=1}^k$ and eigenvalues $\{r_i\}_{i=1}^k$, where typically $k \ll n$. Leveraging the isometric property of the Laplace-Beltrami operator, the kernel function can be approximated using a truncated basis of low-frequency eigenfunctions. In

this case, the vectorized form \mathbf{f} of function f on sample points over \mathcal{X} is:

$$\mathbf{f} = \sum_i v_i w_i = \mathbf{U} \mathbf{w}, \quad (11)$$

where $\mathbf{U} = [v_1, v_2, \dots, v_k]$ is the matrix of k truncated eigenfunctions. The regularization term $\phi(f)$ in Eq.(10) becomes:

$$\phi(f) = \mathbf{w}^\top \mathbf{R} \mathbf{w}, \quad (12)$$

where \mathbf{R} is a diagonal matrix including eigenvalues $\{r_i\}_{i=1}^k$.

Linear Closed-Form Solution According to shape deformation as Eq. (4) of (Eisenberger, Lahner, and Cremers 2020), the deformed positions Γ can be estimated by \mathbf{f} , *i.e.*,

$$\Gamma = \mathbf{X} + \mathbf{f} = \mathbf{X} + \mathbf{U} \mathbf{w}, \quad (13)$$

where $\mathbf{U} \in \mathbb{R}^{n_x \times k}$ denotes k truncated LB eigenvectors, and $\mathbf{w} \in \mathbb{R}^{k \times 3}$ stacks the relevant weights. Applying Eq. (12) to Eq. (8), we obtain a closed-form solution for the unspecified weight matrix \mathbf{w} by:

$$(\mathbf{U}^\top \mathbf{P} \mathbf{U} + \mu \sigma^2 \mathbf{R}) \mathbf{w} = \mathbf{U}^\top \mathbf{P} (\Pi \mathcal{Y} - \mathbf{X}), \quad (14)$$

where \mathbf{P} is a diagonal matrix consisting of posterior probability, *i.e.*, $\mathbf{P} = \text{diag}([p_i])$. Based on linear algebra, other parameters σ^2 and γ can be derived by:

$$\sigma^2 = \frac{\text{tr}((\Pi \mathcal{Y} - \Gamma)^\top \mathbf{P} (\Pi \mathcal{Y} - \Gamma))}{D \cdot \text{tr}(\mathbf{P})}, \quad \gamma = 1 - \frac{\text{tr}(\mathbf{P})}{N}. \quad (15)$$

Hence, the EM algorithm has solved this probabilistic model. Notably, solving Eq. (14) is the most computationally demanding step. Since $k \ll n$, the overall complexity is linear, *i.e.*, $O(n)$, thus avoiding substantial computational cost.

Deformation Consistency Loss

The aforementioned probabilistic model based on shape deformation produces the weight matrix \mathbf{w} and the probability

Train	Geo. error (%)	FAUST		SCAPE		FAUST+SCAPE		
		FAUST	FAUST _a	SCAPE	SCAPE _a	FAUST	SCAPE	SHREC'19
<i>Axiomatic</i>	BCICP (Ren et al. 2018)	6.1	-	11.0	-	6.1	11.0	-
	ZoomOut (Melzi et al. 2019b)	6.1	8.7	7.5	15.0	6.1	7.5	-
	Smooth Shells (Eisenberger, Lahner, and Cremers 2020)	2.5	5.4	4.7	5.0	2.5	4.7	-
	DiscreteOp (Ren et al. 2021)	5.6	6.2	13.1	14.6	5.6	13.1	-
<i>Sup.</i>	FMNet (Litany et al. 2017a)	11.0	43.0	17.0	41.0	-	-	-
	GeomFMaps (Donati, Sharma, and Ovsjanikov 2020)	2.6	3.2	3.0	3.1	2.6	2.9	7.9
<i>Unsupervised</i>	UnsupFMNet (Halimi et al. 2019)	3.8	42.0	16.0	44.0	11.0	13.0	-
	Deep Shells (Eisenberger et al. 2020)	1.7	12.0	2.5	10.0	1.6	2.4	21.1
	NeuroMorph (Eisenberger et al. 2021)	8.5	-	29.9	-	9.1	27.3	25.3
	ConsistFMaps (Cao and Bernard 2022)	1.5	15.3	2.0	4.9	1.7	3.2	17.8
	DUO-FMNet (Donati, Corman, and Ovsjanikov 2022)	2.5	3.2	2.6	2.7	2.5	4.3	6.4
	AttentiveFMaps (Li, Donati, and Ovsjanikov 2022)	1.9	2.4	2.2	2.3	1.9	2.3	5.8
	ULRSSM (Cao, Roetzer, and Bernard 2023)	1.6	1.9	1.9	1.9	1.6	2.1	4.6
	MSSFMapNet (Magnet and Ovsjanikov 2024)	1.9	-	2.4	-	1.9	2.3	3.6
PDCMatch(Ours)	1.5	1.7	1.9	1.9	1.5	2.0	3.2	

Table 1: Evaluation of near-isometric and cross-dataset testing on FAUST, SCAPE, and SHREC'19 datasets. With column-wise best geometric errors (%) boldfaced, our method surpasses state-of-the-art axiomatic, supervised (sup.), and unsupervised baselines under diverse experimental settings.

matrix \mathbf{P} . The weight matrix \mathbf{w} characterizes the deformation function \mathcal{T} between the source set \mathbf{X} and mapped target set \mathbf{IY} , while \mathbf{P} encodes the likelihood probability of each point conforming to this deformation function. To ensure the deformation consistency of \mathbf{P} , we define a deformation discrepancy as \mathcal{D} , which is formulated as:

$$\mathcal{D} = \sum_{i=1}^{n_{\mathcal{X}}} p_i \|y_{i'} - \mathbf{x}_i - \Phi_{\mathcal{X}}^i \mathbf{w}\|_2, \quad (16)$$

where $\Phi_{\mathcal{X}}^i$ denotes the k -dimensional row vector of $\Phi_{\mathcal{X}}$ corresponding to point \mathbf{x}_i . Under ideal conditions, the point-wise map \mathbf{P} perfectly aligns each point pair with the estimated deformation function, causing the probability matrix \mathbf{P} to approximate an identity matrix. In other words, a higher $\text{tr}(\mathbf{P})$ thus reflects greater accuracy. Accordingly, we define an accuracy weight \mathcal{W} , which, together with \mathcal{D} , forms the total deformation consistency loss:

$$\begin{aligned} \mathcal{L}_{deform} &= \mathcal{D} \cdot \mathcal{W} = \mathcal{D} \cdot \frac{n_{\mathcal{X}} - \text{tr}(\mathbf{P})}{n_{\mathcal{X}} \cdot \text{tr}(\mathbf{P})} \\ &= \sum_{i=1}^{n_{\mathcal{X}}} p_i \|y_{i'} - \mathbf{x}_i - \Phi_{\mathcal{X}}^i \mathbf{w}\|_2 \cdot \frac{n_{\mathcal{X}} - \text{tr}(\mathbf{P})}{n_{\mathcal{X}} \cdot \text{tr}(\mathbf{P})}. \end{aligned} \quad (17)$$

The effectiveness of the components \mathcal{W} and \mathcal{D} is demonstrated in the ablation study (Section). The overall unsupervised training loss is formulated as a weighted sum of the above losses, *i.e.*,

$$\mathcal{L} = \mathcal{L}_{spectral} + \lambda_{deform} \mathcal{L}_{deform}, \quad (18)$$

where λ_{deform} is a weighting factor.

Experiments

In this section, we compare our method with other relevant methods across various challenging scenarios, analyzing its advantages and effectiveness. For a comprehensive comparison, the competitors in this section are divided into three categories: *Axiomatic methods*, *Supervised methods*, and *Unsupervised methods*. To ensure fairness, GeofMaps employs DiffusionNet as its feature extractor. Axiomatic methods use parameter initialization as specified by the original authors.

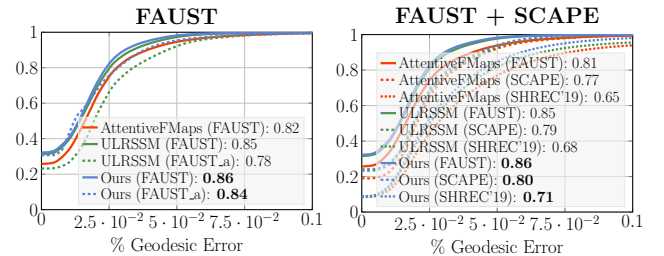


Figure 4: Proportion of correct keypoints (PCK) curves and corresponding area under the curve (AUC) on representative benchmarks. The title of each figure indicates the training dataset, with parenthesized legend entries specifying test datasets. Bold values indicate highest AUC per experimental configuration, with our method dominating these benchmarks.

Near-isometric Shape Matching

Near-isometric shape matching refers to the preservation of geodesic distances between two shapes to be matched, such as different poses of the same object.

Datasets We evaluate on three widely used datasets: **FAUST** (Bogo et al. 2014), **SCAPE** (Angelov et al. 2005), and **SHREC'19** (Melzi et al. 2019a), using their remeshed versions (Donati, Sharma, and Ovsjanikov 2020) for increased challenge. **FAUST** contains 100 high-resolution human scans from 10 subjects in 10 poses. Following standard protocols, we use 80 shapes for training and 20 for testing. **SCAPE** focuses on pose variation of a single subject, with 71 meshes (51 for training, 20 for testing), used to validate pose-invariance. **SHREC'19** provides 44 highly non-rigid human shapes with complex deformations. It includes 430 shape pairs exclusively for testing to evaluate robustness against severe shape distortions.

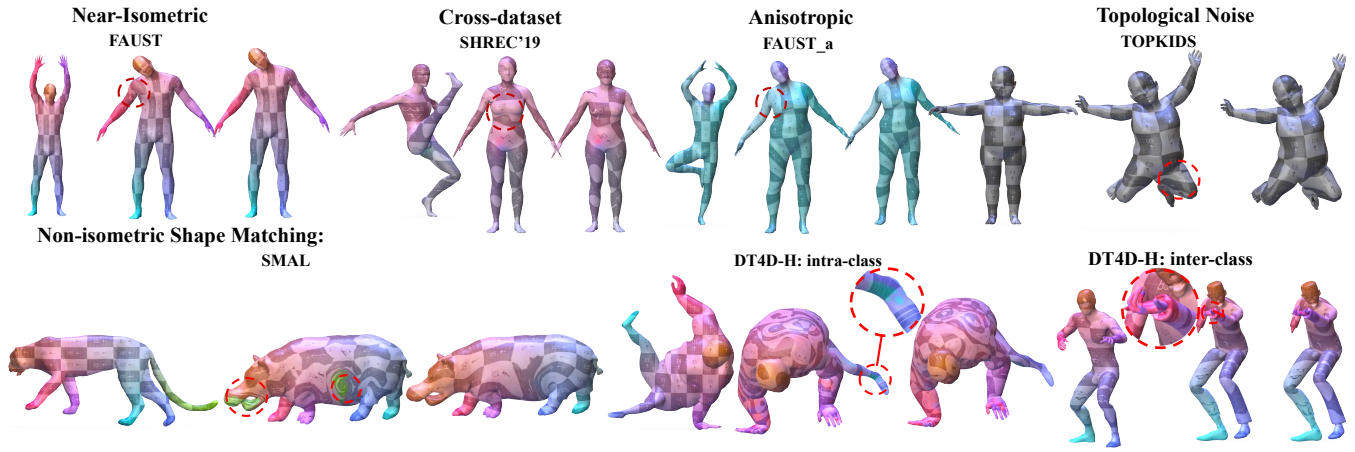


Figure 5: Visualization of texture transfer from source shape to target shape, and each group (from left to right) presents the source shape, the mapping result of ULRSSM (Cao, Roetzer, and Bernard 2023), and that of our PDCMatch. Specifically, the seven groups include near-isometric matching, cross-dataset testing, anisotropic meshing, topological noise, and non-isometric shape matching. It can be observed that our method exhibits strong capability of preserving texture details.

Results Using the mean geodesic error (Kim, Lipman, and Funkhouser 2011) as the evaluation metric, quantitative results are summarized in Table 1, which demonstrates the superior performance of our method over previous axiomatic, supervised, and unsupervised methods in most experimental configurations. Additionally, significant cross-dataset generalization gains of our method are validated in the rightmost column of Table 1. This can be attributed to the strong fitting capacity of the proposed probabilistic deformation model to capture shape transformations, which endows our method with strong generalizability.

The two subplots of Fig. 4 visualize the PCK testing curves for training on **FAUST** and **FAUST+SCAPE**, respectively. Comparative analysis against learning-based state-of-the-arts (Li, Donati, and Ovsjanikov 2022; Cao, Roetzer, and Bernard 2023) reveals that our method achieves the highest AUC. Visual examples are illustrated in the first two groups of Fig. 5, our method preserves more complete texture details compared to ULRSSM.

Matching with Anisotropic Meshing

Datasets To evaluate robustness against mesh discretization, we follow the protocol of (Donati, Corman, and Ovsjanikov 2022) and test on anisotropically remeshed versions of **FAUST** and **SCAPE**, denoted as **FAUST_a** and **SCAPE_a**. These variants feature directionally biased triangle scaling, which challenges methods sensitive to mesh regularity (e.g., fixed edge connectivity).

Results Evaluation results on anisotropically remeshed shapes are reported in the **FAUST_a** and **SCAPE_a** columns of Table 1. Visualizations are shown in the third group of Fig. 5. Compared to ULRSSM, our method better preserves texture details and achieves lower mean geodesic errors. While anisotropic remeshing challenges methods that overfit to mesh connectivity, our method remains robust to triangle meshing and outperforms prior state-of-the-art methods.

Geo. error (%)	TOPKIDS	SMAL	DT4D-H	
			intra-class	inter-class
<i>Axiomatic Methods</i>				
ZoomOut	33.7	38.4	4.0	29.0
Smooth Shells	11.8	36.1	3.1	6.3
DiscreteOp	35.5	38.1	3.6	27.6
<i>Supervised Methods</i>				
FMNet	-	42.0	9.6	38.0
GeomFMaps	-	8.4	2.1	4.1
<i>Unsupervised Methods</i>				
UnsupFMNet	38.5	-	-	-
WSupFMNet	47.6	7.6	3.3	22.6
Deep Shells	13.7	29.3	3.4	31.1
ConsistFMaps	39.3	-	-	-
DUO-FMNet	-	6.7	2.6	15.8
AttenFMaps	23.4	5.4	1.7	11.6
AttenFMaps-Fast	28.5	5.8	1.8	14.6
ULRSSM	9.2	3.9	2.3	4.1
SpiderMatch	-	4.4	2.7	4.1
PDCMatch(Ours)	6.6	4.0	1.5	4.0

Table 2: Evaluation of noisy and non-isometric shape matching on TOPKIDS, SMAL, and DT4D-H benchmarks. With column-wise best results boldfaced, our method shows comparable performance to state-of-the-art.

Matching with Topological Noise

Datasets The **TOPKIDS** (Löhner et al. 2016) dataset serves as a test benchmark to assess robustness against topological noise caused by interactions between disconnected parts in 3D scans. It contains synthetic child-like shapes with prominent topological distortions that challenge geometric isometry. Due to its limited size (26 shapes), we compare only with axiomatic and unsupervised methods.

Results Qualitative and quantitative results are exhibited in Fig. 5 and Table 2, where our method yields signif-




Figure 6: Visualization of the ablation study on SHREC’19. The full version of our method not only reserves fine-grained texture details (e.g., fingers), but also avoids bidirectional ambiguities (e.g., left and right limbs) by leveraging a complete deformation loss.

icantly more accurate results compared to the baseline ULRSSM. This benefit stems from our probabilistic deformation model, which emphasizes global deformation consistency while remaining robust to local topological noise, thus facilitating reliable modeling of non-rigid shapes.

Non-isometric Shape Matching

Datasets For non-isometric shape matching, we evaluate on two benchmarks. The **SMAL** (Zuffi et al. 2017) includes 49 quadruped animal shapes from 8 species. Following (Donati, Corman, and Ovsjanikov 2022), we use 29 shapes (5 species) for training and 20 shapes (3 unseen species) for testing, ensuring no class overlap. This inter-species setup introduces significant non-isometric deformations, posing a major challenge for existing methods. The **DT4D-H** (Magnet et al. 2022), derived from the large-scale DeformingThings4D (Li et al. 2021), contains 9 humanoid shape classes. Following their protocol, we use 198 training and 95 testing shapes to evaluate generalization. Due to SpiderMatch’s (Roetzer and Bernard 2024) high computational cost, we follow its original setup and use downsampled meshes with 1,000 triangles.

Results Table 2 reports the results on **SMAL** and **DT4D-H**. On **SMAL**, our method does not exhibit an improvement over the baseline ULRSSM. This is primarily because the proposed probabilistic deformation model is designed based on the assumption of near-isometric transformations, making it less effective in handling highly non-isometric scenarios such as different species. Nevertheless, as illustrated in Fig. 5, PDCMatch can yield more accurate matching performance in some **SMAL** instances. On **DT4D-H**, including intra-class and inter-class settings, our method achieves the lowest average geodesic error among comparative methods. Qualitative results in Fig. 5 further highlight that our method better preserves fine-grained structures (red circles).

Ablation Study

To validate the components of our proposed probabilistic deformation, we conduct ablation experiments on the challenging cross dataset benchmark, *i.e.*, train on **FAUST** and **SCAPE** while testing on **SHREC’19**. As shown in Fig. 6, *w.o. \mathcal{L}_{deform}* indicates training without the deformation loss, *w.o. \mathcal{W}* means fixing the accuracy weight \mathcal{W} to 1, *w.o. \mathcal{D}* means fixing the deformation discrepancy \mathcal{D} to 1, and *w.o.*

Ablation Setting	SHREC’19
<i>w.o. \mathcal{L}_{deform}</i>	4.63
<i>w.o. accuracy weight \mathcal{W}</i>	training collapsed
<i>w.o. deformation discrepancy \mathcal{D}</i>	3.83
<i>w.o. EM iterations</i>	4.51
Full	3.20

Table 3: Ablation study of our method testing on SHREC’19, evaluated using average geometric errors (%).

Dataset	Faust	SCAPE	TOPKIDS	SMAL	DT4D-H
#Vertices	5000	5000	12000	5213	8000
ULRSSM	0.58	0.55	0.63	0.57	0.64
PDCMatch	0.63	0.61	0.86	0.60	0.71

Table 4: Runtime comparisons on different benchmarks, reporting training propagation time (s) per shape pair.

EM iterations denotes a single-pass version. The ‘Full’ case represents the complete version. Results in Fig. 6 and Table 3 show the positive contribution of each component.

Running Time Analysis

Our method evaluates pointwise maps via a probabilistic deformation model and introduces a deformation consistency loss to guide training, without impacting on inference. For runtime analysis, we measure the training time of both ULRSSM and our method on an RTX3090 GPU and report the average propagation time per shape pair in Table 4. The additional cost, mainly from EM optimization, grows with vertex count but remains moderate and acceptable overall.

Limitations and Conclusion

Our method constructs a probabilistic deformation model based on the near-isometric assumption to estimate consistent deformations, thereby evaluating pointwise correspondences and guiding unsupervised learning of shape features.

However, when faced with pronounced non-isometric deformations, such as those arising from inter-species variations in the **SMAL** dataset, the effectiveness of the deformation regularization diminishes, and the proposed deformation consistency loss may lose its guiding effect. The effectiveness of the deformation consistency loss depends on the reliability of EM optimization, which may be compromised by local optima, leading to suboptimal performance.

In summary, the proposed probabilistic deformation model and designed unsupervised loss provide a deeper understanding of shape matching, with experimental evidence across multiple benchmarks showing improved texture preservation and matching accuracy, especially on cross-dataset generalization and topological noise.

Acknowledgments

This work was supported by the National Natural Science Foundation of China (624B2107, 62276192).

References

- Anguelov, D.; Srinivasan, P.; Koller, D.; Thrun, S.; Rodgers, J.; and Davis, J. 2005. Scape: shape completion and animation of people. In *ACM SIGGRAPH*, 408–416.
- Attaiki, S.; and Ovsjanikov, M. 2022. NCP: Neural correspondence prior for effective unsupervised shape matching. *Adv. Neural Inform. Process. Syst.*, 35: 28842–28857.
- Attaiki, S.; and Ovsjanikov, M. 2023a. Shape non-rigid kinematics (snk): A zero-shot method for non-rigid shape matching via unsupervised functional map regularized reconstruction. *Adv. Neural Inform. Process. Syst.*, 36: 70012–70032.
- Attaiki, S.; and Ovsjanikov, M. 2023b. Understanding and improving features learned in deep functional maps. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 1316–1326.
- Attaiki, S.; Pai, G.; and Ovsjanikov, M. 2021. Dpfm: Deep partial functional maps. In *Int. Conf. 3D Vis.*, 175–185.
- Bastian, L.; Xie, Y.; Navab, N.; and Löhner, Z. 2024. Hybrid functional maps for crease-aware non-isometric shape matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 3313–3323.
- Bernard, F.; Suri, Z. K.; and Theobalt, C. 2020. Mina: Convex mixed-integer programming for non-rigid shape alignment. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 13826–13835.
- Bogo, F.; Romero, J.; Loper, M.; and Black, M. J. 2014. FAUST: Dataset and evaluation for 3D mesh registration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 3794–3801.
- Bogo, F.; Romero, J.; Pons-Moll, G.; and Black, M. J. 2017. Dynamic FAUST: Registering human bodies in motion. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 6233–6242.
- Bronstein, M. M.; and Kokkinos, I. 2010. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 1704–1711.
- Cao, D.; and Bernard, F. 2022. Unsupervised deep multi-shape matching. In *Eur. Conf. Comput. Vis.*, 55–71.
- Cao, D.; and Bernard, F. 2023. Self-supervised learning for multimodal non-rigid 3d shape matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 17735–17744.
- Cao, D.; Roetzer, P.; and Bernard, F. 2023. Unsupervised Learning of Robust Spectral Shape Matching. *ACM Trans. Graph.*, 42(4): 1–15.
- Cao, D.; Roetzer, P.; and Bernard, F. 2024. Revisiting map relations for unsupervised non-rigid shape matching. In *Int. Conf. 3D Vis.*, 1371–1381.
- Deng, B.; Yao, Y.; Dyke, R. M.; and Zhang, J. 2022. A survey of non-rigid 3D registration. In *Comput. Graph. Forum*, volume 41, 559–589.
- Donati, N.; Corman, E.; and Ovsjanikov, M. 2022. Deep orientation-aware functional maps: Tackling symmetry issues in shape matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 742–751.
- Donati, N.; Sharma, A.; and Ovsjanikov, M. 2020. Deep geometric functional maps: Robust feature learning for shape correspondence. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 8592–8601.
- Egger, B.; Smith, W. A.; Tewari, A.; Wuhler, S.; Zollhoefer, M.; Beeler, T.; Bernard, F.; Bolkart, T.; Kortylewski, A.; Romdhani, S.; et al. 2020. 3d morphable face models—past, present, and future. *ACM Trans. Graph.*, 39(5): 1–38.
- Eisenberger, M.; Löhner, Z.; and Cremers, D. 2019. Divergence-free shape correspondence by deformation. In *Comput. Graph. Forum*, volume 38, 1–12.
- Eisenberger, M.; Lahner, Z.; and Cremers, D. 2020. Smooth shells: Multi-scale shape registration with functional maps. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 12265–12274.
- Eisenberger, M.; Novotny, D.; Kerchenbaum, G.; Labatut, P.; Neverova, N.; Cremers, D.; and Vedaldi, A. 2021. Neuronormorph: Unsupervised shape interpolation and correspondence in one go. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 7473–7483.
- Eisenberger, M.; Toker, A.; Leal-Taixé, L.; and Cremers, D. 2020. Deep shells: Unsupervised shape correspondence with optimal transport. *Adv. Neural Inform. Process. Syst.*, 33: 10491–10502.
- Evgeniou, T.; Pontil, M.; and Poggio, T. 2000. Regularization networks and support vector machines. *Adv. Comput. Math.*, 13: 1–50.
- Eynard, D.; Rodola, E.; Glashoff, K.; and Bronstein, M. M. 2016. Coupled functional maps. In *Int. Conf. 3D Vis.*, 399–407.
- Ezuz, D.; Solomon, J.; and Ben-Chen, M. 2019. Reversible harmonic maps between discrete surfaces. *ACM Trans. Graph.*, 38(2): 1–12.
- Fan, A.; Ma, J.; Tian, X.; Mei, X.; and Liu, W. 2022. Coherent point drift revisited for non-rigid shape matching and registration. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 1424–1434.
- Gao, M.; Lahner, Z.; Thunberg, J.; Cremers, D.; and Bernard, F. 2021. Isometric multi-shape matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 14183–14193.
- Girosi, F.; Jones, M.; and Poggio, T. 1995. Regularization theory and neural networks architectures. *Neural Comput.*, 7(2): 219–269.
- Halimi, O.; Litany, O.; Rodola, E.; Bronstein, A. M.; and Kimmel, R. 2019. Unsupervised learning of dense shape correspondence. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4370–4379.
- Holzschuh, B.; Löhner, Z.; and Cremers, D. 2020. Simulated annealing for 3d shape correspondence. In *Int. Conf. 3D Vis.*, 252–260.
- Huang, Q.-X.; Adams, B.; Wicke, M.; and Guibas, L. J. 2008. Non-rigid registration under isometric deformations. In *Comput. Graph. Forum*, volume 27, 1449–1457.
- Jiang, P.; Sun, M.; and Huang, R. 2023. Non-rigid shape registration via deep functional maps prior. *Adv. Neural Inform. Process. Syst.*, 36: 58409–58427.
- Kim, V. G.; Lipman, Y.; and Funkhouser, T. 2011. Blended intrinsic maps. *ACM Trans. Graph.*, 30(4): 1–12.
- Löhner, Z.; Rodola, E.; Bronstein, M. M.; Cremers, D.; Burghard, O.; Cosmo, L.; Dieckmann, A.; Klein, R.;

- Sahillioglu, Y.; et al. 2016. SHREC'16: Matching of deformable shapes with topological noise. In *Eurogr. Workshop 3D Obj. Retr.*, 55–60.
- Li, L.; Donati, N.; and Ovsjanikov, M. 2022. Learning multi-resolution functional maps with spectral attention for robust shape matching. *Adv. Neural Inform. Process. Syst.*, 35: 29336–29349.
- Li, Y.; Takehara, H.; Taketomi, T.; Zheng, B.; and Nießner, M. 2021. 4dcomplete: Non-rigid motion estimation beyond the observable surface. In *Int. Conf. Comput. Vis.*, 12706–12716.
- Litany, O.; Remez, T.; Rodola, E.; Bronstein, A.; and Bronstein, M. 2017a. Deep functional maps: Structured prediction for dense shape correspondence. In *Int. Conf. Comput. Vis.*, 5659–5667.
- Litany, O.; Rodolà, E.; Bronstein, A. M.; and Bronstein, M. M. 2017b. Fully spectral partial shape matching. In *Comput. Graph. Forum*, volume 36, 247–258.
- Magnet, R.; and Ovsjanikov, M. 2024. Memory-scalable and simplified functional map learning. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 4041–4050.
- Magnet, R.; Ren, J.; Sorkine-Hornung, O.; and Ovsjanikov, M. 2022. Smooth non-rigid shape matching via effective dirichlet energy optimization. In *Int. Conf. 3D Vis.*, 495–504.
- Melzi, S.; Marin, R.; Rodolà, E.; Castellani, U.; Ren, J.; Poulencard, A.; Ovsjanikov, P.; et al. 2019a. SHREC'19: matching humans with different connectivity. In *Eurographics Worksh. 3D Obj. Retrieval*, 1–8.
- Melzi, S.; Ren, J.; Rodola, E.; Sharma, A.; Wonka, P.; and Ovsjanikov, M. 2019b. Zoomout: Spectral upsampling for efficient shape correspondence. *arXiv*.
- Munkres, J. 1957. Algorithms for the assignment and transportation problems. *J. Soc. Ind. Appl. Math.*, 5(1): 32–38.
- Myronenko, A.; and Song, X. 2010. Point set registration: Coherent point drift. *IEEE Trans. Pattern Anal. Mach. Intell.*, 32(12): 2262–2275.
- Ovsjanikov, M.; Ben-Chen, M.; Solomon, J.; Butscher, A.; and Guibas, L. 2012. Functional maps: a flexible representation of maps between shapes. *ACM Trans. Graph.*, 31(4): 1–11.
- Qi, C. R.; Yi, L.; Su, H.; and Guibas, L. J. 2017. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *Adv. Neural Inform. Process. Syst.*, 30.
- Ren, J.; Melzi, S.; Wonka, P.; and Ovsjanikov, M. 2021. Discrete optimization for shape matching. In *Comput. Graph. Forum*, volume 40, 81–96.
- Ren, J.; Panine, M.; Wonka, P.; and Ovsjanikov, M. 2019. Structured regularization of functional map computations. In *Comput. Graph. Forum*, volume 38, 39–53.
- Ren, J.; Poulencard, A.; Wonka, P.; and Ovsjanikov, M. 2018. Continuous and orientation-preserving correspondences via functional maps. *ACM Trans. Graph.*, 37(6): 1–16.
- Rodolà, E.; Cosmo, L.; Bronstein, M. M.; Torsello, A.; and Cremers, D. 2017. Partial functional correspondence. In *Comput. Graph. Forum*, volume 36, 222–236.
- Roetzer, P.; Abbas, A.; Cao, D.; Bernard, F.; and Swoboda, P. 2024. DiscoMatch: Fast Discrete Optimisation for Geometrically Consistent 3D Shape Matching. In *Eur. Conf. Comput. Vis.*, 443–460.
- Roetzer, P.; and Bernard, F. 2024. Spidermatch: 3d shape matching with global optimality and geometric consistency. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 14543–14553.
- Roetzer, P.; Swoboda, P.; Cremers, D.; and Bernard, F. 2022. A scalable combinatorial solver for elastic geometrically consistent 3d shape matching. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 428–438.
- Roufousse, J.-M.; Sharma, A.; and Ovsjanikov, M. 2019. Un-supervised deep learning for structured shape matching. In *Int. Conf. Comput. Vis.*, 1617–1627.
- Sahillioglu, Y. 2020. Recent advances in shape correspondence. *The Vis. Comput.*, 36(8): 1705–1721.
- Salti, S.; Tombari, F.; and Di Stefano, L. 2014. SHOT: Unique signatures of histograms for surface and texture description. *Comput. Vis. Image Underst.*, 125: 251–264.
- Sharma, A.; and Ovsjanikov, M. 2020. Weakly supervised deep functional maps for shape matching. *Adv. Neural Inform. Process. Syst.*, 33: 19264–19275.
- Sharp, N.; Attaiki, S.; Crane, K.; and Ovsjanikov, M. 2022. Diffusionnet: Discretization agnostic learning on surfaces. *ACM Trans. Graph.*, 41(3): 1–16.
- Smola, A. J.; Schölkopf, B.; and Müller, K.-R. 1998. The connection between regularization operators and support vector kernels. *Neural Netw.*, 11(4): 637–649.
- Sun, M.; Mao, S.; Jiang, P.; Ovsjanikov, M.; and Huang, R. 2023. Spatially and spectrally consistent deep functional maps. In *Int. Conf. Comput. Vis.*, 14497–14507.
- Thomas, H.; Qi, C. R.; Deschaud, J.-E.; Marcotegui, B.; Goulette, F.; and Guibas, L. J. 2019. Kpconv: Flexible and deformable convolution for point clouds. In *Int. Conf. Comput. Vis.*, 6411–6420.
- 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 *IEEE Conf. Comput. Vis. Pattern Recog.*, 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.
- Zuffi, S.; Kanazawa, A.; Jacobs, D. W.; and Black, M. J. 2017. 3D menagerie: Modeling the 3D shape and pose of animals. In *IEEE Conf. Comput. Vis. Pattern Recog.*, 6365–6373.