Revisiting Map Relations for Unsupervised Non-Rigid Shape Matching



Figure 1. **Qualitative 3D shape matching results of our method.** The leftmost reference shape is matched to the other shapes in the first row. The second row visualises the corresponding functional maps. We observe that the diagonal structure of the functional maps changes significantly, especially under non-isometry or partiality. To better account for the varying structure of the functional map, we propose a self-adaptive functional map solver that can adjust the functional map regularisation strength and structure depending on the input shapes.

Abstract

We propose a novel unsupervised learning approach for non-rigid 3D shape matching. Our approach improves upon recent state-of-the art deep functional map methods and can be applied to a broad range of different challenging scenarios. Previous deep functional map methods mainly focus on feature extraction and aim exclusively at obtaining more expressive features for functional map computation. However, the importance of the functional map computation itself is often neglected and the relationship between the functional map and point-wise map is underexplored. In this paper, we systematically investigate the coupling relationship between the functional map from the functional map solver and the point-wise map based on feature similarity. To this end, we propose a self-adaptive functional map solver to adjust the functional map regularisation for different shape matching scenarios, together with a vertex-wise contrastive loss to obtain more discriminative features. Using different challenging datasets (including non-isometry, topological noise and partiality), we demonstrate that our method substantially outperforms previous state-of-the-art methods.

1. Introduction

3D shape matching is a fundamental problem in shape analysis, computer vision and computer graphics with a broad range of applications, including texture transfer [15], deformation transfer [63] and statistical shape analysis [20, 39, 43]. Even though 3D shape matching is a long-standing problem and has been studied for decades [64, 66], finding correspondences between two non-rigidly deformed 3D shapes is still a challenging problem, especially for shapes with large non-isometric deformation, topological noise, or partiality.

Notably, in the case of 3D shapes represented by triangle meshes, the functional map framework [49] is one of the most dominant pipelines in this area and has been extended by many follow-up works due to its efficiency and well-justified theoretical properties [17, 48, 52, 56].

Meanwhile, with the recent rapid development in deep learning, many learning-based methods for non-rigid 3D shape matching are also based on the functional map framework, including both supervised [4, 16, 41] and unsupervised [10, 12, 18, 23, 30, 38, 58, 61] approaches. Most of them mainly focus on training the feature extraction module to obtain functional maps based on the extracted features and then rely on off-the-shelf post-processing [46] to obtain final point-wise correspondences. In contrast, the recent work by Cao et al. [12] explicitly models the relationship between functional maps and pointwise maps and thus leads to more robust matching in a broad range of challenging scenarios. However, the method only focuses on extracting more expressive features and ignores the importance of the functional map computation itself. Further, it lacks a discussion about insights between the relationship between the functional map and point-wise map.

In this paper, we improve upon the recent work by Cao et al. [12] by proposing a novel functional map solver that is self-adaptive to different shape matching scenarios. Moreover, we systematically analyse the relationship between the functional map and the point-wise map and introduce a vertex-wise contrastive loss to obtain more discriminative features leading to more accurate correspondences. We summarise our main contributions as follows:

- For the first time we propose a functional map solver that is self-adaptive for different challenging matching scenarios.
- We introduce a vertex-wise contrastive loss to obtain more discriminative features that can be used directly for matching via nearest neighbour search.
- We set the new state-of-the-art performance on numerous challenging benchmarks in diverse settings, including non-isometric, topologically noisy and partial shape matching, even compared to recent supervised methods.

2. Related work

3D shape matching is a long-standing problem that has been studied for decades. In the following we focus on reviewing those methods that are most relevant to our work. A more comprehensive overview can be found in [59, 64, 66].

2.1. Axiomatic shape matching methods

Shape matching can be formulated as establishing pointwise correspondences between a given pair of shapes. A simple formulation for doing so is the linear assignment problem (LAP) [47]. However, the LAP cannot take geometric relations into account and thus leads to spatially nonsmooth matchings. To compensate for this, several shape matching approaches [31, 57, 68] establish correspondences by explicitly incorporating geometric constraints. Some methods [6, 21, 26, 33] attempt to solve the problem based on non-rigid shape registration. Overall, directly establishing point-wise correspondences often leads to complex optimisation problems that are difficult to solve.

In contrast, the functional map framework finds correspondences in the functional domain [49]. Here, the correspondence relationship can be encoded with a small matrix, namely the functional map. Due to its simple yet efficient formulation, the functional map framework has been extended by many follow-up works, e.g. in terms of improving the matching accuracy and robustness [25, 53], extending it to more challenging scenarios (e.g. non-isometry [22, 36, 44, 52, 55], partiality [42, 56]), considering multishape matching [13, 28, 32, 34], and matching with nonunique solutions [54]. Nevertheless, axiomatic functional map methods rely on handcrafted features (e.g. HKS [9], WKS [5], SHOT [60]), which limits their performance. In contrast, our method (among others) directly learns discriminative features from training data and achieves more accurate and robust matching performance on challenging settings.

2.2. Deep functional map methods

In contrast to axiomatic approaches, deep functional map methods aim to learn features directly from training data. The supervised FMNet [41] is the pioneer work that learns a non-linear transformation of SHOT feature [60] based on a point-wise MLP. Later works [30, 58] enable unsupervised training of FMNet by introducing isometry regularisation in the spatial and spectral domain, respectively. Instead of using simple point-wise MLPs, follow-up works [16, 61] replace FMNet by point-based networks [51, 65] and lead to better matching performance. More recently, Sharp et al. [62] introduces DiffusionNet with a learnable diffusion process and has set the new state-of-the-art matching performance for a broad range of shape matching scenarios, including near-isometry [3, 10], non-isometry [2, 18, 38], partiality [4, 12], as well as shapes represented as point clouds [11]. Despite the rapid progress of deep functional map methods, existing approaches mostly focus on learning more expressive features for functional map computation, while ignoring the importance of the functional map computation itself. In this work, we systematically investigate the functional map computation process and introduce a self-adaptive functional map solver to better regularise the functional map structure for different kinds of input shapes.

3. Background

In this section we explain the background and introduce the notation used throughout the rest of the paper in Tab. 1.

3.1. Functional map framework

We consider a pair of 3D shapes \mathcal{X} and \mathcal{Y} represented as triangle meshes, with $n_{\mathcal{X}}$ and $n_{\mathcal{Y}}$ vertices, respectively. Here we summarise the common pipeline of the functional map framework.

1. Compute the associated positive semi-definite Laplacian matrices $L_{\mathcal{X}}, L_{\mathcal{Y}}$ [50]. The Laplacian matrix can be computed as $L_{\mathcal{X}} = M_{\mathcal{X}}^{-1}W_{\mathcal{X}}$, where $M_{\mathcal{X}}$ is the diagonal lumped mass matrix and $W_{\mathcal{X}}$ is the cotangent weight matrix.

Symbol	Description
\mathcal{X}, \mathcal{Y}	3D shapes with n_X , n_Y vertices
$L_{\mathcal{X}}$	$\mathbb{R}^{n_{\mathcal{X}} \times n_{\mathcal{X}}}$ Laplacian matrix of shape \mathcal{X}
$\Lambda_{\mathcal{X}}$	$\mathbb{R}^{k \times k}$ eigenvalue matrix of Laplacian $L_{\mathcal{X}}$
$\Phi_{\mathcal{X}}$	$\mathbb{R}^{n_{\mathcal{X}} \times k}$ LBO eigenfunctions of shape \mathcal{X}
$\Phi^{\dagger}_{\mathcal{X}}$	$\mathbb{R}^{k \times n_{\mathcal{X}}}$ Moore-Penrose inverse of $\Phi_{\mathcal{X}}$
$F_{\mathcal{X}}$	$\mathbb{R}^{n_{\mathcal{X}} \times c}$ vertex-wise features of shape \mathcal{X}
$A_{\mathcal{X}}$	$\Phi_{\mathcal{X}}^{\dagger} F_{\mathcal{X}}$ projected feature coefficients of shape \mathcal{X}
C_{XY}	$\mathbb{R}^{k \times k}$ functional map between shapes \mathcal{X} and \mathcal{Y}
$\Pi_{\mathcal{YX}}$	point-wise map between shapes $\mathcal Y$ and $\mathcal X$

Table 1. Summary of the notation used in this paper.

- 2. Compute the first k eigenfunctions $\Phi_{\mathcal{X}}, \Phi_{\mathcal{Y}}$ and the corresponding eigenvalues $\Lambda_{\mathcal{X}}, \Lambda_{\mathcal{Y}}$ of the respective Laplacian matrices (i.e. LBO eigenfunctions/eigenvalues).
- 3. Compute *c*-dimensional features $F_{\mathcal{X}}, F_{\mathcal{Y}}$ defined on each shape either from handcrafted feature descriptors or from a learnable feature extractor.
- 4. Compute the functional map C_{XY} associated with the LBO eigenfunctions by solving (variants of) the least squares problem

$$C_{\mathcal{X}\mathcal{Y}} = \operatorname{argmin}_{C} E_{\text{data}}(C) + \lambda E_{\text{reg}}(C). \quad (1)$$

Here, minimising $E_{\text{data}} = ||CA_{\mathcal{X}} - A_{\mathcal{Y}}||_{F}^{2}$ enforces descriptor preservation, while minimising the regularisation term E_{reg} imposes some form of structural properties (e.g. Laplacian commutativity [49]).

5. Recover the point-wise map $\Pi_{\mathcal{YX}}$ based on the relationship $C_{\mathcal{XY}} = \Phi_{\mathcal{Y}}^{\dagger} \Pi_{\mathcal{YX}} \Phi_{\mathcal{X}}$, e.g. either by nearest neighbour search in the spectral domain or by other postprocessing techniques [27, 46, 67].

We emphasise that most deep functional methods mainly focus on the third step that aims to extract more expressive features for functional map computation while ignoring the importance of the other steps (i.e. the functional map computation and point-wise map conversion). However, we argue that this may lead to sub-optimal performance, since the three interrelated aspects feature learning, functional map computation, and point-wise map conversion are considered in an isolated rather than a joint manner. Therefore, in this paper we systematically investigate the functional map computation step and the relationship between the functional map and the associated point-wise map.

3.2. Deep functional maps

Instead of relying on handcrafted features [5, 9, 60] to compute functional maps, many deep functional map methods [58, 61] have been proposed. The common pipeline of those methods is shown in Fig. 2 (left).

The common deep functional map framework mainly consists of two modules: a feature extractor and a functional map solver. The feature extractor is used to extract vertex-wise features and the functional map solver is used to compute functional maps based on the extracted features. To train the feature extractor, structural regularisation (e.g. orthogonality, bijectivity [58]) is imposed on the computed functional maps, i.e.

$$L_{\rm fmap} = \lambda_{\rm bij} L_{\rm bij} + \lambda_{\rm orth} L_{\rm orth}, \qquad (2)$$

where

$$L_{\text{bij}} = \|C_{\mathcal{X}\mathcal{Y}}C_{\mathcal{Y}\mathcal{X}} - I\|_F^2 + \|C_{\mathcal{Y}\mathcal{X}}C_{\mathcal{X}\mathcal{Y}} - I\|_F^2, \quad (3)$$

$$L_{\text{orth}} = \left\| C_{\mathcal{X}\mathcal{Y}}^{\top} C_{\mathcal{Y}\mathcal{X}} - I \right\|_{F}^{2} + \left\| C_{\mathcal{X}\mathcal{Y}}^{\top} C_{\mathcal{Y}\mathcal{X}} - I \right\|_{F}^{2}.$$
 (4)

After training, off-the-shelf post-processing techniques [46, 67] are used to convert functional maps to point-wise maps.

As pointed out by recent works [3, 12, 55], a major downside of this common pipeline is that the relation between the functional maps and associated point-wise maps is ignored, so that the performance is often sub-optimal, especially in the presence of large non-isometry, topological noise or partiality. To compensate for this, Cao et al. [12] proposed to directly obtain point-wise maps based on the extracted features and introduced a coupling loss L_{couple} to explicitly regularise the relation between the point-wise map $\Pi_{\mathcal{YX}}$ and the corresponding functional map $C_{\mathcal{XY}}$, i.e.

$$L_{\text{couple}} = \left\| C_{\mathcal{X}\mathcal{Y}} - C_{\mathcal{X}\mathcal{Y}}^{\Pi} \right\|_{F}^{2}, \qquad (5)$$

where $C_{\mathcal{X}\mathcal{Y}}^{\Pi} = \Phi_{\mathcal{Y}}^{\dagger} \Pi_{\mathcal{Y}\mathcal{X}} \Phi_{\mathcal{X}}$.

By explicitly modelling the relationship between functional maps and point-wise maps, the method proposed by Cao et al. [12] substantially outperforms existing methods and is robust in different challenging scenarios. However, there are two major limitations:

- In their approach, one functional map $C_{\mathcal{XY}}$ is computed from the functional map solver, while the other one $C_{\mathcal{XY}}^{\Pi}$ is converted from the point-wise map based on the feature similarity. Yet, the underlying relationship between them is not well-understood.
- The coupling loss L_{couple} regularises the functional maps computed from the functional map solver based on the extracted features, while the functional map computation itself is not optimised in a data-driven manner.

In the following, we theoretically analyse the relationship between $C_{\mathcal{X}\mathcal{Y}}$ and $C_{\mathcal{X}\mathcal{Y}}^{\Pi}$, and revisit the map relations by introducing a self-adaptive functional map solver and a vertex-wise contrastive loss.

4. Theoretical analysis of map relations

In this section, we analyse the underlying relation between the functional map computed from the functional map solver and the functional map converted by the pointwise map based on the deep feature similarity.



Figure 2. Left: Common pipeline of deep functional map methods. A Siamese feature extractor computes vertex-wise features for each shape. The extracted features are used for functional map computation. During training, structural regularisation $L_{\rm fmap}$ is imposed on the functional maps. During inference, the computed functional maps are typically converted to point-wise maps via map conversion. **Right: Our proposed shape matching pipeline.** Point-wise correspondences are obtained based on feature similarity. A coupling loss $L_{\rm couple}$ regularises the relation between the point-wise map Π_{VX} and the functional map C_{XY} [12]. To better balance the functional map

 L_{couple} regularises the relation between the point-wise map $\Pi_{\mathcal{YX}}$ and the functional map $C_{\mathcal{XY}}$ [12]. To better balance the functional map regularisation and the coupling relationship between the functional map and the point-wise map, we introduce a self-adaptive functional map solver (with learnable parameters λ and γ) to adjust the regularisation strength and structure, respectively. Additionally, a vertex-wise contrastive loss L_{contrast} is introduced to improve the discriminative power of the features.

W.l.o.g. we assume that $n_{\mathcal{Y}} \leq n_{\mathcal{X}}$. With that, a (partial or full) shape \mathcal{Y} is matched to a full shape \mathcal{X} and thereby the point-wise map $\Pi_{\mathcal{Y}\mathcal{X}}$ should be a (partial) permutation matrix, i.e.

$$\mathcal{P} := \left\{ \Pi \in \{0,1\}^{n_{\mathcal{Y}} \times n_{\mathcal{X}}} : \Pi \mathbf{1}_{n_{\mathcal{X}}} = \mathbf{1}_{n_{\mathcal{Y}}}, \mathbf{1}_{n_{\mathcal{Y}}}^{\top} \Pi \leq \mathbf{1}_{n_{\mathcal{X}}}^{\top} \right\},$$
(6)

where $\Pi_{\mathcal{YX}}(i, j)$ indicates whether the *i*-th point in \mathcal{Y} corresponds to the *j*-th point in \mathcal{X} .

Firstly, we analyse the point-wise map computed based on the feature similarity. We note that

$$\Pi_{\mathcal{Y}\mathcal{X}} = \operatorname{argmin}_{\Pi \in \mathcal{P}} \|\Pi F_{\mathcal{X}} - F_{\mathcal{Y}}\|_{F}^{2}.$$
 (7)

Lemma 4.1. If there exists a unique solution to Eq. (7), then the rows of F_X and F_y have non-repeated rows.

Proof. If $F_{\mathcal{X}}$ has repeated rows, we can find a (full) permutation matrix $\Pi_{\mathcal{X}\mathcal{X}} \neq I$ that satisfies $\Pi_{\mathcal{X}\mathcal{X}}F_{\mathcal{X}} = F_{\mathcal{X}}$. Therefore, any solution $\Pi_{\mathcal{Y}\mathcal{X}}$ has an equivalent solution $\Pi'_{\mathcal{Y}\mathcal{X}} := \Pi_{\mathcal{Y}\mathcal{X}}\Pi_{\mathcal{X}\mathcal{X}}$ and is thus not unique. Due to the orthogonal invariance of the Frobenius norm, an analogous statement can be made for $F_{\mathcal{Y}}$.

Discussion. To obtain a valid point-wise map based on the feature similarity, the features $F_{\mathcal{X}}$, $F_{\mathcal{Y}}$ should have non-repeated rows. To this end, based on Lemma 4.1 we propose a vertex-wise contrastive loss to encourage more discriminative features.

Theorem 4.2. Consider the following conditions:

(i) $\Pi_{\mathcal{YX}}F_{\mathcal{X}} = F_{\mathcal{Y}}, \Pi_{\mathcal{YX}} \in \mathcal{P}$, where $n_{\mathcal{Y}} \leq n_{\mathcal{X}}$.

(ii) $F_{\mathcal{X}}$ is in the span of $\Phi_{\mathcal{X}}$ and $F_{\mathcal{Y}}$ is in the span of $\Phi_{\mathcal{Y}}$. (iii) $\lambda = 0$ in Eq. (1) and $A_{\mathcal{X}} \in \mathbb{R}^{k \times c}$ ($k \leq c$) is full rank. If conditions (i)-(iii) hold, then we have $C_{\mathcal{X}\mathcal{Y}} = C_{\mathcal{X}\mathcal{Y}}^{\Pi}$, and $\|C_{\mathcal{X}\mathcal{Y}}A_{\mathcal{X}} - A_{\mathcal{Y}}\|_{F}^{2} = 0.$ *Proof.* By condition (i), we have $\Pi_{\mathcal{YX}}F_{\mathcal{X}} = F_{\mathcal{Y}}$ and from condition (ii) we know that $F_{\mathcal{X}} = \Phi_{\mathcal{X}}A_{\mathcal{X}}$ (since $A_{\mathcal{X}}$ is the matrix of projected feature coefficients). The same holds for \mathcal{Y} . Putting these together,

$$\Pi_{\mathcal{Y}\mathcal{X}}\Phi_{\mathcal{X}}A_{\mathcal{X}} = \Phi_{\mathcal{Y}}A_{\mathcal{Y}}.$$
(8)

Pre-multiplying Eq. (8) by $\Phi_{\mathcal{V}}^{\dagger}$ we obtain

$$\Phi_{\mathcal{Y}}^{\dagger}\Pi_{\mathcal{Y}\mathcal{X}}\Phi_{\mathcal{X}}A_{\mathcal{X}} = C_{\mathcal{X}\mathcal{Y}}^{\Pi}A_{\mathcal{X}} = A_{\mathcal{Y}}, \qquad (9)$$

where the definition $C_{\mathcal{X}\mathcal{Y}}^{\Pi} = \Phi_{\mathcal{Y}}^{\dagger}\Pi_{\mathcal{Y}\mathcal{X}}\Phi_{\mathcal{X}}$ is used. Thus $\|C_{\mathcal{X}\mathcal{Y}}^{\Pi}A_{\mathcal{X}} - A_{\mathcal{Y}}\|_{F}^{2} = 0$ and $C_{\mathcal{X}\mathcal{Y}} = C_{\mathcal{X}\mathcal{Y}}^{\Pi}$ ($C_{\mathcal{X}\mathcal{Y}} = \operatorname{argmin}_{C} \|CA_{\mathcal{X}} - A_{\mathcal{Y}}\|_{F}^{2}$ achieves 0, and $A_{\mathcal{X}}$ is full rank so that the solution is unique, implying $C_{\mathcal{X}\mathcal{Y}} = C_{\mathcal{X}\mathcal{Y}}^{\Pi}$). \Box

Discussion. Theorem 4.2 builds a connection between the functional map C_{XY} computed from the functional map solver, i.e. Eq. (1), and the functional map C_{XY}^{Π} converted from the point-wise map $\Pi_{\mathcal{Y}\mathcal{X}}$. It explicitly shows that $C_{\mathcal{X}\mathcal{Y}}$ and $C_{\mathcal{X}\mathcal{Y}}^{\Pi}$ are equal under certain conditions. However, in practical situations, the assumptions are too restrictive and often not satisfied. For example, when computing functional maps using Eq. (1), structural regularisation $E_{\rm reg}$ is typically needed to preserve the structure of the functional map (e.g. Laplacian commutativity for isometry). Furthermore, we often do not want to constrain the feature F_{χ} to lie in the span of the corresponding LBO eigenfunctions $\Phi_{\mathcal{X}}$, which limits its discriminative power and expressiveness, since the first k LBO eigenfunctions correspond to the k smoothest orthonormal functions defined on the surface w.r.t. the Dirichlet energy [7]. Even though the conditions of the theoretical results are not strictly met, the results give insights about the relations between variables, which we transfer into soft constraints that approximate the conditions. For instance, we note that the functional map solver plays a crucial role to balance the functional map regularisation and the coupling relation between the functional map and the point-wise map. On the one hand, the regularisation term $E_{\rm reg}$ in Eq. (1) preserves the functional map structure. On the other hand, it may result an invalid functional map (i.e. a functional map without an associated point-wise map). Therefore, it is important to adjust the functional map regularisation in a data-driven manner.

5. Revisiting the map relations

In the previous section, we theoretically analyse the relationship between the functional map computed from the functional map solver and the functional map converted from the point-wise map based on deep feature similarity. Motivated by our analysis, we propose two simple yet efficient extensions from the existing framework, which we introduce in the following. We highlight the different parts in Fig. 2 (right) with red colour, compared to the common deep functional map pipeline shown in Fig. 2 (left).

5.1. Self-adaptive functional map solver

As discussed in Theorem 4.2., we only consider E_{data} and ignore E_{reg} in Eq. (1) with some additional assumptions (i.e. $F_{\mathcal{X}}, F_{\mathcal{Y}}$ in the span of $\Phi_{\mathcal{X}}, \Phi_{\mathcal{Y}}$, and $A_{\mathcal{X}}$ is full rank). With that, the functional map computed by the functional map solver $C_{\mathcal{X}\mathcal{Y}}$ is equal to the functional map $C_{\mathcal{X}\mathcal{Y}}^{\Pi}$ converted by the point-wise map. However, as shown in Fig. 1, the valid functional maps often exhibit certain structures that needs to be imposed by the regularisation term E_{reg} . To this end, we propose a self-adaptive functional map solver that can optimise the regularisation strength and the regularisation structure based on the training data. Specifically, we use the regularisation term proposed by Ren et al. [53], which is an extension of the standard Laplacian commutativity. The standard Laplacian commutativity can be formulated as

$$E_{\text{lap}} = \|C_{\mathcal{X}\mathcal{Y}}\Lambda_{\mathcal{X}} - \Lambda_{\mathcal{Y}}C_{\mathcal{X}\mathcal{Y}}\|_{F}^{2}$$

$$= \sum_{ij} (\Lambda_{\mathcal{Y}}(i,i) - \Lambda_{\mathcal{X}}(j,j))^{2} [C_{\mathcal{X}\mathcal{Y}}]_{ij}^{2}$$

$$= \sum_{ij} [M_{\text{lap}}]_{ij} [C_{\mathcal{X}\mathcal{Y}}]_{ij}^{2}.$$
 (10)

Ren et al. [53] extended the standard mask $M_{\rm lap}$ to a resolvent mask $M_{\rm res}^{\gamma}$ in the form

$$[M_{\rm res}^{\gamma}]_{ij} = [M_{\rm re}^{\gamma}]_{ij} + [M_{\rm im}^{\gamma}]_{ij}, \qquad (11)$$

where

$$[M_{\rm re}^{\gamma}]_{ij} = \left(\frac{\Lambda_{\mathcal{Y}}^{\gamma}(i,i)}{\Lambda_{\mathcal{Y}}^{2\gamma}(i,i)+1} - \frac{\Lambda_{\mathcal{X}}^{\gamma}(j,j)}{\Lambda_{\mathcal{X}}^{2\gamma}(j,j)+1}\right)^2, \quad (12)$$

$$[M_{\rm im}^{\gamma}]_{ij} = \left(\frac{1}{\Lambda_{\mathcal{Y}}^{2\gamma}(i,i)+1} - \frac{1}{\Lambda_{\mathcal{X}}^{2\gamma}(j,j)+1}\right)^2.$$
 (13)

The parameter γ in the resolvent mask controls the regularisation structure of the functional map as shown in Fig. 3. In general, the γ is chosen in the range (0, 1] to keep the funnel-structure regularisation to be similar to the groundtruth functional map, i.e. more diagonal-dominant entries for smaller eigenvalues. Additionally, a larger γ imposes larger penalisation on the non-zero off-diagonal entries and a smaller γ provides more flexibility of the functional map (see Fig. 3).



Figure 3. The resolvent mask $M_{\rm res}^{\gamma}$ for different γ . The red region indicates large penalty, while the blue region indicates small penalty. We notice the funnel-like structure changes w.r.t. the change of γ and it reverses the direction for $\gamma > 1$.

Instead of manually choosing the regularisation strength (i.e. λ in Eq. (1)) and structure (i.e. γ in Eq. (12), Eq. (13)), we propose to learn these parameters during training. To this end, the functional map solver is optimised from the input shapes to find a better balance between the data term E_{data} and regularisation term E_{reg} and thus to better couple $C_{\mathcal{X}\mathcal{Y}}$ and $C_{\mathcal{X}\mathcal{Y}}^{\Pi}$. In the experiment, we show the simple modification leads to better matching performance especially for the most challenging scenarios. We also visualise the regularisation for each evaluated datasets in Fig. 9.

5.2. Vertex-wise contrastive loss

As discussed in Theorem 4.1, a valid point-wise map based on the feature similarity requires $F_{\mathcal{X}}$, $F_{\mathcal{Y}}$ both have distinct rows. To this end, we propose a vertex-wise contrastive loss to encourage more discriminative features. We first compute a point-wise map $\Pi_{\mathcal{X}\mathcal{X}}$ that maps shape \mathcal{X} to itself. To make the computation differentiable, we use the softmax operator to approximate a soft point-wise map, i.e.

$$\Pi_{\mathcal{X}\mathcal{X}} = \operatorname{Softmax}\left(F_{\mathcal{X}}F_{\mathcal{X}}^{T}/\tau\right), \qquad (14)$$

where parameter τ is to determine the softness of the pointwise map. Similar to [12], the computed point-wise map $\Pi_{\mathcal{X}\mathcal{X}}$ is projected to the associated functional map $C_{\mathcal{X}\mathcal{X}}$, i.e.

$$C_{\mathcal{X}\mathcal{X}} = \Phi_{\mathcal{X}}^{\dagger} \Pi_{\mathcal{X}\mathcal{X}} \Phi_{\mathcal{X}}.$$
 (15)

Our vertex-wise contrastive loss regularises the functional map $C_{\mathcal{X}\mathcal{X}}$ to be an identity matrix, i.e.

$$L_{\text{contrast}} = \left\| C_{\mathcal{X}\mathcal{X}} - I \right\|_F^2.$$
(16)

Similarly, we also apply the vertex-wise contrastive loss $L_{\rm contrast}$ for the functional map $C_{\mathcal{Y}\mathcal{Y}}$. There are two main advantages of applying the regularisation on the functional map domain. The first advantage is to make $L_{\rm contrast}$ comparable to other loss terms (i.e. Eq. (2) and Eq. (5)). The second advantage is to make the $L_{\rm contrast}$ discretisation-agnostic. Overall, the total unsupervised loss can be expressed as

$$L_{\text{total}} = L_{\text{fmap}} + \lambda_{\text{couple}} L_{\text{couple}} + \lambda_{\text{contrast}} L_{\text{contrast}}.$$
 (17)

6. Experimental results

In this section we compare our method to previous methods on diverse benchmark shape matching datasets with different settings (including near-isometric, non-isometric, topological noisy, and partial shape matching).

6.1. Near-isometric shape matching

Datasets. We evaluate our method on three standard benchmark datasets, namely the FAUST [8], SCAPE [1] and SHREC'19 [45] datasets. Following prior works, we choose the more challenging remeshed versions from [16, 52]. The FAUST dataset consists of 100 shapes, where the train/test split is 80/20. The SCAPE dataset contains 71 shapes, where the last 20 shapes are used for evaluation. The SHREC'19 dataset is a more challenging dataset with significant variance in the mesh connectivity and shape geometry. It has a total of 430 pairs for evaluation.

Results. The mean geodesic error [35] is used as quantitative measure. We compare our method with state-of-theart axiomatic, supervised and unsupervised methods. The results are summarised in Tab. 2. Our method outperforms the previous state of the art, even in comparison to the supervised methods. Meanwhile, our method achieves substantially better cross-dataset generalisation ability compared to existing learning-based methods.

6.2. Non-isometric shape matching

Datasets. In the context of non-isometric shape matching, we consider the SMAL [69] dataset and the DT4D-H [44] dataset. The SMAL dataset contains 49 animal shapes of eight species. Following Donati et al. [18], five species are used for training and three different species are used for testing (i.e. 29/20 shapes for train/test split). The DT4D-H dataset based on DeformingThings4D [40] is introduced by Magnet et al. [44]. Following Li et al. [38], nine classes of humanoid shapes are used for evaluation, resulting in 198/95 shapes for train/test split.

Train	FAUST	SCAPE	FAUST + SCAPE				
Test	FAUST	SCAPE	SHREC'19				
Axiomatic Methods							
BCICP [52]	6.1	11.0	-				
ZoomOut [46]	6.1	7.5	-				
Smooth Shells [22]	2.5	4.7	-				
DiscreteOp [55]	5.6	13.1	-				
Supervised Methods							
FMNet [41]	11.0	17.0	-				
3D-CODED [29]	2.5	31.0	-				
GeomFMaps [16]	2.6	3.0	7.9				
Unsupervised Methods							
WSupFMNet [61]	3.8	4.4	-				
Deep Shells [23]	1.7	2.5	21.1				
DUO-FMNet [18]	2.5	2.6	6.4				
AttentiveFMaps [38]	1.9	2.2	5.8				
AttentiveFMaps-Fast [38]	1.9	2.1	6.3				
URSSM [12]	1.6	1.9	4.6				
Ours	1.5	1.8	3.4				

Table 2. Near-isometric shape matching and cross-dataset generalisation on FAUST, SCAPE and SHREC'19. The best results in each column are highlighted. Our method outperforms previous axiomatic, supervised and unsupervised methods and demonstrates better cross-dataset generalisation ability.

Geo. error (×100)	SMAL	DT4D-H			
	510H HL	intra-class	inter-class		
Axiomatic Methods					
ZoomOut [46]	38.4	4.0	29.0		
Smooth Shells [22]	36.1	1.1	6.3		
DiscreteOp [55]	38.1	3.6	27.6		
Supervised Methods					
FMNet [41]	42.0	9.6	38.0		
GeomFMaps [16]	8.4	2.1	4.1		
Unsupervised Methods					
WSupFMNet [61]	7.6	3.3	22.6		
Deep Shells [23]	29.3	3.4	31.1		
DUO-FMNet [18]	6.7	2.6	15.8		
AttentiveFMaps [38]	5.4	1.7	11.6		
AttentiveFMaps-Fast [38]	5.8	1.2	14.6		
URSSM [12]	3.9	0.9	4.1		
Ours	3.6	1.0	4.0		

Table 3. Non-isometric matching on SMAL and DT4D-H. Our method outperforms all existing methods for challenging nonisometric inter-class shape matching on both SMAL and DT4D-H datasets and shows comparable performance on intra-class shape matching on DT4D-H dataset.

Results. Tab. 3 summarises the matching results on the SMAL and DT4D-H datasets. In the context of interclass shape matching, our approach outperforms the ex-



Figure 4. Non-isometric matching on SMAL and DT4D-H inter-class datasets. Proportion of correct keypoints (PCK) curves and corresponding area under curve (AUC) of our method compared to the existing state-of-the-art methods.

isting state of the art on both challenging non-isometric datasets, even in comparison to supervised methods. Meanwhile, our method demonstrates comparable and nearperfect matching results for intra-class matching on the DT4D-H dataset. Fig. 4 shows the PCK curves and the corresponding AUC of our method compared to existing stateof-the-art methods. Fig. 5 demonstrates the qualitative results of our method applied on the challenging SHREC'20 dataset [19].



Figure 5. Qualitative results on the challenging SHREC'20 dataset of our method. Our method is capable of finding reliable correspondences even for shapes with extremely large non-isometric deformations.

6.3. Matching with topological noise

Datasets. The mesh topology is often degraded due to selfintersections of separate parts of real-world scanned objects. Such topological noise presents a large challenge to matching methods based on the functional map framework as it distorts the intrinsic shape geometry [37]. To evaluate our method for matching with topologically noisy shapes, we use the TOPKIDS dataset [37]. Due to the small amount of data (26 shapes), we consider only axiomatic and unsupervised methods for comparison.

Results. We compare our method with state-of-the-art axiomatic methods and unsupervised methods. The quantitative results are summarised in Tab. 4. Our method outperforms the existing methods substantially, even in comparison to methods relying on additional extrinsic alignment information. We show the PCK curves of our method in Fig. 6 (left) and qualitative results in Fig. 7.

Geo. error (×100)	TOPKIDS	Fully intrinsic				
Axiomatic Methods						
ZoomOut [46]	33.7	\checkmark				
Smooth Shells [22]	11.8	×				
DiscreteOp [55]	35.5	1				
Unsupervised Methods						
WSupFMNet [61]	47.9	1				
Deep Shells [23]	13.7	×				
NeuroMorph [24]	13.8	×				
AttentiveFMaps [38]	23.4	1				
AttentiveFMaps-Fast [38]	28.5	1				
URSSM [12]	9.2	✓				
Ours	6.6	1				

Table 4. **Quantitative results on the TOPKIDS dataset.** We distinguish fully intrinsic methods (i.e. methods based on the functional map framework) from the methods that rely on additional extrinsic information (e.g. a weak alignment). Our method outperforms all existing methods substantially, even in comparison to methods relying on additional extrinsic information.



Figure 6. Matching with topological noise on TOPKIDS and partial shape matching on SHREC'16. Our method improves the state of the art over existing approaches.



Figure 7. **Qualitative results on TOPKIDS dataset.** Compared to existing fully intrinsic approaches, our method is more robust to topological noise.

6.4. Partial shape matching

Datasets. We evaluate our method on the SHREC'16 partial dataset [14]. The dataset contains 200 training shapes and 400 test shapes, with 8 different classes (humans and animals). Each class has a complete shape to be matched by the other partial shapes. The dataset is divided into two subsets, namely CUTS (missing a large part) with 120/200 train/test split, and HOLES (missing many small parts) with 80/200 train/test split.

Train	CUTS		HOLES			
Test	CUTS	HOLES	CUTS	HOLES		
Axiomatic Methods						
PFM [56]	9.7	23.2	9.7	23.2		
FSP [42]	16.1	33.7	16.1	33.7		
Supervised Methods						
GeomFMaps [16]	12.8	20.6	19.8	15.3		
DPFM [4]	3.2	15.8	8.6	13.1		
Unsupervised Methods						
DPFM-unsup [4]	9.0	22.8	16.5	20.5		
ConsistFMaps [10]	8.4	23.7	15.7	17.9		
URSSM [12]	3.3	13.7	5.2	9.1		
Ours	2.3	15.2	5.1	6.9		

Table 5. **Partial shape matching on SHREC'16 dataset.** Our method substantially outperforms state-of-the-art methods and shows comparable cross-dataset generalisation ability, even in comparison to the supervised approached.



Figure 8. **Qualitative results on SHREC'16 dataset.** Compared to existing methods, our method is more robust to partiality.

Results. We summarise the quantitative results on the SHREC'16 datasets in Tab. 5 and the corresponding PCK curve in Fig. 6 (right). Compared to existing methods, our approach is more robust to partiality. We qualitatively compare our method to existing approaches in Fig. 8.

6.5. Analysis of self-adaptive functional map solver

We summarise the learned parameters of the functional map solver for different kinds of datasets to better understand the learned regularisation strength and structure. Fig. 9 visualises the different regularisation strength (i.e. λ) and different regularisation structure (i.e. γ) for different datasets.

We obverse that the regularisation strength (i.e. λ) for near-isometric shape matching (FAUST, SCAPE) is



Figure 9. Different regularisation strength and structure for different datasets. The self-adaptive functional map solver enables to adjust the regularisation based on the training data.

stronger than the strength for non-isometric shape matching (SMAL, DT4D-H), since in theory functional maps for isometric shape matching are diagonal matrices. In the context of regularisation structure, the funnel-like structure is narrower for topological noisy (TOPKIDS) and partial shapes (CUTS, HOLES).

7. Limitation and future work

We build upon the existing state-of-the-art method [12] by introducing the self-adaptive functional map solver and the vertex-wise contrastive loss, and thereby achieve the new state of the art on a wide range of benchmark datasets. Yet, there are also some limitations that give rise to interesting future researches. Our unsupervised method is applicable in various settings. However, it can not be used for partial-topartial shape matching. Therefore, it is interesting to investigate how to extend the current framework for partial-topartial shape matching. For functional map computation, we optimise the two parameters (i.e. γ, λ) that control the regularisation strength and structure. Meanwhile, the number of LBO eigenfunctions is also an important parameter for functional map computation. How to automatically select the best number of LBO eigenfunctions is thereby an another interesting future work direction.

8. Conclusion

We theoretically analyse the relationship between the functional map from the functional map solver and the functional map from the point-wise map. Based on our theoretical analysis, we extend the current state-of-the-art methods. We evaluate our proposed method on diverse shape matching benchmark datasets with different settings and demonstrate the new state-of-the-art performance. We believe a more accurate and robust non-rigid 3D shape matching method would be beneficial for the shape analysis community to better explore the shape relationship.

References

- Dragomir Anguelov, Praveen Srinivasan, Daphne Koller, Sebastian Thrun, Jim Rodgers, and James Davis. Scape: shape completion and animation of people. In ACM SIGGRAPH. 2005. 6
- [2] Souhaib Attaiki and Maks Ovsjanikov. Ncp: Neural correspondence prior for effective unsupervised shape matching. In *NeurIPS*, 2022. 2
- [3] Souhaib Attaiki and Maks Ovsjanikov. Understanding and improving features learned in deep functional maps. In *CVPR*, 2023. 2, 3
- [4] Souhaib Attaiki, Gautam Pai, and Maks Ovsjanikov. Dpfm: Deep partial functional maps. In *International Conference* on 3D Vision (3DV), 2021. 1, 2, 8
- [5] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In *ICCV*, 2011. 2, 3
- [6] Florian Bernard, Zeeshan Khan Suri, and Christian Theobalt. Mina: Convex mixed-integer programming for non-rigid shape alignment. In CVPR, 2020. 2
- [7] Alexander I Bobenko and Boris A Springborn. A discrete laplace–beltrami operator for simplicial surfaces. *Discrete* & Computational Geometry, 38(4):740–756, 2007. 4
- [8] Federica Bogo, Javier Romero, Matthew Loper, and Michael J Black. Faust: Dataset and evaluation for 3d mesh registration. In *CVPR*, 2014. 6
- [9] Michael M Bronstein and Iasonas Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In *CVPR*, 2010. 2, 3
- [10] Dongliang Cao and Florian Bernard. Unsupervised deep multi-shape matching. In ECCV, 2022. 1, 2, 8
- [11] Dongliang Cao and Florian Bernard. Self-supervised learning for multimodal non-rigid 3d shape matching. In CVPR, 2023. 2
- [12] Dongliang Cao, Paul Roetzer, and Florian Bernard. Unsupervised learning of robust spectral shape matching. ACM Transactions on Graphics (TOG), 2023. 1, 2, 3, 4, 5, 6, 7, 8
- [13] Aharon Cohen and Mirela Ben-Chen. Robust shape collection matching and correspondence from shape differences. In *Computer Graphics Forum*. Wiley Online Library, 2020.
 2
- [14] Luca Cosmo, Emanuele Rodola, Michael M Bronstein, Andrea Torsello, Daniel Cremers, and Y Sahillioglu. Shrec'16: Partial matching of deformable shapes. *Proc. 3DOR*, 2(9): 12, 2016. 7
- [15] Huong Quynh Dinh, Anthony Yezzi, and Greg Turk. Texture transfer during shape transformation. ACM Transactions on Graphics (ToG), 24(2):289–310, 2005. 1
- [16] Nicolas Donati, Abhishek Sharma, and Maks Ovsjanikov. Deep geometric functional maps: Robust feature learning for shape correspondence. In *CVPR*, 2020. 1, 2, 6, 8
- [17] Nicolas Donati, Etienne Corman, Simone Melzi, and Maks Ovsjanikov. Complex functional maps: A conformal link between tangent bundles. In *Computer Graphics Forum*. Wiley Online Library, 2022. 1

- [18] Nicolas Donati, Etienne Corman, and Maks Ovsjanikov. Deep orientation-aware functional maps: Tackling symmetry issues in shape matching. In *CVPR*, 2022. 1, 2, 6
- [19] Roberto M. Dyke, Yu-Kun Lai, Paul L. Rosin, Stefano Zappalà, Seana Dykes, Daoliang Guo, Kun Li, Riccardo Marin, Simone Melzi, and Jingyu Yang. SHREC'20: Shape correspondence with non-isometric deformations. *Computers & Graphics*, 92:28–43, 2020. 7
- [20] Bernhard Egger, William AP Smith, Ayush Tewari, Stefanie Wuhrer, Michael Zollhoefer, Thabo Beeler, Florian Bernard, Timo Bolkart, Adam Kortylewski, Sami Romdhani, et al. 3d morphable face models—past, present, and future. ACM Transactions on Graphics (ToG), 39(5):1–38, 2020. 1
- [21] Marvin Eisenberger, Zorah Lähner, and Daniel Cremers. Divergence-free shape correspondence by deformation. In *Computer Graphics Forum*. Wiley Online Library, 2019. 2
- [22] Marvin Eisenberger, Zorah Lahner, and Daniel Cremers. Smooth shells: Multi-scale shape registration with functional maps. In *CVPR*, 2020. 2, 6, 7
- [23] Marvin Eisenberger, Aysim Toker, Laura Leal-Taixé, and Daniel Cremers. Deep shells: Unsupervised shape correspondence with optimal transport. *NIPS*, 2020. 1, 6, 7
- [24] Marvin Eisenberger, David Novotny, Gael Kerchenbaum, Patrick Labatut, Natalia Neverova, Daniel Cremers, and Andrea Vedaldi. Neuromorph: Unsupervised shape interpolation and correspondence in one go. In CVPR, 2021. 7
- [25] Davide Eynard, Emanuele Rodola, Klaus Glashoff, and Michael M Bronstein. Coupled functional maps. In 2016 Fourth International Conference on 3D Vision (3DV), 2016.
 2
- [26] Danielle Ezuz, Behrend Heeren, Omri Azencot, Martin Rumpf, and Mirela Ben-Chen. Elastic correspondence between triangle meshes. In *Computer Graphics Forum*, pages 121–134. Wiley Online Library, 2019. 2
- [27] Danielle Ezuz, Justin Solomon, and Mirela Ben-Chen. Reversible harmonic maps between discrete surfaces. ACM Transactions on Graphics (ToG), 38(2):1–12, 2019. 3
- [28] Maolin Gao, Zorah Lahner, Johan Thunberg, Daniel Cremers, and Florian Bernard. Isometric multi-shape matching. In CVPR, 2021. 2
- [29] Thibault Groueix, Matthew Fisher, Vladimir G Kim, Bryan C Russell, and Mathieu Aubry. 3d-coded: 3d correspondences by deep deformation. In ECCV, 2018. 6
- [30] Oshri Halimi, Or Litany, Emanuele Rodola, Alex M Bronstein, and Ron Kimmel. Unsupervised learning of dense shape correspondence. In *CVPR*, 2019. 1, 2
- [31] Benjamin Holzschuh, Zorah L\u00e4hner, and Daniel Cremers. Simulated annealing for 3d shape correspondence. In 2020 International Conference on 3D Vision (3DV), 2020. 2
- [32] Qixing Huang, Fan Wang, and Leonidas Guibas. Functional map networks for analyzing and exploring large shape collections. ACM Transactions on Graphics (ToG), 33(4):1–11, 2014. 2
- [33] Qi-Xing Huang, Bart Adams, Martin Wicke, and Leonidas J Guibas. Non-rigid registration under isometric deformations. In *Computer Graphics Forum*. Wiley Online Library, 2008.
 2

- [34] Ruqi Huang, Jing Ren, Peter Wonka, and Maks Ovsjanikov. Consistent zoomout: Efficient spectral map synchronization. In *Computer Graphics Forum*. Wiley Online Library, 2020.
 2
- [35] Vladimir G Kim, Yaron Lipman, and Thomas Funkhouser. Blended intrinsic maps. ACM Transactions on Graphics (ToG), 30(4):1–12, 2011. 6
- [36] Artiom Kovnatsky, Michael M Bronstein, Alexander M Bronstein, Klaus Glashoff, and Ron Kimmel. Coupled quasiharmonic bases. In *Computer Graphics Forum*. Wiley Online Library, 2013. 2
- [37] Zorah Lähner, Emanuele Rodola, Michael M Bronstein, Daniel Cremers, Oliver Burghard, Luca Cosmo, Andreas Dieckmann, Reinhard Klein, and Yusuf Sahillioglu. Shrec'16: Matching of deformable shapes with topological noise. *Proc. 3DOR*, 2(10.2312), 2016. 7
- [38] Lei Li, Nicolas Donati, and Maks Ovsjanikov. Learning multi-resolution functional maps with spectral attention for robust shape matching. *NIPS*, 2022. 1, 2, 6, 7
- [39] Tianye Li, Timo Bolkart, Michael. J. Black, Hao Li, and Javier Romero. Learning a model of facial shape and expression from 4D scans. ACM Transactions on Graphics (ToG), 36(6):194:1–194:17, 2017. 1
- [40] Yang Li, Hikari Takehara, Takafumi Taketomi, Bo Zheng, and Matthias Nießner. 4dcomplete: Non-rigid motion estimation beyond the observable surface. In *ICCV*, 2021. 6
- [41] Or Litany, Tal Remez, Emanuele Rodola, Alex Bronstein, and Michael Bronstein. Deep functional maps: Structured prediction for dense shape correspondence. In *ICCV*, 2017. 1, 2, 6
- [42] Or Litany, Emanuele Rodolà, Alexander M Bronstein, and Michael M Bronstein. Fully spectral partial shape matching. In *Computer Graphics Forum*. Wiley Online Library, 2017.
 2, 8
- [43] Matthew Loper, Naureen Mahmood, Javier Romero, Gerard Pons-Moll, and Michael J Black. Smpl: A skinned multiperson linear model. ACM Transactions on Graphics (ToG), 34(6):1–16, 2015. 1
- [44] Robin Magnet, Jing Ren, Olga Sorkine-Hornung, and Maks Ovsjanikov. Smooth non-rigid shape matching via effective dirichlet energy optimization. In *International Conference* on 3D Vision (3DV), 2022. 2, 6
- [45] Simone Melzi, Riccardo Marin, Emanuele Rodolà, Umberto Castellani, Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Shrec 2019: Matching humans with different connectivity. In *Eurographics Workshop on 3D Object Retrieval*, 2019. 6
- [46] Simone Melzi, Jing Ren, Emanuele Rodolà, Abhishek Sharma, Peter Wonka, and Maks Ovsjanikov. Zoomout: spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (ToG)*, 38(6):1–14, 2019. 1, 3, 6, 7
- [47] James Munkres. Algorithms for the assignment and transportation problems. *Journal of the society for industrial and applied mathematics*, 5(1):32–38, 1957. 2
- [48] Dorian Nogneng and Maks Ovsjanikov. Informative descriptor preservation via commutativity for shape matching. In *Computer Graphics Forum*. Wiley Online Library, 2017. 1

- [49] Maks Ovsjanikov, Mirela Ben-Chen, Justin Solomon, Adrian Butscher, and Leonidas Guibas. Functional maps: a flexible representation of maps between shapes. ACM Transactions on Graphics (ToG), 31(4):1–11, 2012. 1, 2, 3
- [50] Ulrich Pinkall and Konrad Polthier. Computing discrete minimal surfaces and their conjugates. *Experimental mathematics*, 2(1):15–36, 1993. 2
- [51] Charles Ruizhongtai Qi, Li Yi, Hao Su, and Leonidas J Guibas. Pointnet++: Deep hierarchical feature learning on point sets in a metric space. *NIPS*, 2017. 2
- [52] Jing Ren, Adrien Poulenard, Peter Wonka, and Maks Ovsjanikov. Continuous and orientation-preserving correspondences via functional maps. ACM Transactions on Graphics (ToG), 37:1–16, 2018. 1, 2, 6
- [53] Jing Ren, Mikhail Panine, Peter Wonka, and Maks Ovsjanikov. Structured regularization of functional map computations. In *Computer Graphics Forum*. Wiley Online Library, 2019. 2, 5
- [54] Jing Ren, Simone Melzi, Maks Ovsjanikov, and Peter Wonka. Maptree: recovering multiple solutions in the space of maps. ACM Transactions on Graphics (ToG), 39(6):1–17, 2020. 2
- [55] Jing Ren, Simone Melzi, Peter Wonka, and Maks Ovsjanikov. Discrete optimization for shape matching. In *Computer Graphics Forum*. Wiley Online Library, 2021. 2, 3, 6, 7
- [56] Emanuele Rodolà, Luca Cosmo, Michael M Bronstein, Andrea Torsello, and Daniel Cremers. Partial functional correspondence. In *Computer Graphics Forum*. Wiley Online Library, 2017. 1, 2, 8
- [57] Paul Roetzer, Paul Swoboda, Daniel Cremers, and Florian Bernard. A scalable combinatorial solver for elastic geometrically consistent 3d shape matching. In *CVPR*, 2022. 2
- [58] Jean-Michel Roufosse, Abhishek Sharma, and Maks Ovsjanikov. Unsupervised deep learning for structured shape matching. In *ICCV*, 2019. 1, 2, 3
- [59] Yusuf Sahillioğlu. Recent advances in shape correspondence. *The Visual Computer*, 36(8):1705–1721, 2020. 2
- [60] Samuele Salti, Federico Tombari, and Luigi Di Stefano. Shot: Unique signatures of histograms for surface and texture description. *Computer Vision and Image Understanding*, 125:251–264, 2014. 2, 3
- [61] Abhishek Sharma and Maks Ovsjanikov. Weakly supervised deep functional maps for shape matching. *NIPS*, 2020. 1, 2, 3, 6, 7
- [62] Nicholas Sharp, Souhaib Attaiki, Keenan Crane, and Maks Ovsjanikov. Diffusionnet: Discretization agnostic learning on surfaces. arXiv preprint arXiv:2012.00888, 2020. 2
- [63] Robert W Sumner and Jovan Popović. Deformation transfer for triangle meshes. ACM Transactions on Graphics (ToG), 23(3):399–405, 2004. 1
- [64] Gary KL Tam, Zhi-Quan Cheng, Yu-Kun Lai, Frank C Langbein, Yonghuai Liu, David Marshall, Ralph R Martin, Xian-Fang Sun, and Paul L Rosin. Registration of 3d point clouds and meshes: A survey from rigid to nonrigid. *IEEE transactions on visualization and computer graphics*, 19(7):1199– 1217, 2012. 1, 2

- [65] Hugues Thomas, Charles R Qi, Jean-Emmanuel Deschaud, Beatriz Marcotegui, François Goulette, and Leonidas J Guibas. Kpconv: Flexible and deformable convolution for point clouds. In *ICCV*, 2019. 2
- [66] Oliver Van Kaick, Hao Zhang, Ghassan Hamarneh, and Daniel Cohen-Or. A survey on shape correspondence. In *Computer Graphics Forum*. Wiley Online Library, 2011. 1,
 2
- [67] Matthias Vestner, Roee Litman, Emanuele Rodola, Alex Bronstein, and Daniel Cremers. Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In *CVPR*, 2017. 3
- [68] Thomas Windheuser, Ulrich Schlickewei, Frank R Schmidt, and Daniel Cremers. Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In *ICCV*, 2011. 2
- [69] Silvia Zuffi, Angjoo Kanazawa, David W Jacobs, and Michael J Black. 3d menagerie: Modeling the 3d shape and pose of animals. In *CVPR*, 2017. 6