

On Unsupervised Partial Shape Correspondence

Amit Bracha , Thomas Dagès , and Ron Kimmel 

Technion - Israel Institute of Technology
Haifa, Israel
`amit.bracha@cs.technion.ac.il`

Abstract. While dealing with matching shapes to their parts, we often apply a tool known as functional maps. The idea is to translate the shape matching problem into “convenient” spaces by which matching is performed algebraically by solving a least squares problem. Here, we argue that such formulations, though popular in this field, introduce errors in the estimated match when partiality is invoked. Such errors are unavoidable even for advanced feature extraction networks, and they can be shown to escalate with increasing degrees of shape partiality, adversely affecting the learning capability of such systems. To circumvent these limitations, we propose a novel approach for partial shape matching. Our study of functional maps led us to a novel method that establishes direct correspondence between partial and full shapes through feature matching bypassing the need for functional map intermediate spaces. The Gromov Distance between metric spaces leads to the construction of the first part of our loss functions. For regularization we use two options: a term based on the area preserving property of the mapping, and a relaxed version that avoids the need to resort to functional maps. The proposed approach shows superior performance on the SHREC’16 dataset, outperforming existing unsupervised methods for partial shape matching. Notably, it achieves state-of-the-art results on the SHREC’16 HOLES benchmark, superior also compared to supervised methods. We demonstrate the benefits of the proposed unsupervised method when applied to a new dataset PFAUST for part-to-full shape correspondence.

Keywords: Shape correspondence · partial shapes · learning to match

1 Introduction

In recent years, the emergence of deep learning allowed significant leap forwards in efficiently solving computational problems, like finding the point-to-point correspondence between two sampled surfaces. And yet, for unsupervised partial shape matching the contribution of deep learning was somewhat limited. Most state-of-the-art shape correspondence approaches rely on the *functional maps* framework [33], which searches for the mapping between spectral representations of corresponding features on the two input shapes. One popular framework uses a *functional map layer* [26] as a differentiable component in a neural network architecture. We will show that such a *functional map layer* introduces inherent errors in the case of partial shape matching.

Theoretically, we could overcome such unavoidable errors of using functional maps as layers in a neural network. For example, one could first detect the corresponding parts by matching the spectra of an appropriate Hamiltonian of the part with that of the Laplacian of the whole [5, 35]. After detecting the matching parts, the point-to-point correspondence can be solved by resorting to an algebraic problem defined by the relation of projected scalar feature functions onto spectral domains, aka *functional maps*.

State-of-the-art shape correspondence approaches often involve three components [14, 16, 20, 39–41]. Initially, a network computes refined pointwise features. Next, a *functional map layer* uses these features to estimate the correspondence between the shapes. Finally, loss functions based on the resulting functional map are applied during training to further adapt the features to the data at hand.

In this paper, we analyse the use of least squares for estimating the functional map in the *functional map layer* and focus on the resulting errors involved in partiality. While previous studies argue that it should be possible to use least squares in the *functional map layer* to obtain the ground-truth map without errors [4], here, we prove that the projected errors cannot be avoided in such a setting. Furthermore, the magnitude of the error directly relates to the degree of partiality – the smaller the part of the shape the larger the error. The reason for this error stems from the fact that the least squares method considers the inner product of the shape’s basis functions and the feature descriptors. That inner product involves integration over the whole shape.

With this observation in mind, we propose a novel approach for partial shape matching, surpassing all existing unsupervised techniques. The proposed method deviates from the conventional process of generating a functional map from the extracted features prior to the correspondence. We avoid the functional map formulation altogether and establish correspondence between partial and full shapes through direct feature matching. In essence, we demonstrate that for partial shape matching, choosing an appropriate loss, like the Gromov Distance [11, 13, 31], is as important as operating in convenient spaces, like the eigenspaces of the corresponding shape laplacians. Moreover, when treated as a tool implementing intrinsic smoothing [6] rather than a functional map network layer [26], one could harness the power of the Laplace-Beltrami operator (LBO) spectral representation for state-of-the-art solutions for shape matching.

The proposed network optimization process involves a two term loss. The core, inspired by [9–11, 20, 32], utilizes the near isometry property of surfaces in nature as a measure of correspondence that can serve as an unsupervised/semi-supervised way of training our network. This loss measures the distortions of geodesic distance between corresponding points, when mapping one shape to the other. For the second regularization-term we suggest two penalizing options, related to the area preserving mapping [39], and an option to avoid the need to compute the functional map altogether.

We test the proposed framework on the SHREC’16 dataset [15], containing two different types of shape partiality benchmarks, cuts and holes. We also evaluate it on our new dataset, PFAUST, with two different densities of holes. We

show that the proposed method has the lowest error in comparison to previous unsupervised partial shape matching methods.

Contributions.

1. We prove that for partial shape matching a *functional map layer* introduces unavoidable errors into the resulting correspondence function.
2. We introduce PFAUST, a new benchmark that includes two datasets with different densities of holes, dedicated to partial-to-full shape correspondence.
3. We present a novel approach for partial shape matching avoiding the *functional map layer* altogether. The proposed method reaches state-of-the-art results on the SHREC'16 and PFAUST by a large margin.




Fig. 1: Overview of the proposed pipeline. Basic features computed for the full and partial shapes are refined using a Siamese neural diffusion feature extractor. Next, cosine similarity is computed and fed into a softmax layer that produces a soft correspondence. While training, the loss functions are applied to the soft correspondence. At inference, the soft correspondence matrix is binarized for sharp matching.

2 Related efforts

Deep learning was first introduced for dense shape correspondence in [26] proposing a supervised approach known as FM-Net, based on a differentiable functional map layer. Unsupervised versions of the FM-Net appeared in [20, 39]. To transform the supervised approach [26] into an unsupervised one, [20] suggested a loss based on the Gromov-Wasserstein Distance. That distance was suggested as a measure of choice for defining the correspondence of surfaces in nature, see [8, 11, 17–19, 30, 34]. In another example, it was used to align the eigenfunctions of the scale invariant LBO of two shapes, see [8]. Here, we use the Gromov-Wasserstein loss without resorting to the functional map layer.

Later papers suggested new architectures [16, 25, 40, 41] improving the accuracy on networks trained in both supervised and unsupervised settings. In [28], it was argued that learning robust basis functions instead of the LBO eigenfunctions improves correspondence results. In [24], it was suggested to avoid the functional map framework for correspondence of full shapes, relying instead on a supervised contrastive loss. The correspondence is done directly between the

output network features, while a Dirichlet energy is minimized as part of the optimization to enforce smoothness provided by functional map methods. Recently, a method for shape correspondence [14] incorporates a loss [3], penalising the difference between the estimated functional map from the functional map layer, and the functional map extracted from soft correspondences calculated by cosine similarity between the output features of the network. It was noticed in [4] that when matching partial shapes the functional map framework requires some assistance. Motivated by this observation, a network was introduced with an added attention component that makes the feature extractor aware of features on the other shape. The network outputs vanilla features like other methods, those are then fed to the attention layer that refines the features. Theoretically, features of points on the full shape that do not belong to the matched part can be zeroed out with this approach. And yet, in our experiments we noticed that this method is hard to train. Indeed, in [4] only the non-refined features, without the attention part, were used for the functional map on the SHREC’16-CUTS benchmark. In [38], relations between eigenvalues, eigenfunctions, and the functional map matrix for partial shapes are studied. Unlike [38], we explore the error introduced by shape partiality to the functional map matrix produced by the FM-layer, which is a least squares estimate of the matrix extracted from the ground truth correspondence. While the error in [38] relates to the length of the cut, our analysis exhibits a relation to the missing area.

3 Functional maps and shape partiality

In Sec. 3.1 we briefly provide the background for *functional maps*, and in Sec. 3.2 and in the supplementary we analyse its issues with partial shape matching.

3.1 Functional maps

The notion of *functional maps* was introduced to non-rigid shape correspondence in [33]. Given $T : \mathcal{X} \rightarrow \mathcal{Y}$, a bijective mapping between two surfaces \mathcal{X} and \mathcal{Y} , then T induces a natural functional mapping $T_{\mathcal{F}} : \mathcal{F}(\mathcal{X}, \mathbb{R}) \rightarrow \mathcal{F}(\mathcal{Y}, \mathbb{R})$, $f \mapsto g = f \circ T^{-1}$ between scalar functions defined on each surface. Given two orthonormal bases $\{\psi_i^{\mathcal{X}}\}$ and $\{\psi_j^{\mathcal{Y}}\}$ of $\mathcal{F}(\mathcal{X}, \mathbb{R})$ and $\mathcal{F}(\mathcal{Y}, \mathbb{R})$, we can decompose any function $f = \sum_i a_i \psi_i^{\mathcal{X}} \in \mathcal{F}(\mathcal{X}, \mathbb{R})$ and its corresponding function $g = T_{\mathcal{F}}(f) = \sum_j b_j \psi_j^{\mathcal{Y}}$ in these bases. Remarkably, the coefficients (a_i) and (b_j) are related by a linear transformation depending only on the mapping T and the choice of bases, $b_j = \sum_i C_{ji} a_i$, where the coefficients $C_{ij} = \langle \psi_i^{\mathcal{Y}}, T_{\mathcal{F}}(\psi_j^{\mathcal{X}}) \rangle_{\mathcal{Y}}$ define the *functional map*.

When T is unknown, finding $\mathbf{C} = \{C_{ij}\}$ usually requires solving a least squares optimization problem, with constraints imposed by corresponding descriptor functions $(f_i^{\mathcal{X}}, f_i^{\mathcal{Y}})$ defined on the surfaces \mathcal{X} and \mathcal{Y} . Defining $\hat{\mathbf{F}}_x = \langle \psi_i^{\mathcal{X}}, f_j^{\mathcal{X}} \rangle_{\mathcal{X}}$ and $\hat{\mathbf{F}}_y = \langle \psi_i^{\mathcal{Y}}, f_j^{\mathcal{Y}} \rangle_{\mathcal{Y}}$, the matrix \mathbf{C} can be computed by solving

$$\mathbf{C} = \operatorname{argmin}_{\mathbf{C}} \left\| \mathbf{C} \hat{\mathbf{F}}_x - \hat{\mathbf{F}}_y \right\|_F. \quad (1)$$

There is no restriction of which basis to choose. However, it was shown in [1] that selecting the leading eigenfunctions of the LBO, when ordered by their corresponding eigenvalues small to large, optimally and uniquely approximate the family of smooth functions with bounded Dirichlet energy. When mentioning basis functions, eigenfunctions, or eigenvectors in this paper we refer to these low pass subsets of LBO leading basis elements.

The least squares problem Eq. (1) is solved by pseudo-inversion. In deep learning frameworks, this operation is called the functional map layer (FM-layer) [26], and deep learning pipelines for shape correspondence often incorporate it [4, 14, 20, 24, 39]. The FM-layer computes the solution to Eq. (1) as

$$\mathbf{C}_{yx} = \hat{\mathbf{F}}_x \hat{\mathbf{F}}_y^\top (\hat{\mathbf{F}}_y \hat{\mathbf{F}}_y^\top)^{-1}. \tag{2}$$

Here, \mathbf{C}_{yx} is a spectral mapping between \mathcal{Y} and \mathcal{X} . In the discrete setting, denoting \mathbf{F}_x and \mathbf{F}_y the corresponding feature matrices on \mathcal{X} and \mathcal{Y} , this becomes

$$\mathbf{C}_{yx} = \Psi_x^\top \mathbf{A}_x \mathbf{F}_x \mathbf{F}_y^\top \mathbf{A}_y \Psi_y (\Psi_y^\top \mathbf{A}_y \mathbf{F}_y \mathbf{F}_y^\top \mathbf{A}_y \Psi_y)^{-1}, \tag{3}$$

where Ψ_x is the matrix of basis functions on \mathcal{X} and \mathbf{A}_x is the diagonal matrix of the areas about vertices of the triangulated surface \mathcal{X} . For conciseness, we write

$$\mathbf{Q}_y = \Psi_y^\top \mathbf{A}_y \mathbf{F}_y \mathbf{F}_y^\top \mathbf{A}_y \Psi_y. \tag{4}$$

3.2 Shape partiality

Most deep learning approaches for shape correspondence use the same three components. First, a neural network generates features at each point of the surface. Then, these features are fed to an FM-layer to compute the functional map between the surfaces. Note, that this is the first time that the features of the two shapes interact. Finally, the functional map is used to estimate a pointwise correspondence between the two shapes. We prove in this subsection that in the case of partial shapes, the use of an FM-layer introduces unavoidable errors.

We begin our analysis with the continuous case.

Theorem 1 (Continuous case). *Given two surfaces \mathcal{X} and \mathcal{Y} , where $\mathcal{Y} \subset \mathcal{X}$. Denote $\mathcal{Z} = \mathcal{X} \setminus \mathcal{Y}$, and let $\{\psi_i^{\mathcal{X}}\}_{i=1}^k$ and $\{\psi_i^{\mathcal{Y}}\}_{i=1}^k$ be subsets of orthonormal bases of \mathcal{X} and \mathcal{Y} , respectively. Let $\hat{\mathbf{F}}_x$ and $\hat{\mathbf{F}}_y$ be the coefficients of the corresponding feature functions on the corresponding surfaces. Then, the functional map layer between \mathcal{X} and \mathcal{Y} computes a functional map $\hat{\mathbf{C}}_{yx} = \mathbf{C}_{yx} + \mathbf{C}_{yx}^E$, where*

$$\mathbf{C}_{yx} = \hat{\mathbf{F}}^{(\mathcal{Y})} \hat{\mathbf{F}}_y^\top (\hat{\mathbf{F}}_y \hat{\mathbf{F}}_y^\top)^{-1} \quad \text{and} \quad \mathbf{C}_{yx}^E = \hat{\mathbf{F}}^{(\mathcal{Z})} \hat{\mathbf{F}}_y^\top (\hat{\mathbf{F}}_y \hat{\mathbf{F}}_y^\top)^{-1}. \tag{5}$$

Here, the superscripts (\mathcal{Y}) and (\mathcal{Z}) indicate a matrix containing only information related to regions of \mathcal{X} that correspond to the partial surfaces \mathcal{Y} or \mathcal{Z} , respectively.

A proof can be found in the supplementary. The matrix \mathbf{C}_{yx} is the functional map relating $\Psi^{\mathcal{X}}$ and $\Psi^{\mathcal{Y}}$ when considering only descriptors which vanish on

\mathcal{Z} . We show in the supplementary that it is the functional map deduced from a perfect mapping between \mathcal{Y} to \mathcal{X} . At the other end, \mathbf{C}_{yx}^E is an error term that can be interpreted as the functional map relating $\Psi^{\mathcal{X}}$ and $\Psi^{\mathcal{Y}}$ when considering only descriptors that vanish on $\mathcal{X} \setminus \mathcal{Z}$. We readily deduce the following corollary,

Corollary 1. *If $\hat{\mathbf{F}}^{(\mathcal{Z})} \hat{\mathbf{F}}_y^\top \neq \mathbf{0}$, then, the functional map layer injects an error $\mathbf{C}_{yx}^E \neq \mathbf{0}$ into the functional map.*

Generally, $\hat{\mathbf{F}}^{(\mathcal{Z})} \hat{\mathbf{F}}_y^\top \neq \mathbf{0}$, thus, the functional map layer unavoidably introduces a significant error to the estimated functional map. The above theorem also applies to the discrete case, a proof can be found in the supplementary.

Theorem 2 (Discrete case). *Given a sampled surface \mathcal{X} , like a triangle mesh, split into two subsurfaces \mathcal{Y} and \mathcal{Z} . The functional map layer between \mathcal{X} and \mathcal{Y} yields a functional map $\hat{\mathbf{C}}_{yx} = \mathbf{C}_{yx} + \mathbf{C}_{yx}^E$, where*

$$\mathbf{C}_{yx} = (\Psi^{(\mathcal{Y})})^\top \mathbf{A}_y \mathbf{F}^{(\mathcal{Y})} \mathbf{F}_y^\top \mathbf{A}_y \Psi_y \mathbf{Q}_y^{-1} \text{ and } \mathbf{C}_{yx}^E = (\Psi^{(\mathcal{Z})})^\top \mathbf{A}_z \mathbf{F}^{(\mathcal{Z})} \mathbf{F}_y^\top \mathbf{A}_y \Psi_y \mathbf{Q}_y^{-1}. \quad (6)$$

Corollary 2. *If the feature dimensions $d < \min\{n_y, n_z\}$, where n_y and n_z are the number of sampled vertices of surfaces \mathcal{Y} and \mathcal{Z} , and if both $\mathbf{F}^{(\mathcal{Z})}$ and \mathbf{F}_y have d rank, and if both¹ $\mathbf{F}_y \not\perp_{\mathbf{A}_y} \Psi_y$ and $\mathbf{F}^{(\mathcal{Z})} \not\perp_{\mathbf{A}_z} \Psi^{(\mathcal{Z})}$, then, $\mathbf{C}_{yx}^E \neq \mathbf{0}$.*

In practice, $d \ll \min\{n_y, n_z\}$ and the other assumptions are satisfied. This means that for partial shape matching, the FM-layer unavoidably introduces an error into the estimated functional map. Nevertheless, as claimed in [4], it is possible for the error to vanish. Unfortunately, this can only happen when $d = n_x$, that is, in the unreasonable case where the descriptor dimensionality is equal to the total number of sampled vertices. Obviously, taking $\mathbf{F}^{(\mathcal{Z})} \rightarrow \mathbf{0}$ or $\mathbf{F}^{(\mathcal{Z})} \perp_{\mathbf{A}_z} \Psi^{(\mathcal{Z})}$ would eliminate this error; however, any general feature extractor would not be aware of the given shape partiality, that is, the matching region of \mathcal{Y} in \mathcal{X} . As mentioned, the features are extracted independently for each shape without any interaction before the FM-layer. As such, we cannot have a general meaningful feature extractor that provides $\mathbf{0}$ on \mathcal{Z} . See Fig. 2 for an example of the computed matrices \mathbf{C}_{yx} and \mathbf{C}_{yx}^E . Methods such as [38] and those relying on [36] deal with shape partiality by masking the matrix \mathbf{C} . One could easily generalise the proposed analysis along this line of thought, by which one introduces an additive mask matrix to \mathbf{Q} . We explore in detail the error induced by the FM-layer in the supplementary. We show that under simplifying assumptions the error of the FM-layer is proportional to the missing area.

4 Method

We have proven that the FM-layer involves an unavoidable error when applied to partial shape matching. Utilizing this knowledge, we propose to avoid an FM-layer and apply a simple and effective approach, see Figure 1.

¹ Where \perp_M means orthogonality with respect to the metric matrix M .




Fig. 2: Estimating the functional map \hat{C}_{yx} between a full \mathcal{X} and its part \mathcal{Y} using features independently extracted for each yields errors. Recall, $\hat{C}_{yx} = C_{yx} + C_{yx}^E$, where C_{yx} is the ideal functional map given the correct matching, and C_{yx}^E is an error resulting from matching the part \mathcal{Y} to its complementary part in \mathcal{X} . Here, we plot the magnitude of entries in $|C_{yx}|$ and $|C_{yx}^E|$. Indeed, C_{yx} contains an informative structure, whereas C_{yx}^E is a noise-like structure-less matrix.

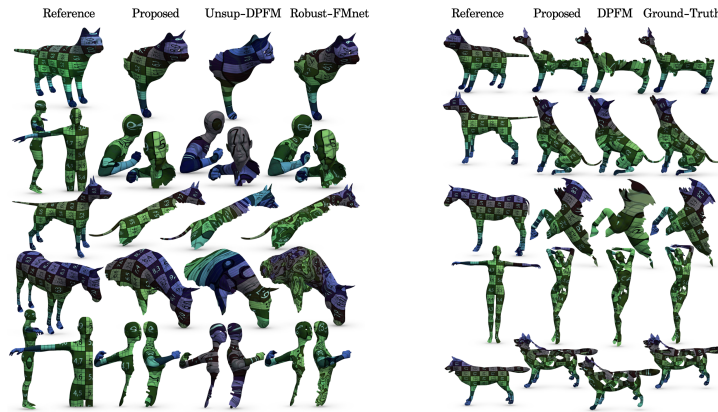


Fig. 3: Qualitative results on SHREC'16 CUTS (*left*) and HOLES (*right*). On SHREC'16 CUTS, we obtain visually appealing results that outperform previous unsupervised methods. On SHREC'16 HOLES, we obtain better matching results both visually and quantitatively than even the best supervised approach (DPFM).

Differentiable softmax of point features' inner product in the context of shape matching were first introduced in [3, 14]. We follow a similar line of thought, and estimate a soft correspondence matrix directly from the feature similarity of the two shapes using the softmax operator. Unlike the methods suggested in [3, 14], we do not feed the features from the feature extractor to an FM-layer. The correspondence is extracted using only feature similarity.

This correspondence matrix is an input to our loss functions. The main component of the proposed loss was first used in [20] in the context of learning surface matching. It penalizes the distortion of surface inter-geodesic distances between corresponding pairs of points. For regularization, we propose two options. The first option [39] is an orthogonality regularization of the estimated functional map extracted from the soft correspondence matrix, but it does not involve an FM-layer. It allows to apply a low pass filter on the computed correspondence in the dual spectral domain. To compute the functional map, we simply multiply the correspondence matrix by the leading LBO basis functions. We also propose

an alternative regularization option. It is based on the assumption that the low frequency or leading eigenfunctions on \mathcal{X} restricted to the regions corresponding to \mathcal{Y} could be well approximated by the low frequency leading eigenfunctions on \mathcal{Y} . This option avoids the need to compute the functional map altogether.

Feature extractor. A feature extractor in our context is a neural network that considers as input a given feature vector $\tilde{\mathbf{F}} \in \mathbb{R}^{n \times \tilde{d}}$ and outputs refined features $\mathbf{F} \in \mathbb{R}^{n \times d}$, where n is the number of vertices of the triangle surface, and \tilde{d} (resp. d) is the number of given (resp. refined) descriptor elements. An example of input features are the coordinates in \mathbb{R}^3 of each vertex of a given surface. Following a Siamese philosophy, the same network, is applied to both shapes.

The correspondence matrix \mathbf{P} . Given the extracted features for each shape, we compute their pointwise cosine similarity

$$S_C(\mathbf{F}_y, \mathbf{F}_x) \equiv \left\langle \frac{\mathbf{F}_y}{\|\mathbf{F}_y\|}, \frac{\mathbf{F}_x}{\|\mathbf{F}_x\|} \right\rangle. \quad (7)$$

We then apply a softmax operator to estimate directly a soft-corresponding matrix \mathbf{P}_{yx} mapping shape \mathcal{Y} to shape \mathcal{X} . Note, that this approach differs from the common practice as we do not use an FM-layer to estimate correspondences. \mathbf{P}_{yx} is calculated as in [6, 14],

$$\mathbf{P}_{yx} = \text{Softmax}(\tau^{-1} S_C(\mathbf{F}_y, \mathbf{F}_x)), \quad (8)$$

where \mathbf{F}_x and \mathbf{F}_y are the refined features of the full shape \mathcal{X} and partial shape \mathcal{Y} respectively, τ is a temperature hyper-parameter, and

$$\text{Softmax}(\mathbf{X})_{ij} = \frac{\exp(\mathbf{X}_{ij})}{\sum_j \exp(\mathbf{X}_{ij})}. \quad (9)$$

Note, that by choosing the softmax operator we only derive the correspondence from the partial shape \mathcal{Y} to its full version \mathcal{X} , as some regions in \mathcal{X} do not have matching regions in \mathcal{Y} . Obviously, the partiality requirement hinders the use of bidirectional loss functions. As elaborated in the next subsection, some published losses require both \mathbf{P}_{xy} and \mathbf{P}_{yx} , see for example, [39]. We discuss further advantages of the softmax operator compared to the FM-layer in the supplementary.

During training, we use regularization based on the functional map computed from the correspondence matrix by

$$\mathbf{C}_{yx} = \Psi_y^\top \mathbf{A}_y \mathbf{P}_{yx} \Psi_x. \quad (10)$$

As in [14], at inference, we only compute the point-to-point matching $\hat{\mathbf{H}}_{yx}^{\text{iso}}$, by

$$\hat{\mathbf{H}}_{yx}^{\text{iso}}(i) = \underset{j}{\text{argmax}} (\Psi_y \Psi_y^\top \mathbf{A}_y \hat{\mathbf{H}}_{yx} \mathbf{A}_x \Psi_x \Psi_x^\top)_{ij}, \quad (11)$$

where $\hat{\mathbf{H}}_{yx} = \text{NN}(\mathbf{F}_y, \mathbf{F}_x)$, and NN is the nearest neighbor operator.

Unsupervised learning to match. The proposed pipeline operates mainly in the spatial domain rather than the spectral one. The core loss, generalising [20] to partial shapes, is based on the Gromov distance [2, 11, 13, 31]. Our loss measures the quality of the partial shape correspondence matrix \mathbf{P}_{yx} . It is given by,

$$\mathcal{L}_G = \left\| \mathbf{A}_y^{\frac{1}{2}} (\mathbf{P}_{yx}^\top \mathbf{D}_x \mathbf{P}_{yx} - \mathbf{D}_y) \mathbf{A}_y^{\frac{1}{2}} \right\|_F^2, \quad (12)$$

where \mathbf{D}_x and \mathbf{D}_y are the inter-geodesic distance matrices of \mathcal{X} and \mathcal{Y} . The (i, j) entry of \mathbf{D}_x is the geodesic distance on \mathcal{X} between its i -th and j -th vertices.

This loss promotes the preservation of inter-point geodesic distances, which should hold if the surfaces are related by an isometric deformation. This property applies for most articulated objects in nature.

Next, we study the unsupervised regularization loss

$$\mathcal{L}_{\text{orth}} = \left\| \hat{\mathbf{C}}_{yx} \hat{\mathbf{C}}_{yx}^\top - \mathbf{J}_r \right\|_F^2. \quad (13)$$

Here, r is the number of chosen eigenfunctions of the LBO of \mathcal{Y} that correspond to smaller than the largest eigenvalue of the chosen eigenbasis of the LBO of \mathcal{X} . Note, by selecting a small finite number of eigenfunctions, we actually keep the low-frequencies. This measure promotes the orthogonality of functional maps, which should hold when the mapping between shapes is area preserving [4, 33].

Plugging Eq. (10) into Eq. (13) yields

$$\mathcal{L}_{\text{orth}} = \left\| \Psi_y^\top \mathbf{A}_y \mathbf{P}_{xy} \Psi_x \Psi_x^\top \mathbf{P}_{xy}^\top \mathbf{A}_y \Psi_y - \mathbf{J}_r \right\|_F^2. \quad (14)$$

Note, that computing $\mathcal{L}_{\text{orth}}$ requires explicit estimation of the functional map from the soft correspondence. Let, $\tilde{\mathcal{L}}_{\text{orth}}$ be the following modified regularization,

$$\begin{aligned} \tilde{\mathcal{L}}_{\text{orth}} &= \left\| \Psi_y (C_{yx} C_{yx}^\top - \mathbf{J}_r) \Psi_y^\top \right\|_F^2 \\ &= \left\| \mathbf{H}_y \mathbf{P}_{xy} \Psi_x \Psi_x^\top \mathbf{P}_{xy}^\top \mathbf{H}_y - \Psi_y \mathbf{J}_r \Psi_y^\top \right\|_F^2 \\ &= \left\| \mathbf{H}_y \mathbf{P}_{xy} \Psi_x \Psi_x^\top \mathbf{P}_{xy}^\top \mathbf{H}_y - \tilde{\Psi}_y \tilde{\Psi}_y^\top \right\|_F^2, \end{aligned} \quad (15)$$

where $\mathbf{H}_y = \Psi_y \Psi_y^\top \mathbf{A}_y$ and $\tilde{\Psi}_y = \Psi_y \mathbf{J}_r$. \mathbf{H}_y is the projection operator onto the set of low frequency basis functions Ψ_y . Thus, \mathbf{H}_y is a geometric low pass filter.

If \mathbf{P}_{xy} is the true mapping matrix from \mathcal{X} to \mathcal{Y} , then the leading low frequency eigenvectors Ψ_x , restricted to \mathcal{Y} , would be well captured by the low frequencies corresponding to leading LBO eigenvectors on \mathcal{Y} , [1]. In that case,

$$\mathbf{H}_y \mathbf{P}_{xy} \Psi_x = \mathbf{H}_y \Psi_x^y \approx \Psi_x^y, \quad (16)$$

where Ψ_x^y is the restriction of Ψ_x to its part corresponding to \mathcal{Y} . Considering these relations, we propose the *low-pass filter* loss

$$\mathcal{L}_{\text{lpf}} = \left\| \mathbf{P}_{xy} \Psi_x \Psi_x^\top \mathbf{P}_{xy}^\top - \tilde{\Psi}_y \tilde{\Psi}_y^\top \right\|_{\mathcal{Y}, \mathcal{Y}}^2, \quad (17)$$

where $\|\mathbf{B}\|_{\mathcal{Y},\mathcal{Y}} = \text{trace}(\mathbf{A}_y \mathbf{B}^T \mathbf{A}_y \mathbf{B})$ as defined in [2]. When the relation in Eq. (16) holds, we readily have $\mathcal{L}_{\text{lpf}} \approx \mathcal{L}_{\text{orth}}$. Note that \mathcal{L}_{lpf} does not require the computation of the functional map matrix.

Our loss functions are the combination of \mathcal{L}_G coupled with a low pass regularization. We propose to use either

$$\mathcal{L} = \mathcal{L}_G + \lambda_{\text{orth}} \mathcal{L}_{\text{orth}}, \quad (18)$$

which requires the computation of the functional map based on the soft correspondence matrix, or

$$\mathcal{L} = \mathcal{L}_G + \lambda_{\text{lpf}} \mathcal{L}_{\text{lpf}}, \quad (19)$$

which bypasses functional maps altogether. The coefficients λ_{orth} and λ_{lpf} are the relative weights of the $\mathcal{L}_{\text{orth}}$ and \mathcal{L}_{lpf} regularization parts of the loss, respectively.

Refinement method. We follow the refinement process of [14] and use test-time adaptation refinement. Unlike other post-processing methods, this refinement adjusts the weights of the network during test-time for each pair of shapes via a few back-propagation iterations. Unlike [14], we do not change the loss functions or add new ones during the refinement process.

Implementation considerations. For a feature extractor we chose the DiffusionNet model [41]. The chosen input features $\tilde{\mathbf{F}}$ of our model are the 3D xyz -coordinates concatenated with the surface normals. The output dimension of the refined features is $d = 256$, and we use four DiffusionNet blocks, where the inner spectral resolution and number of hidden channels are 128. We set τ to be 0.07 and 0.01 for the SHREC’16 HOLES and CUTS benchmarks respectively. For all experiments we set $\lambda_{\text{orth}} = 0.001$ and $\lambda_{\text{lpf}} = 0.02$. Each shape uses a truncated LBO basis of $k = 200$ eigenfunctions. We train the weights of the models using the ADAM optimizer [23] with a learning rate of 10^{-3} . Similar to [14], we set the number of refinement iterations to 15. The distance matrices are calculated using the Fast-Marching Method [22].

5 Experiments

Here, we compare our method to current partial shape matching techniques.²

Datasets and evaluation method. The reference evaluation benchmarks for partial shape matching are the SHREC’16 CUTS and HOLES datasets [15]. Each dataset contains 8 different shapes under various poses, which are non-rigid semi-isometric deformations of a reference pose. The shapes in the CUTS dataset are obtained by cutting full shapes into two parts by a plane and keeping only one

² Code and data can be found at <https://github.com/ABracha/DirectMatchNet> and <https://github.com/ABracha/PFAUST>.

part, whereas those in the HOLES dataset are formed by applying holes and cuts. The CUTS (resp. HOLES) dataset contains 120 (res. 80) shapes in the training-set, and both datasets have 200 shapes for testing. Empirically, HOLES is believed to be a more challenging benchmark. Indeed, its shapes are trickier to handle for partial shape matching; they involve, on average, longer boundaries while the number of triangles and vertices of the partial shape is smaller. When the amount of holes increases, the feature extractor struggles to produce consistent features common to the partial and full shapes. In addition, we further apply the proposed method to a new benchmark we created by introducing holes to shapes in the FAUST remeshed [37] dataset, in a similar way to SHREC’16 HOLES that were generated by introducing holes to shapes from TOSCA [12]. We produce two datasets which we made publicly available, the PFAUST-M and PFAUST-H standing for *medium* and *hard* difficulties, the latter having more holes than the former. Each dataset contains human shapes from 10 different people, each in 20 poses, 10 for partial shapes and 10 for full shapes. More details on PFAUST are provided in the supplementary. We follow the Princeton protocol [21] for evaluation, using the mean geodesic error, and as in other methods [4, 14], we normalize the square root of the area of the shapes to one.

Baselines. Baseline methods we compare to are grouped into three categories. First, axiomatic methods including PFM [38] and FSP [27]. Second, supervised methods including DPFM [4] and GeomFMaps [16]. Third, unsupervised methods including unsupervised DPFM [4] and RobustFmaps [14]. Supervised DPFM is the current state-of-the-art approach. We re-implemented the loss function of the unsupervised DPFM due to source code unavailability. The results presented for RobustFmaps differ from those shown in the original paper, since we did not pre-train on the TOSCA [12] or other external datasets, see supplementary for details. For post-processing refinement, we used the Zoomout method [29] and the time adaptation refinement [14].

Results. We present quantitative comparisons of the proposed method to the baseline methods in Tabs. 1 and 2 and qualitative ones in Fig. 3, see more examples in the supplementary. On the CUTS dataset our method outperforms preexisting unsupervised approaches and almost bridges the gap between the performance of supervised and unsupervised methods for partial shape matching. Our approach also performs state-of-the-art results on the HOLES benchmark, better than previous methods by a large margin. It even surpasses supervised methods such as DPFM. This margin is consistent with our theoretical analysis regarding the error induced by the FM-layer on partial shape correspondence. The reason is that shapes have larger missing regions in the HOLES than in the CUTS dataset, and as shown in our analysis the error is proportional to the area of the missing parts. Furthermore, in the HOLES dataset, more points are closer to the carved-out boundaries than in the CUTS dataset. This complicates the task of the feature extractor to provide consistent features on the full and partial shapes. The reason is that many more points in the partial shape have significantly different local neighbourhood, whereas the feature extractor relies on the

Table 1: Quantitative results on SHREC’16. The numbers correspond to the mean geodesic error (scaled by 100), and the result using post processing refinement. Best performance is marked in bold and second best is underlined. The results demonstrate the superiority of the proposed method in the unsupervised arena. It outperformed both supervised and unsupervised approaches on the HOLES benchmark, which contains extremely challenging scenarios of shape parts.

Test-set	CUTS		HOLES	
Training-set	CUTS	HOLES	CUTS	HOLES
Axiomatic methods				
PFM [38] → Zoomout	9.7 → 9.0		23.2 → 22.4	
FSP [27] → Zoomout	16.1 → 15.2		33.7 → 32.7	
Supervised methods				
GeomFMaps [16] → Zoomout	12.8 → 10.4	19.8 → 16.7	20.6 → 17.4	15.3 → 13.0
DPFM [4] → Zoomout	3.2 → 1.8	8.6 → 7.4	<u>15.8</u> → 13.9	13.1 → 11.9
Unsupervised methods				
Unsupervised-DPFM [4] → Zoomout	11.8 → 12.8	19.5 → 18.7	19.1 → 18.3	17.5 → 16.2
RobustFMnet [14] → Refinement	16.9 → 10.6	22.7 → 16.6	18.7 → 16.2	23.5 → 18.8
Proposed Orthogonal → Refinement	<u>6.9</u> → 5.6	12.2 → 8.0	14.2 → 10.2	11.4 → 7.9
Proposed LPF → Refinement	7.1 → <u>4.7</u>	8.6 → 5.5	16.4 → <u>11.6</u>	<u>12.3</u> → <u>8.6</u>

Table 2: Quantitative results on the new PFAUST benchmark. The numbers correspond to the mean geodesic error (scaled by 100) without post-processing refinement. Best performance is marked in bold. These results are consistent with those obtained on SHREC’16. As expected, the proposed method performs better than previous unsupervised ones and achieves on par results as the supervised one.

		PFAUST-M	PFAUST-H
Supervised	DPFM [4]	3.0	6.8
Unsupervised	Unsupervised-DPFM [4]	9.3	12.7
	RobustFMnet [14]	7.9	12.4
	Proposed method	5.1	7.9

structure of such neighbourhoods to construct consistent point features. Additionally, we show that the LPF approximation, which allows to avoid functional maps, achieves SOTA results in the unsupervised case of the CUTS dataset. For this approximation to hold, an important argument in its derivation was that applying a partial shape low-pass filter to the low eigenfunctions of the full shape does not alter them by much Sec. 4. This fact holds when the boundaries are either short or close to the symmetry axes of the shapes, as discussed in [38]. The LPF approximation is thus better tailored for datasets with such types of boundaries. The parts in the CUTS dataset exhibit such a behaviour, unlike those in the HOLES dataset.

The results on the PFAUST benchmark when compared to unsupervised approaches that are based on functional maps support our theoretical observations and demonstrate the power of the proposed method. As in the SHREC’16 HOLES dataset, the more holes there are, the better is the proposed approach, compared to alternative unsupervised methods.

Ablation study. The major difference between our approach and other methods is the order and type of operators we apply. Another significant change is the use of different loss functions, in particular \mathcal{L}_G , see Eq. (12). To evaluate the contribution of each of these design choices, we enrich our analysis by performing two types of ablation studies to the proposed method. The first focuses on the loss functions, whereas the second is centred about the pipeline as a whole. We perform the ablation studies on the well established SHREC’16 datasets.

Table 3: Loss ablation study results. The results demonstrate that the performance of our method mostly relies on \mathcal{L}_G , whereas $\mathcal{L}_{\text{orth}}$ and \mathcal{L}_{lpf} act mostly as regularisation.

Loss functions	CUTS	HOLES
$\mathcal{L}_{\text{orth}}$	24	14
\mathcal{L}_{lpf}	13	13
\mathcal{L}_G	10	14
$\mathcal{L}_G + \lambda_{\text{orth}}\mathcal{L}_{\text{orth}}$	6.9	11.4
$\mathcal{L}_G + \lambda_{\text{lpf}}\mathcal{L}_{\text{lpf}}$	7.1	12.3

Loss analysis. In this experiment we deactivate, in turns, each part of the loss function and evaluate the outcome. The results, see Tab. 3, indicate that the Gromov part of the loss, Eq. (12), is significant when operating on the CUTS benchmark, while the terms $\mathcal{L}_{\text{orth}}$ and \mathcal{L}_{lpf} serve mostly for regularization. As for the HOLES benchmark, the \mathcal{L}_G part of the loss suffers from severe distortions of the geodesic distances that characterize that benchmark.

Table 4: Pipeline ablation study results. The correspondence matrix is either estimated from the functional map extracted from a regularized FM-layer or, as we propose, directly from the cosine similarity between features. The results demonstrate the benefit of avoiding the use of an FM-layer in partial shape matching. The effect is even more acute when considering more challenging scenarios, like the HOLES dataset.

	CUTS	HOLES
Regularized FM	8.5	15.8
Proposed method	6.9	11.4

FM-layer. In this experiment we compare the proposed pipeline to a network that includes an FM-layer, where the refined features from the feature extractor serve as inputs to an FM-layer, which is regularized as done in recent functional map related publications, see e.g. [36]. This FM-layer applies a weighted mask based on the eigenvalues of the LBO to the functional map matrix. It was shown to give superior results than using the non-regularised version [26].

The loss \mathcal{L}_G can be computed in the FM-layer pipeline using the correspondence matrix based on the estimated functional map $\mathbf{P}_{yx} = \Psi_x \mathbf{C}_{yx} \Phi_y^\top \mathbf{A}_y$. Note,

that both pipelines are trained using \mathcal{L}_G and $\mathcal{L}_{\text{orth}}$, see Eq. (18), but not \mathcal{L}_{lpf} as it is less suited for FM-layer based approaches.

The proposed approach outperforms ones with FM-layers, see Tab. 4, and performs exceptionally well on the HOLES benchmark. This aligns with our analysis as HOLES poses a harder challenge for the FM-layer than CUTS.

Note, that coupling the metric distortion Gromov-loss \mathcal{L}_G with an FM-layer pipeline yields better results than previous approaches. It demonstrates that \mathcal{L}_G is better suited for partial shape matching than alternative traditional losses.

6 Conclusion

In this paper, we explored the problem of unsupervised partial shape matching. Our analysis unveiled a persistent obstacle when using spectral domains as part of neural network architectures rather than just as part of loss functions. Our study showed that using the functional maps framework via least-squares estimation inevitably introduces errors when shape partiality is involved, an issue that escalates with the increase in missing areas of the matched parts.

As a result of our findings, we proposed a novel approach that deviates from conventional schools of thought. Rather than estimating the functional map on the computed features via least squares, and then extract the correspondences from it, we directly compute the correspondence between partial and full shapes through feature similarity. The proposed method achieves superior results compared to previous methods on the SHREC'16 and the new PFAUST datasets, significantly surpassing existing unsupervised methods in the field of partial shape matching of non-rigid shapes. The compelling results endorse the efficacy of our method and proves that spectral representations can be suited for partial shape correspondence when applied properly.

The proposed approach obviously involves limitations. First, in its current form, it is not designed for part-to-part shape correspondence. In that case, \mathcal{L}_G needs to be revisited, since, when defining \mathbf{P} we assume that each point in the partial shape has a corresponding point on the full shape. Another issue is that geodesic distances can be significantly altered by cuts and holes, which affects the contribution of \mathcal{L}_G . And yet, it is far more stable than current unsupervised losses based on functional maps, as demonstrated in the ablation study. This issue is revisited in a follow-up paper [7] where we introduce the *wormhole criterion* that filters out potentially altered geodesic distances. An additional limitation is that partial and full shapes should be of the same scale. This is a common assumption in partial shape matching that is made by most methods, including those based on functional maps. The missing parts also alter the LBO eigenfunctions, which affects DiffusionNet as it uses LBO eigendecomposition to learn heat diffusion on the shape. In practice, DiffusionNet mitigates this by learning to shorten the time scale in the presence of missing parts. Finally, designing a scale-invariant approach, for example, by using a scale-invariant metric, or opting for anisotropic or asymmetric metrics [42], is of great interest that goes beyond the scope of this paper.

References

1. Aflalo, Y., Brezis, H., Kimmel, R.: On the optimality of shape and data representation in the spectral domain. *SIAM Journal on Imaging Sciences* **8**(2), 1141–1160 (2015) [5](#), [9](#)
2. Aflalo, Y., Dubrovina, A., Kimmel, R.: Spectral generalized multi-dimensional scaling. *International Journal of Computer Vision* **118**, 380–392 (2016) [9](#), [10](#)
3. Attaiki, S., Ovsjanikov, M.: Understanding and improving features learned in deep functional maps. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 1316–1326 (2023) [4](#), [7](#)
4. Attaiki, S., Pai, G., Ovsjanikov, M.: DPFM: Deep partial functional maps. In: *2021 International Conference on 3D Vision (3DV)*. IEEE (Dec 2021) [2](#), [4](#), [5](#), [6](#), [9](#), [11](#), [12](#)
5. Bensaïd, D., Bracha, A., Kimmel, R.: Partial shape similarity by multi-metric hamiltonian spectra matching. In: *International Conference on Scale Space and Variational Methods in Computer Vision*. pp. 717–729. Springer (2023) [2](#)
6. Bensaïd, D., Rotstein, N., Goldenstein, N., Kimmel, R.: Partial matching of non-rigid shapes by learning piecewise smooth functions. In: *Computer Graphics Forum*. vol. 42. Wiley Online Library (2023) [2](#), [8](#)
7. Bracha, A., Dagès, T., Kimmel, R.: Wormhole loss for partial shape matching. *Advances in Neural Information Processing Systems* (2024) [14](#)
8. Bracha, A., Halim, O., Kimmel, R.: Shape correspondence by aligning scale-invariant LBO eigenfunctions. In: *Eurographics Workshop on 3D Object Retrieval*. The Eurographics Association (2020) [3](#)
9. Bronstein, A.M., Bronstein, M.M., Bruckstein, A.M., Kimmel, R.: Partial similarity of objects, or how to compare a centaur to a horse. *International Journal of Computer Vision* **84**, 163–183 (2009) [2](#)
10. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Efficient computation of isometry-invariant distances between surfaces. *SIAM Journal on Scientific Computing* **28**(5), 1812–1836 (2006) [2](#)
11. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Generalized multidimensional scaling: a framework for isometry-invariant partial surface matching. *Proceedings of the National Academy of Sciences* **103**(5), 1168–1172 (2006) [2](#), [3](#), [9](#)
12. Bronstein, A.M., Bronstein, M.M., Kimmel, R.: Numerical geometry of non-rigid shapes. Springer Science & Business Media (2008) [11](#)
13. Bronstein, A.M., Bronstein, M.M., Kimmel, R., Mahmoudi, M., Sapiro, G.: A Gromov-Hausdorff framework with diffusion geometry for topologically-robust non-rigid shape matching. *International Journal of Computer Vision* **89**(2-3), 266–286 (2010) [2](#), [9](#)
14. Cao, D., Roetzer, P., Bernard, F.: Unsupervised learning of robust spectral shape matching. *ACM Transactions on Graphics (TOG)* (2023) [2](#), [4](#), [5](#), [7](#), [8](#), [10](#), [11](#), [12](#)
15. Cosmo, L., Rodola, E., Bronstein, M.M., Torsello, A., Cremers, D., Sahillioglu, Y., et al.: Shrec’16: Partial matching of deformable shapes. *Proc. 3DOR* **2**(9), 12 (2016) [2](#), [10](#)
16. Donati, N., Sharma, A., Ovsjanikov, M.: Deep geometric functional maps: Robust feature learning for shape correspondence. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8592–8601 (2020) [2](#), [3](#), [11](#), [12](#)

17. Eisenberger, M., Novotny, D., Kerchenbaum, G., Labatut, P., Neverova, N., Cremers, D., Vedaldi, A.: Neuromorph: Unsupervised shape interpolation and correspondence in one go. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 7473–7483 (2021) [3](#)
18. Elad, A., Kimmel, R.: On bending invariant signatures for surfaces. *IEEE Transactions on pattern analysis and machine intelligence* **25**(10), 1285–1295 (2003) [3](#)
19. Ginzburg, D., Raviv, D.: Cyclic functional mapping: Self-supervised correspondence between non-isometric deformable shapes. In: Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part V 16. pp. 36–52. Springer (2020) [3](#)
20. Halimi, O., Litany, O., Rodola, E., Bronstein, A.M., Kimmel, R.: Unsupervised learning of dense shape correspondence. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition. pp. 4370–4379 (2019) [2](#), [3](#), [5](#), [7](#), [9](#)
21. Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. *ACM transactions on graphics (TOG)* **30**(4), 1–12 (2011) [11](#)
22. Kimmel, R., Sethian, J.A.: Computing geodesic paths on manifolds. *Proceedings of the national academy of Sciences* **95**(15), 8431–8435 (1998) [10](#)
23. Kingma, D.P., Ba, J.: Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014) [10](#)
24. Li, L., Attaiki, S., Ovsjanikov, M.: SRFeat: Learning locally accurate and globally consistent non-rigid shape correspondence. In: International Conference on 3D Vision (3DV). IEEE (2022) [3](#), [5](#)
25. Li, Q., Liu, S., Hu, L., Liu, X.: Shape correspondence using anisotropic chebyshev spectral cnns. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 14658–14667 (2020) [3](#)
26. Litany, O., Remez, T., Rodola, E., Bronstein, A., Bronstein, M.: Deep functional maps: Structured prediction for dense shape correspondence. In: Proceedings of the IEEE international conference on computer vision. pp. 5659–5667 (2017) [1](#), [2](#), [3](#), [5](#), [13](#)
27. Litany, O., Rodolà, E., Bronstein, A.M., Bronstein, M.M.: Fully spectral partial shape matching. In: *Computer Graphics Forum*. vol. 36, pp. 247–258. Wiley Online Library (2017) [11](#), [12](#)
28. Marin, R., Rakotosaona, M.J., Melzi, S., Ovsjanikov, M.: Correspondence learning via linearly-invariant embedding. *Advances in Neural Information Processing Systems* **33**, 1608–1620 (2020) [3](#)
29. Melzi, S., Ren, J., Rodolà, E., Sharma, A., Wonka, P., Ovsjanikov, M.: Zoomout: spectral upsampling for efficient shape correspondence. *ACM Transactions on Graphics (TOG)* **38**(6), 1–14 (2019) [11](#)
30. Mémoli, F.: Some properties of Gromov–Hausdorff distances. *Discrete & Computational Geometry* **48**, 416–440 (2012) [3](#)
31. Mémoli, F., Sapiro, G.: Distance functions and geodesics on submanifolds of \mathbb{R}^d and point clouds. *SIAM Journal on Applied Mathematics* **65**(4), 1227–1260 (2005) [2](#), [9](#)
32. Mémoli, F., Sapiro, G.: A theoretical and computational framework for isometry invariant recognition of point cloud data. *Foundations of Computational Mathematics* **5**, 313–347 (2005) [2](#)
33. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (ToG)* **31**(4), 1–11 (2012) [1](#), [4](#), [9](#)
34. Pazi, I., Ginzburg, D., Raviv, D.: Unsupervised scale-invariant multispectral shape matching. *arXiv preprint arXiv:2012.10685* (2020) [3](#)

35. Rampini, A., Tallini, I., Ovsjanikov, M., Bronstein, A.M., Rodolà, E.: Correspondence-free region localization for partial shape similarity via hamiltonian spectrum alignment. In: 2019 International Conference on 3D Vision (3DV). pp. 37–46. IEEE (2019) [2](#)
36. Ren, J., Panine, M., Wonka, P., Ovsjanikov, M.: Structured regularization of functional map computations. In: Computer Graphics Forum. vol. 38, pp. 39–53. Wiley Online Library (2019) [6](#), [13](#)
37. Ren, J., Poulenard, A., Wonka, P., Ovsjanikov, M.: Continuous and orientation-preserving correspondences via functional maps. ACM Transactions on Graphics (ToG) **37**(6), 1–16 (2018) [11](#)
38. Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D.: Partial functional correspondence. In: Computer Graphics Forum. vol. 36, pp. 222–236. Wiley Online Library (2017) [4](#), [6](#), [11](#), [12](#)
39. Roufousse, J.M., Sharma, A., Ovsjanikov, M.: Unsupervised deep learning for structured shape matching. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 1617–1627 (2019) [2](#), [3](#), [5](#), [7](#), [8](#)
40. Sharma, A., Ovsjanikov, M.: Weakly supervised deep functional maps for shape matching. Advances in Neural Information Processing Systems **33**, 19264–19275 (2020) [2](#), [3](#)
41. Sharp, N., Attaiki, S., Crane, K., Ovsjanikov, M.: Diffusionnet: Discretization agnostic learning on surfaces. ACM Transactions on Graphics (TOG) **41**(3), 1–16 (2022) [2](#), [3](#), [10](#)
42. Weber, S., Dagès, T., Gao, M., Cremers, D.: Finsler-laplace-beltrami operators with application to shape analysis. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 3131–3140 (2024) [14](#)