Synchronous Diffusion for Unsupervised Smooth Non-Rigid 3D Shape Matching

Dongliang Cao¹, Zorah Lähner^{1,2}, and Florian Bernard¹

¹University of Bonn ²Lamarr Institute

Abstract. Most recent unsupervised non-rigid 3D shape matching methods are based on the functional map framework due to its efficiency and superior performance. Nevertheless, respective methods struggle to obtain spatially smooth pointwise correspondences due to the lack of proper regularisation. Inspired by the success of message passing on graphs, in this work we propose a *synchronous diffusion process* which we use as regularisation to achieve smoothness in non-rigid 3D shape matching problems. The intuition of synchronous diffusion is that diffusing the same input function on two different shapes results in consistent outputs. Using different challenging datasets, we demonstrate that our novel regularisation can substantially improve the state-of-the-art in shape matching, especially in the presence of topological noise.

Keywords: 3D shape matching \cdot functional maps \cdot 3D deep learning \cdot heat diffusion \cdot unsupervised learning \cdot deep shape matching



Fig. 1: Left: We present an unsupervised regularisation based on synchronous diffusion for spatially smooth non-rigid 3D shape matching. This is based on the motivation that diffusing a function f on shape \mathcal{M} should lead to comparable results when diffusing $\pi(f)$ (i.e. f transferred to shape \mathcal{N} using the map π) on shape \mathcal{N} . Coloured points are used to illustrate a (subset of) random function values. **Right:** Our approach can be applied in a broad range of challenging scenarios, including topological noise, nonisometry and partiality.

1 Introduction

Finding correspondences between non-rigid 3D shapes is a fundamental problem in computer vision and computer graphics with a diverse range of applications, including texture transfer [23], pose transfer [82,83] and statistical shape analysis [26,54,58], to name just a few. Even though non-rigid 3D shape matching is a long-standing problem and has been studied for decades [22,77], finding accurate pointwise correspondence remains challenging, especially in the presence of topological noise and partiality, which are common in real-world 3D scan data.

Inspiringly, recent deep functional map methods witnessed great improvements for non-rigid 3D shape matching including near-isometric shape matching [4], non-isometric shape matching [24, 53], multi-shape matching [13, 31, 84] as well as partial shape matching [5]. Moreover, a recent deep functional map framework [16,17] achieves state-of-the-art matching performance in diverse settings. Despite the success of deep functional map methods, they map shapes in the low-frequency spectral domain and thus often lead to locally inconsistent pointwise correspondences. To address this limitation, many works [3, 15, 29, 48] attempt to additionally register shapes in the spatial domain to ensure spatial smoothness of the obtained pointwise correspondences. Nevertheless, these methods require that the given shapes are already rigidly aligned and they are sensitive to topological noise and not directly applicable for partial shapes.

In contrast, our objective is to design a universal regularisation that encourages spatially smooth matchings and can be easily integrated into existing non-rigid 3D shape matching methods [4,16,84]. Prior works that ensure smoothness [33,71] typically require computationally expensive iterative optimisation and thus cannot be integrated straightforwardly into learning frameworks. To this end, inspired by message passing mechanisms used in the graph isomorphism test [93] and for deep graph matching [34], we propose *synchronous diffusion* as smoothness regularisation for non-rigid 3D shape matching. In contrast to message passing mechanisms that aggregate information based on local connectivity (e.g. one-hop neighbourhood), the diffusion process [50] utilises smooth communication [21,39], which is more robust to discretisation [80,89,99] and thus more suitable for 3D shape analysis.

In the following we formally define synchronous diffusion:

Definition 1 (Synchronous diffusion). Let \mathcal{M} and \mathcal{N} be two domains (e.g. 3D shapes) and let π be an invertible mapping from functions defined on \mathcal{M} to \mathcal{N} . We say that the diffusion processes $h^{\mathcal{M}}$ (diffusion on \mathcal{M}) and $h^{\mathcal{N}}$ (diffusion on \mathcal{N}) are synchronous (w.r.t. π) if for any function f defined on \mathcal{M} the following property holds: $\pi(h^{\mathcal{M}}(f)) = h^{\mathcal{N}}(\pi(f))$.

Intuitively, the idea of *synchronous diffusion* is to apply a diffusion process to the 'same' function on different domains (3D shapes in our case), see Fig. 1 left for an illustration. Our key observation is that synchronous diffusion favours the mapping to be globally continuous and locally smooth. Therefore, we propose a novel regularisation based on synchronous diffusion to encourage spatially smooth pointwise correspondences.

Table 1: Comparison to related existing methods. Our method is the first approach that is based on multiscale *synchronous* heat diffusion, tailored specifically for unsupervised smooth 3D shape matching, while being more efficient than previous kernel matching methods based on quadratic assignment formulations.

	Heat diffusion	Efficient	Multiscale	Smooth	Synchronous
Heat Kernel Signature [12]	V	1	1	×	×
DiffusionNet [80]	1	×	1	×	×
Dirichlet Energy [33, 59]	×	1	×	×	×
Kernel Matching [91]	×	×	1	×	×
Unsup FMNet [43]	×	×	×	×	×
Heat Kernel UDFMs [7]	×	×	1	1	×
Ours	V	1	1	1	1

Tab. 1 compares our proposed method with the most relevant existing approaches. In contrast to prior works, our method is the first approach that is specifically designed for unsupervised smooth 3D shape matching with a unique combination of desirable properties. In Sec. 5 we analyse our proposed regularisation and build connections with well-established heat kernel matching [91] and the Weisfeiler-Leman graph isomorphism test [93]. In our experiments, we demonstrate that our regularisation can be integrated into the state-of-the-art deep functional map method [16] and substantially improve its performance. We summarise our main contributions as follows:

- For the first time we propose a simple yet efficient regularisation for unsupervised 3D shape matching that encourages spatially smooth pointwise correspondences under a broad range of challenging scenarios.
- To this end, we introduce a synchronous diffusion process that penalises local mismatches without requiring ground truth correspondences.
- We demonstrate that our regularisation substantially improves the state-ofthe-art deep functional map method in various challenging problem settings.

2 Related Work

2.1 Functional Map Methods for Non-Rigid 3D Shape Matching

Non-rigid 3D shape matching is a well-studied problem in computer vision and graphics [86, 90]. Rather than providing an exhaustive literature survey, in the following we will focus on reviewing the most relevant methods that are based on the functional map framework [66]. Instead of directly finding pointwise correspondences, which is often formulated as expensive combinatorial optimisation problem [9, 37, 45, 73, 94], the functional map framework finds correspondences in the functional domain. Here, the correspondence relationship can be encoded into a small matrix, namely the functional map, by using a series of truncated basis functions [66] (typically the first k Laplacian eigenfunctions [52]). Due to its simplicity and efficiency, the functional map framework has been extended in numerous follow-up works, e.g. in terms of improving the

accuracy or robustness [32, 61, 70], as well as extending its application to partial shape matching [57, 72], non-isometric shape matching [65, 69, 71], multishape matching [36, 46, 47] and matching with non-unique solutions [68]. Nevertheless, axiomatic functional map methods typically rely on handcrafted features (e.g. HKS [12], WKS [6]). To further improve the matching performance, many learning-based methods propose to learn the features directly from the training data, either with supervision from ground-truth correspondences [25, 56] or without supervision based on isometry assumptions [43, 75]. However, deep functional map methods often obtain local erroneous pointwise correspondences due to the use of low-frequency basis functions and the lack of proper regularisation. To address this limitation, we propose an efficient regularisation based on the diffusion process that encourages spatially smooth correspondences.

2.2 Message Passing for Graph Analysis

The message passing mechanism is an important tool for graph analysis. Despite its simplicity (i.e. aggregate neighbouring information and update the node features), message passing mechanisms are widely used for graph analysis. Use cases range from traditional graph kernel analysis [74,81] for graph embedding, to recent graph neural networks (GNNs) [96,97] for graph representation learning [40,44,64]. In the context of the graph isomorphism problem (i.e. test whether two graphs have the same topological structure), the Weisfeiler-Leman test (WL test) [93] is a well-known heuristic test based on message passing and has recently gained popularity in the context of graph neural networks [35, 63]. The WL test iteratively checks for node label consistency between two graphs in each message passing step. A more comprehensive review about message passing for graph analysis can be found in [63]. The most relevant application of message passing to our work is deep graph matching. Recent deep graph matching methods [34,38,87,98] find node correspondences based on pairwise node feature similarities. To encourage neighbourhood consensus (i.e. neighbouring nodes are still neighbouring nodes after matching), DGMC [34] proposes a method to iteratively improve the neighbourhood consensus in a subsequent second refinement stage via a differentiable validator (i.e. an additional GNN) for graph isomorphism based on the WL test heuristic. Similar to deep graph matching methods, many recent deep functional map methods [2, 14, 16] also obtain pointwise correspondences based on pairwise feature similarities. To this end, our work is inspired by the idea of using message passing for consensus graph matching in DGMC. However, instead of using the computationally expensive second refinement stage, we propose an efficient regularisation based on synchronous diffusion that is seamlessly integrated into the current deep shape matching framework [16], which in turn significantly accelerates both training and inference, while having better scalability to high-resolution shapes.

2.3 Diffusion for 3D Shape Analysis

The diffusion process, which originates from the analysis of heat transfer, enables various applications in the context of 3D shape analysis, including shape segmen-

tation [76], shape matching [66, 67], multi-resolution shape representation [52], and geometric deep learning [80], among many other applications (see [89] for an overview). In the context of shape matching, heat kernel matching [91] is highly relevant to our work. In this work, a feature-based linear assignment problem (LAP) and a kernel-based quadratic assignment problem (QAP) are solved together in an iterative way. Unlike previous work that uses geodesic distance kernels [92], heat kernels (i.e. the solutions to the heat diffusion equations) are chosen in this method due to their computational efficiency, their bias towards *local* correspondences, and because they are more robust to topological noise [91]. The main difference between our regularisation and heat kernel matching is that heat kernel matching measures the distances in the kernel space, while our regularisation penalises unsmooth correspondences in the feature space, which is more efficient [34]. Moreover, the heat kernel matching method can be considered as a special case of our synchronous diffusion regularisation as shown in Sec. 5.

3 Background

3.1 Diffusion Process

In the continuous case, the diffusion process of a scalar field u is expressed by the heat equation, i.e.

$$\frac{\partial u}{\partial t} = \Delta u,\tag{1}$$

where Δ is the Laplacian operator. The action of diffusion can be represented via the heat operator h_t , which is applied to some initial distribution u_0 and produces the diffused distribution u_t at time t [80]. The solution of the diffusion equation can be expressed as

$$u_t = h_t \left(u_0 \right) = \exp(t\Delta) u_0, \tag{2}$$

where $\exp(\cdot)$ is the operator exponential. For discrete domains, Δ is replaced by the Laplacian matrix $\mathbf{L} \in \mathbb{R}^{n \times n}$. With $\mathbf{M} \in \mathbb{R}^{n \times n}$ being the mass matrix such that $\mathbf{M}^{-1}\mathbf{L} \approx -\Delta$, the diffusion process can be approximated (using the backward Euler method) [80] as

$$h_t(u) := (\mathbf{M} + t\mathbf{L})^{-1}\mathbf{M}u.$$
(3)

However, computing the inverse of the (large) Laplacian matrix is both computationally expensive and numerically unstable [62]. In practice, spectral acceleration is often used [8,80]. Here, the diffusion process is approximated with the first k eigenvectors $\mathbf{\Phi} \in \mathbb{R}^{n \times k}$ and the diagonal eigenvalue matrix $\mathbf{\Lambda} \in \mathbb{R}^{k \times k}$ of the Laplacian matrix, i.e.

$$h_t(u) := \mathbf{\Phi} \exp\left(-t\mathbf{\Lambda}\right) \mathbf{\Phi}^\top u, \quad \text{where} \quad \exp\left(-t\mathbf{\Lambda}\right) = \begin{pmatrix} e^{-\lambda_1 t} & \\ & \ddots & \\ & & e^{-\lambda_k t} \end{pmatrix}. \quad (4)$$

4 Synchronous Diffusion for Smooth Matching



Fig. 2: Illustration of our synchronous diffusion process for smooth matching. First, we transfer the function $\mathbf{F}_{\mathcal{M}}$ from shape \mathcal{M} to shape \mathcal{N} and perform synchronous diffusion on both shapes. Afterwards, the diffused function $\mathbf{F}_{\mathcal{N}}(t)$ from shape \mathcal{N} is transferred back to shape \mathcal{M} . The difference between the two diffused functions is used to penalise spatially unsmooth pointwise correspondences.

We propose a synchronous diffusion regulariser based on Definition 1 that can be applied in any learning-based framework that outputs pointwise correspondences. Typically, initial pointwise correspondences are obtained based on local feature matching [14, 16, 30] without considering *neighbouring information* while using regularisation in the *low-frequency biased* spectral domain [4, 14, 84]. Therefore, the pointwise correspondences inevitably contain many local mismatches because this relies on high-frequency alignment. To address this limitation, we introduce a *synchronous diffusion process* that encourages spatially smooth matchings. Specifically, consider an initial vector-valued function that is sampled from a random distribution $\mathbf{F}_{\mathcal{M}} \in \mathbb{R}^{n_{\mathcal{M}} \times h}$ defined on shape \mathcal{M} . We use the predicted soft pointwise correspondence $\mathbf{\Pi}_{\mathcal{N}\mathcal{M}} \in \mathbb{R}^{n_{\mathcal{N}} \times n_{\mathcal{M}}}$ to transfer the function $\mathbf{F}_{\mathcal{M}}$ from shape \mathcal{M} to shape \mathcal{N} , i.e.

$$\mathbf{F}_{\mathcal{N}} := \mathbf{\Pi}_{\mathcal{N}\mathcal{M}} \mathbf{F}_{\mathcal{M}}.$$
 (5)

Afterwards, the diffusion process is applied synchronously to the initial function defined on both shapes, i.e.

$$\mathbf{F}_{\mathcal{M}}(t) = h_t^{\mathcal{M}}\left(\mathbf{F}_{\mathcal{M}}\right), \quad \mathbf{F}_{\mathcal{N}}(t) = h_t^{\mathcal{N}}\left(\mathbf{F}_{\mathcal{N}}\right) = h_t^{\mathcal{N}}\left(\mathbf{\Pi}_{\mathcal{N}\mathcal{M}}\mathbf{F}_{\mathcal{M}}\right).$$
(6)

Next, we use the predicted pointwise map $\mathbf{\Pi}_{\mathcal{M}\mathcal{N}} \in \mathbb{R}^{n_{\mathcal{M}} \times n_{\mathcal{N}}}$ to transfer the diffused function $\mathbf{F}_{\mathcal{N}}(t)$ from shape \mathcal{N} back to shape \mathcal{M} , i.e.

$$\mathbf{F}_{\mathcal{M}'}(t) := \mathbf{\Pi}_{\mathcal{M}\mathcal{N}} \mathbf{F}_{\mathcal{N}}(t).$$
(7)

Finally, the difference between the function $\mathbf{F}_{\mathcal{M}}(t)$ and the function $\mathbf{F}_{\mathcal{M}'}(t)$ is measured as a regularisation that penalises the non-smoothness of the matching, i.e.

$$E_{\text{diff}} = \left\| \mathbf{F}_{\mathcal{M}}(t) - \mathbf{F}_{\mathcal{M}'}(t) \right\|_{F}^{2} = \left\| h_{t}^{\mathcal{M}}(\mathbf{F}_{\mathcal{M}}) - \mathbf{\Pi}_{\mathcal{M}\mathcal{N}}h_{t}^{\mathcal{N}}(\mathbf{\Pi}_{\mathcal{N}\mathcal{M}}\mathbf{F}_{\mathcal{M}}) \right\|_{F}^{2}.$$
 (8)

Since diffusion exchanges information along neighbourhoods but without relying on exact discretisation of the mesh, the term penalises correspondence that do not preserve local neighbourhoods while indicating at which areas respective violations occur. In practice, the initial function $\mathbf{F}_{\mathcal{M}}$ is random sampled with unit norm along the feature dimension h, which ensures that $\mathbf{F}_{\mathcal{M}}$ is uniformly distributed on the entire shape \mathcal{M} (i.e. with the same amount of heat in each vertex). Instead of manually fixing the diffusion time t [91], we randomly sample different diffusion times for different initial functions individually, i.e.

$$t_i \sim \text{Uniform}(0, T), \quad i \in \{1, \dots h\},$$
(9)

where T is a pre-defined maximum diffusion time that controls the local/global information aggregation. To this end, our regularisation can be expressed as

$$L_{\text{diff}} = \sum_{i=1}^{h} \left\| h_{t_i}^{\mathcal{M}} \left(\mathbf{F}_{\mathcal{M}}^i \right) - \mathbf{\Pi}_{\mathcal{M}\mathcal{N}} h_{t_i}^{\mathcal{N}} \left(\mathbf{\Pi}_{\mathcal{N}\mathcal{M}} \mathbf{F}_{\mathcal{M}}^i \right) \right\|_F^2, \tag{10}$$

where $\mathbf{F}_{\mathcal{M}}^{i}$ denotes the *i*-th column of matrix $\mathbf{F}_{\mathcal{M}}$. As such, we have for each time t_{i} a different scalar-valued initial function. In this case, our loss L_{diff} enables multiscale regularisation by using different diffusion times for each function. The diffusion process on 3D shapes with a large diffusion time is analogous to performing several message passing operations in the graph and vice versa.

5 Theoretical Justification

We will motivate the diffusion loss from two different perspectives: an approximation of a quadratic assignment problem and a continuous version of the Weisfeiler-Leman test.

A common formulation for non-rigid shape correspondence is as a quadratic assignment problem of the form

$$\underset{\mathbf{\Pi}_{\mathcal{N}\mathcal{M}}\in\mathcal{P}_{n}}{\arg\min} \|\mathbf{D}_{\mathcal{M}}-\mathbf{\Pi}_{\mathcal{N}\mathcal{M}}^{\top}\mathbf{D}_{\mathcal{N}}\mathbf{\Pi}_{\mathcal{N}\mathcal{M}}\|$$
(11)

where \mathbf{D}_{\bullet} might be chosen as the heat kernel [91]. In that case, the formulation penalises the differences between heat diffusion based on *initial Dirac functions* on all points. In [85] this is solved using a series of alternating diffusion steps with the objective of promoting neighbourhood preservation. The concept of synchronous diffusion is closely related: instead of using Dirac deltas, we sample from a random distribution in each iteration, which can be interpreted as a form of *matrix sketching* [95]. In this case we decrease the dimension from n (for

Dirac functions in Eq. (11)) to $h \ll n$. Additionally, we show experimentally that using random functions instead of Dirac functions increases both the empirical performance and runtime, see Section 7.

Our loss can also be understood as a continuous version of the Weisfeiler-Leman test. In the continuous domain, manifolds do not have the same discrete notion of neighbourhood (the meshing can only provide a noisy replacement) and the notion of discrete labels in each node is less meaningful. However, conceptually, heat diffusion provides both a weighted notion of neighbourhood as well as a message passing mechanism. In [34] the graph matching problem is solved through synchronous message passing aggregation of permuted random features. Our approach can be seen as replacing the neighbourhood message passing with heat diffusion, i.e.

$$\mathbf{F}_{\mathcal{M}}^{(t)}(i) = \operatorname{aggregate}(\{\{\mathbf{F}_{\mathcal{M}}^{(t-1)}(j) | j \in \mathcal{N}(i)\}\}) \stackrel{\text{becomes}}{\leadsto} \mathbf{F}_{\mathcal{M}}^{(t)}(i) = h_t(\mathbf{F}_{\mathcal{M}}^{(0)})(i).$$
(12)

Instead of a discrete number of message-passing steps, the influence of distant points is determined by the diffusion time parameter t, as was also proposed in [8]; the update step becomes a simple integration. Additionally, the Weisfeiler-Leman distance has been shown to be a lower bound for the Gromov-Wasserstein distance [18, 19] which in turn is related to solving Eq. (11).

6 Experimental Results

In this section we compare our method to previous methods on various shape matching datasets with different settings (including near-isometric, topologically noisy, non-isometric, and partial shape matching). To this end, we integrate our regularisation into the state-of-the-art unsupervised functional map framework of [16]. The details of [16] can be found in our supplementary material.

6.1 Near-isometric shape matching

Datasets. Following prior works [25,79], we evaluate our method on three standard benchmark datasets, namely the FAUST [11], SCAPE [1] and SHREC'19 [60] datasets. To evaluate the robustness against different discretisation, we choose the more challenging remeshed versions from [71]. The FAUST dataset contains 100 shapes, where the train/test split is 80/20. The SCAPE dataset consists of 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 [60] and used only for evaluation not training.

Results. We compare our method to state-of-the-art axiomatic, supervised and unsupervised methods. We use the mean geodesic error [49] as our evaluation metric. The quantitative and qualitative results are summarised in Fig. 3 and Fig. 4. Our method outperforms the previous state of the art, even in comparison to supervised methods. Meanwhile, our method achieves substantially better cross-dataset generalisation compared to existing learning-based methods (see last column in Fig. 3 left).

Train	FAUST	SCAPE F	$\mathbf{FAUST} + \mathbf{SCAPE}$
Test	FAUST	SCAPE	SHREC'19
	Axiomati	c Methods	
BCICP [71]	6.1	11.0	-
ZoomOut [61]	6.1	7.5	-
Smooth Shells [28]	2.5	4.7	-
DiscreteOp [69]	5.6	13.1	-
	Supervise	d Methods	
FMNet [56]	11.0	17.0	-
3D-CODED [41]	2.5	31.0	-
GeomFMaps [25]	2.6	3.0	7.9
TransMatch [88]	1.7	12.0	10.9
1	Unsupervis	sed Methods	
Deep Shells [30]	1.7	2.5	21.1
DUO-FMNet [24]	2.5	2.6	6.4
AttnFMaps [53]	1.9	2.2	5.8
AttnFMaps-Fast [53]	1.9	2.1	6.3
SSCDFM [84]	1.7	2.6	3.8
URSSM [16]	1.6	1.9	4.6
Ours	1.5	1.8	3.4

Fig. 3: Left: Near-isometric shape matching and cross-dataset generalisation on FAUST, SCAPE and SHREC'19 datasets. The mean geodesic error [49] is used as quantitative evaluation metric. The **best** results in each column are highlighted. Right: Qualitative results of our method on the challenging SHREC'19 dataset trained on FAUST and SCAPE datasets. Our method outperforms existing state-of-the-art methods and demonstrates superior cross-dataset generalisation ability.



Fig. 4: Near-isometric shape matching and cross-dataset generalisation on FAUST, SCAPE and SHREC'19 datasets where we evaluate the performance on the SHREC'19 dataset trained on FAUST and SCAPE datasets. Proportion of correct keypoints (PCK) curves and corresponding area under curve (scores in the legend) of our method in comparison to the existing unsupervised state-of-the-art methods.

6.2 Matching with topological noise

Datasets. The topology of meshes can undergo significant inconsistencies, e.g. due to self-intersections of separate parts of real-world scanned objects. Such topological noise presents a great challenge, both for matching methods based on functional maps as it distorts the intrinsic shape geometry [51], as well as methods solely based on spatial registrations [31]. To evaluate our method for matching with topologically noisy shapes, we use the TOPKIDS dataset [51].

			Source	AttnFMaps
Geo. error (×100)	TOPKIDS	Fully intrinsic		
Axio	matic Methods	3		126
ZoomOut [61]	33.7	✓	7 7 5 F4 F5 B 7 7 5	
Smooth Shells [28]	11.8	×		K BD
DiscreteOp [69]	35.5	\checkmark		
Unsup	ervised Metho	ds		- 🦉 🌾
UnsupFMNet [43]	38.5	\checkmark		
SURFMNet [75]	48.6	✓	URSSM	Ours
WSupFMNet [79]	47.9	✓		
Deep Shells [30]	13.7	×		
NeuroMorph [29]	13.8	×	2 22 23	T55 P*
AttnFMaps [53]	23.4	✓	03 17 23	14 15 ED
AttnFMaps-Fast [53]	28.5	✓		84 G5 C
URSSM [16]	9.2	\checkmark		
Ours	5.4	\checkmark	. 🔰 🌄	- 🚺 🌄
			- X	7

Fig. 5: Results on the TOPKIDS dataset. We distinguish fully intrinsic methods (i.e. methods solely based on functional maps) from the methods that rely on additional extrinsic information (e.g. rigid alignment). Our method substantially outperforms all existing methods, even compared to methods relying on extrinsic information.

Results. We compare our method with prominent axiomatic methods and unsupervised methods. The quantitative results are summarised in Fig. 5. Our method substantially outperforms existing methods, even in comparison to methods relying on additional extrinsic information. Fig. 6 left shows PCK curves of our method in comparison to existing state-of-the-art unsupervised methods and demonstrates the superior performance of our method even under large topological noise. The main reason is the relative robustness of the diffusion process w.r.t. small topological noise [91] when choosing small diffusion times t and the multiscale regularisation based on randomly sampled diffusion time (see comparison to fixed diffusion time in Sec. 7). By choosing a small maximum diffusion time T in Eq. (9), we ensure our regularisation to focus more on the local neighbourhood and thus to be less effected by the global distortion caused from topological noisy parts. More details about the choice of diffusion time can be found in the supplementary materials.

6.3 Non-isometric shape matching

Datasets. In the context of non-isometric shape matching, we consider the SMAL [101] dataset and the DT4D-H [59] dataset. The SMAL dataset contains 49 animal shapes of 8 species. Similarly to previous work [16,24], five species are used for training and three different species are used for testing (i.e. 29/20 shapes for the train/test split). The DT4D-H dataset based on DeformingThings4D [55] is firstly introduced in [59]. Following prior evaluations [16,53], nine classes of shapes are used for train/test split.



Fig. 6: Left: Shape matching with topological noise on the TOPKIDS dataset. Our method significantly outperforms existing unsupervised methods. **Right**: Nonisometric shape matching on SMAL and DT4D-H datasets. Our method demonstrates state-of-the-art performance in both challenging scenarios.



Fig. 7: Results on SMAL and DT4D-H datasets. Our method outperforms existing methods for challenging non-isometric shape matching on both SMAL and DT4D-H inter-class datasets and shows comparable performance on near-isometric shape matching on the DT4D-H intra-class dataset.

Results. Fig. 7 summarises the qualitative and quantitative matching results on the SMAL and DT4D-H datasets. In the context of non-isometric inter-class shape matching, our approach outperforms the existing state-of-the-art methods on both challenging datasets, even in comparison to supervised methods. Mean-while, our method demonstrates comparable and near-perfect matching results for intra-class matching on the DT4D-H dataset. Fig. 6 right shows the PCK curves and the corresponding AUC of our method compared to existing state-of-the-art methods. Compared to the results on topological noisy dataset, the performance improvement is less significant – one of the main reasons is that large non-isometry leads to inconsistent diffusion processes on two shapes.

6.4 Partial shape matching

Datasets. We evaluate our method on the SHREC'16 partial dataset [20]. The dataset contains 200 training shapes and 400 test shapes, with eight different

classes (including humans and animals). Each class has a complete template shape to which the other partial shapes are matched. The dataset is divided into two subsets, namely CUTS (missing a large part) with a 120/200 train/test split and HOLES (missing many small parts) with a 80/200 train/test split.



Fig. 8: Qualitative comparison between our method and existing state-of-the-art methods on the SHREC'16 dataset. Our method leads to spatially smoother matchings.



Fig. 9: Results on the SHREC'16 partial dataset. Our method substantially outperforms state-of-the-art methods, even in comparison to the supervised approaches. Meanwhile, our method demonstrates better cross-dataset generalisation ability.

Results. Fig. 8 and Fig. 9 provide the qualitative and quantitative results of our method in comparison to previous axiomatic, supervised and unsupervised methods that are suitable for partial shape matching. Similar to prior works [13, 16], we pre-train the feature extractor on complete shapes due to the limited amount of training shapes [27] and use vertex position as input features [5]. Compared to existing methods, our approach is more robust to partiality and enables spatially smoother pointwise correspondences.

7 Ablation Study

In this section we conduct ablative experiments to analyse the effectiveness of our regularisation. For all ablation experiments, we choose the most challenging TOPKIDS dataset for evaluation. We study a total of six different ablative settings and explain these in detail in the following.

(a) Removal of our regularisation. We remove our regularisation L_{diff} defined in Eq. (10) for training. In this case, no spatial smoothness regularisation is taken into consideration.

(b) Fixed single diffusion time. We replace our randomly sampled diffusion time defined in Eq. (9) by a heuristically manually selected single diffusion time applied to the whole distribution $\mathbf{F}_{\mathcal{M}}$ as defined in Eq. (8).

(c) Pre-defined initial distributions. In Sec. 4, we choose randomly sampled initial distributions that we diffuse. In this ablative setting, we consider a predefined initial distributions. To this end, we choose to use the first h Laplacian eigenfunctions $\Phi_{\mathcal{M}} \in \mathbb{R}^{n_{\mathcal{M}} \times h}$, which corresponds to the smoothest orthogonal basis w.r.t. the Dirichlet energy [10].

(d) Comparison to heat kernel matching. In Sec. 5 we conceptually relate our regularisation to heat kernel matching [91], which we now study empirically. To this end, we replace L_{diff} defined in Eq. (10) by the quadratic assignment formulation defined in Eq. (11). Similarly, we consider a multiscale setting, i.e.

$$L_{\text{kernel}} = \sum_{i=1}^{h} \left\| \mathbf{D}_{\mathcal{M}}^{t_i} - \mathbf{\Pi}_{\mathcal{N}\mathcal{M}}^{\top} \mathbf{D}_{\mathcal{N}}^{t_i} \mathbf{\Pi}_{\mathcal{N}\mathcal{M}} \right\|_F^2,$$
(13)

where $\mathbf{D}_{\bullet}^{t_i} = \mathbf{\Phi}_{\bullet} \exp\left(-t_i \mathbf{\Lambda}_{\bullet}\right) \mathbf{\Phi}_{\bullet}^{\top}$ is the heat kernel, and t_i is the randomly sampled diffusion time defined in Eq. (9). We emphasise that heat kernel matching is much more computationally expensive, since the matrix size in L_{kernel} is $n_{\mathcal{M}} \times n_{\mathcal{N}}$, opposed to $n_{\mathcal{M}} \times h$ in our regularisation.

(e) Comparison to Dirichlet energy. Many prior works [16, 33, 59] utilise the Dirichlet energy to prompt smooth pointwise correspondences. The Dirichlet energy is expressed as

$$L_{\rm D} = \left\| \mathbf{\Pi}_{\mathcal{N}\mathcal{M}} \mathbf{V}_{\mathcal{M}} \right\|_{\mathbf{L}_{\mathcal{N}}}^{2} = \operatorname{tr}((\mathbf{\Pi}_{\mathcal{N}\mathcal{M}} \mathbf{V}_{\mathcal{M}})^{\top} \mathbf{L}_{\mathcal{N}}(\mathbf{\Pi}_{\mathcal{N}\mathcal{M}} \mathbf{V}_{\mathcal{M}})),$$
(14)

where $\mathbf{V}_{\mathcal{M}}$ is the matrix of vertex locations of shape \mathcal{M} . However, the optimal solution of $L_{\rm D}$ leads to a degenerate result, i.e. all vertices of shape \mathcal{M} are matched to the same vertex of shape \mathcal{N} . In contrast, our regularisation encourages bijective/high-coverage correspondences.

(f) Comparison to cycle-consistency. We point out that if $t_i = 0$ for all distributions (i.e. no diffusion), our regularisation degrades to the commonly used cycle-consistency loss [42, 100], i.e.

$$L_{\text{diff}} = \|\mathbf{F}_{\mathcal{M}} - \mathbf{\Pi}_{\mathcal{M}\mathcal{N}} (\mathbf{\Pi}_{\mathcal{N}\mathcal{M}} \mathbf{F}_{\mathcal{M}})\|_{F}^{2}, \qquad (15)$$

since $h_0(u) = u$ in Eq. (3).

Results. We summarise the quantitative results in Tab. 2. Compared to the use of a fixed diffusion time (see column (a)), we notice that our random diffusion time sample strategy enables multiscale regularisation, and thus obtains better performance. Compared to the pre-defined $\mathbf{F}_{\mathcal{M}}$ (see column (b)), we observe

random initial distributions is better. In comparison to kernel matching (KM), Dirichlet energy, and cycle-consistency (column (d), column (e), and column (f)), our regularisation demonstrates better matching performance and significantly improves the baseline (i.e. compared to column (a)).

Table 2: Qualitative results of our ablation study on TOPKIDS dataset. Compared to different settings, our proposed regularisation obtains the best matching performance.

TOPKIDS (a) w/o $L_{ m diff}$	$_{\rm f}$ (b) Fixed t (e	c) Pre-defined l	$\mathbf{F}_{\mathcal{M}}$ (d) KM (e)	Dirichlet	(f) Cycle-con.	Ours
Initial $F_{\mathcal{M}}$ Geo. error	${ m n/a} {9.2}$	Rand. 7.8	$\Phi_{\mathcal{M}}$ 10.7	Dirac 8.9	${n/a} 9.5$	Rand. 9.8	Rand 5.4

8 Discussion and Limitations

For the first time, we propose a simple yet efficient regularisation via synchronous diffusion for unsupervised smooth non-rigid 3D shape matching. Yet, there are some limitations that give rise to further investigations.

One limitation of our work is its scope: although conceptually our synchronous diffusion idea can be applied to other settings (including image keypoint matching [78], graph matching [34] and point cloud matching [14]), in this work we specifically focus on non-rigid 3D shape matching based on deep functional maps. As such, we cannot handle partial-to-partial shape matching [5], which is a limitation that we inherit from the functional map framework [72]. Analysing and extending our regularisation to the challenging and highly relevant case of partial-to-partial shape matching is an interesting direction for future work.

Another limitation of our work is that – from a theoretical point of view – the diffusion process is inconsistent for pairs of shapes with severe non-isometries and topological differences. Yet, experimentally we demonstrate that our regularisation works well in such settings. However, this requires a careful choice of the maximum diffusion time T to adequately balance local and global information exchange (see supplementary for more details).

9 Conclusion

In this work we propose the first unsupervised regularisation based on a synchronous diffusion process for smooth non-rigid 3D shape matching. We demonstrate that our simple regularisation can easily be integrated into a recent stateof-the-art shape matching method and that this leads to a notable performance improvement in different challenging cases, including non-isometric, topologically noisy and partial shapes. Our work is the first attempt to design a diffusion-based smoothness regularisation for unsupervised learning of matching problems and we hope that it will inspire follow-up works across different domains and variants of matching problems.

References

- 1. Anguelov, D., Srinivasan, P., Koller, D., Thrun, S., Rodgers, J., Davis, J.: Scape: shape completion and animation of people. In: ACM SIGGRAPH (2005) 8
- 2. Attaiki, S., Ovsjanikov, M.: Ncp: Neural correspondence prior for effective unsupervised shape matching. In: NeurIPS (2022) 4
- 3. Attaiki, S., Ovsjanikov, M.: Shape non-rigid kinematics (snk): A zero-shot method for non-rigid shape matching via unsupervised functional map regularized reconstruction. In: NIPS (2023) 2
- 4. Attaiki, S., Ovsjanikov, M.: Understanding and improving features learned in deep functional maps. In: CVPR (2023) 2, 6
- Attaiki, S., Pai, G., Ovsjanikov, M.: Dpfm: Deep partial functional maps. In: International Conference on 3D Vision (3DV) (2021) 2, 12, 14
- Aubry, M., Schlickewei, U., Cremers, D.: The wave kernel signature: A quantum mechanical approach to shape analysis. In: ICCV (2011) 4
- Aygün, M., Lähner, Z., Cremers, D.: Unsupervised dense shape correspondence using heat kernels. In: Conference on 3D Vision (3DV) (2020) 3
- Behmanesh, M., Krahn, M., Ovsjanikov, M.: Tide: Time derivative diffusion for deep learning on graphs. In: ICML. PMLR (2023) 5, 8
- 9. Bernard, F., Suri, Z.K., Theobalt, C.: Mina: Convex mixed-integer programming for non-rigid shape alignment. In: CVPR (2020) 3
- Bobenko, A.I., Springborn, B.A.: A discrete laplace-beltrami operator for simplicial surfaces. Discrete & Computational Geometry 38(4), 740–756 (2007) 13
- 11. Bogo, F., Romero, J., Loper, M., Black, M.J.: Faust: Dataset and evaluation for 3d mesh registration. In: CVPR (2014) 8
- 12. Bronstein, M.M., Kokkinos, I.: Scale-invariant heat kernel signatures for non-rigid shape recognition. In: CVPR (2010) 3, 4
- Cao, D., Bernard, F.: Unsupervised deep multi-shape matching. In: ECCV (2022) 2, 12
- Cao, D., Bernard, F.: Self-supervised learning for multimodal non-rigid 3d shape matching. In: CVPR (2023) 4, 6, 14
- 15. Cao, D., Eisenberger, M., El Amrani, N., Cremers, D., Bernard, F.: Spectral meets spatial: Harmonising 3d shape matching and interpolation. In: CVPR (2024) 2
- Cao, D., Roetzer, P., Bernard, F.: Unsupervised learning of robust spectral shape matching. ACM Transactions on Graphics (ToG) (2023) 2, 3, 4, 6, 8, 9, 10, 11, 12, 13
- Cao, D., Roetzer, P., Bernard, F.: Revisiting map relations for unsupervised nonrigid shape matching. In: 2024 International Conference on 3D Vision (3DV). IEEE (2024) 2
- Chen, S., Lim, S., Mémoli, F., Wan, Z., Wang, Y.: Weisfeiler-lehman meets gromov-wasserstein. In: International Conference on Machine Learning (ICML) (2022) 8
- Chen, S., Lim, S., Mémoli, F., Wan, Z., Wang, Y.: The weisfeiler-lehman distance: Reinterpretation and connection with gnns. In: Annual Workshop on Topology, Algebra and Geometry in Machine Learning (TAG-ML) (2023) 8
- Cosmo, L., Rodola, E., Bronstein, M.M., Torsello, A., Cremers, D., Sahillioglu, Y.: Shrec'16: Partial matching of deformable shapes. Proc. 3DOR 2(9), 12 (2016) 11
- 21. Defferrard, M., Bresson, X., Vandergheynst, P.: Convolutional neural networks on graphs with fast localized spectral filtering. NIPS (2016) 2

- 16 Dongliang Cao et al.
- Deng, B., Yao, Y., Dyke, R.M., Zhang, J.: A survey of non-rigid 3d registration. In: Computer Graphics Forum. vol. 41, pp. 559–589. Wiley Online Library (2022)
 2
- Dinh, H.Q., Yezzi, A., Turk, G.: Texture transfer during shape transformation. ACM Transactions on Graphics (ToG) 24(2), 289–310 (2005) 2
- 24. Donati, N., Corman, E., Ovsjanikov, M.: Deep orientation-aware functional maps: Tackling symmetry issues in shape matching. In: CVPR (2022) 2, 9, 10, 11
- Donati, N., Sharma, A., Ovsjanikov, M.: Deep geometric functional maps: Robust feature learning for shape correspondence. In: CVPR (2020) 4, 8, 9, 11, 12
- Egger, B., Smith, W.A., Tewari, A., Wuhrer, S., Zollhoefer, M., Beeler, T., Bernard, F., Bolkart, T., Kortylewski, A., Romdhani, S., et al.: 3d morphable face models—past, present, and future. ACM Transactions on Graphics (ToG) 39(5), 1–38 (2020) 2
- Ehm, V., Gao, M., Roetzer, P., Eisenberger, M., Cremers, D., Bernard, F.: Partialto-partial shape matching with geometric consistency. In: CVPR (2024) 12
- Eisenberger, M., Lahner, Z., Cremers, D.: Smooth shells: Multi-scale shape registration with functional maps. In: CVPR (2020) 9, 10, 11
- Eisenberger, M., Novotny, D., Kerchenbaum, G., Labatut, P., Neverova, N., Cremers, D., Vedaldi, A.: Neuromorph: Unsupervised shape interpolation and correspondence in one go. In: CVPR (2021) 2, 10
- Eisenberger, M., Toker, A., Leal-Taixé, L., Cremers, D.: Deep shells: Unsupervised shape correspondence with optimal transport. NIPS (2020) 6, 9, 10, 11
- Eisenberger, M., Toker, A., Leal-Taixé, L., Cremers, D.: G-msm: Unsupervised multi-shape matching with graph-based affinity priors. In: CVPR (2023) 2, 9
- Eynard, D., Rodola, E., Glashoff, K., Bronstein, M.M.: Coupled functional maps. In: 2016 Fourth International Conference on 3D Vision (3DV) (2016) 4
- Ezuz, D., Solomon, J., Ben-Chen, M.: Reversible harmonic maps between discrete surfaces. ACM Transactions on Graphics (ToG) 38(2), 1–12 (2019) 2, 3, 13
- Fey, M., Lenssen, J.E., Morris, C., Masci, J., Kriege, N.M.: Deep graph matching consensus. In: ICLR (2020) 2, 4, 5, 8, 14
- Franks, B.J., Morris, C., Velingker, A., Geerts, F.: Weisfeiler–leman at the margin: When more expressivity matters. arXiv:2402.07568 (2024) 4
- Gao, M., Lahner, Z., Thunberg, J., Cremers, D., Bernard, F.: Isometric multishape matching. In: CVPR (2021) 4
- Gao, M., Roetzer, P., Eisenberger, M., Lähner, Z., Moeller, M., Cremers, D., Bernard, F.: Sigma: Scale-invariant global sparse shape matching. In: ICCV (2023) 3
- Gao, Q., Wang, F., Xue, N., Yu, J.G., Xia, G.S.: Deep graph matching under quadratic constraint. In: CVPR (2021) 4
- Gasteiger, J., Weißenberger, S., Günnemann, S.: Diffusion improves graph learning. NIPS (2019) 2
- Gilmer, J., Schoenholz, S.S., Riley, P.F., Vinyals, O., Dahl, G.E.: Neural message passing for quantum chemistry. In: ICML. PMLR (2017) 4
- Groueix, T., Fisher, M., Kim, V.G., Russell, B.C., Aubry, M.: 3d-coded: 3d correspondences by deep deformation. In: ECCV (2018) 9
- Groueix, T., Fisher, M., Kim, V.G., Russell, B.C., Aubry, M.: Unsupervised cycle-consistent deformation for shape matching. In: Computer Graphics Forum. vol. 38, pp. 123–133. Wiley Online Library (2019) 13
- Halimi, O., Litany, O., Rodola, E., Bronstein, A.M., Kimmel, R.: Unsupervised learning of dense shape correspondence. In: CVPR (2019) 3, 4, 10

- Hamilton, W., Ying, Z., Leskovec, J.: Inductive representation learning on large graphs. NIPS (2017) 4
- Holzschuh, B., Lähner, Z., Cremers, D.: Simulated annealing for 3d shape correspondence. In: 2020 International Conference on 3D Vision (3DV) (2020) 3
- 46. Huang, Q., Wang, F., Guibas, L.: Functional map networks for analyzing and exploring large shape collections. ACM Transactions on Graphics (ToG) 33(4), 1–11 (2014) 4
- Huang, R., Ren, J., Wonka, P., Ovsjanikov, M.: Consistent zoomout: Efficient spectral map synchronization. In: Computer Graphics Forum. Wiley Online Library (2020) 4
- 48. Jiang, P., Sun, M., Huang, R.: Non-rigid shape registration via deep functional maps prior. In: NIPS (2023) 2
- Kim, V.G., Lipman, Y., Funkhouser, T.: Blended intrinsic maps. ACM Transactions on Graphics (ToG) 30(4), 1–12 (2011) 8, 9
- 50. Kondor, R.I., Lafferty, J.: Diffusion kernels on graphs and other discrete structures. In: ICML (2002) 2
- Lähner, Z., Rodola, E., Bronstein, M.M., Cremers, D., Burghard, O., Cosmo, L., Dieckmann, A., Klein, R., Sahillioglu, Y.: Shrec'16: Matching of deformable shapes with topological noise. Proc. 3DOR 2(10.2312) (2016) 9
- Lévy, B.: Laplace-beltrami eigenfunctions towards an algorithm that" understands" geometry. In: IEEE International Conference on Shape Modeling and Applications 2006 (SMI'06). pp. 13–13. IEEE (2006) 3, 5
- 53. Li, L., Donati, N., Ovsjanikov, M.: Learning multi-resolution functional maps with spectral attention for robust shape matching. NIPS (2022) 2, 9, 10, 11
- Li, T., Bolkart, T., Black, M.J., Li, H., Romero, J.: Learning a model of facial shape and expression from 4D scans. ACM Transactions on Graphics (ToG) 36(6), 194:1–194:17 (2017) 2
- Li, Y., Takehara, H., Taketomi, T., Zheng, B., Nießner, M.: 4dcomplete: Non-rigid motion estimation beyond the observable surface. In: ICCV (2021) 10
- Litany, O., Remez, T., Rodola, E., Bronstein, A., Bronstein, M.: Deep functional maps: Structured prediction for dense shape correspondence. In: ICCV (2017) 4, 9, 11
- Litany, O., Rodolà, E., Bronstein, A.M., Bronstein, M.M.: Fully spectral partial shape matching. In: Computer Graphics Forum. Wiley Online Library (2017) 4, 12
- Loper, M., Mahmood, N., Romero, J., Pons-Moll, G., Black, M.J.: Smpl: A skinned multi-person linear model. ACM Transactions on Graphics (ToG) 34(6), 1–16 (2015) 2
- Magnet, R., Ren, J., Sorkine-Hornung, O., Ovsjanikov, M.: Smooth non-rigid shape matching via effective dirichlet energy optimization. In: International Conference on 3D Vision (3DV) (2022) 3, 10, 13
- Melzi, S., Marin, R., Rodolà, E., Castellani, U., Ren, J., Poulenard, A., Wonka, P., Ovsjanikov, M.: Shrec 2019: Matching humans with different connectivity. In: Eurographics Workshop on 3D Object Retrieval (2019) 8
- 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) 4, 9, 10, 11
- Moler, C., Van Loan, C.: Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. SIAM review 45(1), 3–49 (2003) 5

- 18 Dongliang Cao et al.
- Morris, C., Lipman, Y., Maron, H., Rieck, B., Kriege, N.M., Grohe, M., Fey, M., Borgwardt, K.: Weisfeiler and leman go machine learning: The story so far. Journal of Machine Learning Research (JMLR) (2023) 4
- 64. Morris, C., Ritzert, M., Fey, M., Hamilton, W.L., Lenssen, J.E., Rattan, G., Grohe, M.: Weisfeiler and leman go neural: Higher-order graph neural networks. In: AAAI (2019) 4
- Nogneng, D., Ovsjanikov, M.: Informative descriptor preservation via commutativity for shape matching. In: Computer Graphics Forum. Wiley Online Library (2017) 4
- 66. 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) **3**, 5
- Ovsjanikov, M., Mérigot, Q., Mémoli, F., Guibas, L.: One point isometric matching with the heat kernel. In: Computer Graphics Forum. Wiley Online Library (2010) 5
- Ren, J., Melzi, S., Ovsjanikov, M., Wonka, P.: Maptree: recovering multiple solutions in the space of maps. ACM Transactions on Graphics (ToG) 39(6), 1–17 (2020) 4
- Ren, J., Melzi, S., Wonka, P., Ovsjanikov, M.: Discrete optimization for shape matching. In: Computer Graphics Forum. Wiley Online Library (2021) 4, 9, 10, 11
- Ren, J., Panine, M., Wonka, P., Ovsjanikov, M.: Structured regularization of functional map computations. In: Computer Graphics Forum. Wiley Online Library (2019) 4
- Ren, J., Poulenard, A., Wonka, P., Ovsjanikov, M.: Continuous and orientationpreserving correspondences via functional maps. ACM Transactions on Graphics (ToG) 37, 1–16 (2018) 2, 4, 8, 9
- Rodolà, E., Cosmo, L., Bronstein, M.M., Torsello, A., Cremers, D.: Partial functional correspondence. In: Computer Graphics Forum. Wiley Online Library (2017) 4, 12, 14
- 73. Roetzer, P., Swoboda, P., Cremers, D., Bernard, F.: A scalable combinatorial solver for elastic geometrically consistent 3d shape matching. In: CVPR (2022) 3
- Rogers, D., Hahn, M.: Extended-connectivity fingerprints. Journal of chemical information and modeling 50(5), 742–754 (2010) 4
- Roufosse, J.M., Sharma, A., Ovsjanikov, M.: Unsupervised deep learning for structured shape matching. In: ICCV (2019) 4, 10
- Rustamov, R.M., et al.: Laplace-beltrami eigenfunctions for deformation invariant shape representation. In: Symposium on geometry processing (2007) 5
- 77. Sahillioğlu, Y.: Recent advances in shape correspondence. The Visual Computer ${\bf 36}(8),\,1705{-}1721$ (2020) ${\bf 2}$
- Sarlin, P.E., DeTone, D., Malisiewicz, T., Rabinovich, A.: Superglue: Learning feature matching with graph neural networks. In: CVPR (2020) 14
- Sharma, A., Ovsjanikov, M.: Weakly supervised deep functional maps for shape matching. NIPS (2020) 8, 10, 11
- Sharp, N., Attaiki, S., Crane, K., Ovsjanikov, M.: Diffusionnet: Discretization agnostic learning on surfaces. arXiv preprint arXiv:2012.00888 (2020) 2, 3, 5
- 81. Shervashidze, N., Borgwardt, K.: Fast subtree kernels on graphs. NIPS (2009) 4
- 82. Song, C., Wei, J., Li, R., Liu, F., Lin, G.: 3d pose transfer with correspondence learning and mesh refinement. NIPS (2021) 2

- Song, C., Wei, J., Li, R., Liu, F., Lin, G.: Unsupervised 3d pose transfer with cross consistency and dual reconstruction. IEEE Transactions on Pattern Analysis and Machine Intelligence (2023) 2
- Sun, M., Mao, S., Jiang, P., Ovsjanikov, M., Huang, R.: Spatially and spectrally consistent deep functional maps. In: ICCV (2023) 2, 6, 9
- Talmon, R., Wu, H.T.: Latent common manifold learning with alternating diffusion: Analysis and applications. Applied and Computational Harmonic Analysis 47(3), 848–892 (2019) 7
- Tam, G.K., Cheng, Z.Q., Lai, Y.K., Langbein, F.C., Liu, Y., Marshall, D., Martin, R.R., Sun, X.F., Rosin, P.L.: 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) 3
- Tourani, S., Rother, C., Khan, M.H., Savchynskkyy, B.: Unsupervised deep graph matching based on cycle consistency. arXiv preprint arXiv:2307.08930 (2023) 4
- Trappolini, G., Cosmo, L., Moschella, L., Marin, R., Melzi, S., Rodolà, E.: Shape registration in the time of transformers. NIPS (2021) 9
- Vallet, B., Lévy, B.: Spectral geometry processing with manifold harmonics. In: Computer Graphics Forum. vol. 27, pp. 251–260. Wiley Online Library (2008) 2, 5
- Van Kaick, O., Zhang, H., Hamarneh, G., Cohen-Or, D.: A survey on shape correspondence. In: Computer Graphics Forum. Wiley Online Library (2011) 3
- 91. Vestner, M., Lähner, Z., Boyarski, A., Litany, O., Slossberg, R., Remez, T., Rodola, E., Bronstein, A., Bronstein, M., Kimmel, R., et al.: Efficient deformable shape correspondence via kernel matching. In: 2017 international conference on 3D vision (3DV). IEEE (2017) 3, 5, 7, 10, 13
- Vestner, M., Litman, R., Rodola, E., Bronstein, A., Cremers, D.: Product manifold filter: Non-rigid shape correspondence via kernel density estimation in the product space. In: CVPR (2017) 5
- Weisfeiler, B., Leman, A.: The reduction of a graph to canonical form and the algebra which appears therein. Nauchno-Technicheskaya Informatsia 2(9), 12–16 (1968) 2, 3, 4
- Windheuser, T., Schlickewei, U., Schmidt, F.R., Cremers, D.: Geometrically consistent elastic matching of 3d shapes: A linear programming solution. In: ICCV (2011) 3
- 95. Woodruff, D.P., et al.: Sketching as a tool for numerical linear algebra. Foundations and Trends[®] in Theoretical Computer Science 10(1–2), 1–157 (2014) 7
- Wu, Z., Pan, S., Chen, F., Long, G., Zhang, C., Philip, S.Y.: A comprehensive survey on graph neural networks. IEEE transactions on neural networks and learning systems 32(1), 4–24 (2020) 4
- 97. Xu, K., Hu, W., Leskovec, J., Jegelka, S.: How powerful are graph neural networks? In: ICLR (2018) 4
- 98. Zanfir, A., Sminchisescu, C.: Deep learning of graph matching. In: CVPR (2018) $\underline{4}$
- Zhang, H., Van Kaick, O., Dyer, R.: Spectral mesh processing. In: Computer graphics forum. vol. 29, pp. 1865–1894. Wiley Online Library (2010) 2
- 100. Zhou, T., Krahenbuhl, P., Aubry, M., Huang, Q., Efros, A.A.: Learning dense correspondence via 3d-guided cycle consistency. In: CVPR (2016) 13
- 101. Zuffi, S., Kanazawa, A., Jacobs, D.W., Black, M.J.: 3d menagerie: Modeling the 3d shape and pose of animals. In: CVPR (2017) 10