

Coupled Laplacian Eigenmaps for Locally-Aware 3D Rigid Point Cloud Matching

Matteo Bastico^{1*}, Etienne Decencière², Laurent Corté¹, Yannick Tillier³, David Ryckelynck¹
 Mines Paris, Université PSL

¹Centre des Matériaux (MAT), UMR7633 CNRS, 91003 Evry, France

²Centre de Morphologie Mathématique (CMM), 77300 Fontainebleau, France

³Centre de Mise en Forme des Matériaux (CEMEF), UMR7635 CNRS, 06904 Sophia Antipolis, France

Abstract

Point cloud matching, a crucial technique in computer vision, medical and robotics fields, is primarily concerned with finding correspondences between pairs of point clouds or voxels. In some practical scenarios, emphasizing local differences is crucial for accurately identifying a correct match, thereby enhancing the overall robustness and reliability of the matching process. Commonly used shape descriptors have several limitations and often fail to provide meaningful local insights about the paired geometries. In this work, we propose a new technique, based on graph Laplacian eigenmaps, to match point clouds by taking into account fine local structures. To deal with the order and sign ambiguity of Laplacian eigenmaps, we introduce a new operator, called Coupled Laplacian¹, that allows to easily generate aligned eigenspaces for multiple registered geometries. We show that the similarity between those aligned high-dimensional spaces provides a locally meaningful score to match shapes. We firstly evaluate the performance of the proposed technique in a point-wise manner, focusing on the task of object anomaly localization on the MVTec 3D-AD dataset. Additionally, we define a new medical task, called automatic Bone Side Estimation (BSE), which we address through a global similarity score derived from coupled eigenspaces. In order to test it, we propose a benchmark collecting bone surface structures from various public datasets. Our matching technique, based on Coupled Laplacian, outperforms other methods by reaching an impressive accuracy on both tasks.

1. Introduction

Point cloud matching, or more generally 3D shape matching, is a fundamental task in computer vision. It involves

finding the closest matching geometry to a target shape within a set of reference shapes [65]. In addition, if the task involves finding rigid transformations that best align the target shape with the reference, it is often part of a registration process. In particular, point-set *rigid registration* determines the relative transformation needed to align two point clouds without altering their internal structures [41]. This problem is essential for many practical computer vision tasks, such as medical image analysis [3, 32, 48, 61], intelligent vehicles [21, 34], human pose estimation [22] and objects retrieval and tracking [46, 64]. Traditional [8, 19, 57] and probabilistic registration and matching methods [15, 20, 30, 45], while robust, often struggle to optimally align complex geometries, especially in cases with intricate local structures or slight deformations.

Over the years, several methods have been proposed to tackle the challenge of accurate and efficient 3D shape matching and retrieval [5, 9, 11, 51, 53, 65, 73]. Data-driven 3D shape descriptors [54], capturing underlying properties of the shapes under study, are the common denominator of early shape matching techniques. Global descriptors, such as volume and areas descriptors [78], describe the entirety of the shape in one compact representation, often failing to capture local fine details of complex geometries. On the other hand, local descriptors [38, 57] aim to tackle this issue but they generally are sensitive to noise, based on landmarks and they might not capture semantic information [63]. More recently, deep-learned shape descriptors [6, 72] and neural networks for shape matching, based on auto-encoders [73] or transformers [59, 68], have been proposed. Despite their good performances, these methods require a huge amount of annotated data for training, which are hard to collect in fields such as medical imaging [37]. Furthermore, non-rigid point cloud matching and retrieval methods [35, 36, 71] are designed to handle shape deformations and, therefore, they might be excessively flexible ignoring fine local details that are not due to deformations, such as anomalies.

*Corresponding author: matteo.bastico@minesparis.psl.eu

¹Code: <https://github.com/matteo-bastico/CoupLap>

In this study, we introduce a novel method for rigid 3D point cloud matching, based on spectral Laplacian eigenmaps [4, 23], which focus on local details. This technique is designed to overcome the limitations of shape descriptors and, although it can be seen as a linear graph neural network, it does not require training. One of the main reasons that combinatorial and geometric Laplacians are often considered for spectral graph processing is that their eigenvectors possess properties like the classical Fourier basis functions [79]. We leverage these characteristics to perform locally-aware shape comparison of point clouds, equipped with k-nearest neighbor graph, without relying on additional descriptors. Nevertheless, eigenspaces alignment, including both eigenvalue ordering and sign disambiguity of the eigendecomposition, is required to correctly match shapes in the spectral space. Current State-of-The-Art methods for such alignment are based on Laplacian eigenfunctions matching [42, 60], i.e. matching the histograms of the eigenvectors. The latter is not robust and frequently fails in possible scenarios with highly symmetric geometries, such as bones, with small eigenvalue separation. In this work, we introduce a novel operator, called *Coupled Laplacian*, designed to simultaneously produce aligned eigenspaces for multiple registered geometries. Namely, we show that when two or more shapes, each with their corresponding graphs, are merged into a single graph using artificial cross-edges, the eigendecomposition of the Laplacian derived from such a combined graph yields aligned spectral spaces for each individual component. Furthermore, this method is order-invariant, count-agnostic and landmarks-free, enabling it to handle point clouds without being influenced by their specific arrangement or the number of points contained. Finally, we utilize the distance between these aligned higher-dimensional spectral spaces, which accounts for intricate local structures, as a global or point-wise score for different applications of 3D shape matching, as shown in Fig. 1.

One natural application of the proposed technique, is 3D object anomaly detection through the identification of local differences between two shapes. Hence, we tested our method on the MVTec 3D-AD dataset [7], recently proposed for unsupervised point cloud anomaly detection and localization. Furthermore, in the context of medical imaging, correctly identifying the side of a bone (left or right) is crucial for diagnosis, treatment planning, or skeletal analysis [12]. We refer to this task as *Bone Side Estimation* (BSE). Recent studies on bone landmark detection [18] and statistical shape modeling [31] have highlighted the limitations of manual side identification for ensuring proper functioning. To the best of our knowledge, manual bone markings [2] are currently the only technique used for this purpose. Consequently, automatic BSE arises as an interesting, yet challenging, task to assist the development of fully auto-




Figure 1. Overview of the proposed technique applied to different tasks. The global similarity between aligned eigenmaps of multiple geometries, generated from the Coupled Laplacian, is used to match bones and predict their body side. While local similarity is exploited for accurate 3D anomaly detection.

mated pipelines for medical image analysis, such as patient-specific preoperative planning for Total Knee Arthroplasty (TKA) [33] or Anterior Cruciate Ligament Reconstruction (ACLR) [17, 44]. Due to the bilateral symmetry of the animal and human body [27, 67], i.e. right and left sides are mirror shapes of one another, the BSE can be defined as a non-trivial chiral shape matching problem. Then, given a known reference or source bone, assessing the side of a target bone involves finding the best match between the source-target pair and its mirrored counterpart. The complexity of this task arises from subtle local differences in mirrored bones, which makes it a suitable application of the proposed shape matching technique utilizing coupled Laplacian eigenmaps as local geometry descriptors. We propose a benchmark for human BSE by extracting surface point clouds of different bones, e.g. femur, hip and tibia, from public datasets [18, 31, 47]. We also discuss, in the Supplementary Material, a non-rigid correspondence application of the proposed matching algorithm, the cross-species BSE (human to animal), on an internal dataset [40].

Our contributions are summarized as follows:

- We propose a new method to perform locally-aware 3D rigid point cloud matching, considering fine local structures, based on graph Laplacian eigenmaps.
- We define the Coupled Laplacian operator for aligned graphs to tackle the order and sign ambiguity issue of the eigendecomposition.
- Based on the proposed technique, we introduce a new method for automatic BSE to assist fully automated medical pipelines and we propose a benchmark to test it.

- We extensively evaluate our method on two tasks: anomaly detection and BSE, outperforming previously proposed techniques.

2. Related Works

Early efforts in developing hand-crafted 3D local features for point cloud matching, registration and retrieval have typically drawn inspiration from 2D descriptors. Several methods, such as Signature of Histograms of Orientations (SHOT) [58, 66], Rotational Projection Statistics (RoSP) [25] and Unique Shape Context (USC) [66] rely on the estimation of a unique Local Reference Frame (LRF). The latter is usually based on the eigendecomposition of the covariance matrix of the neighbours of a point of interest, which are then projected into the LRF to analyze their geometric properties. For instance, SHOT [58] captures local shape and orientation information by computing histograms of surface normals and point distribution around keypoints. In contrast, LRF-free approaches [10, 56, 57] try to rely just on features that are intrinsically invariant. The most common LRF-free features are Fast Point Feature Histograms (FPFH) [57] which, similarly to SHOT, calculates histograms of surface normals considering their relationships in a local region around a keypoint and generating 33-dimensional descriptors. Despite the progress made with hand-crafted 3D local features, they encounter difficulties when it comes to deal with issues like point cloud resolutions, noisy data, and occlusions [26].

On the contrary, spectral-based point cloud descriptors [1, 28, 52, 55, 62, 69, 70] are a category of feature extraction methods that leverage spectral analysis techniques from graph theory to capture the underlying structure and intrinsic geometric properties of point clouds. Among them, shape-DNA [52] is a surface descriptor based on eigenvalue analysis of the Laplace-Beltrami operator. They propose to use the sequence of eigenvalues (spectrum) of the Laplace operator as a fingerprint characterizing the intrinsic geometry of 3D shapes represented as point clouds. This method has been used for shape retrieval, classification, and correspondence. Weinmann *et al.* [70] also proposed to extract a features set consisting of 8 eigenvalues-based indices for each 3D point of a cloud. Furthermore, Heat Kernel Signature (HKS) [62] and Wave Kernel Signatures (WKS) [1] are descriptors measuring how heat and wave propagate across a shape, having the eigendecomposition as leading element of the computation. Scaled eigenvectors evaluated at each point are instead directly exploited by the Global Point Signature (GPS) [55] to represent a point cloud as a set of infinite-dimensional vectors, characterizing each point within the global context of the surface it belongs. Nevertheless, the vast majority of these works assumes that the eigenvalues of a shape are distinct and, therefore, can be ordered. Indeed, in practice, due to numerical approxima-

tions, we cannot guarantee that the eigenvalues of the Laplacian are all distinct and, possible symmetries in the shapes may cause some of them to have multiplicity greater than one. As shown by Mateus *et al.* [42], when dealing with shape matching, an elegant way to overcome this problem is to use the Laplacian eigenmaps scheme [4], which can be seen as a reduced GPS, and perform a-posteriori alignment of the resulting point cloud embeddings. Matching the eigenfunctions histograms, i.e. their signatures, is the only reliable method for such embeddings alignment [42, 60]. Recently, Ma *et al.* [39] proposed a canonization algorithm for sign and basis invariance, called Maximal Axis Projection (MAP), that adopts the permutation-invariant axis projection functions to determine the canonical directions. Unfortunately, when the eigenvalues separation is small due to symmetries in the geometries, these methods become unstable and sensitive to noise. To overcome this issue, in this work we propose the Coupled Laplacian operator to produce a-priori aligned eigenmaps for several registered shapes, and perform locally-aware shape matching.

Deep-learned point-cloud descriptors [6, 24, 72, 77] arose as an alternative to generate local features for 3D surfaces. We can distinguish three main categories depending on the backbone architecture employed. Convolutional Neural Networks (CNNs) are often used on point clouds projected into 2D depth images [16, 29] or directly on 3D voxels [24, 29, 77]. Secondly, to work directly on raw point cloud data, PointNet and PointNet++ have been proposed aiming to learn rotation and permutation invariant features [49, 50]. Based on PointNet, several learned descriptors have been introduced [13, 14, 75, 76]. Among them, PPFNet [14] and PPF-FoldNet [13] try to improve the feature representation of PointNet by incorporating global context and point-pair features. Nevertheless, the lack of convolutional layers in these models limits the learning of local geometries. Finally, transformers-based descriptors, such as Deep Closest Point [68], have been recently proposed trying to exploit the attention mechanisms to capture shapes and surfaces intrinsic characteristics. Arguably, training a deep-learning model, especially if based on transformers, is not always feasible in terms of training samples required to achieve good performances. Furthermore, in supervised algorithms [13, 14, 24, 77], paired 3D patches, such as in the 3DMatch dataset [77], are needed to train the models. For this reason, the 3D rigid point cloud matching method we propose does not need training and utilizes only the similarity among properly aligned spectral embeddings to relate an unseen target shape to the given references.

3. Method

Graph Laplacian. A 3D point-cloud $\{\mathbf{x}_i\}_{i=1}^n$ can be treated as a connected undirected weighted graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, where $\mathcal{V} = \{\mathbf{x}_i\}_{i=1}^n$ is the nodes set and $\mathcal{E} = \{\mathbf{e}_{ij}\}$ is the

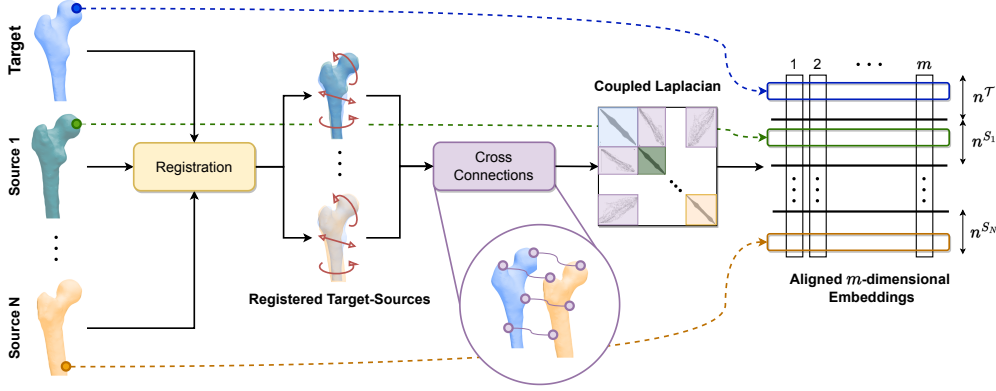


Figure 2. Proposed workflow of the Coupled Laplacian applied to proximal femur shapes. The N sources are registered to the target using a rigid or affine registration. After that, cross-connection are added between each target-source pair (for simplicity, in the zoom the shapes are not overlapping) and the Coupled Laplacian is computed on the global graph. Its eigendecomposition leads to aligned spectral embeddings of the input geometries that can be used for shape matching.

edge set. The latter is generally obtained through the construction of a nearest neighbors graph [4]. To that purpose, in this work, we consider the k -Nearest Neighbour (k -NN) approach, that is, node j is connected to node i if \mathbf{x}_i is among the k nearest neighbours of \mathbf{x}_j . Hence, we can build a weighted adjacency matrix, $\mathbf{W} = \{w_{ij}\}$, which stores the connections between nodes. In particular, in spectral graph theory, a Radial Basis Function (RBF) is commonly used as weight function for the edges between \mathbf{x}_i and \mathbf{x}_j , as

$$w_{ij} = \exp\left(-\frac{d^2(\mathbf{x}_i, \mathbf{x}_j)}{\sigma^2}\right) \quad (1)$$

where $d^2(\cdot, \cdot)$ is the Euclidean distance between two vertices and σ^2 a free parameter which, for simplicity, we set to the maximum distance between connected nodes, $\max_{ij} d^2(\mathbf{x}_i, \mathbf{x}_j)$. The *Laplacian matrix*, $\mathbf{L} \in \mathbb{R}^{n \times n}$, of a graph constructed in such way, is defined as $\mathbf{L} = \mathbf{D} - \mathbf{W}$ where $\mathbb{R}^{n \times n} \ni \mathbf{D} = \text{diag}([d_1, \dots, d_n])$ is the degree matrix with diagonal elements $d_i = \sum_{j=1}^n w_{ij}$ [43]. The Laplacian eigenvalues, $\{\lambda_i\}_{i=0}^n$, and eigenvectors, $\{\phi_i\}_{i=0}^n$, can be computed by solving the generalized eigenproblem

$$\mathbf{L}\phi_i = \lambda_i \mathbf{B}\phi_i, \quad i = 0, \dots, n. \quad (2)$$

where $\mathbf{B} \in \mathbb{R}^{n \times n}$ is generally set as \mathbf{D} , and $\lambda_i \leq \lambda_{i+1}$. Finally, eigenmaps are simply eigenvectors sub-spaces, generated by leaving out ϕ_0 , corresponding to $\lambda_0 = 0$, and using the next m eigenvectors for embedding graph nodes in an m -dimensional space, $\mathbf{x}_i \rightarrow [\phi_1(i), \dots, \phi_m(i)]$

Coupled Laplacian. The complete pipeline to generate aligned Laplacian embeddings through graph coupling is illustrated in Fig. 2. Let \mathcal{G}^T be the graph of a target 3D point cloud and $\mathcal{G}^{S_1}, \dots, \mathcal{G}^{S_N}$ the ones of N different sources. We construct a global graph, or coupled graph, \mathcal{G}^C , by adding *cross-connections* separately between the

vertices of the source, $\mathcal{V}^T = \{\mathbf{x}_i^T\}_{i=1}^{n^T}$, and each of the reference shapes, $\mathcal{V}^{S_k} = \{\mathbf{x}_i^{S_k}\}_{i=1}^{n^{S_k}}$ with $k = 1, \dots, N$. To include meaningful cross-connections, we first perform a rigid [19, 20, 45] or affine [45] registration of the source geometries to align to the target. These methods are preferred to a non-rigid registration because, when dealing with the identification of fine local variations, they do not change the relative position of the points inside a point-set and, therefore, the deformations are kept unchanged. After that, a sub-set of vertices, $F^T \subset \mathcal{V}^T$, is stochastically extracted from the target geometry and their nearest correspondences are searched in each of the aligned reference point clouds as

$$F^{S_k} = \{\mathbf{f}_j^{S_k} : \arg \min_{\mathbf{x}_i^{S_k} \in \mathcal{V}^{S_k}} d^2(\mathbf{x}_i^{S_k}, \mathbf{f}_j^T) | \mathbf{f}_j^T \in F^T\} \quad (3)$$

with $k = 1, \dots, N$. In this way, there is no constraint on the original number points comprising each shape nor on their initial coordinates systems. Furthermore, the cardinality of the target sub-set, is chosen as a fraction of its total points, $|F^T| = l \cdot n^T$, where $0 < l \leq 1$. The cross-connection are therefore added for each pair $(\mathbf{f}_i^{S_k}, \mathbf{f}_i^T)$, $1 \leq i \leq |F^T|$, between the target and the k -th source. Finally, we arrange the vertex indices of the coupled graph such that they are grouped for each individual geometry, and the whole weighted adjacency matrix can be computed as in Eq. (1).

We define the Laplacian matrix of a global graph, constructed as describe above, as the Coupled Laplacian, $\mathbf{L}^C \in \mathbb{R}^{n \times n}$, where $n = n^T + \sum_{k=1}^N n^{S_k}$. Thanks to the vertex ordering of the coupled graph, the solution of the generalized eigenproblem of Eq. (2) applied to the Coupled Laplacian, with $\mathbf{B}^C = \text{diag}([\mathbf{D}^T, \mathbf{D}^{S_1}, \dots, \mathbf{D}^{S_N}])$, yields to *coupled eigenvectors* $\{\phi_i^C\}_{i=0}^n$. These eigenvectors can then be split into each single component of the coupled graph as

$$\phi_i^C = [\phi_i^T, \phi_i^{S_1}, \dots, \phi_i^{S_N}]^T. \quad (4)$$

The eigenspaces restricted to the single geometries, obtained through the Coupled Laplacian, are intrinsically aligned up to a certain component, depending on the factor l , and the eigenvalues ordering issue is automatically solved since the split eigenvectors are associated to the same eigenvalues. The proof of this property for the ideal case of perfect match and more theoretical aspects of the Coupled Laplacian are reported in the Supplementary Material.

Shape Matching with Eigenmaps. The Coupled Laplacian allows the generation of m -dimensional embeddings aligned for the graph vertices of the target and reference shapes, $\Phi^T = (\phi_1^T, \dots, \phi_m^T) \in \mathbb{R}^{n_T \times m}$ and $\Phi^{S_k} = (\phi_1^{S_k}, \dots, \phi_m^{S_k}) \in \mathbb{R}^{n_{S_k} \times m}$ for $1 \leq k \leq N$, respectively. Therefore, a comparison of multiple geometries in this higher dimensional spectral space yields to a proper consideration of local structures. We consider in the following two possible applications of the coupled eigenmaps: (1) a global shape matching score and (2) local similarity scores. The first one can be obtained by measuring the similarity between the aligned eigenspaces through the Grassmann distance, $d_G(\cdot, \cdot)$ [74]. Nevertheless, it cannot be computed directly, since we consider point clouds of arbitrary size and we do not have point-to-point correspondences. Hence, we restrict the distance computation only to the set of cross-connected vertices, F^T for the target and F^{S_k} for the k -th reference. In order to get the reduced basis restricted only to those coupled points, we perform a QR decomposition of the restriction of the eigenmodes to the cross-points. The best matching source is then given by

$$\arg \min_k d_G(Q^T, Q^{S_k}). \quad (5)$$

where Q^T and Q^{S_k} are obtained from the QR factorization of $\Phi^T(F^T, \cdot)$ and $\Phi^{S_k}(F^{S_k}, \cdot)$, respectively. This first approach is used in the following to solve the BSE task by using two source shapes, i.e. the reference bone of known side and its contralateral mirrored version.

On the other hand, to obtain point-wise similarity scores we propose to compare the m -dimensional embeddings of the cross-connection vertices by using the cosine distance function. In this case, we interpret the distance value as the probability of local structural difference, where 0 means that the two compared points have the same local structure.

4. Experiments

4.1. Experiment settings

Bone Side Estimation. The proposed global shape matching score is used to perform automatic BSE. Assuming general unknown initial frames, we used the Principal Component Analysis (PCA) to generate a mirrored version of the reference surface, i.e. a synthetic contralateral bone. Namely, we identified that the bilateral symmetry of human and animal bodies is equivalent to a mirroring around

Table 1. Summary of the point cloud bone structures collected for the BSE benchmark. L and R stands for Left and Right, respectively, and S for Sheep.

Dataset		Femur		Hip		Tibia		Fibula	
		L	R	L	R	L	R	L	R
Human	Fisher et al. [18]	18	19	20	20	-	-	-	-
	SSM-Tibia [31]	-	-	-	-	-	30	-	30
	ICL [47]	35	35	-	-	35	35	35	35
S	Internal [40]	18	18	-	-	18	18	-	-
	Total	71	72	20	20	53	83	35	65

the second principal component on the vast majority of bones, such as the ones considered in this study. Hence, given the three point clouds, target, source and mirrored source, the Coupled Laplacian of Fig. 2 can be applied. More specifically, we selected the Random Sample Consensus (RANSAC) [19], preceded by a spectral scaling, as rigid registration method to handle varying bone lengths without altering local structures. The scaling is performed using only the information carried by the Fiedler vector, i.e. the eigenvector ϕ_1 corresponding to the first non-zero eigenvalue λ_1 . A rough estimation of the bones length, independently of the Euclidean frame, is given by the distance between the points corresponding to minimum and maximum values of the Fiedler vectors. Thanks to that, one of the two bones can be scaled in order to match the length of the other and improve the RANSAC registration. The three aligned m -dimensional embeddings, derived from the eigendecomposition of the Coupled Laplacian, are then used as in Eq. (5) to retrieve the best matching source and, consequently, the target side. Note that, the proposed BSE method is fully independent of the target and reference frames and therefore can be applied directly on the segmentation obtained from a medical image, without any previous knowledge, in a fully automated pipeline. In alternative, one can use two different bone references, one left and one right, avoiding the mirroring step, or apply the mirroring on the target shape while letting the source unchanged. Full details of the described algorithm are provided in the Supplementary Material.

We generated a benchmark to test the BSE task by collecting several bone surface point cloud data from public datasets [18, 31, 47]. Namely, from Fisher et al. [18] we extracted Femur and Hip structures, from SSM-Tibia [31] only right Tibia and Fibula, and from the Imperial College London (ICL) [47] we collected Femur, Tibia and Fibula point clouds. The detailed composition of the benchmark is reported in Tab. 1. Such a variety of shapes, which are acquired using different methods, makes the task more challenging, including some intra-clinic variability. In our experiments, for a given bone class, we performed cross-testing by selecting each shape once as source and the others

Table 2. Anomaly localization results. The area under the PRO curve is reported for an integration limit of 0.3 for each evaluated method and dataset category. GAN, AE and VM results are provided by *Bergmann et al.* [7]. Moreover, we include the results obtained by restricted GPS [55], eigenfunctions (Hist) [42, 60] and Euclidean matching and MAP [39]. All the matching methods are applied after affine CPD registration, when not specified, or CMM + CPD Non-Rigid (*NR*) registration. The subscript on a method indicates the number of eigenmaps used. The overall best performing methods are highlighted in boldface, while the bests for each category are underlined.

Method		Bagel	Cable Gland	Carrot	Cookie	Dowel	Foam	Peach	Potato	Rope	Tire	Mean \uparrow
3D RGB	VM	0.388	0.321	0.194	0.570	0.408	0.282	0.244	0.349	0.268	0.331	0.335
	GAN	0.421	<u>0.422</u>	0.778	<u>0.696</u>	0.494	0.252	<u>0.285</u>	<u>0.362</u>	0.402	<u>0.631</u>	0.474
	AE	<u>0.432</u>	<u>0.158</u>	0.808	0.491	0.841	<u>0.406</u>	<u>0.262</u>	<u>0.216</u>	<u>0.716</u>	0.478	<u>0.481</u>
3D Only	GAN	0.111	0.072	0.212	0.174	0.160	0.128	0.003	0.042	0.446	0.075	0.143
	AE	0.147	0.069	0.293	0.217	0.207	0.181	0.164	0.066	0.545	0.142	0.203
	VM	0.280	0.374	0.243	0.526	0.485	0.314	0.199	0.388	0.543	0.385	0.374
	Euclidean (<i>NR</i>)	0.404	0.623	0.731	0.366	0.771	0.303	0.590	0.772	0.697	0.583	0.584
	GPS ₁₀₀ [55]	0.452	0.616	0.695	0.364	0.738	0.471	0.659	0.844	0.647	0.651	0.613
	GPS ₂₀₀ [55]	0.465	0.621	0.690	0.363	0.739	0.480	0.672	0.833	0.653	0.670	0.619
	Hist ₁₀₀ [60]	0.476	0.629	0.703	0.365	0.744	0.473	0.661	0.840	0.647	0.693	0.623
	MAP ₂₀₀ [39]	0.481	0.630	0.694	0.399	0.742	0.497	0.653	0.832	0.649	0.675	0.625
	Hist ₂₀₀ [60]	0.491	0.629	0.698	0.351	0.746	0.501	0.663	0.841	0.652	0.695	0.627
	Euclidean	0.655	0.631	0.743	0.615	0.803	0.528	0.726	0.875	0.762	0.695	0.703
	Ours ₁₀₀	0.669	0.642	0.808	0.714	<u>0.812</u>	0.582	0.748	0.897	0.750	0.733	0.736
	Ours ₂₀₀	0.702	0.630	0.728	0.735	<u>0.812</u>	0.701	0.780	0.914	0.767	0.713	0.748

Table 3. Accuracy [%] of human BSE. All the matching methods are applied after RANSAC registration with spectral scaling, when not specified, or CMM + CPD Non-Rigid (*NR*) registration. The overall best performing methods are highlighted in boldface.

Method	Femur	Hip	Fibula	Tibia	Mean \uparrow
Chamfer (<i>NR</i>)	56.58	86.58	57.89	74.74	68.95
Hausdorff (<i>NR</i>)	57.37	87.63	59.47	73.68	69.54
Hausdorff	73.52	97.95	59.30	72.52	75.82
Chamfer	71.64	98.65	64.62	74.47	77.35
FPFH [57]	68.88	96.67	66.69	78.35	77.65
MAP ₂₀ [39]	76.05	98.54	71.32	74.21	80.13
Hist ₂₀ [60]	77.63	98.42	69.47	77.37	80.72
Ours ₂₀	78.79	98.78	71.28	78.46	81.83
Ours ₁₀	79.68	97.76	73.47	78.66	82.39

as targets on which the side has to be inferred. Therefore, the average accuracy of all the experiments, achieved when predicting the correct body side, i.e. binary classification, takes into account the variability of the source shape. The robustness of the method with respect to the reference is discussed in the Supplementary Material.

Anomaly Localization. We tested the proposed approach for the 3D anomaly detection task in the MVTEC 3D-AD dataset [7]. This recent dataset was designed for the unsupervised detection of anomalies in point clouds of industrially manufactured products. It contains over 4000 high-resolution 3D scans of 10 object categories. In our experiments, we chose one anomaly-free training sample for each class as source and we used the point-wise distances of the aligned eigenmaps between this shape and each target in the test set as anomaly scores. We performed the pre-




Figure 3. Average BSE accuracy with respect to the fraction of cross-edges used to build the Coupled Laplacian and the number of eigenmaps used for matching.

registration using affine Coherent Point Drift (CPD) [45], which is more suited for objects like the ones included in the MVTEC 3D-AD, e.g. length and thickness variations are better captured by an affine registration rather than rigid. In order to speed-up the computation, we pre-processed the point clouds by removing their flat backgrounds with a threshold on the z -axis after a 3-dimensional PCA and, without losing generality, we sub-sampled the number of foreground points to a maximum of 13000. Finally, to compare the obtained anomaly scores with the Ground Truth (GT) image, we projected back the points to the original 2D plane and, if any sub-sampling was performed, we applied a dilation with a structuring element of size the inverse of the sampling factor.

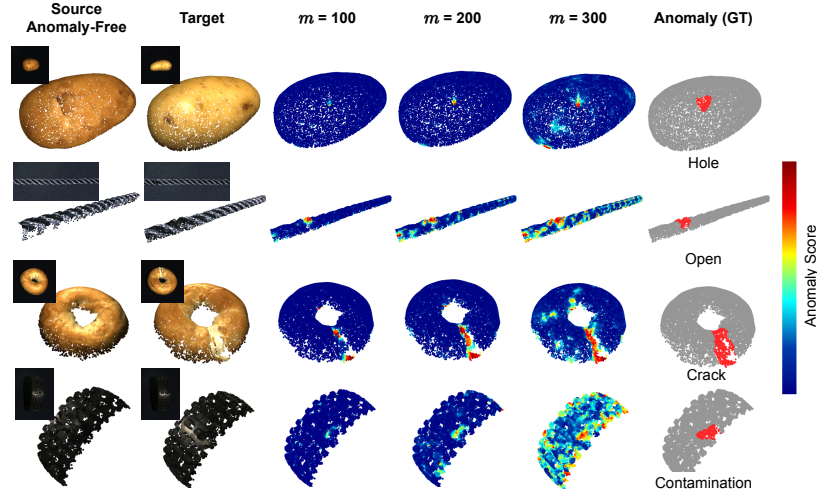


Figure 4. Graphical comparison of 3D anomaly localization using different numbers of aligned eigenmaps, m . Tuning the dimension of the embeddings computed with the Coupled Laplacian, it is possible to decide the extent and size of local surface differences that are detected.

4.2. Experiment results

Bone Side Estimation. In Fig. 3, the average accuracy of BSE is depicted based on the fraction of artificial cross-edges added to create the coupled graph and the dimension of the spectral embeddings, l and m , respectively. We can observe that the more coupled eigenmaps we aim to utilize for computing a global similarity score, the greater the need for additional cross-connections to ensure the reliability of the results. This occurs because, when graphs are weakly coupled with only a few cross-connections, the Coupled Laplacian captures limited intra-shape characteristics and predominantly emphasizes individual local geometries. In fact, when $l = 0$ and therefore the coupled graph is not fully connected, the eigenmaps obtained using the Coupled Laplacian are equivalent to the ones computed independently on each single graph, if the eigenvalues are the same (proof in the Supplementary Material). On the other hand, not too many cross-edges, nor eigenmaps, are needed to obtain meaningful global matching scores and achieve good performances in the BSE task. In Tab. 3 a quantitative comparison with other methods of human BSE accuracy is reported. For a fair comparison, all methods are performed after RANSAC registration of the sources. Hausdorff and Chamfer discrepancies are calculated on the Euclidean coordinates of the points, for completeness, also in case of 2-Step non-rigid GMM + CPD registration (2000x slower than RANSAC). Our method, using 10 and 20 eigenmodes with $l = 0.5$ fraction of cross-connections, outperforms the other techniques achieving higher accuracy both on the single bones and in average. Hence, the description provided by the aligned eigenmaps, obtained with Coupled Laplacian, is more aware of local details than the other methods.

Anomaly Localization. Tab. 2 lists quantitative results

of each evaluated method for the localization of anomalies. For each category, the normalized area under the Per-Region Overlap (PRO) curve with an upper integration limit of 0.3 [7], as well as, the mean performance, are reported. Performance of Generative Adversarial Network (GAN), Autoencoder (AE) and Variation Model (VM) on the same test set are provided by the dataset authors [7]. Furthermore, we tested restricted GPS [55], eigenmaps histogram matching [42, 60], MAP [39] and our method, all with the same pre-processing and source shapes, using 100 and 200 eigenmaps. In order to obtain denser and more precise anomaly localization, the coupled graph is built using the whole set of points as cross-connections, i.e. $l = 1$. Furthermore, we include the result obtained using as anomaly score for each target point the normalized Euclidean distance of the nearest point in the source geometry, both using rigid and non-rigid registration. Our method, using only 3D information, outperforms all the other techniques. Moreover, we obtained better results than deep-learning methods having RGB and Depth in combination as input. Interestingly, the point-wise similarity computed on eigenmaps not properly aligned (GPS and Hist) is worst than just considering Euclidean distances between points of registered shapes, making worthless the computation of spectral embeddings. In Fig. 4 we compare the qualitative results obtained using different numbers of eigenmaps to score the point-wise similarities. Using a smaller m , the proposed technique is prone to individuate only small regions with highly dissimilar local geometries. By increasing the size of the spectral embeddings, we can identify larger and more subtle surface differences, even if they are not necessarily classified as anomalies in the GT. For instance, the potato and bagel surfaces have some natural irregularities, with re-

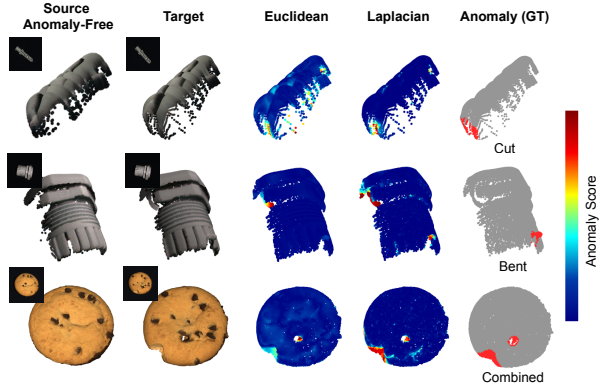


Figure 5. Comparison of the anomaly score obtained using the Coupled Laplacian technique, with $m = 200$ and $l = 1$, and Euclidean distance, both after affine CPD alignment.

spect to the source, that are not highlighted using 100 and 200-dimensional embeddings, but are instead detected with 300 eigenmaps. This concept is linked to the *modal length*, which we define in the Supplementary Material. The trade-off between number of maps and extent of differences detected is interesting to tune the method to other tasks requiring specific attention to identify surface dissimilarities.

Ablation Study. To further motivate the preference for Coupled Laplacian spectral space over the Euclidean one, in Fig. 5, we display comparisons of point-wise distances (anomaly scores) between source and target in both spaces, after the same registration. In the first and second row the better resistance to noise and outliers in the spectral space is highlighted. The bending of the cable gland, i.e. the anomaly, is barely detected in the Euclidean space because of noisy points on the back. Instead, with the aligned eigenmaps comparison, the anomaly, as well as the noise, is correctly localized. This is because both are local differences, with respect to the reference, and, therefore, they can be easily captured with coupled eigenmaps. In addition, in the third row, we show that our technique provides a sharper anomaly split of the cookie point cloud, which simplifies the anomaly detection task. Nevertheless, the usage of the proposed landmarks-free Coupled Laplacian is limited to tasks in which a correspondence between target and source shapes can be defined by mean of a registration process. Problems including graph isomorphism, such as articulated shape matching [42, 60], or non-rigid shape matching, can be still solved using our method, without relying on the rigid alignment, if some cross-connection are defined a-priori. In this case, landmarks and their neighbouring points would be most likely the easiest option to connect separate graphs and obtain aligned spectral embeddings for further processing. Moreover, a wrong registration of the shapes, due to randomness and errors in the selected registration algorithm, may affect the alignment of the m -dimensional embeddings

derived from the Coupled Laplacian. We report here some statistics on registration failures in BSE (failure rate [%] / BSE accuracy on failures [%]), using different references, for femur $6.84 \pm 7.82 / 96.67 \pm 11.30$ (high symmetry) and hip 0.0 ± 0.0 (low symmetry). Moreover, to test robustness, we added white Gaussian rotation and translation noise, i.i.d. in the 3 axis post-registration, obtaining (Rot. SD [$^\circ$] / Trans. SD [mm] / Acc. Drop [%]): $2.5 / 5 / 8.16 - 5 / 5 / 15.13 - 7.5 / 10 / 18.16$. Given the low registration failure rates and high robustness to noise, we can rely on the proposed coupling algorithm, making the method independent of the acquisition frames. Moreover, the method can well generalize to any chiral shape matching problem and to tasks like identification of topological noise in graphs and partial matching (see example in Supplementary Material).

5. Conclusions

We presented a versatile method to perform 3D rigid point cloud matching, both globally and locally, by using aligned eigenspaces for two or more similar shapes without prior knowledge on the local frames or specific markers on them. The introduction of the Coupled Laplacian operator enables the generation of aligned eigenmaps relying on cross-connections added between graphs of registered geometries. We introduced a novel task consisting in the automatic detection of a bone side, i.e. Bone Side Estimation (BSE), and we proposed a benchmark to test it. Leveraging global similarities of eigenmaps derived from the Coupled Laplacian, we define a full pipeline to perform BSE on arbitrary bone surfaces segmented from a generic medical image. Moreover, we tested the ability of the proposed method in capturing local surface differences by performing 3D anomaly detection on the MVTec 3D-AD dataset. The proposed technique outperforms other methods on the two tasks, therefore capturing better both global and local similarities, thanks to the matching performed on correctly aligned spectral spaces. Beside the applications showcased here, we believe that our work can help several tasks in many fields in which the localization of local difference is crucial for matching, detection or retrieval. Future extensions could explore generalizations of the technique, not only for rigid matching but also to address challenges such as graph isomorphism. This approach opens new avenues for advancing the State-of-The-Art in 3D shape analysis and matching techniques.

6. Acknowledgements

This project has received funding from the European Union’s Horizon 2020 research and innovation programme under the Marie Skłodowska-Curie grant agreement No 945304-Cofund AI4theSciences hosted by PSL University. This work was granted access to the HPC/AI resources of IDRIS under the allocation 2022-AD011013902 made by GENCI.

References

- [1] Mathieu Aubry, Ulrich Schlickewei, and Daniel Cremers. The wave kernel signature: A quantum mechanical approach to shape analysis. In 2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops), pages 1626–1633, Barcelona, Spain, 2011. IEEE.
- [2] Ivan Bandovic, Matthew R. Holme, Asa C. Black, and Bennett Futterman. Anatomy, Bone Markings. In StatPearls. StatPearls Publishing, Treasure Island (FL), 2023.
- [3] Zachary M C Baum, Yipeng Hu, and Dean C Barratt. Real-time multimodal image registration with partial intraoperative point-set data. Medical Image Analysis, 74:102231, 2021.
- [4] Mikhail Belkin and Partha Niyogi. Laplacian Eigenmaps for Dimensionality Reduction and Data Representation. Neural Computation, 15(6):1373–1396, 2003.
- [5] S. Belongie, J. Malik, and J. Puzicha. Shape matching and object recognition using shape contexts. IEEE Transactions on Pattern Analysis and Machine Intelligence, 24(4):509–522, 2002. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [6] Paul Bergmann and David Sattlegger. Anomaly Detection in 3D Point Clouds using Deep Geometric Descriptors. In 2023 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV), pages 2612–2622, Waikoloa, HI, USA, 2023. IEEE.
- [7] Paul Bergmann, Xin Jin, David Sattlegger, and Carsten Steger. The MVTEC 3D-AD Dataset for Unsupervised 3D Anomaly Detection and Localization. In Proceedings of the 17th International Joint Conference on Computer Vision, Imaging and Computer Graphics Theory and Applications, pages 202–213, 2022. arXiv:2112.09045 [cs].
- [8] P.J. Besl and Neil D. McKay. A method for registration of 3-D shapes. IEEE Transactions on Pattern Analysis and Machine Intelligence, 14(2):239–256, 1992. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [9] S. Bickel, B. Schleich, and S. Wartzack. A Novel Shape Retrieval Method for 3D Mechanical Components Based on Object Projection, Pre-Trained Deep Learning Models and Autoencoder. Computer-Aided Design, 154:103417, 2023.
- [10] Tolga Birdal and Slobodan Ilic. Point Pair Features Based Object Detection and Pose Estimation Revisited. In 2015 International Conference on 3D Vision, pages 527–535, 2015.
- [11] Michael M. Bronstein and Iasonas Kokkinos. Scale-invariant heat kernel signatures for non-rigid shape recognition. In 2010 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, pages 1704–1711, 2010. ISSN: 1063-6919.
- [12] Michael C. Corballis. Bilaterally Symmetrical: To Be or Not to Be? Symmetry, 12(3):326, 2020. Number: 3 Publisher: Multidisciplinary Digital Publishing Institute.
- [13] Haowen Deng, Tolga Birdal, and Slobodan Ilic. PPF-FoldNet: Unsupervised Learning of Rotation Invariant 3D Local Descriptors, 2018. arXiv:1808.10322 [cs].
- [14] Haowen Deng, Tolga Birdal, and Slobodan Ilic. PPFNet: Global Context Aware Local Features for Robust 3D Point Matching, 2018. arXiv:1802.02669 [cs].
- [15] Ben Eckart, Kihwan Kim, and Jan Kautz. Fast and Accurate Point Cloud Registration using Trees of Gaussian Mixtures, 2018. arXiv:1807.02587 [cs].
- [16] Gil Elbaz, Tamar Avraham, and Anath Fischer. 3D Point Cloud Registration for Localization Using a Deep Neural Network Auto-Encoder. In 2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR), pages 2472–2481, Honolulu, HI, 2017. IEEE.
- [17] Francisco Figueroa, David Figueroa, Rodrigo Guiloff, Sven Putnis, Brett Fritsch, and Minerva Itriago. Navigation in anterior cruciate ligament reconstruction: State of the art. Journal of ISAKOS, 8(1):47–53, 2023.
- [18] Maximilian C. M. Fischer, Sonja A. G. A. Grothues, Juliana Habor, Matías de la Fuente, and Klaus Radermacher. A robust method for automatic identification of femoral landmarks, axes, planes and bone coordinate systems using surface models. Scientific Reports, 10(1):20859, 2020. Number: 1 Publisher: Nature Publishing Group.
- [19] Martin A. Fischler and Robert C. Bolles. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. Communications of the ACM, 24(6):381–395, 1981.
- [20] Wei Gao and Russ Tedrake. FilterReg: Robust and Efficient Probabilistic Point-Set Registration using Gaussian Filter and Twist Parameterization, 2019. arXiv:1811.10136 [cs].
- [21] Yang Gao, Honglin Yuan, Tao Ku, Remco C. Veltkamp, Georgios Zamanakos, Lazaros Tsochatzidis, Angelos Amanatiadis, Ioannis Pratikakis, Aliko Panou, Ioannis Romanelis, Vlassis Fotis, Gerasimos Arvanitis, and Konstantinos Moustakas. SHREC 2023: Point cloud change detection for city scenes. Computers & Graphics, 115:35–42, 2023.
- [22] Song Ge and Guoliang Fan. Articulated Non-Rigid Point Set Registration for Human Pose Estimation from 3D Sensors. Sensors, 15(7):15218–15245, 2015. Number: 7 Publisher: Multidisciplinary Digital Publishing Institute.
- [23] Benyamin Ghogh, Ali Ghods, Fakhri Karray, and Mark Crowley. Laplacian-Based Dimensionality Reduction Including Spectral Clustering, Laplacian Eigenmap, Locality Preserving Projection, Graph Embedding, and Diffusion Map: Tutorial and Survey, 2022. arXiv:2106.02154 [cs, stat].
- [24] Zan Gojcic, Caifa Zhou, Jan D. Wegner, and Andreas Wieser. The Perfect Match: 3D Point Cloud Matching With Smoothed Densities. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 5540–5549, Long Beach, CA, USA, 2019. IEEE.
- [25] Yulan Guo, Ferdous Sohel, Mohammed Bennamoun, Min Lu, and Jianwei Wan. Rotational Projection Statistics for 3D Local Surface Description and Object Recognition. International Journal of Computer Vision, 105(1):63–86, 2013. arXiv:1304.3192 [cs].
- [26] Yulan Guo, Mohammed Bennamoun, Ferdous Sohel, Min Lu, Jianwei Wan, and Ngai Ming Kwok. A Comprehen-

- sive Performance Evaluation of 3D Local Feature Descriptors. International Journal of Computer Vision, 116(1):66–89, 2016.
- [27] Gábor Holló and Mihály Novák. The manoeuvrability hypothesis to explain the maintenance of bilateral symmetry in animal evolution. Biology Direct, 7(1):22, 2012.
- [28] Jiayi Hu and Jing Hua. Salient spectral geometric features for shape matching and retrieval. The Visual Computer, 25(5):667–675, 2009.
- [29] Haibin Huang, Evangelos Kalogerakis, Siddhartha Chaudhuri, Duygu Ceylan, Vladimir G. Kim, and Ersin Yumer. Learning Local Shape Descriptors from Part Correspondences With Multi-view Convolutional Networks, 2017. arXiv:1706.04496 [cs].
- [30] Bing Jian and Baba C. Vemuri. Robust Point Set Registration Using Gaussian Mixture Models. IEEE Transactions on Pattern Analysis and Machine Intelligence, 33(8):1633–1645, 2011. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [31] Meghan Keast, Jason Bonacci, and Aaron Fox. Geometric variation of the human tibia-fibula: a public dataset of tibia-fibula surface meshes and statistical shape model. PeerJ, 11:e14708, 2023. Publisher: PeerJ Inc.
- [32] Kazuma Kobayashi, Lin Gu, Ryuichiro Hataya, Takaaki Mizuno, Mototaka Miyake, Hirokazu Watanabe, Masamichi Takahashi, Yasuyuki Takamizawa, Yukihiko Yoshida, Satoshi Nakamura, Nobuji Kouno, Amina Bolatkan, Yusuke Kurose, Tatsuya Harada, and Ryuji Hamamoto. Sketch-based Medical Image Retrieval, 2023. arXiv:2303.03633 [cs].
- [33] Adriaan Lambrechts, Roel Wirix-Speetjens, Frederik Maes, and Sabine Van Huffel. Artificial Intelligence Based Patient-Specific Preoperative Planning Algorithm for Total Knee Arthroplasty. Frontiers in robotics and AI, 9:840282, 2022.
- [34] Liang Li, Ming Yang, Chunxiang Wang, and Bing Wang. Rigid Point Set Registration Based on Cubature Kalman Filter and Its Application in Intelligent Vehicles. IEEE Transactions on Intelligent Transportation Systems, 19(6):1754–1765, 2018. Conference Name: IEEE Transactions on Intelligent Transportation Systems.
- [35] Zhouhui Lian, Afzal Godil, Benjamin Bustos, Mohamed Daoudi, Jeroen Hermans, Shun Kawamura, Yukinori Kurita, Guillaume Lavoué, Hien Van Nguyen, Ryutarou Ohbuchi, Yuki Ohkita, Yuya Ohishi, Fatih Porikli, Martin Reuter, Ivan Sipiran, Dirk Smeets, Paul Suetens, Hedi Tabia, and Dirk Vandermeulen. A comparison of methods for non-rigid 3D shape retrieval. Pattern Recognition, 46(1):449–461, 2013.
- [36] Z. Lian, J. Zhang, S. Choi, H. ElNaghy, J. El-Sana, T. Furuya, A. Giachetti, R. A. Guler, L. Lai, C. Li, H. Li, F. A. Limberger, R. Martin, R. U. Nakanishi, A. P. Neto, L. G. Nonato, R. Ohbuchi, K. Pevzner, D. Pickup, P. Rosin, A. Sharf, L. Sun, X. Sun, S. Tari, G. Unal, and R. C. Wilson. Non-rigid 3D shape retrieval. In Proceedings of the 2015 Eurographics Workshop on 3D Object Retrieval, pages 107–120, Goslar, DEU, 2015. Eurographics Association.
- [37] Eyal Lotan, Charlotte Tschider, Daniel K. Sodickson, Arthur L. Caplan, Mary Bruno, Ben Zhang, and Yvonne W. Lui. Medical Imaging and Privacy in the Era of Artificial Intelligence: Myth, Fallacy, and the Future. Journal of the American College of Radiology : JACR, 17(9):1159–1162, 2020.
- [38] David G. Lowe. Distinctive Image Features from Scale-Invariant Keypoints. International Journal of Computer Vision, 60(2):91–110, 2004.
- [39] Jiangyan Ma, Yifei Wang, and Yisen Wang. Laplacian Canonicalization: A Minimalist Approach to Sign and Basis Invariant Spectral Embedding, 2024. arXiv:2310.18716 [cs].
- [40] Deyo Maeztu Redin, Julien Caroux, Pierre-Yves Rohan, Hélène Pillet, Alexia Cermolacce, Julien Trnka, Mathieu Manassero, Véronique Viateau, and Laurent Corté. A wear model to predict damage of reconstructed ACL. Journal of the Mechanical Behavior of Biomedical Materials, page 105426, 2022.
- [41] Baraka Maiseli, Yanfeng Gu, and Huijun Gao. Recent developments and trends in point set registration methods. Journal of Visual Communication and Image Representation, 46:95–106, 2017.
- [42] Diana Mateus, Radu Horaud, David Knossow, Fabio Cuzzolin, and Edmond Boyer. Articulated Shape Matching Using Laplacian Eigenfunctions and Unsupervised Point Registration. In 2008 IEEE Conference on Computer Vision and Pattern Recognition, pages 1–8, 2008. arXiv:2012.07340 [cs].
- [43] Russell Merris. Laplacian matrices of graphs: a survey. Linear Algebra and its Applications, 197-198:143–176, 1994.
- [44] Kento Morita, Syoji Kobashi, Kaori Kashiwa, Hiroshi Nakayama, Shunichiro Kambara, Masakazu Morimoto, Shinichi Yoshiya, and Satoru Aikawa. Computer-aided Surgical Planning of Anterior Cruciate Ligament Reconstruction in MR Images. Procedia Computer Science, 60:1659–1667, 2015.
- [45] Andriy Myronenko and Xubo Song. Point-Set Registration: Coherent Point Drift. IEEE Transactions on Pattern Analysis and Machine Intelligence, 32(12):2262–2275, 2010. arXiv:0905.2635 [cs].
- [46] Thao Nguyen, Nakul Gopalan, Roma Patel, Matt Corsaro, Ellie Pavlick, and Stefanie Tellex. Robot Object Retrieval with Contextual Natural Language Queries, 2020. arXiv:2006.13253 [cs].
- [47] Daniel Nolte, Chui Kit Tsang, Kai Yu Zhang, Ziyun Ding, Angela E. Kedgley, and Anthony M. J. Bull. Non-linear scaling of a musculoskeletal model of the lower limb using statistical shape models. Journal of Biomechanics, 49(14):3576–3581, 2016.
- [48] Abdol Hamid Pilevar. CBMIR: Content-based Image Retrieval Algorithm for Medical Image Databases. Journal of Medical Signals and Sensors, 1(1):12–18, 2011.
- [49] Charles R. Qi, Hao Su, Kaichun Mo, and Leonidas J. Guibas. PointNet: Deep Learning on Point Sets for 3D Classification and Segmentation, 2017. arXiv:1612.00593 [cs].
- [50] Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. PointNet++: Deep Hierarchical Feature Learning on Point Sets in a Metric Space, 2017. arXiv:1706.02413 [cs].

- [51] Martin Reuter, Franz-Erich Wolter, and Niklas Peinecke. Laplace-spectra as fingerprints for shape matching. In Proceedings of the 2005 ACM symposium on Solid and physical modeling, pages 101–106, New York, NY, USA, 2005. Association for Computing Machinery.
- [52] Martin Reuter, Franz-Erich Wolter, and Niklas Peinecke. Laplace–Beltrami spectra as ‘Shape-DNA’ of surfaces and solids. Computer-Aided Design, 38(4):342–366, 2006.
- [53] Martin Reuter, Silvia Biasotti, Daniela Giorgi, Giuseppe Patanè, and Michela Spagnuolo. Discrete Laplace–Beltrami operators for shape analysis and segmentation. Computers & Graphics, 33(3):381–390, 2009.
- [54] R. Rostami, F. S. Bashiri, B. Rostami, and Z. Yu. A Survey on Data-Driven 3D Shape Descriptors. Computer Graphics Forum, 38(1):356–393, 2019. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13536](https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13536).
- [55] Raif M. Rustamov. Laplace-Beltrami Eigenfunctions for Deformation Invariant Shape Representation. The Eurographics Association, 2007. Accepted: 2014-01-29T09:43:15Z ISSN: 1727-8384.
- [56] Radu Bogdan Rusu, Nico Blodow, Zoltan Csaba Marton, and Michael Beetz. Aligning point cloud views using persistent feature histograms. In 2008 IEEE/RSJ International Conference on Intelligent Robots and Systems, pages 3384–3391, 2008. ISSN: 2153-0866.
- [57] Radu Bogdan Rusu, Nico Blodow, and Michael Beetz. Fast Point Feature Histograms (FPFH) for 3D registration. In 2009 IEEE International Conference on Robotics and Automation, pages 3212–3217, Kobe, 2009. IEEE.
- [58] Samuele Salti, Federico Tombari, and Luigi Di Stefano. SHOT: Unique signatures of histograms for surface and texture description. Computer Vision and Image Understanding, 125:251–264, 2014.
- [59] Dimple A. Shajahan, Mukund Varma T, and Ramanathan Muthuganapathy. Point Transformer for Shape Classification and Retrieval of 3D and ALS Roof PointClouds, 2021. [arXiv:2011.03921 \[cs\]](https://arxiv.org/abs/2011.03921).
- [60] Avinash Sharma, Radu Horaud, and Diana Mateus. 3D Shape Registration Using Spectral Graph Embedding and Probabilistic Matching, 2021. [arXiv:2106.11166 \[cs, stat\]](https://arxiv.org/abs/2106.11166).
- [61] Martin Sinko, Patrik Kamencay, Robert Hudec, and Miroslav Benco. 3D registration of the point cloud data using ICP algorithm in medical image analysis. In 2018 ELEKTRO, pages 1–6, 2018.
- [62] Jian Sun, Maks Ovsjanikov, and Leonidas Guibas. A Concise and Provably Informative Multi-Scale Signature Based on Heat Diffusion. Computer Graphics Forum, 28(5):1383–1392, 2009.
- [63] Sarah Tang and Afzal Godil. An evaluation of local shape descriptors for 3D shape retrieval. In Three-Dimensional Image Processing (3DIP) and Applications II, pages 217–231. SPIE, 2012.
- [64] Te Tang and Masayoshi Tomizuka. Track deformable objects from point clouds with structure preserved registration. The International Journal of Robotics Research, 41(6):599–614, 2022. Publisher: SAGE Publications Ltd STM.
- [65] Johan W. H. Tangelder and Remco C. Veltkamp. A survey of content based 3D shape retrieval methods. Multimedia Tools and Applications, 39(3):441–471, 2008.
- [66] Federico Tombari, Samuele Salti, and Luigi Di Stefano. Unique shape context for 3d data description. In Proceedings of the ACM workshop on 3D object retrieval, pages 57–62, New York, NY, USA, 2010. Association for Computing Machinery.
- [67] Søren Toxvaerd. The Emergence of the Bilateral Symmetry in Animals: A Review and a New Hypothesis. Symmetry, 13(2):261, 2021. Number: 2 Publisher: Multidisciplinary Digital Publishing Institute.
- [68] Yue Wang and Justin M. Solomon. Deep Closest Point: Learning Representations for Point Cloud Registration, 2019. [arXiv:1905.03304 \[cs\]](https://arxiv.org/abs/1905.03304) version: 1.
- [69] Yiqun Wang, Jianwei Guo, Dong-Ming Yan, Kai Wang, and Xiaopeng Zhang. A Robust Local Spectral Descriptor for Matching Non-Rigid Shapes With Incompatible Shape Structures. In 2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), pages 6224–6233, Long Beach, CA, USA, 2019. IEEE.
- [70] Martin Weinmann, Boris Jutzi, and Clément Mallet. Semantic 3D scene interpretation: A framework combining optimal neighborhood size selection with relevant features. ISPRS Annals of the Photogrammetry, Remote Sensing and Spatial Information Sciences, II-3:181–188, 2014.
- [71] Hao Wu, Lincong Fang, Qian Yu, and Chengzhan Yang. Learning Robust Point Representation for 3D Non-Rigid Shape Retrieval. IEEE Transactions on Multimedia, pages 1–15, 2023. Conference Name: IEEE Transactions on Multimedia.
- [72] Jin Xie, Guoxian Dai, Fan Zhu, Edward K. Wong, and Yi Fang. DeepShape: Deep-Learned Shape Descriptor for 3D Shape Retrieval. IEEE Transactions on Pattern Analysis and Machine Intelligence, 39(7):1335–1345, 2017. Conference Name: IEEE Transactions on Pattern Analysis and Machine Intelligence.
- [73] Guoqing Xu and Weiwei Fang. Shape retrieval using deep autoencoder learning representation. In 2016 13th International Computer Conference on Wavelet Active Media Technology and Information Processing (ICCWAMTIP), pages 227–230, 2016.
- [74] Ke Ye and Lek-Heng Lim. Schubert Varieties and Distances between Subspaces of Different Dimensions. SIAM Journal on Matrix Analysis and Applications, 37(3):1176–1197, 2016.
- [75] Zi Jian Yew and Gim Hee Lee. 3DFeat-Net: Weakly Supervised Local 3D Features for Point Cloud Registration. In Computer Vision – ECCV 2018, pages 630–646, Cham, 2018. Springer International Publishing.
- [76] Kwang Moo Yi, Eduard Trulls, Vincent Lepetit, and Pascal Fua. LIFT: Learned Invariant Feature Transform, 2016. [arXiv:1603.09114 \[cs\]](https://arxiv.org/abs/1603.09114).
- [77] Andy Zeng, Shuran Song, Matthias Nießner, Matthew Fisher, Jianxiong Xiao, and Thomas Funkhouser. 3DMatch: Learning Local Geometric Descriptors from RGB-D Reconstructions, 2017. [arXiv:1603.08182 \[cs\]](https://arxiv.org/abs/1603.08182).

- [78] Cha Zhang and Tsuhan Chen. Efficient feature extraction for 2D/3D objects in mesh representation. In Proceedings 2001 International Conference on Image Processing (Cat. No.01CH37205), pages 935–938 vol.3, 2001.
- [79] H. Zhang, O. Van Kaick, and R. Dyer. Spectral Mesh Processing. Computer Graphics Forum, 29(6):1865–1894, 2010.