

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

Supplementary Material

7. Coupled Laplacian

We present here some theoretical aspects of the Coupled Laplacian, introduced in Sec. 3, which we proposed to generate aligned spectral embeddings for multiple registered point clouds.

Zero Cross-Connections. The coupled Laplacian matrix has the following structure:

$$\mathbf{L}^C = \mathbf{L}^U + \mathbf{L}^+ \quad (6)$$

where \mathbf{L}^U is the Laplacian matrix of the global graph without any cross-connection between \mathcal{G}^T and \mathcal{G}^{S_k} , for $k = 1, \dots, N$. Hence, \mathbf{L}^+ is the Laplacian matrix of the global graph restricted to its cross-connections between sub-graphs. Moreover, \mathbf{L}^+ enforces the matching constraint in the global eigenproblem.

By construction of the matrix \mathbf{B}^C (see Sec. 3), eigenvectors related to $(\mathcal{G}^{S_k})_{k=1, \dots, N}$ and \mathcal{G}^T , that have the same eigenvalue λ , can be stacked into a global eigenvector ϕ^U such that the following property holds

$$\mathbf{L}^U \phi^U = \lambda \mathbf{B}^C \phi^U \quad (7)$$

where the above system is block diagonal. Therefore, solving the Coupled Laplacian eigenproblem without cross-connections, i.e. $l = 0$ and then $\mathbf{L}^+ = \mathbf{0}$, it is equivalent of solving separate eigenproblems, when the eigenvalues are the same, for each single component of the global graph. Otherwise, if eigenvalues are not the same, the solution of the coupled eigenproblems with $\mathbf{L}^+ = \mathbf{0}$ will likely give separate eigenvectors for each sub-graph, that is, vectors with non-zero values only for indices corresponding to a single component.

Ideal Matching. In the ideal case of perfect matching, when source graphs bare copies of the target graph, whether or not vertices are randomly reordered, the set of n^T coupled eigenvectors contains $N + 1$ copies of target's eigenvectors. For the sake of simplicity, in the following proof, the vertices in the copies of the target graph are not reordered. Therefore, $\mathcal{G}^{S_k} = \mathcal{G}^T$ and $F^{S_k} = F^T$, for $k = 1, \dots, N$. Let $\eta_i^C \in \mathbb{R}^{(N+1)n^T}$ be the vector in which $N + 1$ copies of an eigenvector ϕ_i^T are stacked. Then, from Eq. (6) and Eq. (7) we have,

$$\mathbf{L}^C \eta_i^C - \lambda_i^T \mathbf{B}^C \eta_i^C = \mathbf{L}^+ \eta_i^C \quad (8)$$

where non-zero rows in the right hand side term are related to cross-connections between a vertex of \mathcal{G}^T and its copies

on the references, at the same location. Hence, the weight function is equal to one, for each of the cross-connections, and

$$\forall j \in F^T \quad (\mathbf{L}^+ \eta_i^C)_j = \sum_{k=1}^N \phi_{ji}^T - \phi_{ji}^{S_k}. \quad (9)$$

Since $\phi_i^{S_k} = \phi_i^T$ for $\mathcal{G}^{S_k} = \mathcal{G}^T$, the right hand side term of Eq. (8) is equal to zero and $(\eta_i^C)_{i=1, \dots, n^T}$ are eigenvectors of the Coupled Laplacian containing $N+1$ copies of a target graph.

In addition to the above property, when target graph has eigenvalues with multiplicity higher than one, the coupled term $\mathbf{L}^+ \phi_i^C = 0$ enforces the shape matching of eigenvectors that do not belong to the kernel of the matrix \mathbf{L}^+ .

Penalization Term. The computation of matching vectors can be penalized in the initial eigenvectors by incorporating a coefficient $\alpha > 1$ into the Laplacian matrix for the coupled graph, as follows:

$$\mathbf{L}^{C\alpha} = \mathbf{L}^U + \alpha \mathbf{L}^+ \quad (10)$$

The larger α , the higher the eigenvalues of non matching eigenvectors such that $\|\mathbf{L}^+ \phi_i^C\| > 0$.

Modal Length. Let ϕ_k be the k -th eigenvector of a single graph $\mathcal{G}(\mathcal{V}, \mathcal{E})$, its gradient on the graph can be expressed as

$$\|\nabla \phi_k\|_{\mathcal{G}}^2 = \sum_{(i,j) \in \mathcal{E}} \left(\frac{\phi_{ik} - \phi_{jk}}{d(\mathbf{x}_i, \mathbf{x}_j)} \right)^2 \quad (11)$$

where the sum is over the set of edges \mathcal{E} and $d(\cdot, \cdot)$ is the Euclidean distance. We define the modal length as the ratio between the norm of the eigenvector and its gradient

$$L_k = \frac{\|\phi_k\|}{\|\nabla \phi_k\|_{\mathcal{G}}} \quad (12)$$

with $L_0 = +\infty$, because ϕ_0 is constant, and $L_k > L_{k+1}$ for $k \neq 0$. Note that the unit of L_k is the same as the unit given by the Euclidean distance of points in the graph. Therefore, the modal length can give an insight about the number of eigenmpas, m , to produce with the Coupled Laplacian, based on the extent of surface differences we want to detect. Intuitively, representing differences at small scales is easier than at large ones. Therefore, the more modes are used, the larger the range of defect sizes that can be represented.

Algorithm 1: Spectral-Based Bone Side Estimation (BSE) using aligned embeddings from the Coupled Laplacian.

Data: Source surface and side - (\mathcal{V}^S, S^S)
 Target surface - \mathcal{V}^T
 Neighbours - k
 Cross-Fraction - l
 Eigenmaps - m

Result: Target Side - S^T

$\mathbf{W}^S \leftarrow \text{Adjacency}(\mathcal{V}^S, k);$
 $\mathbf{W}^T \leftarrow \text{Adjacency}(\mathcal{V}^T, k);$
 $L^S, L^T \leftarrow \text{Bone lengths from Fiedler vectors};$
 $\alpha \leftarrow L^S/L^T; \quad \triangleright \text{Scale factor}$
 $\mathcal{V}^T \leftarrow \alpha \times \mathcal{V}^T;$
 $\mathcal{V}_M^S \leftarrow \text{Mirror Source using PCA};$
 $\mathcal{V}^S, \mathcal{V}_M^S \leftarrow \text{RANSAC registrations to } \mathcal{V}^T;$
 Select target cross-connections set $F^T \subset \mathcal{V}^T$
 Find sets $F^S \subset \mathcal{V}^S$ and $F_M^S \subset \mathcal{V}_M^S$ as in Eq. (3)
 $L^C \leftarrow \text{Coupled Laplacian}(\mathbf{W}^T, \mathbf{W}^S, F^T, F^S, F_M^S);$
 $\{\lambda_i^C\}_{i=0}^m, \{\phi_i^C\}_{i=0}^m \leftarrow \text{Eig}(L^C, m);$
 $\Phi^T, \Phi^S, \Phi_M^S \leftarrow \text{Split}(\{\phi_i^C\}_{i=0}^m);$
 QR decompose $\Phi^T(F^T, :), \Phi^S(F^S, :), \Phi_M^S(F_M^S, :);$
if $d_G(\mathbf{Q}^T, \mathbf{Q}^S) \leq d_G(\mathbf{Q}^T, \mathbf{Q}_M^S)$ **then**
 | $S^T = S^S; \quad \triangleright \text{Same side}$
else
 | $S^T = -S^S; \quad \triangleright \text{Opposite side}$
end

8. Spectral-Based BSE

In Algorithm 1 we provide full details on the spectral-based algorithm for BSE introduced in Sec. 4.1. It takes as input the set of points, i.e. graph vertices, representing source and target surfaces, \mathcal{V}^S and \mathcal{V}^T , respectively, and the target side, S^S . Additional input parameters are the number of neighbours used to create the k -NN graphs from the point clouds, the fraction of target nodes used to create cross-connection, l , and the dimension of the embeddings generated through the Coupled Laplacian, m .

As first step, the single adjacency matrices are built from the set of vertices using an RBF kernel, as in Eq. (1). After that, the length of each shape is computed only with the information carried by the Fiedler vector, f , as follows

$$L = \|x_{\max} - x_{\min}\|^2 \quad (13)$$

where $x_{\max} = x_{\arg \max f}$ and $x_{\min} = x_{\arg \min f}$ are the points corresponding to the maximum and minimum values of the Fiedler vector, respectively. In this way, the input surfaces do not need to lie on the same Euclidean frame in order to have comparable length measurements. Even if Fiedler vectors of different shapes may have opposite sign,




Figure 6. Graphical representation of the structure of the coupled adjacency matrix. \mathbf{W}^S and \mathbf{W}^T are the source and target adjacency matrices, while F^T , F^S and F_M^S are the set of nodes for cross-connections of target, source and mirrored source. Note that the nodes are sorted in groups such that the eigenmaps derived from the Coupled Laplacian can be easily split into the single components.

the inversion of maximum and minimum does not affect the bone length calculation. Note that, as extension for a more precise length computation, it is also possible to select the M maximum and minimum points of the Fiedler vector and compute x_{\max} and x_{\min} as the barycenters of those points. The ratio between source and target lengths gives a scaling factor, α , that we use to scale the target surface in order to match the length of the reference. The mirrored version of the source bone, \mathcal{V}_M^S , representing its contralateral, is then generated by flipping it around the point cloud second principal component and both versions of the reference are registered to the target with RANSAC algorithm [19]. Interestingly, the adjacency matrix of the two references, before and after registration, is the same as the original one and, therefore, we can avoid its computation multiple times.

After the described pre-processing, the core of the algorithm based on the Coupled Laplacian takes place. The subset of target nodes used for cross-connections is stochastically selected and the corresponding points on the two versions of the reference are queried as in Eq. (3). Alternatively, assuming not flat surfaces, other methods for cross-connections can be used instead of NN, e.g. spectral clustering medoids. Hence, the three graphs can now be connected and the Coupled Laplacian computed from the whole weighted adjacency matrix. For more clarity, the structure of the coupled graph adjacency matrix is depicted in Fig. 6. Therefore, the first m coupled eigenmaps are computed with a given eigensolver and split into the components of each of the three single shapes, as in Eq. (4). Finally, the Grassman distances, $d_G(\cdot, \cdot)$, of the QR normalized aligned eigenmaps restricted to the cross-connection sets are com-




Figure 7. Box plot of the accuracy of various methods for BSE depending on the source shape. Each measurement is the accuracy obtained using a different source shape among the ones available in the proposed benchmark. The median is highlighted with the notch while the mean is represented by the dashed line.

Table 4. Minimum and maximum accuracy of human BSE for each bone structures using different sources. All the matching methods are applied after RANSAC registration. The overall best performing methods are highlighted in boldface.

	Method	Femur	Hip	Fibula	Tibia
Best	Hausdorff	89.62	100.0	75.76	86.87
	Chamfer	85.85	100.0	86.87	90.91
	FPFH [57]	83.96	100.0	83.84	93.94
	Ours₂₀	93.40	100.0	89.90	94.95
	Ours₁₀	93.40	100.0	91.92	93.94
Worst	Hausdorff	51.89	89.74	39.39	51.51
	Chamfer	54.72	92.31	35.35	38.38
	FPFH [57]	34.34	87.18	34.34	41.41
	Ours ₂₀	49.06	94.87	38.38	39.39
	Ours₁₀	57.55	84.62	43.43	45.45

pared in order to predict the target side. The latter is equal to S^S if the Grassman distance between target and reference is lower than the distance between target and mirrored reference, otherwise it will be the opposite body side.

9. Robustness to Source Variation

In this section we briefly discuss the robustness with respect to the selected source shape of the methods tested in Sec. 4.2.

Bone Side Estimation. Fig. 7 shows the variation of the BSE accuracy, with respect to the source bone, for different bones and different methods. We recall that in each experiment one bone is chosen as source and the side is inferred on all the other bones of the same category. Therefore, we automatically have N_b accuracy measurements for each class, where N_b is the number of samples of the bone class b . The average accuracy of these experiments is reported in Tab. 3 of the main manuscript. While median and average accuracy are generally higher using our technique in all the bones, the values spread is similar between methods. Additionally, we report in Tab. 4 the maximum and the




Figure 8. Area under the PRO curve (AUC-PRO) with an upper integration limit of 0.3 for each category of the MVTEC 3D-AD dataset. Each score, represented by a dot, is computed with a different source shape selected from the train set. The median is highlighted with the notch while the mean is represented by the dashed line.

minimum accuracies obtained using different sources. We can observe that the maximum accuracy is always achieved with our method, both using 10 or 20 eigenmaps for matching. Moreover, it is also better in the worst case for all the bones classes except for the tibia, for which the Hausdorff distance seems to be better. Nevertheless, as also shown in Fig. 7, there is no much difference in the span between best and worst cases among different methods. This suggests that, to improve the performance of our BSE technique, and also of other methods, a proper choice of source bone is essential.

Anomaly Localization. We tested our method on the 3D anomaly localization task using 20 different anomaly-free sources randomly drawn from the train set of the MVTEC 3D-AD dataset. For each train sample, and for each category, we computed the normalized area under the PRO




Figure 9. Graphical comparison of 3D anomaly localization using different source shapes on different objects. The target example with RGB colors and the GT anomaly are shown on the left and right, respectively. In the middle 6 different results obtained changing the anomaly-free source are reported.

curve with an upper integration limit of 0.3 and their distribution is shown in Fig. 8. All the experiments are performed using the Coupled Laplacian with $m = 200$ and $l = 1$. We can observe that the results are much more stable and less spread than in the BSE task. This is likely due to the fact that different samples of the same industrially manufactured product are very similar and, therefore, all equally good as source shape. On the other hand, bones may vary in size and slightly in form from patient to patient. Hence, the selection of a good representative source is essential to enhance the capabilities of the proposed matching technique.

Moreover, in Fig. 9 we compare graphically the results obtained on some test samples using different anomaly-free sources. We can observe the strength of the proposed method which always correctly localizes the anomalies, even with different sources. In opposition, some false positives are sometimes present, especially on the borders of the shapes. This weakness is due to fact that the consid-

ered point clouds are not closed surfaces, since they are acquired with objects placed on a flat surface. Therefore, the points on the edges between object and surface are likely to contain noise which may produce different coupled embeddings leading to a detection of local surface differences. Interestingly, in the peach sample, some anomalies are located in the area of the stem end because only some anomaly-free sources has it. The peach test example shown in Fig. 9 does not have the stem and therefore, when compared with a source that has it, a surface difference is highlighted in that specific area.

10. Graphical Comparison

In this section we first provide more graphical comparisons of different methods on the anomaly localization task. After that, we show the difference between eigenmodes produced with the Coupled Laplacian with respect to the ones generate independently on the single geometries.




Figure 10. Graphical comparison of 3D anomaly localization using different methods on different objects. The target example with RGB colors and the GT anomaly are shown on the left and right, respectively. In the middle the results obtained with FPFH, Euclidean distance and distance between aligned spectral spaces are reported.

In Fig. 10 the anomaly scores obtained using FPFH features comparison [57], Euclidean distance and Coupled Laplacian, with $m = 200$ and $l = 1$, on different samples are compared. For a fair comparison, all the techniques are performed with the same anomaly-free source sample and after affine CPD [45] source to target registration. We observed that FPFH tends to overestimate anomalies, whereas the Euclidean distance underestimates them, leading to misleading detections. In contrast, our method achieves a better trade-off and improves the localization of anomalies. With reference to the modal length of Eq. (12), decreasing the number of modes used for spectral comparison we will obtain similar results to the Euclidean method, while increasing it, the localization will tend to the one of FPFH. Once again, the parameter m plays a crucial role in our method, helping us select surface differences that are tailored to the

specific task at hand.

In Fig. 11 and Fig. 12 the first 6 eigenmodes of a femur and hip BSE problem produced with the Coupled Laplacian are shown and compared with the ones obtained through independent eigendecomposition of target and source, i.e. uncoupled eigenmaps. In both cases the target side is opposite to the source and therefore the best matching coupled eigenmaps are with the mirrored source. This is because its RANSAC registration to the target shape is more precise and so the added cross-connections. On the contrary, the cross-edges between source and target are weaker leading to less precise coupled modes and larger Grassmann distance, which result in a correct BSE. In the uncoupled eigenmaps we can instead observe both the eigenvalue ordering and sign disambiguity of the eigendecomposition. In the femur example, eigenmode 3 has flipped sign in source and target,

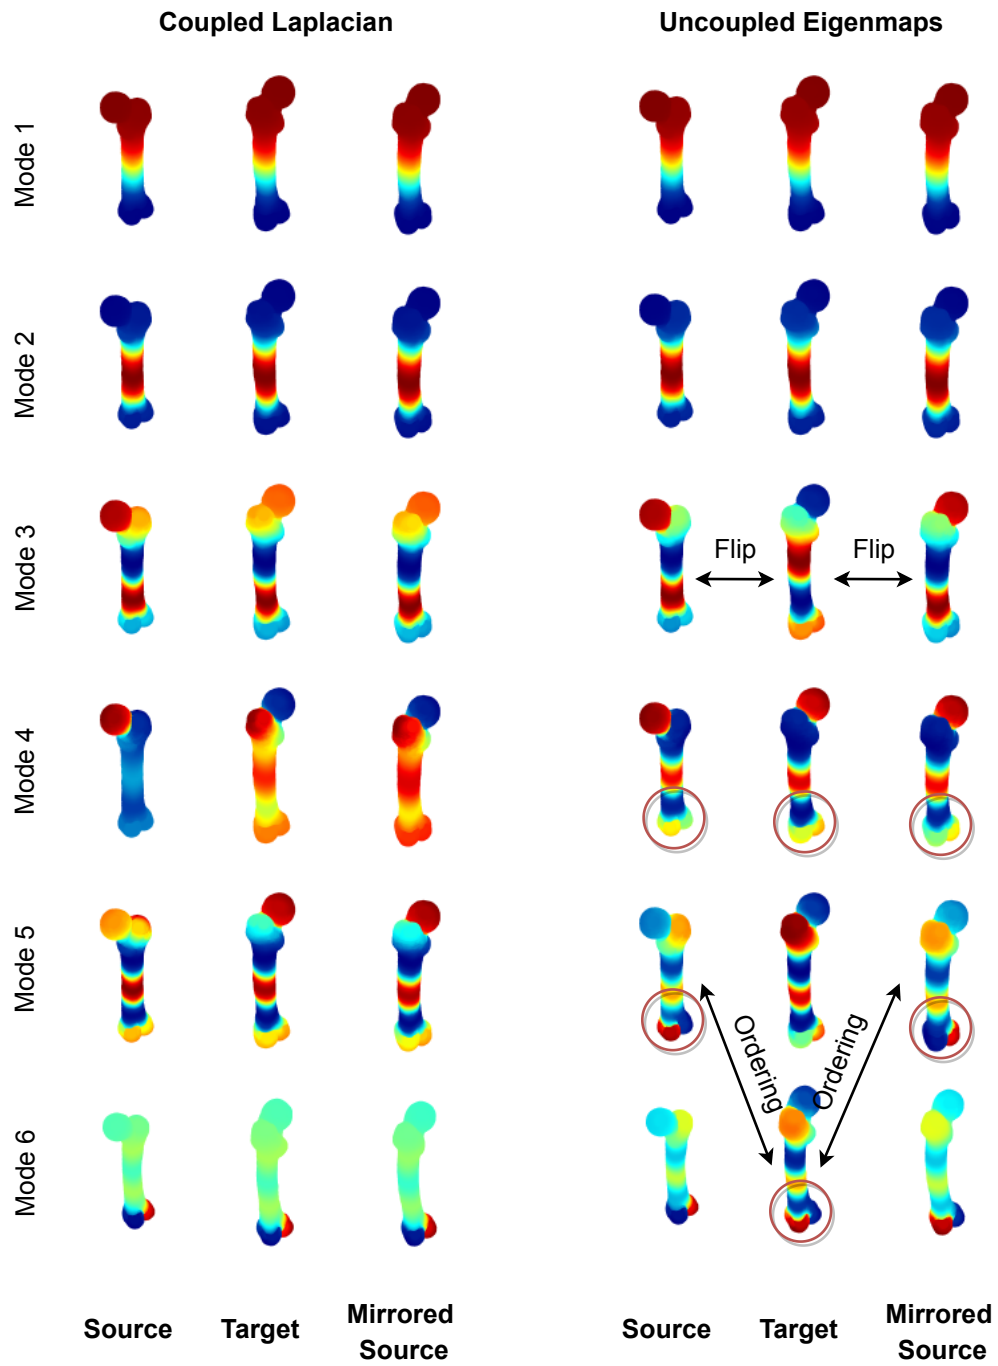


Figure 11. Graphical comparison of the eigenmaps obtained through Coupled Laplacian and with independent eigendecomposition of point cloud graphs, i.e. uncoupled eigenmaps. The first 6 eigenmodes are shown in both cases for source, target and mirrored target in the case of femur bones.

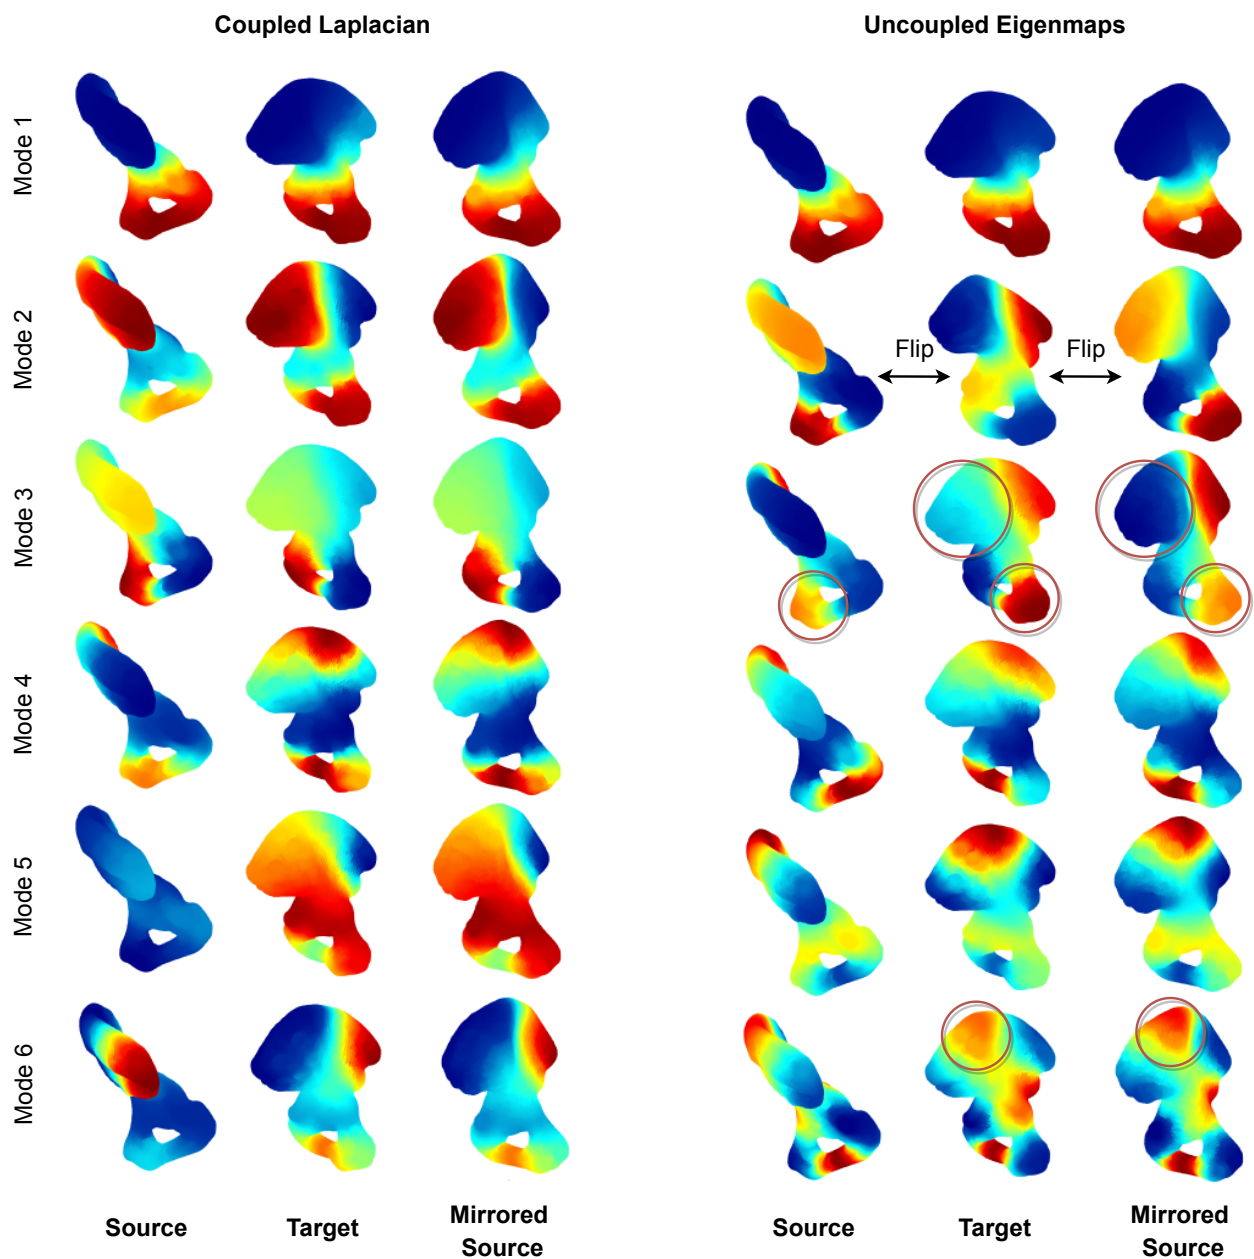


Figure 12. Graphical comparison of the eigenmaps obtained through Coupled Laplacian and with independent eigendecomposition of point cloud graphs, i.e. uncoupled eigenmaps. The first 6 eigenmodes are shown in both cases for source, target and mirrored target in the case of hip bones.

while eigenvalues 5 and 6 are inverted. Moreover, we can observe slight local variations of modes in the condyles area of the distal femur (highlighted with red circles). Similar observations can be also derived from the uncoupled modes of the hip bones in Fig. 12.

A comparison of higher modes is instead shown for two samples of the anomaly detection task in Fig. 13. In relation to the modal length defined in Eq. (12), it is evident that lower modes represent larger scales compared to higher modes, enabling them to better capture differences. More-

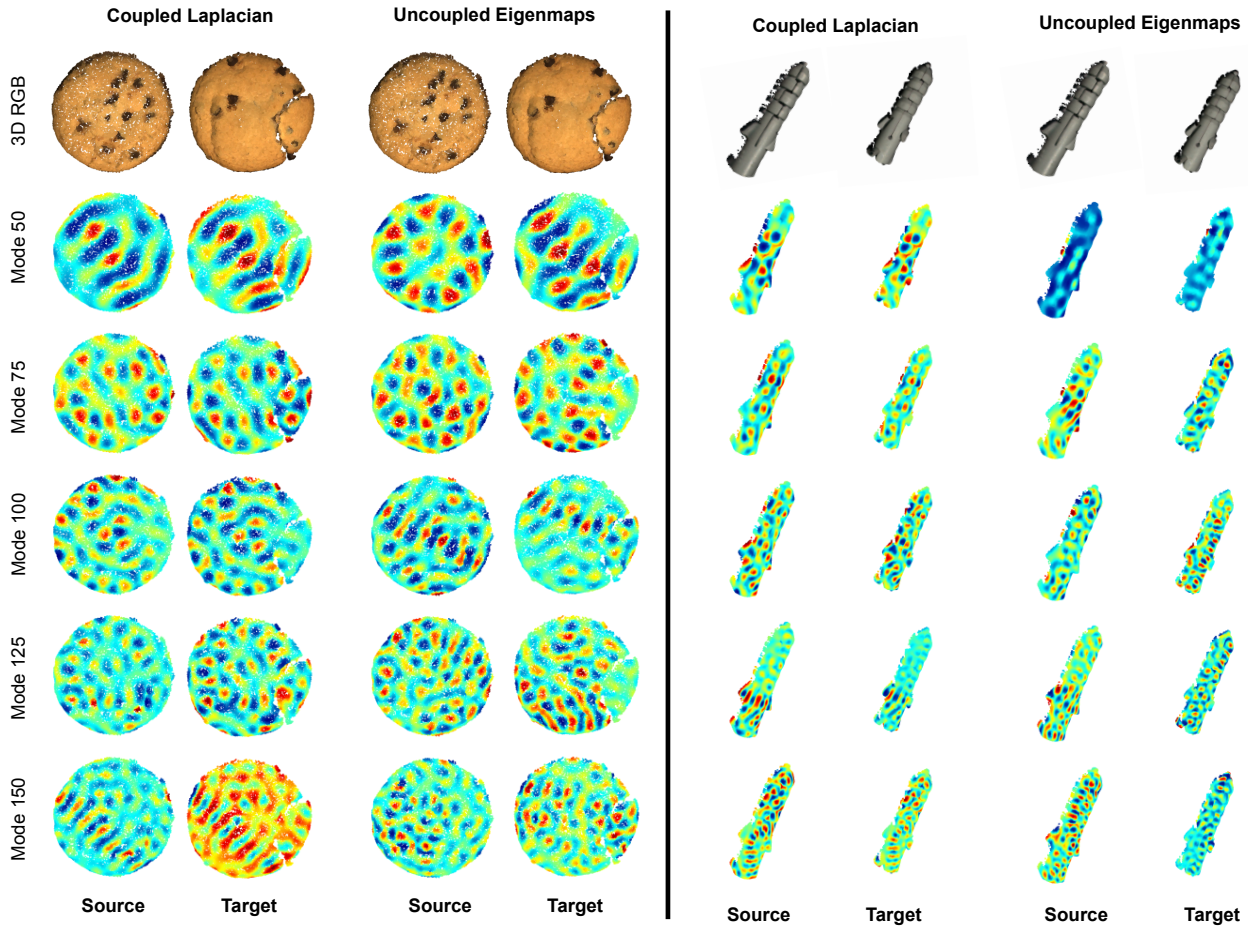


Figure 13. Graphical comparison of the eigenmaps obtained through Coupled Laplacian and with independent eigendecomposition of point cloud graphs, i.e. uncoupled eigenmaps. Some eigenmodes of increasing number are shown for two samples of the MVtec 3D-AD dataset.

over, the proposed coupling process leads to aligned eigenmaps for source and target shapes. Uncoupled eigenmodes exhibit completely different patterns, rendering them impractical for matching purposes without a reordering process, which complexity increases as the represented scale decreases. On the other hand, modes derived from the Coupled Laplacian exhibit the exact same pattern on both source and target point clouds, making them feasible for direct comparison. However, it is worth noting that as we increase the mode number, there might be an observable weakening in the coupling of the intensities of the eigenvectors, despite the preservation of the underlying patterns. In the computation of similarity scores using aligned embeddings, the observed coupling weakening with increasing mode numbers may not significantly impact the results. In fact, this characteristic could be leveraged to our advantage when determining the truncation number for the eigenspaces, highlighting

the adaptability and effectiveness of our method.

11. Cross-Species BSE

We test here the generalization of the proposed BSE algorithm on an internal dataset of sheep femur and tibia bones and we discuss its cross-species capabilities, i.e. when source and target bones are from different species. The latter poses an interesting challenge in the medical field, as tests are often performed on animals before being expanded to humans. The experiments only on sheep bones are performed as described in Sec. 4.1 for the public human benchmark. While, for the cross-species experiments, each bone of one species is chosen once as source to infer the side of the bones of the same category from the other species. The results of these sets of experiments are reported in Tab. 5. We can observe that, the sheep-to-sheep BSE is much more effective than the human-to-human comparison reported in

Table 5. Accuracy of cross-species BSE for femur and tibia using different methods. All the matching methods are applied after RANSAC registration. **S** and **H** stands for Sheep and Human, respectively, and the overall best performing methods are highlighted in boldface.

	Method	Femur	Tibia	Mean
S → S	Hausdorff	94.37	97.14	90.48
	Chamfer	94.29	88.57	94.88
	FPFH [57]	97.14	97.14	97.14
	Ours₂₀	97.30	97.30	97.30
	Ours₁₀	97.30	97.30	97.30
H → S	Hausdorff	72.22	69.44	70.83
	Chamfer	69.44	75.00	72.22
	FPFH [57]	77.78	75.00	76.39
	Ours₂₀	69.44	72.22	70.83
	Ours₁₀	83.33	77.78	80.56
S → H	Hausdorff	57.94	63.00	60.47
	Chamfer	62.62	66.00	64.31
	FPFH [57]	63.89	63.00	63.45
	Ours₂₀	63.55	60.00	61.78
	Ours₁₀	63.55	66.00	64.78

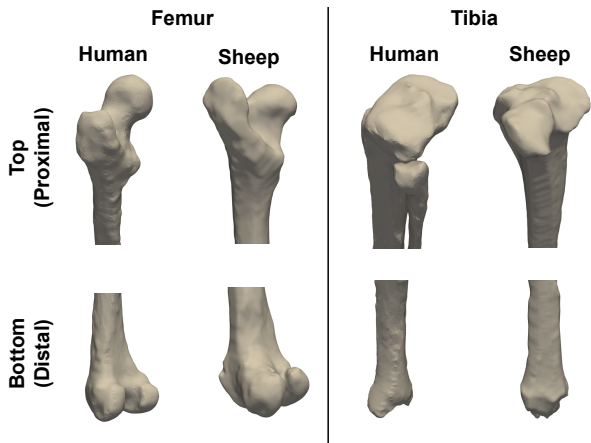


Figure 14. Comparison of human and sheep femur and tibia bones. For both of them, the top and lower parts are shown separately to highlight their differences.

the main manuscript, achieving almost correct estimation in all the sample. This is due to the fact that sheep femur and tibia have less symmetric shapes than their human counterparts, hence making easier the side estimation. Moreover, all the samples in our internal datasets are acquired with the same settings and therefore there is small sample variability, unlike the human benchmark we proposed to collect.

However, despite our method maintaining its superiority, performance drops are noted for cross-species tasks. A clearer understanding of this phenomenon can be obtained from Fig. 14, where human and sheep femur and tibia bones are graphically compared. While the overall shapes are similar, some major differences can be noticed between species.




Figure 15. Average computational time of the eigendecomposition using different combinations of number of shapes (1, 3 or 5) and modes (10 or 20).

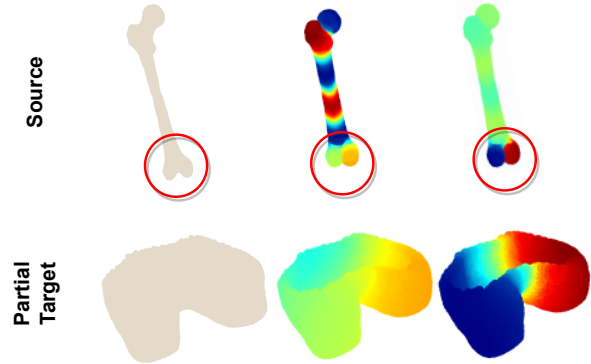


Figure 16. Coupled modes correctly aligned in a partial matching scenario.

For instance, the sheep proximal femur is similar to the human distal part. Such variations, and other artifacts, may affect the result of the initial registration, leading to wrong cross-connections and miss-aligned eigenmaps. Furthermore, while bones of different species are globally similar, local differences introduce complexities in the BSE task. The global similarity score derived from the Coupled Laplacian is therefore affected by these differences, leading to a more challenging task of aligning and comparing bones across species. Addressing this challenge will contribute to advancing the field of BSE and refining the understanding of global similarities and localized distinctions in diverse bone structures.

12. Scalability

The efficiency trade-off of decomposing larger Laplacians is illustrated in Fig. 15. The average computational time is reported for the eigendecomposition, using 10 and 20 maps, coupling the same shape with itself 3 or 5 times. As expected, larger, i.e. number of shapes, and denser, i.e. number of cross-connections, is the Laplacian, higher is the time needed for the computation. Scalability might be a concern in tasks like shape retrieval with numerous classes. However, simple measures, e.g. downsampling shapes and interpolating eigenmaps, can reduce the complexity. Given the

intrinsic coupling of modes, faster, but randomized, solvers, e.g. AMG, are still precise and have been used, cutting the need of resources.

13. Partial Matching

In Fig. 16 we show a graphical example of coupled eigenmaps derived from partially matching shapes. More in detail, an entire femur bone is matched with the distal part, i.e. lower part, of another femur. In this case, artificial cross-connections are added only on the partial match of the source shapes, leading to correctly aligned eigenmaps. Namely, the embeddings are successfully coupled in the distal area of the femur while they are independent on the remaining surface of the source. This example shows the potential of the Coupled Laplacian operator in different, and more complex, scenarios.