Spectral Meets Spatial: Harmonising 3D Shape Matching and Interpolation

Supplementary Material

In this supplementary document we first provide implementation details of our method. Next, we provide more ablative experiments to demonstrate the advantages of our method. Afterwards, we explain the preparation process of the medical data for our statistical shape analysis. Eventually, we show more qualitative results of our method.

10. Implementation details

Firstly, we provide definitions of our spectral regularisation L_{struct} in Eq. (2). The L_{bij} is the bijectivity loss to encourage the functional map from $\mathcal{X}$ through $\mathcal{Y}$ back to $\mathcal{X}$ is an identity map, and vice versa. It can be expressed in the form

$$L_{\text{bij}} = \|\mathbf{C}_{\mathcal{X}\mathcal{Y}}\mathbf{C}_{\mathcal{Y}\mathcal{X}} - \mathbf{I}\|_F^2 + \|\mathbf{C}_{\mathcal{Y}\mathcal{X}}\mathbf{C}_{\mathcal{X}\mathcal{Y}} - \mathbf{I}\|_F^2 . \quad (20)$$

The L_{orth} is the orthogonality loss to prompt a locally area-preserving matching, see [73] for more details. It can be expressed in the form

$$L_{\text{orth}} = \left\|\mathbf{C}_{\mathcal{X}\mathcal{Y}}^{\top}\mathbf{C}_{\mathcal{Y}\mathcal{X}} - \mathbf{I}\right\|_F^2 + \left\|\mathbf{C}_{\mathcal{X}\mathcal{Y}}^{\top}\mathbf{C}_{\mathcal{Y}\mathcal{X}} - \mathbf{I}\right\|_F^2 . \quad (21)$$

In the following we explain each component in our framework in detail. Following prior works [13, 49], we use DiffusionNet [76] as our feature extractor and wave kernel signature (WKS) [4] with 128 dimensions as input features, since it is agnostic to shape discretisation and orientation. The dimension of the output features is 256 for non-isometric datasets and 384 for near-isometric datasets. In the context of functional map computation, we use the regularised functional map solver [67] based on the resolvent mask $\mathbf{M}$, in which the regularisation term E_{reg} in Eq. (1) can be expressed in the form

$$E_{\text{reg}} = \sum_{ij} \mathbf{C}_{ij}^2 \mathbf{M}_{ij}, \quad (22)$$

where

$$\mathbf{M}_{ij} = \left(\frac{\mathbf{\Lambda}_{\mathcal{Y}}(i)^{\gamma}}{\mathbf{\Lambda}_{\mathcal{Y}}(i)^{2\gamma}+1} - \frac{\mathbf{\Lambda}_{\mathcal{X}}(j)^{\gamma}}{\mathbf{\Lambda}_{\mathcal{X}}(j)^{2\gamma}+1}\right)^2 + \left(\frac{1}{\mathbf{\Lambda}_{\mathcal{Y}}(i)^{2\gamma}+1} - \frac{1}{\mathbf{\Lambda}_{\mathcal{X}}(j)^{2\gamma}+1}\right)^2 . \quad (23)$$

The regularisation strength λ in Eq. (1) is 100. The number of eigenfunctions $\mathbf{\Phi}$ and eigenvalues $\mathbf{\Lambda}$ used for functional map computation is 200 for non-isometric datasets and 300 for near-isometric datasets. In terms of our spectral loss L_{spectral}, we empirically set $\lambda_{\text{bij}} = 1, \lambda_{\text{orth}} = 1$ in Eq. (2), $\lambda_{\text{struct}} = 1, \lambda_{\text{couple}} = 1$ in Eq. (8). For our spatial loss L_{spatial}, we empirically set $\lambda_{\text{align}} = 5, \lambda_{\text{arap}} = 100, \lambda_{\text{sym}} = 1, \lambda_{\text{var}} = 1$ in Eq. (15). We use Adam [45] optimiser with learning rate equal to 10^{-3}. In the context of test-time adaptation, the shape-dominant deformation field $\mathbf{\Delta}_s(t)$ is initialised as all zeros. The Dirichlet energy in Eq. (16) can be expressed in the form

$$L_{\text{D}} = \|\mathbf{\Delta}\|_{\mathbf{L}}^2 , \quad (24)$$

where $\|\mathbf{X}\|_L = \text{Trace}(\mathbf{X}^{\text{T}}\mathbf{L}\mathbf{X})$. We empirically set $\lambda_{\text{D}} = 0.1$. The number of iterations for optimisation is 2×10^3.

11. Additional evaluations

11.1. Smoothness of point-wise maps

In Sec. 5.1 we demonstrate that our method substantially outperforms existing state-of-the-art shape matching methods in terms of matching accuracy based on mean geodesic error. In this experiment, we evaluate the matching smoothness based on the conformal distortion metric [41] that is also used in [24, 27]. We compare our method with the

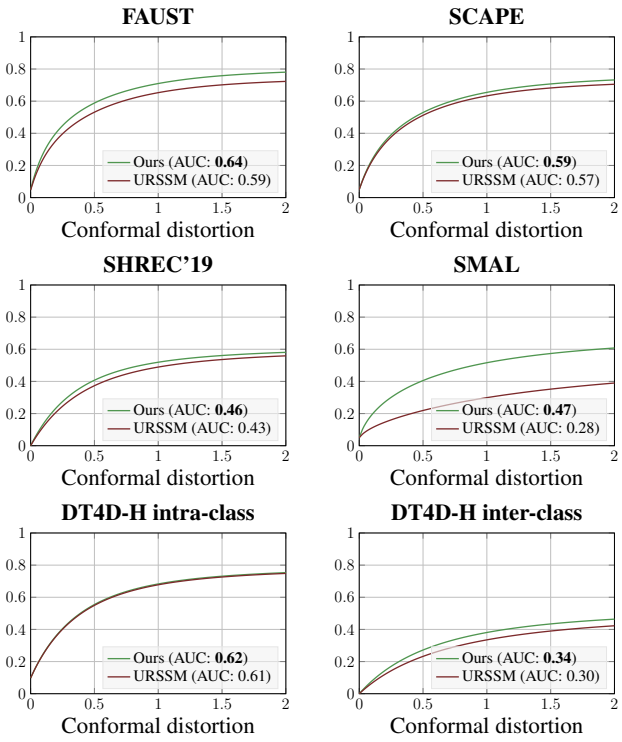

Figure 11. **Matching smoothness on all datasets in our matching experiments.** Our method obtains smoother point-wise correspondences based on the combination of spectral and spatial maps.

state-of-the-art matching method (i.e. URSSM [13]) based on deep functional map framework. Fig. 11 summarises the results on all datasets in our matching experiments. Our method obtains smoother point-wise correspondences by harmonising spectral and spatial maps.

11.2. Matching on topologically noisy data

We evaluate the performance of our method on topologically noisy data. Such topological noise presents a great challenge to shape matching methods, especially for matching methods with spatial regularisation [30]. To this end, we use the TOPKIDS dataset [46], which contains synthetic shapes of children with topological merging. Due to the small amount of training data, we only consider axiomatic and unsupervised methods for comparison, similar to [13, 30]. Tab. 6 summarises the matching results. We observe that our method obtains comparable but slightly worse performance compared to state-of-the-art, due to the incorporation of explicit spatial regularisation. Meanwhile, our method outperforms existing shape matching methods with explicit spatial regularisation. Moreover, our spatial regularisation ensures that our method does not suffer from symmetry flip as shown in Fig. 12.

Geo. error (×100)	**TOPKIDS**	**Spatial reg.**
Axiomatic Methods		
ZoomOut [61]	33.7	✗
Smooth Shells [27]	11.8	✓
DiscreteOp [68]	35.5	✗
Unsupervised Methods		
WSupFMNet [75]	47.9	✗
Deep Shells [28]	13.7	✓
NeuroMorph [29]	13.8	✓
AttnFMaps [49]	23.4	✗
URSSM [13]	**9.2**	✗
Ours	9.4	✓

Table 6. **Matching on TOPKIDS dataset.** Our method achieves comparable but slightly worse performance compared to the state-of-the-art method, while outperforming existing shape matching methods with explicit spatial regularisation.

Figure 12. **Qualitative comparison on TOPKIDS dataset.** By incorporating spatial regularisation, our method does not suffer from front-back flips that occurs for functional map methods.

12. Preparation of lung shapes

To generate the lung meshes used in our statistical shape analysis experiment (see Sec. 5.3), we randomly select 22 CT images from the 9th subset of the LUNA dataset [74]. Each CT image is accompanied by segmentation masks of the lung structures. We first use the Marching Cubes algorithm [57] to extract triangular meshes for the left lungs. After extraction, we remove outliers by clustering triangle meshes and retaining only the biggest triangle mesh in terms of the number of triangles. Next, we employ a Laplacian smoothing algorithm to smooth the triangle mesh. Finally, we apply the Quadric Error Metric Decimation [34], resulting in simplified mesh structures comprising approximately 8000 triangles per mesh.

13. More qualitative results

In this section we provide more qualitative results obtained by our method for shape matching and interpolation.

13.1. Shape matching

Figure 13. **Qualitative shape matching results of our method on the SHREC'19 dataset.** The top-left shape is the source shape that is matched to the other shapes. Our method obtains accurate matchings for human shapes with diverse poses and appearances.

Figure 14. **Qualitative shape matching results of our method on the SMAL dataset.** The top-left shape is the source shape that is matched to the other shapes. Our method obtains accurate correspondences for shapes in different classes.

Figure 15. **Qualitative shape matching results of our method on the DT4D-H inter-class dataset.** The top-left shape is the source shape that is matched to the other shapes. Our method obtains accurate correspondences for non-isometric deformed shapes.

13.2. Shape interpolation

In this subsection we demonstrate the shape interpolation results of our method on different datasets. Additionally, we show the matching results by texture transfer from the source shape to the target shape. We observe that our method can obtain accurate point-wise correspondences and realistic shape interpolation trajectories even under large non-isometry and pose variations.

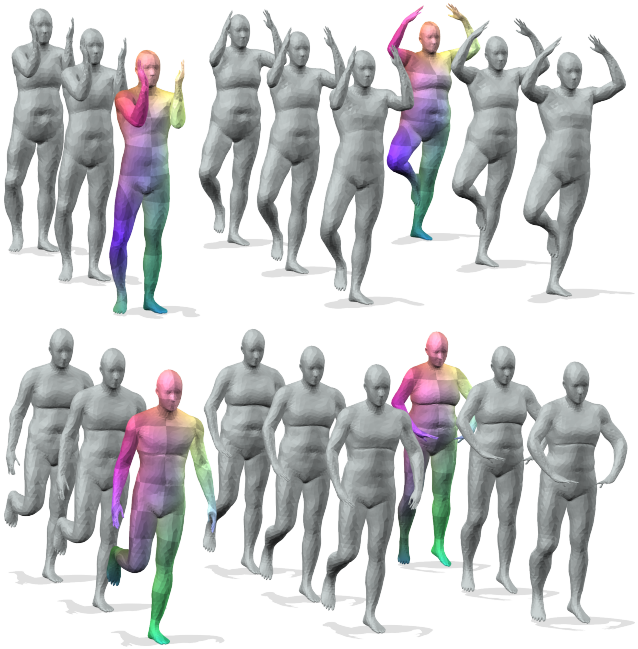

Figure 16. **Qualitative shape interpolation results of our method on the FAUST dataset.** Our method obtains realistic interpolation trajectories that capture both pose-dominant (horizontal) and shape-dominant (vertical) deformations.

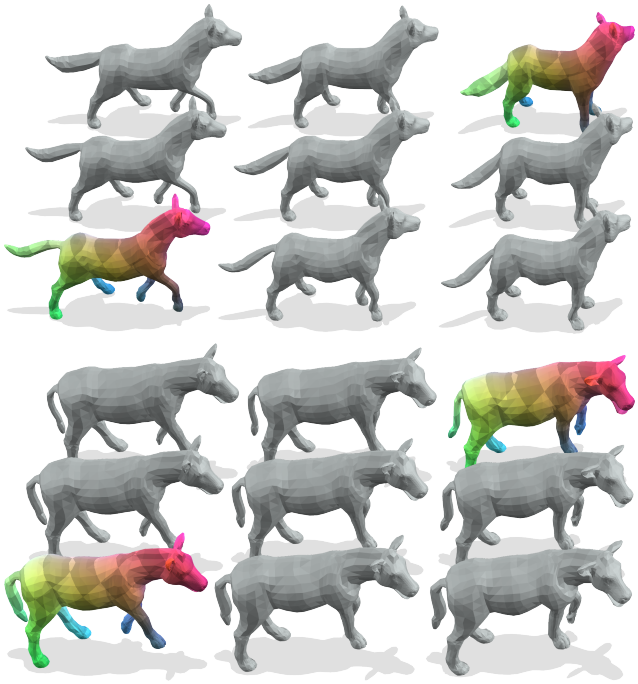

Figure 17. **Qualitative shape interpolation results of our method on the SMAL dataset.** Our method obtains realistic interpolation trajectories between different shape categories.

Figure 18. **Qualitative shape interpolation results of our method on the MANO dataset.** Our method obtains realistic interpolation trajectories between hands in different poses.

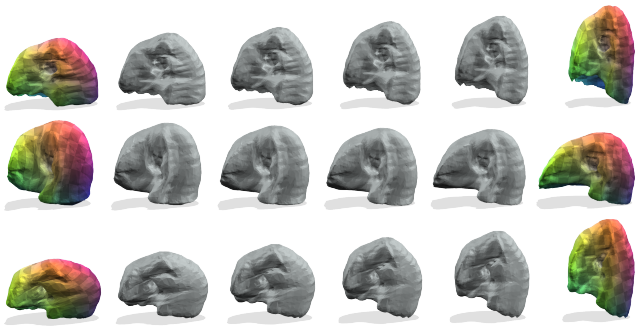

Figure 19. **Qualitative shape interpolation results of our method on the LUNA dataset.** Our method obtains realistic interpolation trajectories between lungs despite large non-isometries.