


A Universal and Flexible Framework for Unsupervised Statistical Shape Model Learning

Nafie El Amrani , Dongliang Cao, and Florian Bernard

University of Bonn, Bonn, Germany
n.elamrani@uni-bonn.de

Abstract. We introduce a novel unsupervised deep learning framework for constructing statistical shape models (SSMs). Although unsupervised learning-based 3D shape matching methods have made a major leap forward in recent years, the correspondence quality of existing methods does not meet the demanding requirements necessary for the construction of SSMs of complex anatomical structures. We address this shortcoming by proposing a novel *deformation coherency loss* to effectively enforce smooth and high-quality correspondences during neural network training. We demonstrate that our framework outperforms existing methods in creating high-quality SSMs by conducting extensive experiments on five challenging datasets with varying anatomical complexities. Our proposed method sets the new state of the art in unsupervised SSM learning, offering a universal solution that is both flexible and reliable. Our source code is publicly available at <https://github.com/NafieAmrani/FUSS>.

Keywords: Unsupervised 3D Correspondence Matching · 3D Shape Matching · Statistical Shape Modelling.

1 Introduction

Statistical shape models (SSM) are an essential tool in medical image analysis and computational anatomy, facilitating a deeper understanding of anatomical variability across populations. Its applications are diverse and range from diagnosis [24], pathology detection [22] to treatment planning [6].

Despite the evident utility of SSMs, they often come with the drawback of depending on (some form of) human supervision for their construction. This process is labour-intensive, prone to errors and can also introduce bias, affecting the reliability of the resulting SSMs [28]. With that, a large body of research has evolved that studies the automated construction of SSMs. Traditional methods for the automated construction of SSMs rely on axiomatic optimisation formalisms to establish correspondences across a collection of 3D shapes [4, 8, 11]. In recent years, several deep learning methods for SSM construction were proposed, some of them even without relying on labelled data [2, 13]. Despite these remarkable achievements, existing methods do not meet the demanding requirements necessary for the construction of SSMs of complex anatomical structures. For example, recent unsupervised learning-based methods require post-processing

heuristics to mend inaccurate, non-smooth and non-bijective point-wise matchings [3, 23] – while this improves the matching quality to some extent, the two-stage processing is suboptimal since respective heuristics are generic and not tailored towards the need of SSMs (e.g. having spatially smooth and coherent deformations between shapes).

To address these issues, we introduce a novel unsupervised SSM learning framework based on a *deformation coherency regularisation*. Our approach handles a wide spectrum of datasets, encompassing anatomical structures with considerable variability in shape and size, such as the liver, pancreas, lung, spleen, and hippocampus. We summarise our main contributions as follows:

1. We propose a fully unsupervised framework for SSM learning that offers a universal and flexible solution for a broad and diverse range of shapes.
2. From a technical perspective, we achieve this by introducing a deformation coherency loss to ensure smooth and accurate deformations, which we fuse into a recent state-of-the-art shape matching method to enable accurate unsupervised SSM construction.
3. We set new state-of-the-art performance on five challenging datasets characterised by high shape variability and different anatomical sizes, demonstrating our method’s strong ability to produce high-quality SSMs.

2 Related Work

Recent studies have proposed the direct generation of SSMs from volumetric images [1, 5]. These techniques typically rely on supervised learning paradigms and require pre-existing Point Distribution Models (PDMs) for training. In contrast, our work focuses on constructing SSMs from mesh data. Therefore, this section will explore methodologies for creating SSMs from mesh data, which can be categorised into correspondence-based and deformation-based approaches.

Correspondence-based SSM. The construction of SSMs typically begins with establishing correspondences across a population of shapes. This process has evolved from traditional pair-wise registration methods [12, 21] to more sophisticated group-wise optimisation techniques [4, 8, 11]. However, they use time-consuming iterative optimisation strategies, hindering their application to large medical datasets. Recently, several methods [3, 15] proposed the use of functional maps [20] to abstract the notion of point-to-point correspondences into finding correspondences in the spectral domain. However, the map in the spectral domain does not guarantee smooth bijective point-wise correspondences. Therefore, both μ Match [15] and S3M [3] require extra post-processing steps to fix non-smooth point-wise correspondences – this two-stage processing may produce suboptimal point-wise correspondences that impairs SSM quality (which we experimentally demonstrate in Sec. 5). In contrast, our method combines functional and spatial maps to estimate *smooth* point-to-point correspondences which are then used in learning the deformation field between training shapes.

Deformation-based SSM. Several methods [10, 17, 19] have been proposed to estimate deformation fields that map a template shape to shapes in the population. However, they are computationally expensive and constrained by a finite set of momentum vectors that parameterise the deformation field [9]. To overcome these issues, FlowSSM [16] and ShapeFlow [14] proposed to use an encoder-free setup and parameterise the deformation field within the low-dimensional latent space of neural networks. Nonetheless, this methodology complicates the reconstruction of unseen shapes, necessitating the optimisation of a randomly initialised latent representation to model the optimal deformation. More recently, Mesh2SSM [13] proposed to estimate the deformation of a template point cloud to subject-specific meshes. However, it embeds the shape variations into the latent space of a variational autoencoder, thereby complicating interpretability of the resulting SSM. Unlike existing methods, our method parameterises a smooth and non-rigid deformation with the help of accurate and smooth point-to-point correspondences, thereby resulting in a continuous deformation field that reliably captures biological deformations. We experimentally demonstrate that this combination of point-wise and deformation-based correspondence representation leads to high-quality SSMs.

3 Background: Spectral Meets Spatial Shape Matching

Our methodology builds upon the recent state-of-the-art *Spectral meets Spatial* (SmS) shape matching framework [7]. SmS is designed to predict point-wise correspondences while simultaneously interpolating between 3D shapes. To this end, it leverages a combination of spectral and spatial self-supervision. The core of SmS lies in its ability to handle *articulated* objects (e.g. humans, or animals) by utilising an As-Rigid-As-Possible (ARAP) deformation energy [27] to capture pose-dominant transformations effectively. Additionally, SmS relies on a post-processing strategy (test-time adaptation) to better capture shape-dominant deformations, further enhancing its interpolation capabilities.

Despite its efficacy for articulated objects, the SmS framework exhibits limitations when constructing SSMs of anatomical structures. While the ARAP deformation energy is well-suited to model articulation (e.g. of human or animal motion), it is based on locally rigid transformations and is thus inadequate for modelling anatomical shape variability across a population (e.g. of organs), see our experimental results in Sec. 5.

To compensate for these shortcomings, in this work, we build upon SmS and replace their ARAP loss with a novel *deformation coherency regularisation* designed for the construction of SSMs. Our deformation coherency term combines the Chamfer distance and Dirichlet energy [18] to capture smooth and geometrically plausible shape deformations. This dual-faceted loss eliminates the need for test-time post-processing (as used in SmS) by directly capturing shape-dominant deformations during training.

4 Shape Model Learning with Deformation Coherency

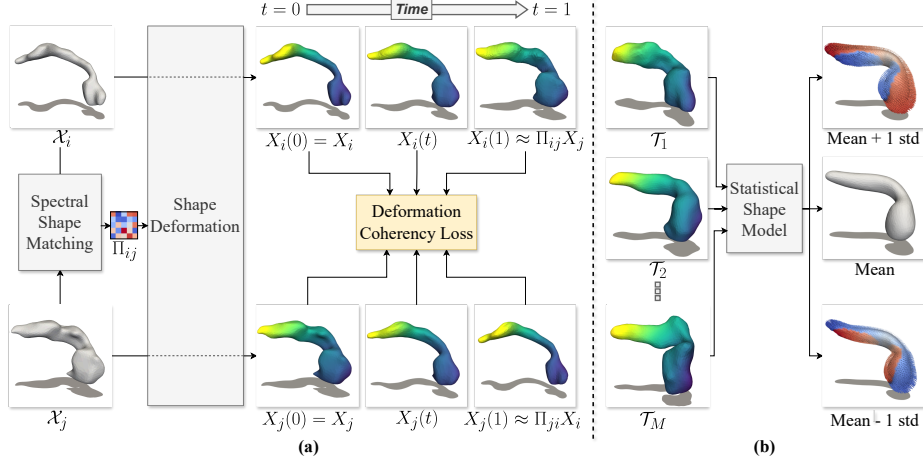


Fig. 1. Method overview. **(a)** First, the *Spectral Shape Matching* module computes point-wise maps Π_{ij} between shapes \mathcal{X}_i and \mathcal{X}_j . Subsequently, the *Shape Deformation* network predicts two deformation trajectories from \mathcal{X}_i to \mathcal{X}_j and vice versa. During training, our novel *Deformation Coherency Loss* is used to regularise the deformation fields. **(b)** At test-time, a chosen template \mathcal{T} is deformed to all M training shapes forming the set of shapes $\{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_M\}$ (in point-wise correspondence), which are then used to construct an SSM.

Given is a set of M shapes $\mathbf{X} = \{\mathcal{X}_1, \mathcal{X}_2, \dots, \mathcal{X}_M\}$, where each shape \mathcal{X}_i is represented as a 3D triangular mesh with n_i vertices encoded as matrix $X_i \in \mathbb{R}^{n_i \times 3}$. Note that the input shapes \mathbf{X} need to be spatially aligned. Using an unsupervised learning approach, our method aims to establish accurate point-wise correspondences and realistic shape deformations between pairs of shapes. Eventually, we use these correspondences to construct an SSM (in the space of a template shape \mathcal{T}) that captures shape variation across \mathbf{X} . Our overall processing pipeline is summarised in Fig. 1. We use a *Spectral Shape Matching* module to generate a point-wise map Π_{ij} between two input shapes \mathcal{X}_i and \mathcal{X}_j . This map serves as an input for the *Shape Deformation* network, which predicts the deformation field $\Delta_i(t) \in \mathbb{R}^{n_i \times 3}$ that smoothly transforms the vertices X_i of \mathcal{X}_i along the trajectory $X_i(t) = X_i + \Delta_i(t)$ for $t \in [0, 1]$. This trajectory deforms the original 3D positions $X_i(0) = X_i$ to new positions $X_i(1) \approx \Pi_{ij}X_j$, approximating the corresponding vertices of \mathcal{X}_j (analogous for transforming \mathcal{X}_j to \mathcal{X}_i). We refer to SmS [7] for more details about both modules.

Deformation Coherency Regularisation. We propose a novel *deformation coherency regularisation* to ensure smooth and geometrically plausible deforma-

tions. To this end, we uniformly sample $K+1$ timesteps in the time interval $[0, 1]$, and consider the sequence of deformed shapes $X_i^k := X_i(k/K)$ for $k = 0, \dots, K$. To ensure precise spatio-temporal alignment of shapes throughout the deformation, we penalise differences between consecutive shapes in the deformation sequence using

$$L_{\text{cd}} = \sum_{k=0}^{K-1} f_{\text{cd}}(X_i^k, X_i^{k+1}), \quad (1)$$

where $f_{\text{cd}}(\cdot, \cdot)$ is the bi-directional Chamfer distance. Furthermore, we use the Dirichlet energy [18] to promote spatio-temporally smooth deformations, i.e.

$$L_{\text{dir}} = \sum_{k=0}^{K-1} \text{Tr}(\delta_i(k)^T L_i \delta_i(k)), \quad (2)$$

where $\text{Tr}(\cdot)$ stands for the matrix trace, L_i is the Laplace-Beltrami operator of \mathcal{X}_i , and $\delta_i(k) = X_i^k - X_i^{k+1}$ is the vertex offset between \mathcal{X}_i^k and \mathcal{X}_i^{k+1} . The losses in Eqns. (1) and (2) are analogously applied to the deformation of shape \mathcal{X}_j .

To further promote local area-preserving deformations, we introduce L_{edge} , which maintains the proximity of neighbouring vertices, i.e.

$$L_{\text{edge}} = \|E_i X_i^K\|_F^2 + \|E_j X_j^K\|_F^2, \quad (3)$$

where E_i and E_j are the incidence matrices of (the mesh graphs of) shapes \mathcal{X}_i and \mathcal{X}_j , respectively. Our overall deformation coherency loss is

$$L_{\text{dc}} = \lambda_{\text{cd}} L_{\text{cd}} + \lambda_{\text{dir}} L_{\text{dir}} + \lambda_{\text{edge}} L_{\text{edge}}. \quad (4)$$

L_{dc} effectively regularises the shape deformations, while eliminating the need for test-time post-processing (required by SmS).

Training, Inference and SSM Construction. During training, pairs of shapes are sampled from the dataset. At test time, our method produces pair-wise deformed shapes. To accumulate these over the entire dataset of M shapes, we select a template shape \mathcal{T} (with $n_{\mathcal{T}}$ vertices) as the sample with the smallest total loss to all other shapes. Then, we use our method to deform \mathcal{T} to all other shapes, forming a set of M shapes $\mathbf{T} = \{\mathcal{T}_1, \mathcal{T}_2, \dots, \mathcal{T}_M\}$ that share the same topology as \mathcal{T} and are in point-wise correspondence, i.e. the s -th vertex in shape \mathcal{T}_i corresponds to the s -th vertex of all other shapes in \mathbf{T} .

To build the SSM we use PCA, i.e. we calculate the mean shape $\mu \in \mathbb{R}^{3n_{\mathcal{T}}}$ of the (vectorised vertex positions of) shapes in \mathbf{T} , and the covariance matrix $S \in \mathbb{R}^{3n_{\mathcal{T}} \times 3n_{\mathcal{T}}}$ with eigenvectors $\{v_i\}$ and eigenvalues $\{\lambda_i\}$. We can then model the distribution of the shapes as $\mu + \sum_i \alpha_i \lambda_i v_i$, with $\alpha_i \sim \mathcal{N}(0, 1)$.

5 Experiments

Datasets. We evaluate our method on five publicly available datasets with diverse anatomical and morphological characteristics: the pancreas, spleen, liver

and hippocampus datasets from the Medical Segmentation Decathlon (MSD) challenge [26], and the lung dataset from the LUNA16 challenge [25] (see Table S2 in the supplementary material for more details). We use ShapeWorks [8] to smoothen, centre and convert the (segmented) CT/MRI images to triangular meshes of around 2000 vertices (our method can handle meshes with a varying number of vertices across samples). All datasets are split into 80% training, 10% validation and 10% test datasets.

Experimental Setup. We evaluate our method through comparative analysis with S3M [3] and FlowSSM [16]. Additionally, we benchmark against SmS [7] (with test-time adaptation), as our method builds and improves upon SmS by substituting the ARAP loss with our newly introduced *deformation coherency loss*, while maintaining all other loss functions inherent to the SmS framework. For further implementation details, refer to Table S1 in the supplementary material. It is worth noting that FlowSSM lacks a latent space analogous to those in our method, S3M, and SmS. Therefore, the deformed mesh vertices from FlowSSM are used as correspondences, and we perform PCA for evaluation.

Metrics. The performance of all methods is evaluated using three standardised metrics: generalisation, specificity, and compactness. Generalisation is measured by computing the average point-to-surface distance between test shapes and their reconstructions, indicating the SSM’s ability to represent unseen shapes. Specificity is evaluated by generating 1000 random samples from the SSM and calculating each sample’s minimum Chamfer distance to the training data, reflecting how well the random samples represent the training data. Compactness is determined by the explained variance in PCA, which quantifies the SSM’s ability to reconstruct new instances with fewer parameters.

Results. Our approach outperforms existing methods in terms of performance across all five datasets. As depicted in Fig. 2, our approach achieves low generalisation and specificity errors, indicating a robust ability to accurately represent unseen shapes and generate diverse samples that are representative of the underlying dataset distribution.

In contrast, FlowSSM and SmS produce low specificity and high generalisation errors (lung and liver datasets for FlowSSM; pancreas and spleen datasets for SmS), suggesting a lack of diversity in the generated shapes and an inability to capture the variance present in the data. The qualitative results in Fig. 3 (highlighted with red arrows) further corroborate this observation, where these methods introduce non-existent small structures or result in over-smoothed meshes. Although S3M avoids these specific issues, it displays the poorest compactness among the compared methods and is prone to generating outlier points that do not accurately match the surface of the test mesh, as indicated by the blue arrows in Fig. 3. This causes S3M to produce worse generalisation and specificity than our method.

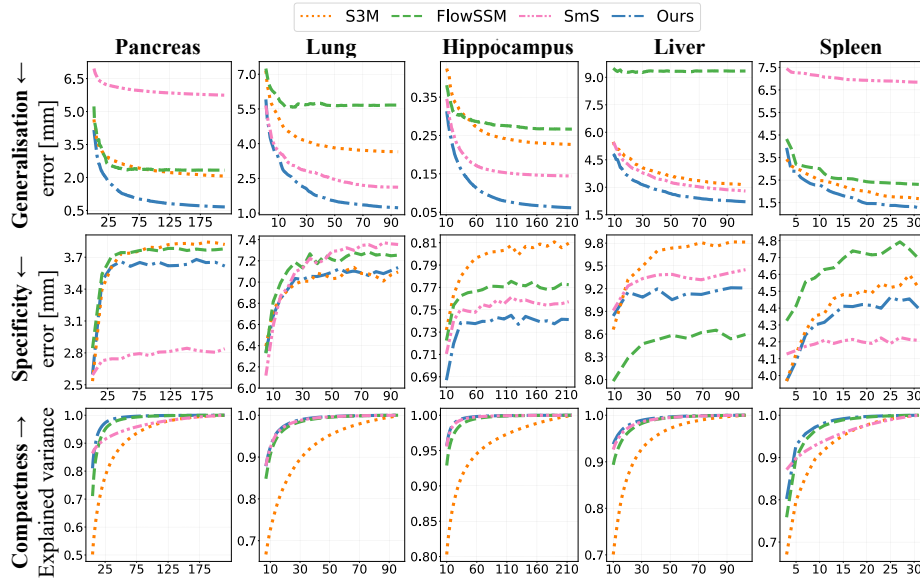


Fig. 2. Quantitative plots of generalisation, specificity and compactness (rows) on five datasets (columns). Horizontal axes in each plot show the number of modes (basis functions in PCA). Our method outperforms previous methods on all datasets by producing low generalisation and specificity errors while exhibiting higher compactness.

Notably, SmS is prone to finding suboptimal solutions (black arrows in Fig. 3), indicating its inability to adequately handle the variability of complex anatomical structures. In contrast, our method successfully addresses these shortcomings, rendering our method the sole technique to accurately capture complex anatomical variations, as shown with green arrows in Fig. 3. This is further corroborated in our ablation study (see Fig. S1 in supplementary material), where all three terms in Eq. (4) are shown to reduce generalisation and specificity errors.

6 Discussion and Limitations

Our novel approach sets the new state of the art in constructing SSM, but it also has some limitations. Firstly, the SSM construction is not learned in an end-to-end fashion. Incorporating end-to-end learning, potentially extending it with a *nonlinear* shape model, could enhance the quality of the resulting SSM. Secondly, our method is template-based. Despite our efforts to mitigate bias by selecting the most representative template shape (defined as the one closest to all shapes on average), our approach may still introduce a degree of bias. Future research could benefit from developing a technique that learns the template directly from the population, thereby generalising the approach and potentially reducing template-related biases.




Fig. 3. Reconstructions of a test mesh for 5 datasets (S3M, FlowSSM, SmS and our method). (Green arrows) Our method outputs high-quality reconstructions compared to other methods. (Black arrows) SmS can produce distorted meshes. (Red arrows) FlowSSM and SmS hallucinate parts non-existent in the test mesh. (Blue arrows) S3M produces a significant amount of outlier points.

7 Conclusion

We have introduced an unsupervised framework for statistical shape model learning that effectively combines correspondence and deformation-based paradigms. Our framework is versatile and accommodates a broad spectrum of anatomical variations in shape and size. By integrating a deformation coherency loss into the state-of-the-art shape matching method, we have successfully enforced smooth and high-quality correspondences necessary for accurate SSM construction. Comparative experiments on five datasets show that our approach surpasses existing methods, striking a good balance between generalisation to new shapes and accurate sample generation. This advancement in SSMs provides a flexible, universal solution for diverse anatomical structures and sets a new standard for SSM construction.

Acknowledgments. Funded by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – 458610525.

Disclosure of Interests. The authors have no competing interests to declare that are relevant to the content of this article.

References

1. Adams, J., Elhabian, S.: From images to probabilistic anatomical shapes: A deep variational bottleneck approach. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*. pp. 474–484. Springer Nature Switzerland, Cham (2022)
2. Adams, J., Khan, N., Morris, A., Elhabian, S.: Spatiotemporal cardiac statistical shape modeling: A data-driven approach. In: *STACOM MICCAI (2022)*
3. Bastian, L., Baumann, A., Hoppe, E., Bürgin, V., Kim, H.Y., Saleh, M., Busam, B., Navab, N.: S3m: scalable statistical shape modeling through unsupervised correspondences. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 459–469. Springer (2023)
4. Bernard, F., Vlassis, N., Gemmar, P., Husch, A., Thunberg, J., Goncalves, J., Hertel, F.: Fast correspondences for statistical shape models of brain structures. In: *Medical Imaging 2016: Image Processing*. vol. 9784, pp. 197–204. SPIE (2016)
5. Bhalodia, R., Elhabian, S.Y., Kavan, L., Whitaker, R.T.: Deepssm: A deep learning framework for statistical shape modeling from raw images. In: Reuter, M., Wachinger, C., Lombaert, H., Paniagua, B., Lüthi, M., Egger, B. (eds.) *Shape in Medical Imaging*. pp. 244–257. Springer International Publishing, Cham (2018)
6. Bruse, J.L., McLeod, K., Biglino, G., Ntsinjana, H.N., Capelli, C., Hsia, T.Y., Serresant, M., Pennec, X., Taylor, A.M., Schievano, S., Taylor, A., Giardini, A., Khambadkone, S., de Leval, M., Hsia, T.Y., Bove, E., Dorfman, A., Baker, G.H., Hlavacek, A., Migliavacca, F., Pennati, G., Dubini, G., Marsden, A., Feinstein, J., Vignon-Clementel, I., Figliola, R., McGregor, J., of Congenital Hearts Alliance (MOCHA) Collaborative Group, f.t.M.: A statistical shape modelling framework to extract 3d shape biomarkers from medical imaging data: assessing arch morphology of repaired coarctation of the aorta. *BMC Medical Imaging* **16**(1), 40 (May 2016)
7. Cao, D., Eisenberger, M., El Amrani, N., Cremers, D., Bernard, F.: Spectral meets spatial: Harmonising 3d shape matching and interpolation. In: *CVPR (2024)*
8. Cates, J., Elhabian, S., Whitaker, R.: Shapeworks: particle-based shape correspondence and visualization software. In: *Statistical Shape and Deformation Analysis*, pp. 257–298. Elsevier (2017)
9. Cury, C., Glaunès, J.A., Toro, R., Chupin, M., Schumann, G., Frouin, V., Poline, J.B., Colliot, O., t.I.C.: Statistical shape analysis of large datasets based on diffeomorphic iterative centroids. *Frontiers in Neuroscience* **12** (2018)
10. Durrleman, S., Prastawa, M., Charon, N., Korenberg, J.R., Joshi, S., Gerig, G., Trounev, A.: Morphometry of anatomical shape complexes with dense deformations and sparse parameters. *NeuroImage* **101**, 35–49 (2014)
11. Heimann, T., Wolf, I., Williams, T., Meinzer, H.P.: 3d active shape models using gradient descent optimization of description length. In: *Biennial International Conference on Information Processing in Medical Imaging*. pp. 566–577. Springer (2005)

12. Heitz, G., Rohlfing, T., Jr., C.R.M.: Statistical shape model generation using non-rigid deformation of a template mesh. In: Fitzpatrick, J.M., Reinhardt, J.M. (eds.) *Medical Imaging 2005: Image Processing*. vol. 5747, pp. 1411 – 1421. International Society for Optics and Photonics, SPIE (2005)
13. Iyer, K., Elhabian, S.Y.: Mesh2ssm: From surface meshes to statistical shape models of anatomy. In: Greenspan, H., Madabhushi, A., Mousavi, P., Salcudean, S., Duncan, J., Syeda-Mahmood, T., Taylor, R. (eds.) *Medical Image Computing and Computer Assisted Intervention – MICCAI 2023*. pp. 615–625 (2023)
14. Jiang, C., Huang, J., Tagliasacchi, A., Guibas, L.: Shapeflow: Learnable deformations among 3d shapes. In: *Advances in Neural Information Processing Systems* (2020)
15. Klatzow, J., Dalmasso, G., Martínez-Abadías, N., Sharpe, J., Uhlmann, V.: μ match: 3d shape correspondence for biological image data. *Frontiers in Computer Science* **4** (2022)
16. Lüdke, D., Amiranashvili, T., Ambellan, F., Ezhov, I., Menze, B.H., Zachow, S.: Landmark-free statistical shape modeling via neural flow deformations. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022*. pp. 453–463. Springer Nature Switzerland (2022)
17. Lüthi, M., Gerig, T., Jud, C., Vetter, T.: Gaussian process morphable models. *IEEE Transactions on Pattern Analysis and Machine Intelligence* **40**(8), 1860–1873 (2018)
18. Magnet, R., Ren, J., Sorkine-Hornung, O., Ovsjanikov, M.: Smooth non-rigid shape matching via effective dirichlet energy optimization. In: *3DV* (2022)
19. Miller, M.I., Trounev, A., Younes, L.: Geodesic shooting for computational anatomy. *Journal of Mathematical Imaging and Vision* **24**(2), 209–228 (Mar 2006)
20. Ovsjanikov, M., Ben-Chen, M., Solomon, J., Butscher, A., Guibas, L.: Functional maps: a flexible representation of maps between shapes. *ACM Transactions on Graphics (ToG)* **31**(4), 1–11 (2012)
21. Paulsen, R., Larsen, R., Nielsen, C., Laugesen, S., Ersbøll, B.: Building and testing a statistical shape model of the human ear canal. In: *MICCAI* (2002)
22. Peiffer, M., Burssens, A., De Mits, S., Heintz, T., Van Waeyenberge, M., Buedts, K., Victor, J., Audenaert, E.: Statistical shape model-based tibiofibular assessment of syndesmotic ankle lesions using weight-bearing ct. *Journal of Orthopaedic Research* **40**(12), 2873–2884 (2022)
23. Roufousse, J.M., Sharma, A., Ovsjanikov, M.: Unsupervised deep learning for structured shape matching. In: *ICCV* (2019)
24. Schaufelberger, M., Kühle, R., Wachter, A., Weichel, F., Hagen, N., Ringwald, F., Eisenmann, U., Hoffmann, J., Engel, M., Freudlsperger, C., Nahm, W.: A radiation-free classification pipeline for craniosynostosis using statistical shape modeling. *Diagnostics* **12**(7) (2022)
25. Setio, A.A.A., Traverso, A., De Bel, T., Berens, M.S., Van Den Bogaard, C., Cerello, P., Chen, H., Dou, Q., Fantacci, M.E., Geurts, B., et al.: Validation, comparison, and combination of algorithms for automatic detection of pulmonary nodules in computed tomography images: the luna16 challenge. *Medical image analysis* **42**, 1–13 (2017)
26. Simpson, A.L., Antonelli, M., Bakas, S., Bilello, M., Farahani, K., van Ginneken, B., Kopp-Schneider, A., Landman, B.A., Litjens, G., Menze, B., Ronneberger, O., Summers, R.M., Bilic, P., Christ, P.F., Do, R.K.G., Gollub, M., Golia-Pernicka, J., Heckers, S.H., Jarnagin, W.R., McHugo, M.K., Napel, S., Vorontsov, E., Maier-Hein, L., Cardoso, M.J.: A large annotated medical image dataset for the development and evaluation of segmentation algorithms (2019)

27. Sorkine, O., Alexa, M.: As-rigid-as-possible surface modeling. In: Symposium on Geometry processing. vol. 4, pp. 109–116. Citeseer (2007)
28. Zhang, L., Tanno, R., Xu, M.C., Jin, C., Jacob, J., Ciccarelli, O., Barkhof, F., Alexander, D.C.: Disentangling human error from the ground truth in segmentation of medical images. In: Proceedings of the 34th International Conference on Neural Information Processing Systems. NIPS’20, Curran Associates Inc., Red Hook, NY, USA (2020)

1 Supplementary Material

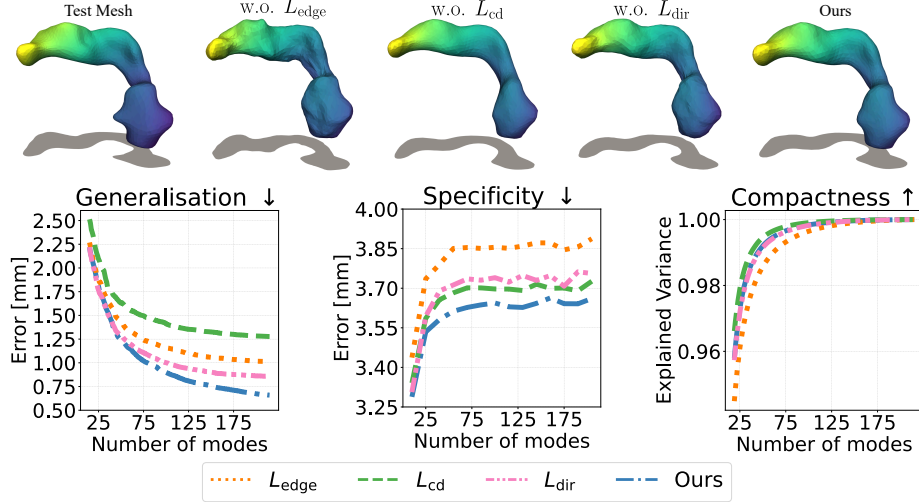


Fig. S1. Ablation study on the pancreas dataset. (Top) Reconstruction of a test mesh using our method trained without individuals terms of our deformation coherency loss introduced Eq. (4) (i.e. L_{edge} , L_{cd} , L_{dir}). The reconstruction using our full method (all three losses) achieves the best qualitative results. (Bottom) Generalisation, specificity of compactness of our method trained with and without L_{edge} , L_{cd} , L_{dir} .

Table S1. Implementation details for our method. The last four rows (denoted by *) are details from the SmS framework. See [7] for more details.

Implementation details	Value	Loss	Weight
Epochs	15	L_{cd}	10^{+4}
Time steps K	6	L_{dir}	10^{+2}
GPU	NVIDIA A100	L_{edge}	10^{+5}
Optimizer for DiffusionNet*	Adam with lr 10^{-3}	L_{struct} *	1
Optimizer for ResnetECPos*	Adam with lr 10^{-3}	L_{couple} *	1
Number of eigenfunctions*	40	L_{symm} *	1
τ for PMap computation*	0.07	L_{align} *	10

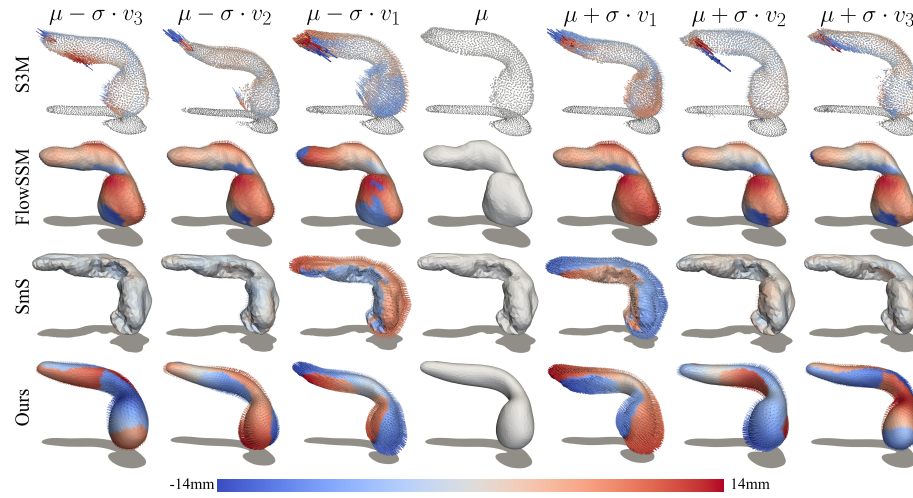


Fig. S2. Top three PCA modes of variations identified by S3M [3], FlowSSM [16], SmS [7] and our proposed method. Our method can operate on the head, neck, and body of the pancreas using just the first three modes of variation. The colour map and arrows show the signed distance and direction from the mean shape.




Fig. S3. Random samples created using our method. The shape diversity of the samples showcases our method’s ability to capture the biological variability of the training data (low specificity errors).

Table S2. Additional information on the datasets [25, 26] used in our paper.

Organ	Scans	Curated from
Pancreas	273 segmented CT	patients undergoing resection of pancreatic masses
Spleen	41 segmented CT	patients with metastatic liver disease
Liver	131 segmented CT	patients receiving chemotherapy for liver metastases
Hippocampus	260 segmented MRI	a combination of healthy adults and individuals with non-affective psychiatric disorders
Lung	subset of 126 segmented CT	patients with pulmonary nodules