



HHS Public Access

Author manuscript

Med Image Comput Assist Interv. Author manuscript; available in PMC 2024 October 01.

Published in final edited form as:

Med Image Comput Assist Interv. 2023 October ; 14220: 615–625.

doi:10.1007/978-3-031-43907-0_59.

Mesh2SSM: From Surface Meshes to Statistical Shape Models of Anatomy

Krithika Iyer^{1,2}, Shireen Elhabian^{1,2}

¹Scientific Computing and Imaging Institute, University of Utah, SLC, UT, US

²Kahlert School of Computing, University of Utah, Salt Lake City, UT, USA

Abstract

Statistical shape modeling is the computational process of discovering significant shape parameters from segmented anatomies captured by medical images (such as MRI and CT scans), which can fully describe subject-specific anatomy in the context of a population. The presence of substantial non-linear variability in human anatomy often makes the traditional shape modeling process challenging. Deep learning techniques can learn complex non-linear representations of shapes and generate statistical shape models that are more faithful to the underlying population-level variability. However, existing deep learning models still have limitations and require established/optimized shape models for training. We propose Mesh2SSM, a new approach that leverages unsupervised, permutation-invariant representation learning to estimate how to deform a template point cloud to subject-specific meshes, forming a correspondence-based shape model. Mesh2SSM can also learn a population-specific template, reducing any bias due to template selection. The proposed method operates directly on meshes and is computationally efficient, making it an attractive alternative to traditional and deep learning-based SSM approaches.

Keywords

Statistical Shape Modeling; Representation Learning; Point Distribution Models

1 Introduction

Statistical shape modeling (SSM) is a powerful tool in medical image analysis and computational anatomy to quantify and study the variability of anatomical structures within populations. SSM has shown great promise in medical research, particularly in diagnosis [12,23], pathology detection [19,25], and treatment planning [27]. SSM has enabled researchers to better understand the underlying biological processes, leading to the development of more accurate and personalized diagnostic and treatment plans [17,3,14,9].

Over the years, several SSM approaches have been developed that implicitly represent the shapes (deformation fields [8], level set methods [22]) or explicitly represent them as a ordered set of landmarks or *correspondence points* (aka point distribution models, PDMs).

Here, we focus on the automated construction of PDMs because, compared to deformation fields, point correspondences are easier to interpret by clinicians, are computationally efficient for large datasets, and less sensitive to noise and outliers than deformation fields [5].

SSM performance depends on the underlying process used to generate shape correspondences and the quality of the input data. Various correspondence generation methods exist, including non-optimized landmark estimation and parametric and non-parametric correspondence optimization. Non-optimized methods manually label a reference shape and warp the annotated landmarks using registration techniques [18,10,16]. Parametric methods use fixed geometrical bases to establish correspondences [26], while group-wise non-parametric approaches find correspondences by considering the variability of the entire cohort during the optimization process. Examples of non-parametric methods include particle-based optimization [4] and Minimum Description Length (MDL) [7].

Traditional SSM methods assume that population variability follows a Gaussian distribution, which implies that a linear combination of training shapes can express unseen shapes. However, anatomical variability can be far more complex than this linear approximation, in which case nonlinear variations normally exist (e.g., bending fingers, soft tissue deformations, and vertebrae with different types). Furthermore, conventional SSM pipelines are computationally intensive, where inferring PDMs on new samples entail an optimization process. Deep learning-based approaches for SSM have emerged as a promising avenue to overcoming these limitations. Deep learning models can learn complex non-linear representations of the shapes, which can be used to generate shape models. Moreover, they can efficiently perform inference on new samples without computation overhead or re-optimization. Recent works such as FlowSSM [15], ShapeFlow [11], DeepSSM [2], and VIB-DeepSSM [1] have incorporated deep learning to generate shape models. FlowSSM [15] and ShapeFlow [11] operate on surface meshes and use neural networks to parameterize the deformations field between two shapes in a low dimensional latent space and rely on an encoder-free setup. Encoder-free methods randomly initialize the latent representations for each sample that are then optimized to produce the optimal deformations. One major caveat of an encoder-free setup is that inference on new meshes is no longer straightforward; the latent representation has to be re-optimized for every new sample. On the other hand, DeepSSM [2], TL-DeepSSM [2], and VIB-DeepSSM [1] learn the PDM directly from unsegmented CT/MRI images, and hence alleviate the need for PDM optimization given new samples and can bypass anatomy segmentation by operating directly on unsegmented images. However, these methods rely on supervised losses and require volumetric images, segmented images, and established/optimized PDMs for training. This reliance on supervised losses introduces linearity assumptions in generating ground truth PDMs. TL-DeepSSM [2], a variant of DeepSSM [2], differs from the others by not utilizing PCA scores as shape descriptors. Instead, it adopts an established correspondence model hence, similar to the vanilla DeepSSM [2] learns a linear model.

In this paper, we introduce Mesh2SSM³, a deep learning method that addresses the limitations of traditional and deep learning-based SSM approaches. Mesh2SSM leverages *unsupervised, permutation-invariant representation learning* to learn the low dimensional

nonlinear shape descriptor directly from mesh data and uses the learned features to generate a correspondence model of the population. Mesh2SSM also includes an analysis network that operates on the learned correspondences to obtain a data-driven template point cloud (i.e., template point cloud), which can replace the initial template, and hence reducing the bias that could arise from template selection. Furthermore, the learned representation of meshes can be used for predicting related quantities that rely on shape. Our main contributions are:

1. We introduce Mesh2SSM, a fully unsupervised correspondence generation deep learning framework that operates directly on meshes. Mesh2SSM uses an autoencoder to extract the shape descriptor of the mesh and uses this descriptor to transform a template point cloud using IM-Net [6].
2. The proposed method uses an autoencoder that combines geodesic distance features and EdgeConv [28] (dynamic graph convolution neural network) to extract meaningful feature representation of each mesh that is permutation-invariant.
3. Mesh2SSM also includes a variational autoencoder (VAE) [13,21] operating on the learned correspondence points and trained end-to-end with correspondence generation network. This VAE branch serves two purposes: (a) serves as a shape analysis module for the non-linear shape variations and (b) learns a data-specific template from the latent space of the correspondences that is fed back to the correspondence generation network.

To motivate the need for the mesh feature encoder and study the effect of the template selection, we considered the box-bump dataset, a synthetic dataset of 3D shapes of boxes with a moving bump. In Figure 1, we compare Mesh2SSM (sans the VAE analysis branch) with FlowSSM [15] since this approach is the closest to Mesh2SSM. We performed experiments with three templates: medoid, sphere, and box without the bump. Although both methods show some sensitivity to the choice of template, FlowSSM is more sensitive toward the choice of the template than Mesh2SSM. Moreover, FlowSSM fails to identify the correct mode of variation, the horizontal movement of the bump as the primary variation, which can also be inferred by comparing the compactness curves in Figure 1.c. Mesh2SSM performs best when the template is a medoid shape, which makes the case for learning a data-specific template. Since Mesh2SSM model uses an autoencoder, inference on unseen meshes only requires a single forward pass (1 second per sample); FlowSSM requires re-optimization, increasing the inference time drastically and require a convergence criteria to determine the best number of iterations per sample (0.15 seconds for one iterations per sample).

2 Method

The overview of the proposed pipeline is provided in Figure 2. This section provides a brief description of each module.

³Source code: https://github.com/iyerkrithika21/mesh2SSM_2023

2.1 Correspondence Generation

Given a set of N aligned surface meshes $\mathcal{X} = \{X_1, X_2, \dots, X_N\}$, each mesh $X_i = (V_i, E_i)$, where V_i and E_i represent the vertices and edge connectivity, respectively. The goal of the model is to predict a set C_i of M 3D correspondence points that fully describe each surface X_i and are anatomically consistent across all meshes. This goal is achieved by learning a low dimensional representation of the surface mesh $\mathbf{z}^m \in \mathbb{R}^L$ using the mesh autoencoder and then \mathbf{z}^m is used to transform the template point cloud via the implicit field decoder (IMNet) [6]. The network optimization is driven primarily by point-set to point-set two-way Chamfer distance between the learned correspondence point sets C_i and the vertex locations V_i of the original meshes. To ensure that the encoder learns useful features for the task, we regularize the optimization using the vertex reconstruction loss of the autoencoder between the input V_i and the predicted \hat{V}_i . The correspondence loss function is given by:

$$\mathcal{L}_C = \sum_{i=1}^N \left[\mathcal{L}_{L_2\text{Chamfer}}(V_i, C_i) + \alpha \mathcal{L}_{L_1\text{Chamfer}}(V_i, C_i) + \gamma \mathcal{L}_{MSE}(V_i, \hat{V}_i) \right] \quad (1)$$

where α, γ , are the hyperparameters. We consider a combination of L_1 and L_2 two-way Chamfer distance for numerical stability as the magnitude of L_2 loss can be low over epochs and L_1 can compensate for it. The correspondence generation uses two networks:

Mesh Autoencoder (M-AE): We use EdgeConv [28] blocks, which are dynamic graph convolution neural network (DGCNN) blocks in the encoder and decoder to capture local geometric features of the mesh. The model takes vertices as input, computes an edge feature set of size k (using nearest neighbors) for each vertex at an EdgeConv layer, and aggregates features within each set to compute EdgeConv responses. The output features of the last EdgeConv layer are then globally aggregated to form a 1D global descriptor \mathbf{z}^m of the mesh. The first EdgeConv block uses geodesic distance on the surface of the mesh to calculate the k features. The dynamic feature creation property of EdgeConv and the global pooling make this autoencoder permutation invariant.

Implicit field decoder (IM-NET): The IM-NET [6] architecture consists of fully connected layers with non-linearity and skip-layer connections. This network enforces the notion of correspondence across the samples. The network takes in two inputs, the latent representation of the mesh \mathbf{z}^m and a template point cloud (a set of unordered points). IM-NET estimates the deformation of each point in the template required to deform the template to each sample, conditioned on \mathbf{z}^m . Based on the learned deformation, IM-NET directly produces the resultant displaced template point without the computational complexity of the deformation fields. Correspondence is established since the same template is deformed to all the samples.

2.2 Analysis

The Mesh2SSM model also consists of an analysis branch that acts as a shape analysis module to capture non-linear shape variations identified by the learned correspondences $\{C_i\}_{i=1}^N$ and also learns a data-informed template from the latent space of correspondences to be fed back into the correspondence generation network during training. This branch uses one network module:

Shape Variation Autoencoder (SP-VAE): The VAE [13,21] is a latent variable model parameterized by an encoder ϕ , decoder θ , and the prior $p(\mathbf{z}^p) \sim \mathcal{N}(0, \mathbf{I})$. The encoder maps the shape represented by the learned correspondence points C to the latent space and the decoder reconstructs the correspondences from the latent representation \mathbf{z}^p . By capturing the underlying structure of the PDM through a low-dimensional representation, SP-VAE allows for the estimation of the mean shape of the learned correspondences. The SP-VAE is trained using the loss function given by:

$$\mathcal{L}(\theta, \phi) = -\mathbb{E}_{q_\phi(\mathbf{z}_i^p | C_i)}[\log p_\theta(C_i | \mathbf{z}_i^p)] + KL(q_\phi(\mathbf{z}_i^p | C_i) \parallel p(\mathbf{z}_i^p)) \quad (2)$$

The main difference between M-AE and a SP-VAE lies in the input and output representations they handle. SP-VAE operates directly on sets of landmarks or correspondences, aiding in the analysis of shape models. It takes a set of correspondences describing a shape as input and aims to learn a compressed latent representation of the shape. Importantly, the SP-VAE maintains the same ordering of correspondences at the input and output, so it does not use permutation-invariant layers or operations like pooling. .

2.3 Training

We begin with a burn-in stage, where only the correspondence generation module is trained while the analysis module is frozen. After the burn-in stage, alternate optimization of the correspondence and analysis module begins. During the alternate optimization phase, we generate the data-informed template from the latent space of SP-VAE at regular intervals. The learned data-informed template is used in the correspondence generation module in the subsequent epochs. For the learned template, we sample 500 samples from the prior $p(\mathbf{z}^p) \sim \mathcal{N}(0, \mathbf{I})$ and pass it through the decoder of SP-VAE to get the reconstructed correspondence point set. The mean template is defined by taking the average of these generated samples. Inference with unseen meshes is straight forward; the meshes are passed through the mesh encoder and IM-NET of the correspondence generation module to get the predicted correspondences. All hyperparameters and network architecture details are mentioned in the supplementary material.

3 Experiments and Discussion

Dataset:

We use the publicly available Decath-Pancreas dataset of 273 segmentations from patients who underwent pancreatic mass resection [24]. The shapes of the pancreas are highly

variable and have thin structures, making it a good candidate for non-linear SSM analysis. The segmentations were isotropically resampled, smoothed, centered, and converted to meshes with roughly 2000 vertices. Although the DGCNN mesh autoencoder used in Mesh2SSM does not require the same number of vertices, uniformity across the dataset makes it computationally efficient; hence, we pad the smallest mesh by randomly repeating the vertices (akin to padding image for convolutions). The samples were randomly divided, with 218 used for training, 26 for validation, and 27 for testing.

3.1 Results

We perform experiments with two templates: sphere and medoid. We compare the performance of FlowSSM [15] with Mesh2SSM with the template feedback loop. For Mesh2SSM template, we use 256 points uniformly spread across the surface of the sample. Mesh2SSM and FlowSSM do not have a equivalent latent space for comparison of the shape models, hence, we consider the deformed mesh vertices of FlowSSM as correspondences and perform PCA analysis. Figure 3 shows the top three PCA modes of variations identified by Mesh2SSM and FlowSSM. Similar to the observations made box-bump dataset, FlowSSM is affected by the choice of the template, and the modes of variation differ as the template changes. On the other hand, PDM predicted by Mesh2SSM identifies the same primary modes consistently. Pancreatic cancer mainly presents itself on the head of the structure [20] and for the Decath dataset, we can see the first mode identifies the change in the shape of the head. We evaluate the models based on compactness, generalization, and specificity. Compactness measures the ability of the model to reconstruct new shape instances with fewer parameters using PCA explained variance. Generalization measures the average surface distance between all test shapes and their reconstructions, and specificity measures the distance between randomly generated PCA samples. Figure 4.a shows the metrics for the pancreas dataset. Mesh2SSM outperforms FlowSSM in all three metrics, despite using only 256 correspondence points compared to FlowSSM's ~2000 vertices. Mesh2SSM correspondence generation module efficiently parameterizes the surface of the pancreas with a minimum number of parameters. Mesh2SSM template, shown in Figure 4.b, becomes more detailed as optimization continues, regardless of the starting template. The model can learn correct deformations in the correspondence generation module and identify the correct mean shape in the latent space of SP-VAE in the analysis module. Using the analysis module of Mesh2SSM, we visualized the top three modes of variation identified by sorting the latent dimensions of SP-VAE based on the standard deviations of the latent embeddings of the training dataset. Variations are generated by perturbing the latent representation of a sample in three directions, resulting in non-linear modes such as changes in the size and shape of the pancreas head and narrowing of the neck and body. This is shown in Figure 4.c for MeshSSM model with medoid starting template. The distance metrics for the reconstructions of the testing samples were also computed. The results of the metrics are summarized in Table 1. The calculation involved the L_1 Chamfer loss between the predicted points (correspondences in the case of Mesh2SSM and the deformed mesh vertices in the case of FlowSSM) and the original mesh vertices. Additionally, the surface to surface distance of the mesh reconstructions (using the correspondences in Mesh2SSM and deformed meshes in FlowSSM) was included. For the pancreas dataset with the medoid as the initial template, Mesh2SSM with the template feedback produced more precise models.

3.2 Limitations and Future Scope

As SSM is included a part of diagnostic clinical support systems, it is crucial to address the drawbacks of the models. Like most deep learning models, performance of Mesh2SSM could be affected by small dataset size, and it can produce overconfident estimates. An augmentation scheme and a layer uncertainty calibration are could improve its usability in medical scenarios. Additionally, enforcing disentanglement in the latent space of SP-VAE can make the analysis module interpretable and allow for effective non-linear shape analysis by clinicians.

4 Conclusion

The paper presents a new systematic approach of generating non-linear statistical shape models using deep learning directly from meshes, which overcomes the limitations of traditional SSM and current deep learning approaches. The use of an autoencoder for meaningful feature extraction of meshes to learn the PDM provides a versatile and scalable framework for SSM. Incorporating template feedback loop via VAE [13,21] analysis module helps in mitigating bias and capturing non-linear characteristics of the data. The method is demonstrated to have superior performance in identifying shape variations using fewer parameters on synthetic and clinical datasets. To conclude, our method of generating highly accurate and detailed models of complex anatomical structures with reduced computational complexity has the potential to establish statistical shape modeling from non-invasive imaging as a powerful diagnostic tool.

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgements

This work was supported by the National Institutes of Health under grant numbers NIBIB-U24EB029011, NIAMS-R01AR076120, and NHLBI-R01HL135568. We thank the University of Utah Division of Cardiovascular Medicine for providing left atrium MRI scans and segmentations from the Atrial Fibrillation projects and the ShapeWorks team.

References

1. Adams J, Elhabian S: From images to probabilistic anatomical shapes: A deep variational bottleneck approach. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part II. pp. 474–484. Springer (2022)
2. Bhalodia R, Elhabian SY, Kavan L, Whitaker RT: Deepssm: a deep learning framework for statistical shape modeling from raw images. In: Shape in Medical Imaging: International Workshop, ShapeMI 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 20, 2018, Proceedings. pp. 244–257. Springer (2018)
3. Bruse JL, McLeod K, Biglino G, Ntsinjana HN, Capelli C, Hsia TY, Sermesant M, Pennec X, Taylor AM, Schievano S, et al. : 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, 1–19 (2016) [PubMed: 26746426]

4. 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)
5. Cerrolaza JJ, Picazo ML, Humbert L, Sato Y, Rueckert D, Ballester MÁG, Linguraru MG: Computational anatomy for multi-organ analysis in medical imaging: A review. *Medical Image Analysis* 56, 44–67 (2019) [PubMed: 31181343]
6. Chen Z: Im-net: Learning implicit fields for generative shape modeling (2019)
7. Davies RH: Learning shape: optimal models for analysing natural variability. The University of Manchester (United Kingdom) (2002)
8. Durrleman S, Prastawa M, Charon N, Korenberg JR, Joshi S, Gerig G, Trouvé A: Morphometry of anatomical shape complexes with dense deformations and sparse parameters. *NeuroImage* 101, 35–49 (2014) [PubMed: 24973601]
9. Faghieh Roohi S, Aghaeizadeh Zoroofi R: 4d statistical shape modeling of the left ventricle in cardiac mr images. *International journal of computer assisted radiology and surgery* 8, 335–351 (2013) [PubMed: 22893114]
10. Heitz G, Rohlfing T, Maurer CR Jr: Statistical shape model generation using nonrigid deformation of a template mesh. In: *Medical Imaging 2005: Image Processing*, vol. 5747, pp. 1411–1421. SPIE (2005)
11. Jiang C, Huang J, Tagliasacchi A, Guibas LJ: Shapeflow: Learnable deformation flows among 3d shapes. *Advances in Neural Information Processing Systems* 33, 9745–9757 (2020)
12. Khan RA, Luo Y, Wu FX: Machine learning based liver disease diagnosis: A systematic review. *Neurocomputing* 468, 492–509 (2022)
13. Kingma DP, Welling M: Auto-encoding variational bayes. arXiv preprint arXiv:1312.6114 (2013)
14. Lindberg O, Walterfang M, Looi JC, Malykhin N, Östberg P, Zandbelt B, Styner M, Paniagua B, Velakoulis D, Örndahl E, et al. : Hippocampal shape analysis in alzheimer’s disease and frontotemporal lobar degeneration subtypes. *Journal of Alzheimer’s disease* 30(2), 355–365 (2012)
15. Lüdke D, Amiranashvili T, Ambellan F, Ezhov I, Menze BH, Zachow S: Landmark-free statistical shape modeling via neural flow deformations. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part II*, pp. 453–463. Springer (2022)
16. McInerney T, Terzopoulos D: Deformable models in medical image analysis. In: *Proceedings of the workshop on mathematical methods in biomedical image analysis*, pp. 171–180. IEEE (1996)
17. Merle C, Innmann MM, Waldstein W, Pegg EC, Aldinger PR, Gill HS, Murray DW, Grammatopoulos G: High variability of acetabular offset in primary hip osteoarthritis influences acetabular reaming—a computed tomography–based anatomic study. *The Journal of Arthroplasty* 34(8), 1808–1814 (2019) [PubMed: 31122846]
18. Paulsen R, Larsen R, Nielsen C, Laugesen S, Ersbøll B: Building and testing a statistical shape model of the human ear canal. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 373–380. Springer (2002)
19. 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) [PubMed: 35249244]
20. Ralston SH, Penman ID, Strachan MW, Hobson R: *Davidson’s Principles and Practice of Medicine E-Book*. Elsevier Health Sciences (2018)
21. Rezende DJ, Mohamed S, Wierstra D: Stochastic backpropagation and approximate inference in deep generative models. In: *International conference on machine learning*, pp. 1278–1286. PMLR (2014)
22. Samson C, Blanc-Féraud L, Aubert G, Zerubia J: A level set model for image classification. *International journal of computer vision* 40(3), 187–197 (2000)
23. Schaufelberger M, Köhle R, Wachter A, Weichel F, Hagen N, Ringwald F, Eisenmann U, Hoffmann J, Engel M, Freudlsperger C, et al. : A radiation-free classification pipeline for craniosynostosis using statistical shape modeling. *Diagnostics* 12(7), 1516 (2022) [PubMed: 35885422]

24. Simpson AL, Antonelli M, Bakas S, Bilello M, Farahani K, Van Ginneken B, Kopp-Schneider A, Landman BA, Litjens G, Menze B, et al. : A large annotated medical image dataset for the development and evaluation of segmentation algorithms. arXiv preprint arXiv:1902.09063 (2019)
25. Sophocleous F, Bône A, Shearn AI, Forte MNV, Bruse JL, Caputo M, Biglino G: Feasibility of a longitudinal statistical atlas model to study aortic growth in congenital heart disease. *Computers in Biology and Medicine* 144, 105326 (2022) [PubMed: 35245697]
26. Styner M, Oguz I, Xu S, Brechböhler C, Pantazis D, Levitt JJ, Shenton ME, Gerig G: Framework for the statistical shape analysis of brain structures using spharm-pdm. *The insight journal* (1071), 242 (2006) [PubMed: 21941375]
27. Vicory J, Herz C, Allemang D, Nam HH, Cianciulli A, Vigil C, Han Y, Lasso A, Jolley MA, Paniagua B: Statistical shape analysis of the tricuspid valve in hypoplastic left heart syndrome. In: *Statistical Atlases and Computational Models of the Heart. Multi-Disease, Multi-View, and Multi-Center Right Ventricular Segmentation in Cardiac MRI Challenge: 12th International Workshop, STACOM 2021, Held in Conjunction with MICCAI 2021, Strasbourg, France, September 27, 2021, Revised Selected Papers*. pp. 132–140. Springer (2022)
28. Wang Y, Sun Y, Liu Z, Sarma SE, Bronstein MM, Solomon JM: Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)* 38(5), 1–12 (2019)

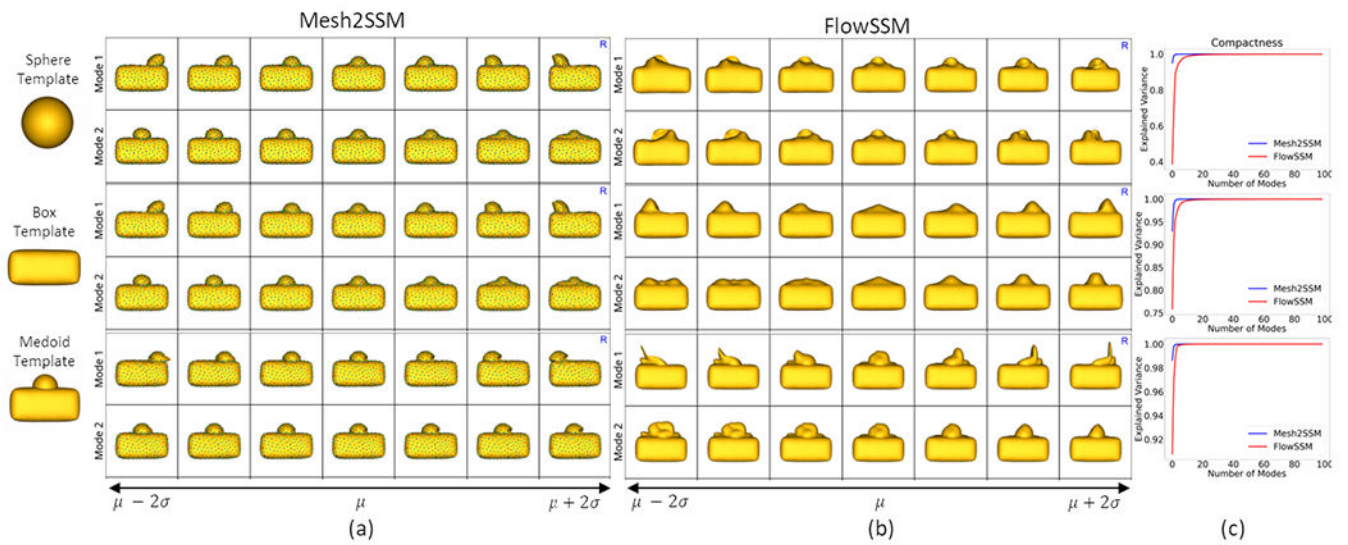


Fig. 1. Top two PCA modes of variations identified by (a) Mesh2SSM and (b) FlowSSM [15] with three templates: sphere, box without a bump, and medoid shape. FlowSSM fails to capture the horizontal movement as the primary mode of variation. (c) The compactness curves for both models with different templates.

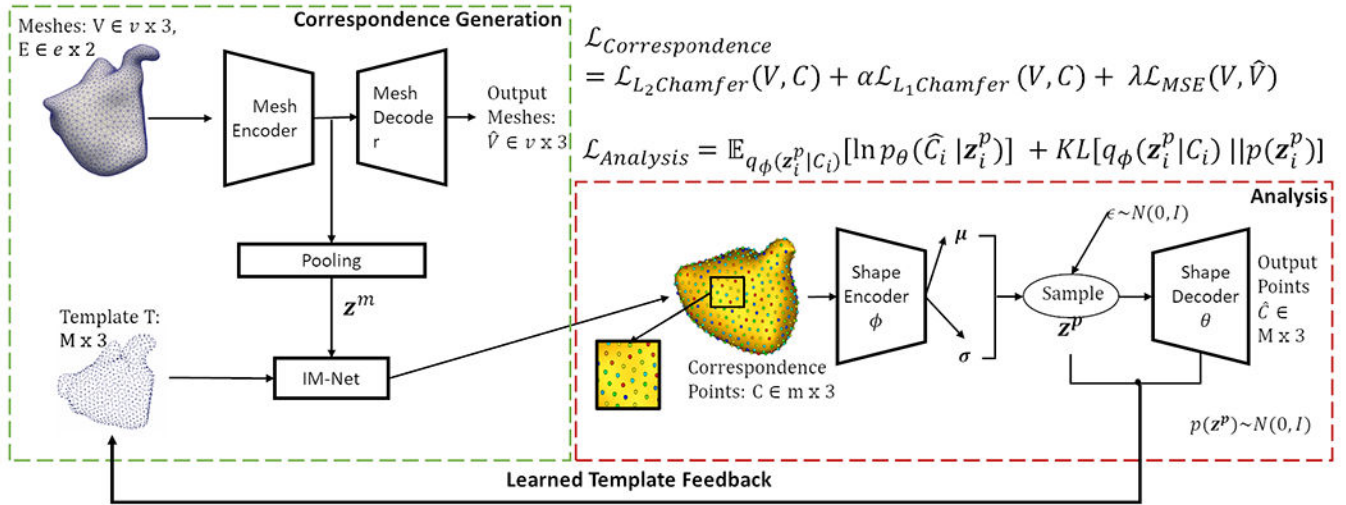


Fig. 2. Mesh2SSM:
Architecture and loss of the proposed method.

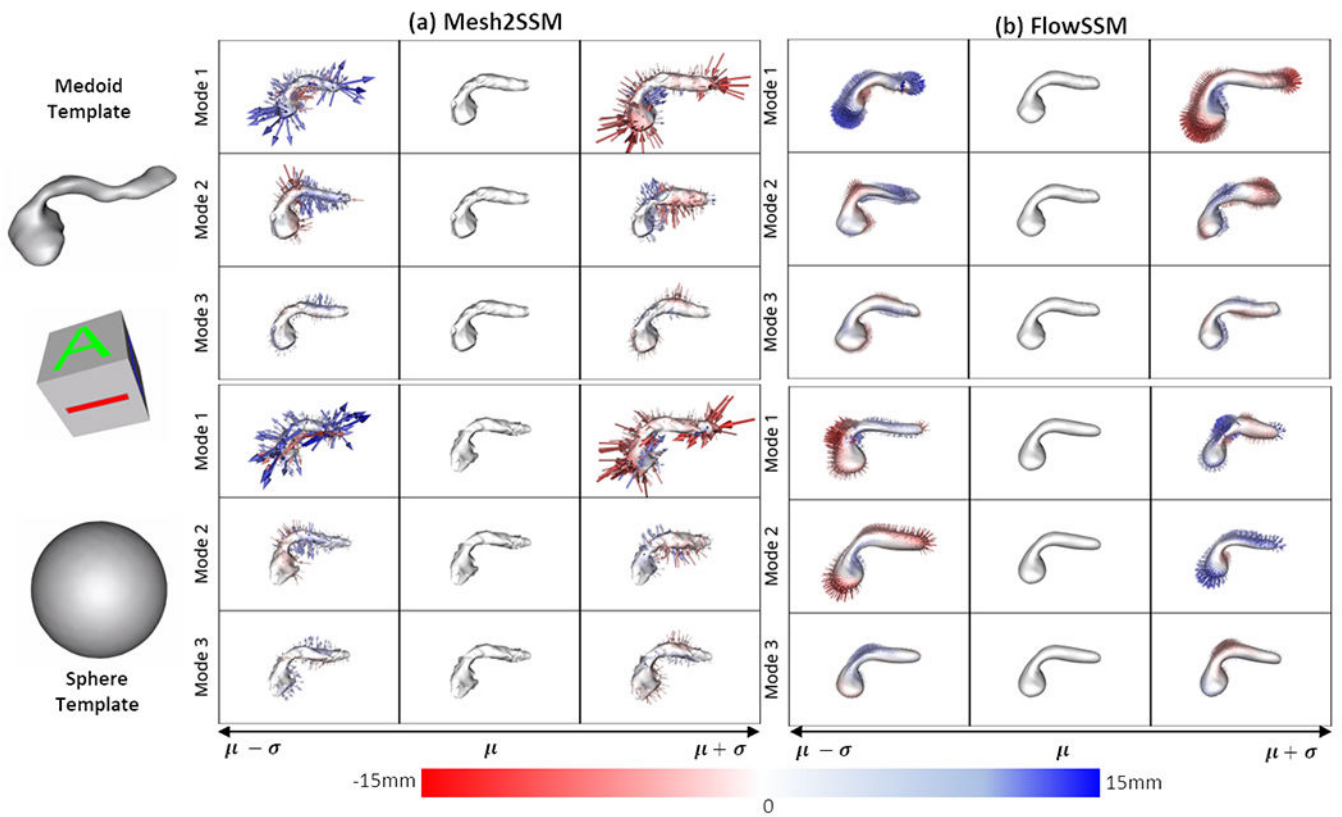


Fig. 3. Top three PCA modes of variations identified by (a) Mesh2SSM and (b) FlowSSM [15] with two templates: sphere, medoid. The color map and arrows show the signed distance and direction from the mean shape.

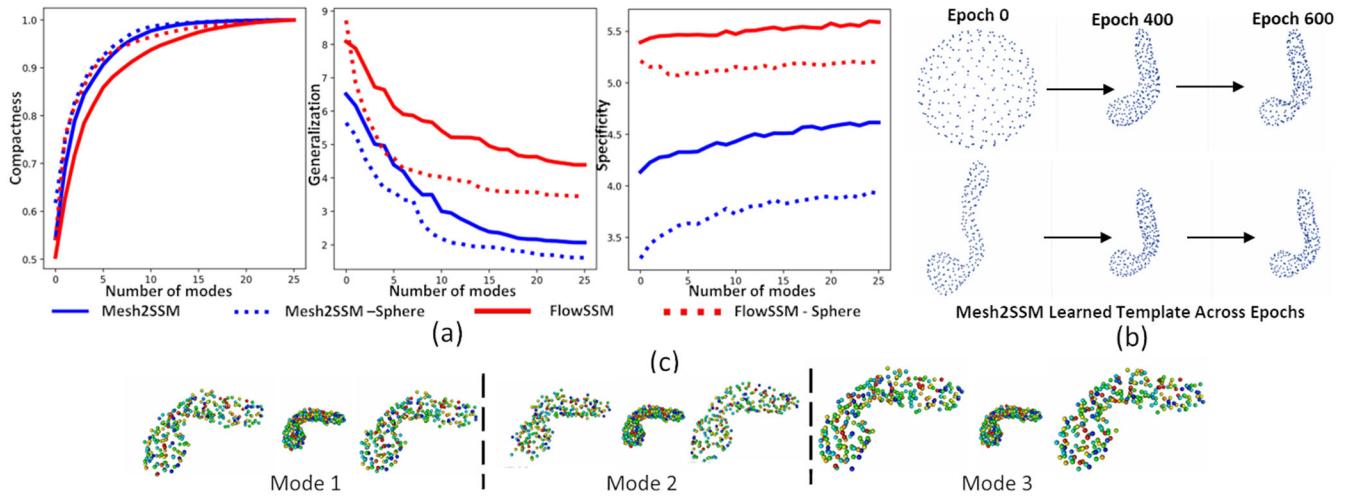


Fig. 4. (a) Shape statistics of pancreas dataset: compactness (higher is better), generalization (lower is better), and specificity (lower is better). (b) Mesh2SSM Learned template across epochs for pancreas dataset. (c) Non-linear modes of variations identified by Mesh2SSM.

Table 1.

Distance metrics (measured in mm) of the testing samples and their reconstructions for the pancreas dataset

| Metrics | Mesh2SSM | | FlowSSM [15] | |
|--------------------|--------------------------------------|-------------------|-------------------|-------------------|
| | Template | | | |
| | Medoid | Sphere | Medoid | Sphere |
| L_1 Chamfer | 0.033 ± 0.002 | 0.035 ± 0.002 | 0.391 ± 0.162 | 1.91 ± 0.687 |
| Surface-to-Surface | 2.378 ± 0.7325 | 5.436 ± 2.232 | 5.918 ± 2.026 | 4.918 ± 1.925 |

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript