

Can point cloud networks learn statistical shape models of anatomies?

Jadie Adams^{1,2} and Shireen Elhabian^{1,2}

¹ Scientific Computing and Imaging Institute, University of Utah, UT, USA

² School of Computing, University of Utah, UT, USA

jadie.adams@utah.edu, shireen@sci.utah.edu

Abstract. Statistical Shape Modeling (SSM) is a valuable tool for investigating and quantifying anatomical variations within populations of anatomies. However, traditional correspondence-based SSM generation methods require a time-consuming re-optimization process each time a new subject is added to the cohort, making the inference process prohibitive for clinical research. Additionally, they require complete geometric proxies (e.g., high-resolution binary volumes or surface meshes) as input shapes to construct the SSM. Unordered 3D point cloud representations of shapes are more easily acquired from various medical imaging practices (e.g., thresholded images and surface scanning). Point cloud deep networks have recently achieved remarkable success in learning permutation-invariant features for different point cloud tasks (e.g., completion, semantic segmentation, classification). However, their application to learning SSM from point clouds is to-date unexplored. In this work, we demonstrate that existing point cloud encoder-decoder-based completion networks can provide an untapped potential for SSM, capturing population-level statistical representations of shapes while reducing the inference burden and relaxing the input requirement. We discuss the limitations of these techniques to the SSM application and suggest future improvements. Our work paves the way for further exploration of point cloud deep learning for SSM, a promising avenue for advancing shape analysis literature and broadening SSM to diverse use cases.

Keywords: Statistical Shape Modeling · Point Cloud Deep Networks · Morphometrics.

1 Introduction

Statistical Shape Modeling (SSM) enables population-based morphological analysis, which can reveal patterns and correlations between shape variations and clinical outcomes. SSM can help researchers understand the differences between healthy and pathological anatomy, assess the effectiveness of treatments, and identify biomarkers for diseases (e.g., [4,7,15,9,8]). The traditional pipeline for constructing SSM entails the segmentation of 3D images to acquire either binary volumes or meshes and then aligning these shapes. SSM can then be constructed *explicitly* via finding surface-to-surface correspondences across the cohort or *implicitly* via deforming a predefined atlas to each shape sample. Correspondence-based shape models are widely used due to their intuitive interpretation [20];

they comprise of sets of landmarks or *correspondence points* that are defined consistently using invariant points across populations that vary in their form. Historically, correspondence points were established manually to capture biologically significant features. This cumbersome, subjective process has since been replaced via automated *optimization*, which defines dense sets of correspondence points, aka a Point Distribution Model (PDM).

PDM optimization schemes have been defined using metrics such as entropy and minimum description length [11,12] and via parametric representations [16,22]. However, a significant drawback of these methods is that PDM optimization must be performed on the entire shape cohort of interest, which is time-consuming and hinders inference. To evaluate a new patient scan using PDM, the new scan must undergo segmentation and alignment, and the PDMs must be optimized again for the entire population of shapes. Moreover, current approaches require a complete, high-resolution mesh or binary volume representation of the shape that is free from noise and artifacts. Therefore, lightweight shape acquisition methods, such as thresholding clinical images, anatomical surface scanning, and shape acquired from stacked or orthogonal 2D slices, cannot be directly used for SSM [24,25]. Deep learning solutions have been proposed to mitigate these limitations by predicting PDMs directly from unsegmented 3D images using convolutional neural networks [6,1,2]. However, these frameworks require PDM supervision and, hence, need the traditional optimization-based workflow to acquire training data.

Effective SSM from point clouds would be widely applicable in clinical research, from artery disease progression from point clouds acquired via biplane angiographic and intravascular ultrasound image data fusion [24], to orthopedics implant design from point clouds acquired via 3D body scanning [25]. Applying existing methods for SSM generation from point clouds requires converting them to meshes or rasterizing them into segmentations, which is nontrivial given that point clouds are unordered and do not retain surface normals. SSM directly from point clouds would enable many clinical studies. Recently point cloud deep learning has gained attention, with significant efforts focused on effective point completion networks (PCNs) for generating complete point clouds from partial observations. Most point-completion methods use order-invariant feature encoding and two-stage coarse-to-fine decoding. An important outcome of such networks, which has been overlooked and not reported, is that the learned coarse point clouds are ordered and provide correspondence. In this work, we acknowledge this missed potential and explore the use of PCNs to predict PDM from 3D point clouds in an unsupervised manner. We investigate state-of-the-art PCNs as potential solutions for generating PDMs that (1) accurately represent shapes via uniformly distributed points constrained to the surface and (2) provide good correspondences that capture population-level statistics. We discuss the benefits of this approach, its robustness to missingness and training size, current limitations, and possible improvements. This discussion will bring awareness to the community about the potential for learning SSM from point clouds and ultimately make SSM a more accessible, viable option in future clinical research.

2 Background

Point Distribution Models The goal of SSM is to capture the inherent characteristics or underlying parameters of a shape class that remain when global geometrical information is removed. Given a PDM, correspondence points can be averaged across subjects to provide a mean shape, and principal component analysis (PCA) can be performed to compute the modes of shape variation, which can then be visualized and used in downstream medical tasks. Furthermore, if a PDM contains sub-populations, such as disease versus control, the differences in mean shapes can be quantified and visualized, providing group characterization.

Point Cloud Deep Learning Deep learning from 3D point clouds is an emerging area of research with numerous applications in computer vision, robotics, and medicine (e.g., classification, object tracking, segmentation, registration, pose estimation) [3]. PointNet [19] pioneered a Multi-Layer Perceptron (MLP) and max-pooling-based approach for permutation invariant feature learning from raw point clouds. FoldingNet [29] proposed a point cloud auto-encoder with a folding-based decoder that utilizes 2D grid deformation for reconstruction. Several convolutional approaches have been proposed, including mapping point clouds to voxel grids to directly apply 3D CNNs [14] and graph-based methods [27].

The initial *point cloud completion* network, PCN [31], utilized PointNet [19] and FoldingNet [29] with a coarse-to-fine decoder. Since then, numerous point cloud completion approaches have been proposed, including point MLP and folding-based extensions, 3D convolution approaches, graph-based methods, generative modeling approaches (including generative adversarial network GAN-based and variational autoencoder VAE-based), and transformer-based methods. See [13] for a recent survey. Many approaches follow the general framework of first encoding the point cloud into a permutation-invariant feature representation, then decoding the encoded shape feature to acquire a coarse or sparse point cloud, and finally, refining the coarse point cloud to acquire the dense complete prediction. The general architecture of these methods is shown in Figure 1.

3 Methods

3.1 Point Completion Networks for SSM

Our experiments demonstrate that when coarse-to-fine point completion networks are trained on anatomical shapes, the bottleneck captures a population-specific shape prior. Directly decoding the shape feature representation results in a consistent ordering of the intermediate coarse point cloud across samples, providing a PDM. This phenomenon can be intuitively understood as an application of Occam’s razor, where the model prefers to learn the simplest solution, resulting in consistent output ordering. Many point completion networks contain skip connections from the feature and/or input space to the refinement network. In the case where the unordered input point cloud is fed to the refinement network, the ordering of the output dense point cloud is understandably lost.




Fig. 1: General coarse-to-fine architecture of the considered PCNs.

To study whether point completion networks can learn anatomical SSM, we first extract point clouds from mesh vertices, then train point completion models, and finally evaluate the effectiveness of the predicted coarse point clouds as PDMs. Note this approach is not restricted to input point clouds obtained from meshes; point clouds from any acquisition process can be used. Global geometric information is factored out by aligning all shapes via iterative closest points [5] to a reference shape. We utilized the open-source toolkit ShapeWorks [10] for this step. In our experiments, the aligned, unordered mesh vertices serve as ground truth complete point clouds. The ground truth points are randomly downsampled to 2048 points and permuted to provide input point clouds. As is standard, the point clouds are uniformly scaled to be between -1 and 1 to assist network training. We consider a state-of-the-art model from the major point completion approach categories:

- **PCN**[31]: Point Completion Network (MLP-based)
- **ECG**[17]: Edge-aware Point Cloud Completion with Graph Convolution (convolution-based)
- **VRCnet**[18]: Variational Relational Point Completion Network (generative-modeling based)
- **SFN**[28]: SnowflakeNet (transformer-based)
- **PointAttN**[26]: Point Attention Network (attention-based)

Point completion network loss is based on the $L1$ **Chamfer Distance (CD)** [13], which defines the minimum distance between two sets of points. The loss is typically defined as a combination of coarse and dense loss with weighting parameter α , which we consistently set across models. All model’s hyperparameters are set to the original implementation values, and training is run until convergence (as assessed by training CD).

3.2 Evaluation Metrics

In addition to CD, **Fscore** [23] is typically used in point completion to quantify the percentage of points that are reconstructed correctly. In analyzing accurate surface sampling for SSM, we quantify the **point-to-face distance (P2F)** from each point in the predicted point cloud to the closest surface of the ground truth mesh. To analyze point **uniformity**, we quantify the variance in the distance of each point to its six nearest neighbors. A uniform PDM would result in a small variance in the point nearest neighbor distance.

In addition, we consider PDM correspondence analysis metrics. An ideal PDM is compact, meaning that it represents the distribution of the training data using the minimum number of parameters. We quantify **compactness** as the

number of PCA modes required to capture 99% of the variation in the correspondences. Moreover, a good PDM should generalize well from training examples to unseen examples. The **generalization** metric is defined as the reconstruction error ($L2$) between a predicted set of correspondences of a held-out point cloud and the correspondences reconstructed via the training PDM. Finally, effective SSM is specific, generating only valid instances of the shape class presented in the training set. The average distance between correspondences sampled from the training PDM and the closest existing training correspondences provides the **specificity** metric.

4 Experiments

We use five datasets in experiments – one synthetic ellipsoid dataset for a proof-of-concept and four real anatomical shapes: proximal femurs, left atrium of the heart, spleen, and pancreas. These datasets vary greatly in size (see Table 1, Column 1) and shape variation. Details and visualization of these cohorts are provided in the supplementary materials. In all experiments, the input point cloud size is 2048, the coarse output size is 512, and the dense output size is 2048. The real datasets were split (90%/10%) into training and testing sets. Stratified sampling via clustering was used to define the test set to ensure it is representative, given the low sample size.

Table 1: **Results:** Evaluation metrics across all models and datasets. All metrics are quantified on coarse predictions, with the exception of the gray column, which reports the dense prediction CD. The CD values are scaled by 1000 for reporting and Fscore is calculated at 1% threshold. Best test values are shown in bold.

Data	Model	Point Accuracy Metrics (train/test)					SSM Evaluation Metrics		
		Dense CD ↓	CD ↓	Fscore ↑	P2F (mm) ↓	Uniformity ↓	Comp. ↓	Gen. ↓	Spec. ↓
Ellipsoids	PCN[31]	0.908/0.917	1.80/1.80	0.561/ 0.582	0.0142/ 0.0168	0.325/0.354	2	0.140	0.348
	ECG[17]	0.929 /0.923	1.80/1.78	0.554/0.579	0.0171/0.0196	0.308/0.335	2	0.136	0.348
	VRcnet[18]	0.923/0.919	1.78/ 1.77	0.563/ 0.582	0.0384/0.0367	0.295/0.327	2	0.143	0.348
	SFN[28]	0.842/ 0.838	1.92/1.91	0.559/0.566	0.0960/0.0949	0.237/0.256	2	0.103	0.332
	PointAttN[26]	0.912/0.911	1.78/ 1.77	0.554/0.571	0.0255/0.0315	0.217/ 0.243	2	0.137	0.356
Femur	PCN[31]	1.27/1.80	2.30/ 2.91	0.501/ 0.386	0.235/ 0.712	1.77/1.63	18	0.28	1.30
	ECG[17]	1.59/2.29	2.37/3.02	0.485/0.377	0.297/0.750	1.71/1.63	15	0.255	1.23
	VRcnet[18]	1.78/2.33	2.88/3.27	0.399/0.362	0.676/0.876	1.62/1.60	5	0.259	0.786
	SFN[28]	1.04/ 1.29	2.36/3.13	0.464/0.365	0.341/0.774	1.01/ 0.967	17	0.297	1.37
	PointAttN[26]	1.27/1.50	3.25/3.94	0.352/0.295	0.759/1.03	2.32/2.29	9	0.199	0.905
Left Atrium	PCN[31]	0.405/0.773	0.707/1.09	0.941/0.865	0.245/1.02	1.87/2.01	82	0.932	5.85
	ECG[17]	0.571/0.868	0.814/1.16	0.919/0.848	0.440/1.07	1.85/2.04	57	0.929	5.22
	VRcnet[18]	0.070/ 0.085	0.886/1.62	0.904/ 0.768	0.632/1.53	1.87/2.19	81	1.04	6.09
	SFN[28]	0.310/0.311	0.822/ 0.948	0.927/0.902	0.511/ 0.831	1.30/ 1.46	49	0.783	5.07
	PointAttN[26]	0.332/0.360	0.875/1.20	0.909/0.840	0.523/1.08	1.65/1.79	82	0.807	5.46
Spleen	PCN[31]	3.59/7.88	4.51/9.77	0.326/0.155	1.67/3.73	15.3/12.0	7	1.47	4
	ECG[17]	1.84/6.46	3.0/9.69	0.456/0.174	1.07/3.94	10.2/7.98	14	1.62	5.57
	VRcnet[18]	0.408/ 0.583	5.03/14.3	0.318/ 0.117	1.81/4.59	16.1/12.5	6	1.73	4.48
	SFN[28]	0.986/1.28	4.57/ 7.23	0.265/0.189	1.64/ 3.07	17.8/14.7	6	1.08	4.18
	PointAttN[26]	1.13/3.70	3.49/11.5	0.380/0.158	1.33/4.22	7.67/ 6.72	15	1.86	6.52
Pancreas	PCN[31]	0.571/1.63	0.869/2.02	0.895/0.710	0.526/1.92	3.44/3.99	66	1.12	5.31
	ECG[17]	1.11/2.51	1.10/2.08	0.843/0.700	0.883/1.99	3.53/4.02	48	1.01	4.7
	VRcnet[18]	0.100/ 0.138	2.40/3.79	0.614/ 0.507	2.13/3.17	4.75/5.25	18	0.904	3.49
	SFN[28]	0.329/0.557	0.95/ 1.85	0.880/0.736	0.764/ 1.83	3.34/ 3.61	55	0.955	5.16
	PointAttN[26]	0.407/0.837	1.07/2.55	0.849/0.629	0.897/2.31	3.52/3.65	96	1.05	5.55

Proof-of-Concept: Ellipsoids As a proof-of-concept, we generate 3D axis-aligned ellipsoid shapes with fixed z -radius and random x and y -radius. A testing set of 30 and a training set of just 50 were randomly defined to emulate the scarce data scenario. The results in Table 1 show that all of the model variants performed well with low S2F distance ($<0.1\text{mm}$), and all PDMs correctly captured just two modes of variation (x and y-radius). Results from the PCN[31] model are shown in Figure 2.




Fig. 2: **Ellipsoid PDM from PCN[31]** Left: Example test predictions are shown on the ground truth mesh with color-denoting point correspondence. Points are reasonably uniformly spread with good correspondence, shown via zoomed-in boxes. Right: The primary and secondary modes (training and testing combined) are shown with ± 1 standard deviation from the mean shape. Color and vectors denote the difference from the mean. The shape model correctly characterizes the x and y-radius as the only source of variation.

Femur The femur dataset is comprised of 56 femoral heads segmented from CT scans, nine of which have the cam-FAI pathology characterized by an abnormal bone growth lesion that causes hip osteoarthritis [4]. We utilize this pathology to analyze if PDM from point completion networks can correctly characterize group differences. Table 1 shows the predicted coarse particles are close to the surfaces (P2F distance of 0.1mm), and the PCN[31] model performs best in this regard. The PCN[31] predictions were used to analyze the difference between the normal and CAM pathology mean shapes. Figure 3 shows the pathology is correctly characterized, and the *Linear Discrimination of Variation* (LDA) plot shows the difference in normal and CAM distributions captured by the PDM.

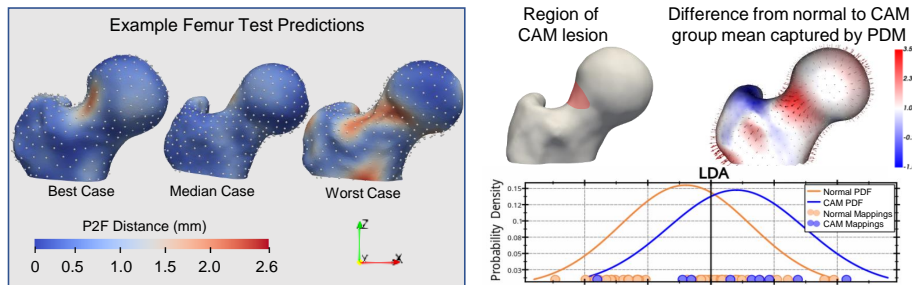


Fig. 3: **Femur PDM from PCN[31]** Left: Test examples with P2F distance displayed as a heatmap. Right: The expected region of CAM lesion correlates with the difference found from the normal to CAM group means.

Left Atrium The left atrium dataset comprises of 1096 shapes segmented from cardiac LGE MRI images from unique atrial fibrillation patients. This cohort contains significant morphological variation in overall size, the size of the left atrium appendage, and the number and arrangement of the pulmonary veins. This variation is reflected in the large compactness values in Table 1. Despite this variation, the models achieve reasonable CD and P2F scores due to the large training size. Figure 4 shows the PDM predicted by SFN[28], which performed best. From the prediction examples, we can see the model represents the training data very well, even in the worst case, which has an extremely enlarged left atrium appendage. The size of the appendage is appropriately captured in the modes of variation and the performance on the test examples is reasonable with the exception of those that are not well represented by the training data, such as the test set worst case. This highlights the importance of a large, representative training set. To further illustrate this importance, we perform an ablation experiment evaluating the performance of the PCN[31] model with respect to the left atrium training test size (Figure 4).

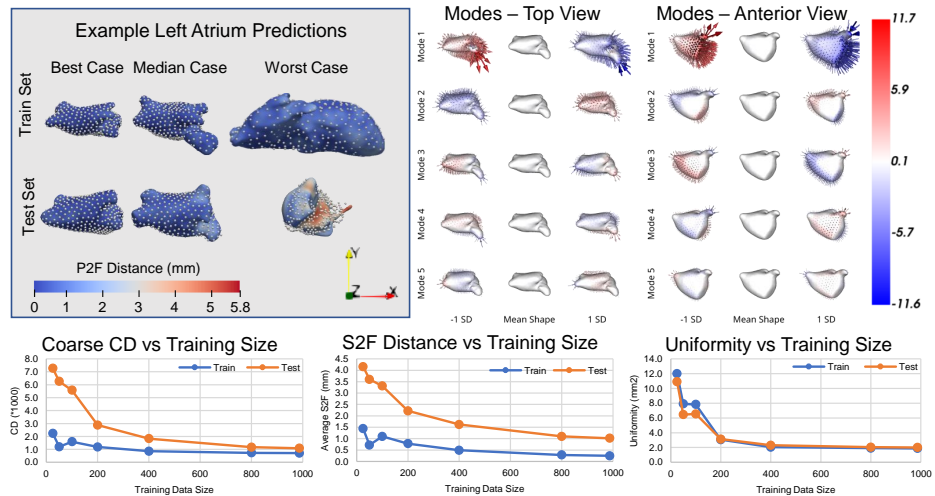


Fig. 4: **Top: Left Atrium PDM from SFN[28]** Left: Examples from the train and test sets are shown from top view with P2F distance. Right: The first 5 modes of variation are displayed from top and anterior view. **Bottom: Left Atrium PCN[31] results vs Training Size**

Pancreas We utilize the pancreas dataset [21] to analyze the impact of incomplete input point clouds as the point completion networks were designed to address. Cases of incomplete observations frequently arise in clinical research. For example, in the analysis of bones where some are clipped due to scan size or in cases where 3D shape is interpreted from stacked or orthogonal 2D observations. Traditional methods of SSM generation are unable to handle such cases,

but the point cloud learning approach has the potential to. Figure 5 shows how the test set error increase as the percentage of missing points increases. The SFN[28] model provides the best results given partial input.




Fig. 5: **Effect of Partial Input on Pancreas Prediction Accuracy**

Spleen The spleen dataset [21] is included to provide an example of a small dataset with a large amount of variation to stress test the point completion models. Table 1 shows the models perform the worst on this dataset with regards to CD, Fscore, P2F, and Uniformity. This example illustrates the limitations of this approach to SSM generation.

5 Discussion and Conclusion

Our experiments demonstrate that point cloud networks can learn accurate SSM of anatomy when provided with a sufficiently large and representative training dataset. The transformer-based SFN[28] architecture provided the best overall results among the models we explored. Interestingly, the simplest model, PCN[31], achieved similar, and sometimes better, SSM accuracy than more current state-of-the-art methods, despite its inferior performance in point completion benchmarks [17,18,28,26]. This may be attributed to point completion benchmarks involving multiple object class point completion, which is a more challenging task. Another significant difference between our experiments and point completion benchmarks is the PCN datasets have tens of thousands of examples [31,18,30], while we worked with limited training data – the typical scenario in shape analysis. Our experiments demonstrate that PCNs can effectively predict SSM under limited training data when shape variation is minimal, as in the case of proximal femurs. However, they struggle when there is significant shape variation, such as in the spleen cohort.

This work indicates a promising potential for adapting characteristics of point completion architectures and learning schemes to tailor to the task of predicting SSMs from point clouds. Potential improvements could be made to the training objective, such as penalization for non-uniformity and bottleneck regularization for compact population-statistical learning. Additionally, improvements could be made to address the scarce training data scenario, such as model-based data augmentation and probabilistic transfer learning. This work establishes the groundwork for future research into the potential of point cloud deep learning for SSM,

which offers significant benefits such as reducing the inference burden, increasing SSM accessibility, and potentially accelerating its adoption as a widespread clinical tool. Moreover, it reduces the input shape restrictions required by current SSM generation techniques, vastly expanding the potential applications of SSM.

Acknowledgements This work was supported by the National Institutes of Health under grant numbers NIBIB-U24EB029011, NIAMS-R01AR076120, NHLBI-R01HL135568, and NIBIB-R01EB016701. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health. The authors would like to thank the University of Utah Division of Cardiovascular Medicine for providing left atrium MRI scans and segmentations from the Atrial Fibrillation projects as well as the Orthopaedic Research Laboratory (Andrew Anderson, PhD) at the University of Utah for providing femur CT scans and corresponding segmentations.

References

1. Adams, J., Bhalodia, R., Elhabian, S.: Uncertain-deepssm: From images to probabilistic shape models. *Shape in Medical Imaging : International Workshop, ShapeMI 2020, Held in Conjunction with MICCAI 2020, Lima, Peru, October 4, 2020, Proceedings* **12474**, 57–72 (2020)
2. 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)
3. Akagic, A., Krivić, S., Dizdar, H., Velagić, J.: Computer vision with 3d point cloud data: Methods, datasets and challenges. In: *2022 XXVIII International Conference on Information, Communication and Automation Technologies (ICAT)*. pp. 1–8. IEEE (2022)
4. Atkins, P.R., Elhabian, S.Y., Agrawal, P., Harris, M.D., Whitaker, R.T., Weiss, J.A., Peters, C.L., Anderson, A.E.: Quantitative comparison of cortical bone thickness using correspondence-based shape modeling in patients with cam femoroacetabular impingement. *Journal of Orthopaedic Research* **35**(8), 1743–1753 (2017)
5. Besl, P.J., McKay, N.D.: Method for registration of 3-d shapes. In: *Sensor fusion IV: control paradigms and data structures*. vol. 1611, pp. 586–606. Spie (1992)
6. Bhalodia, R., Elhabian, S.Y., Kavan, L., Whitaker, R.T.: Deepssm: A deep learning framework for statistical shape modeling from raw images. In: *Shape In Medical Imaging at MICCAI. Lecture Notes in Computer Science*, vol. 11167, pp. 244–257. Springer (2018)
7. Bischoff, J.E., Dai, Y., Goodlett, C., Davis, B., Bandi, M.: Incorporating population-level variability in orthopedic biomechanical analysis: a review. *Journal of biomechanical engineering* **136**(2), 021004 (2014)
8. Bruse, J.L., McLeod, K., Biglino, G., Ntsinjana, H.N., Capelli, C., Hsia, T.Y., Sermesant, M., Pennec, X., Taylor, A.M., 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)

9. Carriere, N., Besson, P., Dujardin, K., Duhamel, A., Defebvre, L., Delmaire, C., Devos, D.: Apathy in parkinson's disease is associated with nucleus accumbens atrophy: a magnetic resonance imaging shape analysis. *Movement disorders* **29**(7), 897–903 (2014)
10. 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)
11. Cates, J., Fletcher, P.T., Styner, M., Shenton, M., Whitaker, R.: Shape modeling and analysis with entropy-based particle systems. In: *IPMI*. pp. 333–345. Springer (2007)
12. Davies, R.H., Twining, C.J., Cootes, T.F., Waterton, J.C., Taylor, C.J.: A minimum description length approach to statistical shape modeling. *IEEE transactions on medical imaging* **21**(5), 525–537 (2002)
13. Fei, B., Yang, W., Chen, W.M., Li, Z., Li, Y., Ma, T., Hu, X., Ma, L.: Comprehensive review of deep learning-based 3d point cloud completion processing and analysis. *IEEE Transactions on Intelligent Transportation Systems* (2022)
14. Le, T., Duan, Y.: Pointgrid: A deep network for 3d shape understanding. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 9204–9214 (2018)
15. Merle, C., Inmann, M.M., Waldstein, W., Pegg, E.C., Aldinger, P.R., Gill, H.S., Murray, D.W., 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)
16. 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)
17. Pan, L.: Ecg: Edge-aware point cloud completion with graph convolution. *IEEE Robotics and Automation Letters* **5**(3), 4392–4398 (2020)
18. Pan, L., Chen, X., Cai, Z., Zhang, J., Zhao, H., Yi, S., Liu, Z.: Variational relational point completion network. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 8524–8533 (2021)
19. Qi, C.R., Su, H., Mo, K., Guibas, L.J.: Pointnet: Deep learning on point sets for 3d classification and segmentation. In: *Proceedings of the IEEE conference on computer vision and pattern recognition*. pp. 652–660 (2017)
20. Sarkalkan, N., Weinans, H., Zadpoor, A.A.: Statistical shape and appearance models of bones. *Bone* **60**, 129–140 (2014)
21. Simpson, A.L., Antonelli, M., Bakas, S., Bilello, M., Farahani, K., Van Ginneken, B., Kopp-Schneider, A., Landman, B.A., 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)
22. Styner, M., Oguz, I., Xu, S., Brechbühler, C., Pantazis, D., Levitt, J.J., Shenton, M.E., Gerig, G.: Framework for the statistical shape analysis of brain structures using spharm-pdm. *The insight journal* p. 242 (2006)
23. Tatarchenko, M., Richter, S.R., Ranftl, R., Li, Z., Koltun, V., Brox, T.: What do single-view 3d reconstruction networks learn? In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 3405–3414 (2019)
24. Timmins, L.H., Samady, H., Oshinski, J.N.: Effect of regional analysis methods on assessing the association between wall shear stress and coronary artery disease progression in the clinical setting. In: *Biomechanics of Coronary Atherosclerotic Plaque*, pp. 203–223. Elsevier (2021)

25. Treleaven, P., Wells, J.: 3d body scanning and healthcare applications. *Computer* **40**(7), 28–34 (2007). <https://doi.org/10.1109/MC.2007.225>
26. Wang, J., Cui, Y., Guo, D., Li, J., Liu, Q., Shen, C.: Pointattn: You only need attention for point cloud completion. *arXiv preprint arXiv:2203.08485* (2022)
27. Wang, Y., Sun, Y., Liu, Z., Sarma, S.E., Bronstein, M.M., Solomon, J.M.: Dynamic graph cnn for learning on point clouds. *Acm Transactions On Graphics (tog)* **38**(5), 1–12 (2019)
28. Xiang, P., Wen, X., Liu, Y.S., Cao, Y.P., Wan, P., Zheng, W., Han, Z.: Snowflakenet: Point cloud completion by snowflake point deconvolution with skip-transformer. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 5499–5509 (2021)
29. Yang, Y., Feng, C., Shen, Y., Tian, D.: Foldingnet: Point cloud auto-encoder via deep grid deformation. In: *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)* (June 2018)
30. Yu, X., Rao, Y., Wang, Z., Liu, Z., Lu, J., Zhou, J.: Pointtr: Diverse point cloud completion with geometry-aware transformers. In: *Proceedings of the IEEE/CVF international conference on computer vision*. pp. 12498–12507 (2021)
31. Yuan, W., Khot, T., Held, D., Mertz, C., Hebert, M.: Pcn: Point completion network. In: *2018 international conference on 3D vision (3DV)*. pp. 728–737. *IEEE* (2018)

A Dataset Visualization

Figures 6-10 display example shapes from each of the datasets used in experiments from two viewpoints.

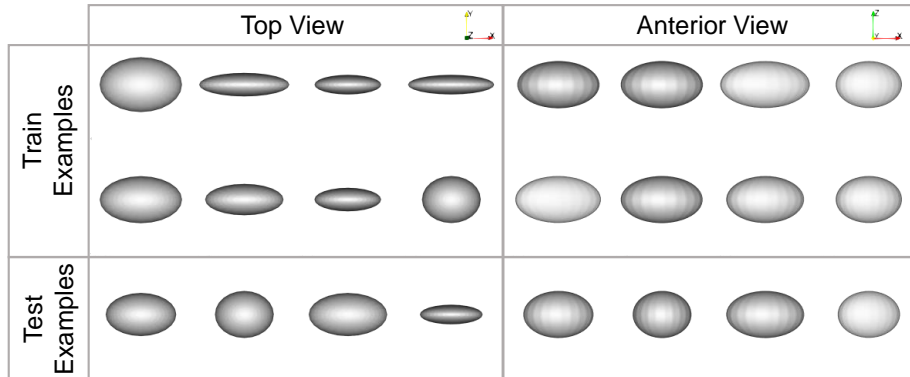


Fig. 6: **Ellipsoid Dataset** Eight of the 50 training examples and 4 of the 30 test examples are shown from both the top and anterior view. Ellipsoids have random x and y -radius and fixed z radius.

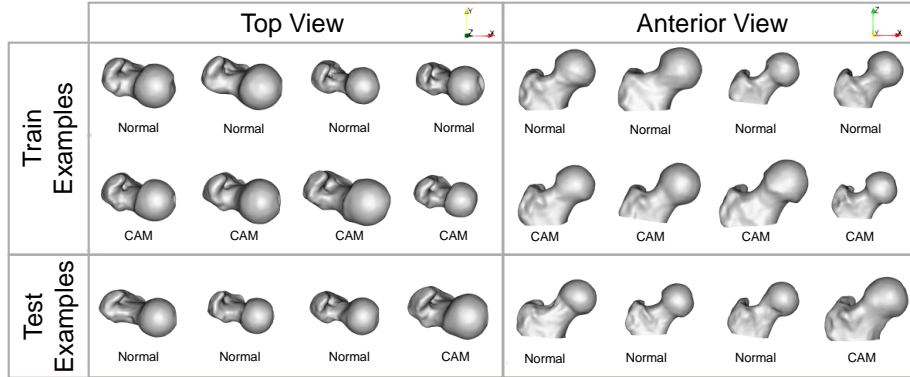


Fig. 7: **Femur Dataset** Eight of the 51 training examples and 4 of the 5 test examples are shown from both top and anterior view. The group label is displayed, denoting which femurs have the CAM pathology. Overall 47 of the femurs were normal and 9 were CAM, 8 were in the train set, and 1 was in the test.

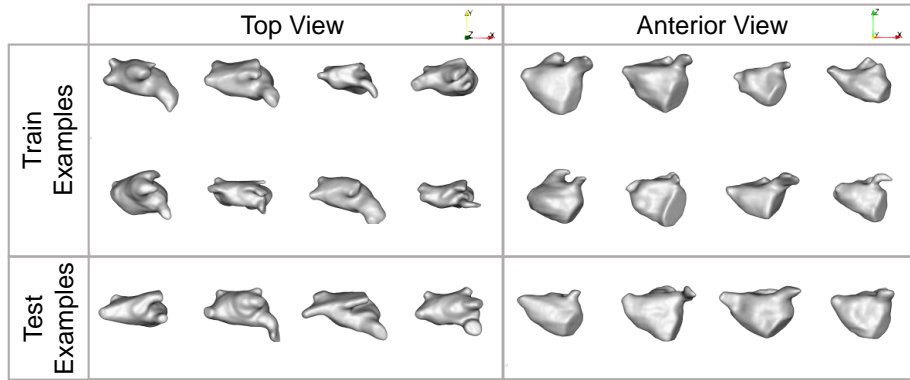


Fig. 8: **Left Atrium Dataset** Eight of the 987 training examples and 4 of the 109 test examples are shown from both top and anterior view.

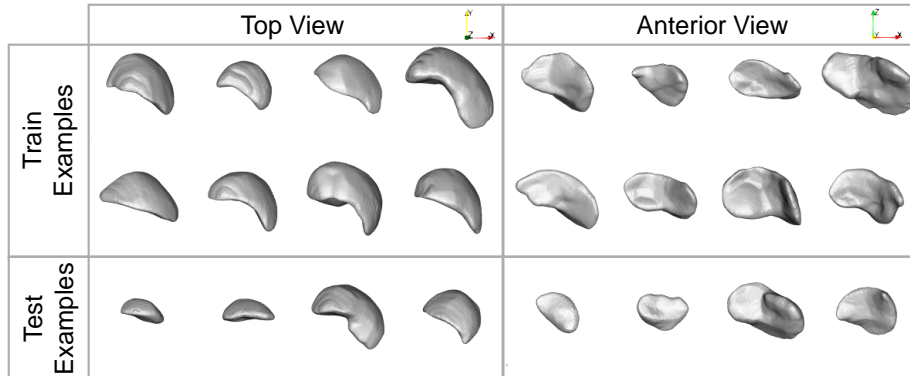


Fig. 9: **Spleen Dataset [21]** Eight of the 36 training examples and all 4 of the test examples are shown from both top and anterior view.




Fig. 10: **Pancreas Dataset [21]** Eight of the 245 training examples and 4 of the 28 test examples are shown from both top and anterior view.