Beyond the Left Atrium: Joint Generative Modeling of Bi-Atrial Anatomy

Jule Bender[1 [0009−0000−7013−4344]], Christian Goetz[1,2 [0009−0006−5222−4390]], Pascal Maierhofer[1 [0009−0002−6308−9252]], Constanze Schmidt[2,3 [0000−0001−5897−237X]], and Axel Loewe[1 [0000−0002−2487−4744]]

[1] Institute of Biomedical Engineering, Karlsruhe Institute of Technology (KIT), Fritz-Haber-Weg 1, 76131 Karlsruhe, Germany
publications@ibt.kit.edu

[2] Department of Cardiology, Angiology and Pneumology, Heidelberg University Hospital, Im Neuenheimer Feld 410, 69120 Heidelberg, Germany

[3] Department of Cardiology and Pneumology, University Medical Center Goettingen, Robert-Koch-Str. 40, Goettingen, 37075, Lower Saxony, Germany

Abstract. Atrial shape has been linked to atrial fibrillation (AF), yet many aspects remain poorly understood. Recent evidence suggests that the right atrium (RA) affects AF vulnerability in both chambers, highlighting the need for bi-atrial analysis. However, open bi-atrial cohorts are scarce. Statistical shape models (SSMs) address this gap, but model variability only linearly and require surface correspondence. Recently, deep learning (DL)-based generative models have emerged as an alternative. We present a DL-based bi-atrial generative model, extending an existing signed distance function based auto-decoder framework for the left atrium (LA) to a dual-head network with shared trunk, trained on 344 patients from a public CT dataset. We investigated latent dimensionality effects on reconstruction and generation performance. The model produced anatomically realistic geometries with separated atria, variable shapes, and diverse LA appendage morphologies. Volumes and volume ratios aligned with clinically observed ranges. Reconstruction and generation improved with latent dimension, with dimension 64 best balancing complexity and quality where mean Chamfer distances were 1.2 mm for reconstruction and 2.5 mm (LA) and 2.2 mm (RA) for generation. LA reconstruction distances exceeded those in the RA, likely due to greater topological complexity and the comparatively lower RA data fidelity. Compared to the existing bi-atrial SSM, the DL model required higher dimension, possibly confounded by dataset differences. Overall, loss plateaued within 400 epochs for all models, suggesting early stopping for future work. In conclusion, the model can support virtual cohort generation for AF *in silico* studies; allowing conditioning on clinical metadata in future work.

Keywords: Generative DL · Bi-atrial model · Atrial geometry · Signed distance functions.

1 Introduction

Atrial fibrillation (AF) is the most prevalent cardiac arrhythmia and a leading cause of stroke [25]. While atrial shape characteristics have been linked to AF pathology and thrombogenesis, many aspects of anatomical shape variation and their clinical implications remain insufficiently characterized [10,21]. Notably, the right atrium (RA) has been shown to affect AF vulnerability in both atria [12], underscoring the need for three-dimensional, bi-atrial analysis. However, large open bi-atrial cohorts remain scarce, and computational studies are therefore often limited to small, private datasets [17].

Synthetic virtual patient cohorts can address these limitations, mitigating ethical and logistical constraints [17]. Statistical shape models (SSMs) have been established as a methodological basis for cohort modeling approaches [21,22]. For the atria, Nagel et al. proposed and publicly released the first bi-atrial SSM along with simulation-ready geometries [14,16]. This model has since been successfully applied for large-scale electrocardiogram (ECG) dataset generation [7], machine learning tasks [13,15], and for evaluating *in silico* ablation and pharmacological treatment strategies [4]. However, traditional point distribution SSMs rely on dense surface correspondence, model variability only linearly, and are typically derived from small datasets.

Deep learning (DL)-based generative models have emerged as a compelling alternative, allowing additionally conditioning on clinical metadata such as age and sex [2,20,26,27]. By learning a compact, non-linear latent space, these models represent anatomical variability without requiring explicit shape correspondence. Recently, Sørensen et al. [24] proposed a spatio-temporal model of the left atrium (LA). However, the model is restricted to a single chamber, preventing its application in studies requiring bi-atrial geometries, such as P-wave simulations [7,13,15] or the assessments of RA contributions to AF vulnerability [12]. While whole-heart generative models have been proposed [5,11], they suffer from limited representation of the atria, rely on data correspondence or were designed for specific pathologies with a small training set. A high fidelity deep-learning based bi-atrial generative model has not yet been proposed.

In this work, we present a DL-based generative model for bi-atrial anatomy, extending a signed distance function (SDF)-based auto-decoder framework [24] into a dual-head architecture that jointly represents left and right atrial geometry. We systematically evaluate the model's ability to reconstruct and generate anatomically realistic shapes and investigate the effect of latent space dimensionality and training duration on model performance. Specifically, we hypothesize that the non-linear latent representation requires fewer dimensions than traditional SSMs to achieve comparable fidelity. We further hypothesize that the optimal model configuration for geometry reconstruction aligns with that for shape generation. Thus, although RA data fidelity is currently limited by the dataset, this study proposes a foundational framework for high-fidelity, DL-based bi-atrial generative modeling to create synthetic bi-atrial cohorts. This framework can be extended in future work by conditioning on clinical metadata such as age, sex, and AF history.

2 Materials and Methods

2.1 Data and Preprocessing

This study used the open CT dataset of 685 patients by Hansen et al. [9] with in-plane resolution of $0.29 - 0.43\,\mathrm{mm}^2$ and a spacing of $0.25 - 0.45\,\mathrm{mm}$. All individuals were older than 18 years and had documented history of ischemic stroke, transient ischemic attack and/or peripheral artery disease [9]. Notably, RA data fidelity in this dataset was comparatively lower than for the LA, particularly near the venae cavae. The provided blood pool segmentation masks were automatically refined to resolve intersections between left atrial appendages (LAAs) and pulmonary veins (PVs), and three-dimensional surface models were extracted using 3D Slicer [6]. Seed points for the PVs and LAA were extracted automatically from the provided anatomical labels, restricting this step to cases with exactly four PVs. PV clipping was performed using DIVAID [8] based on the seed points. Surfaces were subsequently remeshed to a target edge length of 1 mm. The resulting geometries underwent visual inspection, and models with visible segmentation artifacts or incomplete anatomy were excluded, yielding a final cohort of 344 bi-atrial models. Main reasons for exclusion were failure of automatic vein clipping ($\sim$ 100 meshes) and incorrect clipping of the PVs or LAA ($\sim$ 200 meshes) due to automatic seed placement. Included models had slightly smaller LA body volumes (excluding PVs and LAA) than excluded models (69.9 ± 16.3 ml vs. 72.4 ± 26.8 ml), and similarly smaller RA volumes (80.1 ± 17.8 ml vs. 81.1 ± 34.5 ml). All models were rigidly aligned at the interatrial septum and the joint envelope of all individual bounding boxes was normalized to the unit cube, leaving rotation and size as learnable parameters. $K = 210,000$ points were sampled per geometry, using the importance sampling strategy proposed by Sørensen et al. [23]. While favoring samples at geometrically complex structures via the shape diameter, also random samples throughout the unit cube were generated to ensure global coverage. SDFs to the LA and RA were computed separately.

2.2 Dual-Head Auto-Decoder Framework

Our model[1] builds on the auto-decoder framework of Sørensen et al. [24], which models left atrial geometry conditioned on clinical metadata via spatio-temporal SDFs following the DeepSDF framework [18]. Briefly, a sampling point $s_{n,k}$ and its corresponding SDF value $d_{n,k}$, i.e. the distance to the left atrial surface, with sign indicating interior or exterior, are combined with a shape-specific latent vector $\mathbf{z}_n$ and passed to a multi-layer perceptron that learns the mapping $d_{n,k} = f_\theta(\mathbf{z}_n \oplus \mathbf{s}_{n,k})$. Thereby, each latent vector $\mathbf{z}_n$ is a learnable embedding representing an individual training shape $n \in \{1, \ldots, N\}$, optimized jointly with the network parameters θ. The atrial surface is recovered as the zero-level isosurface of the predicted SDF, extracted via Marching Cubes on a uniform grid.

[1] Code is available at https://gitlab.kit.edu/kit/ibt-public/sdf4ba.git and data is available at https://doi.org/10.5281/zenodo.21900735.

We adopted the same multi-layer perceptron architecture as Sørensen et al., removing the temporal and clinical conditioning components. When generating bi-atrial surfaces, two failure modes can occur: either the two atria fuse, i.e., only a single closed surface is predicted, or the two chambers intersect. To adapt the architecture for bi-atrial model creation, we introduced a dual-head output layer. Specifically, the network shares all layers except the last, which is split into two heads that independently predict the SDF for the LA and RA, yielding f_θ^{LA} and f_θ^{RA} from a single forward pass. Thus, instead of a single SDF value for the combined bi-atrial mesh, the network learns an SDF vector $d_{n,k} = [d_{n,k}^{\mathrm{LA}}, d_{n,k}^{\mathrm{RA}}]$, comprising the distances to the left and right atrium, respectively. In contrast to separately trained decoders for each chamber, this design encourages a joint latent representation of bi-atrial anatomy to learn bi-atrial relationships while still allowing each chamber's geometry to be decoded independently. Accordingly, the new training objective minimizes the sum of the chamber-wise clamped L1 loss $\mathcal{L}$ and a newly added intersection penalty P, while keeping the regularization loss $\|\mathbf{z}_n\|_2^2$ with weight $\sigma = 10^{-4}$. For all N shapes, K sampling points $\mathbf{s}_{n,k}$, and the corresponding SDF vector $d_{n,k}$, this yields

$$\underset{\theta, \{\mathbf{z}_n\}_{n=1}^N}{\arg\min} \sum_{n=1}^{N} \sum_{k=1}^{K} \Big(\mathcal{L}\big(f_\theta^{\mathrm{LA}}\left(\mathbf{z}_n \oplus \mathbf{s}_{n,k}\right), d_{n,k}^{\mathrm{LA}}\big) + \mathcal{L}\big(f_\theta^{\mathrm{RA}}\left(\mathbf{z}_n \oplus \mathbf{s}_{n,k}\right), d_{n,k}^{\mathrm{RA}}\big) + \frac{1}{\sigma^2}\|\mathbf{z}_n\|_2^2 + \frac{1}{2}P \Big). \tag{1}$$

The intersection penalty P discourages overlapping chamber surfaces by

$$P = \frac{1}{|I|} \sum_{i \in I} \min\left(-f_\theta^{\mathrm{LA}}\left(\mathbf{z}_n \oplus \mathbf{s}_{n,k}\right), -f_\theta^{\mathrm{RA}}\left(\mathbf{z}_n \oplus \mathbf{s}_{n,k}\right)\right) \tag{2}$$

for the set I of points that could in principle be predicted to lie inside both chambers simultaneously by the two independent network heads.

The dataset was split randomly into 70 % training, 10 % validation, and 20 % testing. The model was implemented in PyTorch and trained for 1000 epochs on an Nvidia RTX 2080 Ti (11 GB). Validation was performed every 25 epochs by freezing the decoder and optimizing latent codes for the validation set. This optimization ran for 200 epochs at a fixed learning rate of 0.001 and a grid resolution of 128^3, after which the average Chamfer and Hausdorff distances were computed. To investigate the effect of latent space dimensionality, we trained five model variants with latent dimensions $8, 16, 32, 64, 128$, keeping all other parameters as in Sørensen et al. [24].

2.3 Evaluation

Training and validation loss curves were monitored to assess convergence behavior and detect overfitting. Additionally, validation set generalizability and specificity were evaluated across training epochs for dimension 64. To evaluate generalizability, i.e., the ability to reconstruct unseen geometries, reconstruction was evaluated on a hold-out validation set. To evaluate specificity, i.e., how closely generated shapes resemble training set shapes as an indicator of shape realism, $N = 100$ generated shapes were compared to the training set.

Reconstruction For each test shape, a latent code was optimized for 800 epochs by freezing the decoder parameters and minimizing the SDF reconstruction loss as in Sørensen et al. [24]. Reconstructed surfaces from a grid with resolution 256 were compared to the original meshes using the symmetric, normalized, bi-directional Chamfer distance

$$CD_{\text{sym, norm}}(S_1, S_2) = \frac{1}{2}\left(\frac{1}{|S_1|}\sum_{x\in S_1}\min_{y\in S_2}\|x-y\|_2 + \frac{1}{|S_2|}\sum_{y\in S_2}\min_{x\in S_1}\|x-y\|_2\right) \quad (3)$$

and the Hausdorff distance

$$HD(S_1, S_2) = \max\left(\max_{x\in S_1}\min_{y\in S_2}\|x-y\|, \max_{y\in S_2}\min_{x\in S_1}\|x-y\|\right) \quad (4)$$

separately for the LA and RA, where S_1 and S_2 denote the original and reconstructed surface points. Both metrics capture complementary aspects; while Chamfer distance assesses the average distance between the meshes to evaluate the overall shape fidelity, Hausdorff distances evaluates the maximum distance to analyze the worst-case error, making it more sensitive to localized artifacts. To assess the qualitative fidelity of a reconstructed anatomy, a subset of test shapes spanning the range of known LAA morphologies (chicken wing, cactus, cauliflower, windsock) was identified through manual classification, and representative reconstructions were inspected visually [3].

Generation For each model, $N = 100$ novel latent vectors were sampled from a multivariate Gaussian distribution fitted to the training latent codes. Specificity was measured to assess anatomical plausibility using Chamfer (3) and Hausdorff (4) distance to the closest training set mesh, defined by the smallest average Chamfer distance for RA and LA. Additionally, coverage was measured by computing the fraction of training shapes that serve as the nearest neighbor of at least one generated shape to assess the diversity. Furthermore, generated shapes were inspected qualitatively to assess realism and morphological variety. To assess bi-atrial relationships and clinical representativeness, LA (including PV trunks and LAA) and RA volumes, as well as LA-RA volume ratios for the generated and training meshes were compared to a clinical cohort of 641 individuals (38% men) free of prevalent cardiovascular disease, obesity, and hypertension [19]. To account for differences between 3D volumetric and area–length volume measurements, volumes were converted assuming the same correlation applies for both atria [1].

3 Results

All generated and reconstructed geometries were anatomically plausible, smooth, closed, and free of self-intersections, with clearly separated LA and RA. Notably, the generated surfaces were smoother than the original meshes, effectively suppressing artifacts such as small holes in the input data. To assess convergence

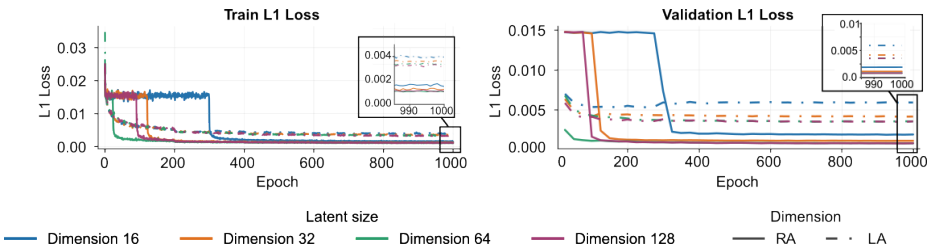

Fig. 1: Loss curves for different latent space dimensions for 1000 epochs. Training L1 loss (left) and validation L1 loss, evaluated every 25th epoch (right).

behavior, we evaluated training and validation loss curves (Figure 1). All models converged over the course of 1000 training epochs, with loss curves reaching stable plateaus. The epoch at which the primary loss drop occurred varied with latent dimension: the model with dimension 64 converged first, followed by dimensions 128, 32, and 16. Across all models, an asymmetry in chamber learning dynamics was observed, with LA loss decreasing prior to RA loss. Validation loss curves mirrored the training behavior, reaching plateaus at similar epochs. The model with dimension 8 failed to learn a meaningful RA representation and was therefore excluded from all subsequent evaluations. We additionally performed an ablation study to assess the benefit of the dual-head structure, replacing it with a single shared output layer predicting SDF values for the entire bi-atrial surface. This yielded geometries that were fused at the septum, confirming the benefit of the dual-head design. While an explicit intersection penalty was added as a safeguard against chamber overlap since separately predicted SDFs of the dual-head model could in theory yield intersecting chambers, ablating it did not lead to intersecting chambers in practice.

Reconstruction To evaluate reconstruction quality, we evaluated distances between reconstructed and original meshes in the test set (Figure 2). Overall, distances decreased with increasing latent space dimension (Figure 2A). Notably, dimensions 64 and 128 yielded similar distances, and RA distances were generally smaller than LA distances. For dimension 64, mean Chamfer distances were 1.2 ± 0.5 mm for both atria while mean Hausdorff distances were 8.9 ± 4.5 mm (LA) and 5.1 ± 2.2 mm (RA) respectively. Furthermore, for dimension 64, reconstruction quality plateaued across all evaluated epochs beyond epoch 199 for the validation set (Figure 2C). To assess variability qualitatively, we additionally inspected representative geometries for each LAA type visually (Figure 3). While this confirmed that the diversity of known LAA types is largely captured, fine structures were not fully reconstructed including the lateral parts in the chicken wing type LAA. While reconstruction generally improved with increasing dimension, this was not true for some individual models, for example, the cactus type LAA reconstruction in Figure 3 degraded at higher dimensions.

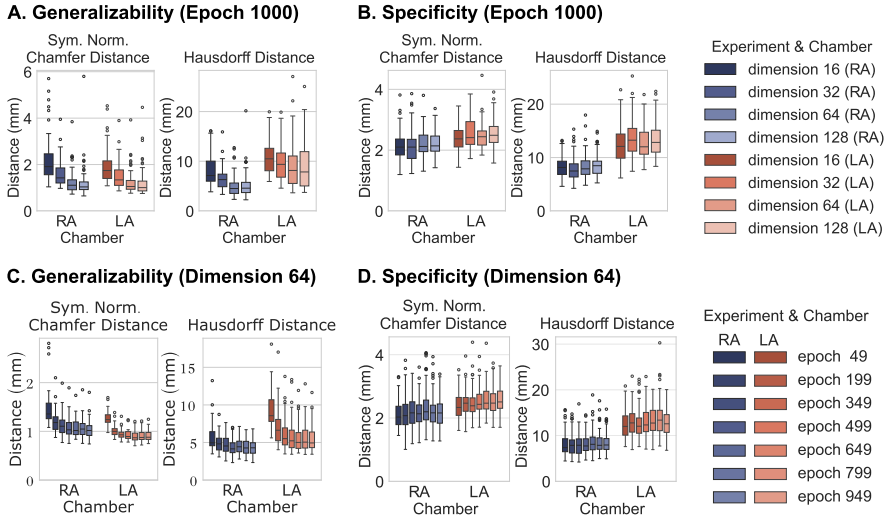

Fig. 2: **A.** Generalizability for the 69 test set models and **B.** specificity across 100 randomly generated models for different latent dimensions at epoch 1000. **C.** Generalizability for the 34 validation models and **D.** specificity across 100 randomly generated models for different epochs for dimension 64.

Generation To assess anatomical plausibility, we evaluated specificity across randomly generated geometries compared to the training set (Figure 2). Overall, distances were generally larger for the LA than for the RA (Figure 2B). While for the RA, dimension 32 yielded the lowest distance to the training set, no clear trend was observed for the LA. For specificity based on Chamfer distance at dimension 64, mean Chamfer distances were 2.5 ± 0.4 mm (LA) and 2.2 ± 0.4 mm (RA) while mean Hausdorff distances were 12.7 ± 2.9 mm (LA) and 8.3 ± 2.2 mm (RA). Similar to other validation metrics, specificity for dimension 64 remained stable across all evaluated epochs (Figure 2D). To assess variability qualitatively, we additionally visually evaluated the geometries (Figure 4). Across all latent dimensions, random samples drawn from the learned multivariate Gaussian distribution exhibited variable anatomy. Generation coverage, i.e., the fraction of training samples that serve at least once as the nearest neighbor of a generated sample, was approximately 30% across all dimensions. Training coverage, i.e., the fraction of samples that were the nearest neighbor of at least one other training sample, was 63% for the RA and 53% for the LA. To assess whether the model learned a meaningful bi-atrial relationship and evaluate clinical plausibility, RA and LA volumes were compared to a clinical cohort [19] (Figure 5). Overall, the majority of training and generated meshes fell within the reported reference ranges. The LA/RA volume ratio was similar for training and generated meshes but smaller than in the healthy reference cohort.

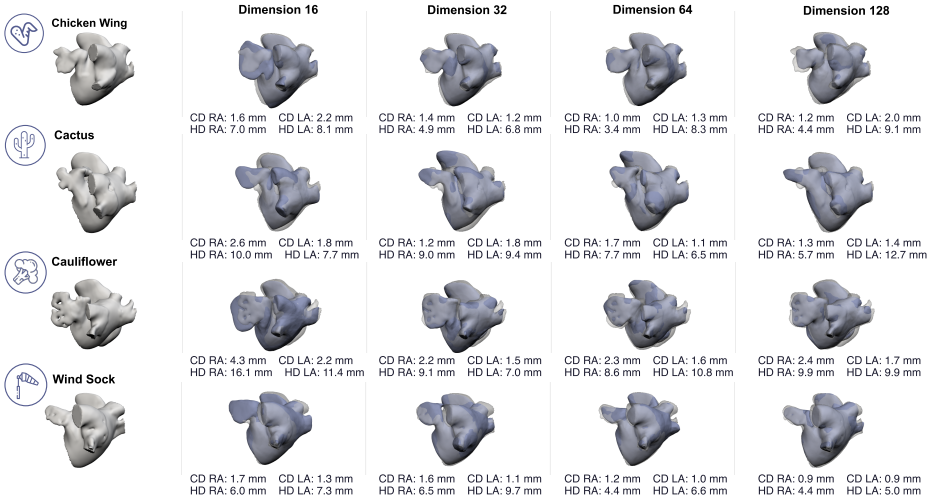

Fig. 3: Reconstruction (gray color) vs. original (white color) for different LAA morphology types across different latent space dimensions with Chamfer and Hausdorff distances at epoch 1000. Note that for windsock type in dimension 32, this is one of the rare cases where the LAA and PVs are connected. Icons in the figure created with flaticon.com.

4 Discussion and Conclusion

This work presents a DL-based generative model for joint left and right atrial modeling using SDFs. Building on an existing dataset [9] and a spatio-temporal LA model [24], we developed a dual-head network with a shared trunk but separate RA and LA output layers, similar to [11], who proposed a model using separated SDF values for each cardiac structure for type and shape-disentangled modeling of congenital heart disease. Our design enables a shared latent space that captures both chamber-specific variability and bi-atrial relationships while producing realistic geometries without intersecting chambers, unlike a single-head architecture predicting a combined SDF which yielded geometries fused at the septum. Furthermore, SDFs circumvent the need for correspondence between samples, which is advantageous for the complex and variable bi-atrial anatomy.

Optimal performance for both reconstruction and generation, was achieved with latent dimension 64, following Occam's razor by selecting the simplest model among those with comparable performance. This dimension is commonly used in related work but rarely explored in detail [20,24,26]. Furthermore, LA/RA volume ratios of the generated meshes closely matched those of the training set, suggesting that the model learned a meaningful LA-RA coupling rather than generating each chamber in isolation. Importantly, atrial volumes and volume ratios also fell within clinically observed ranges [19], although the volume ratios in both the training and generated sets were smaller than clinically observed. This remaining discrepancy may be attributable to differences in population

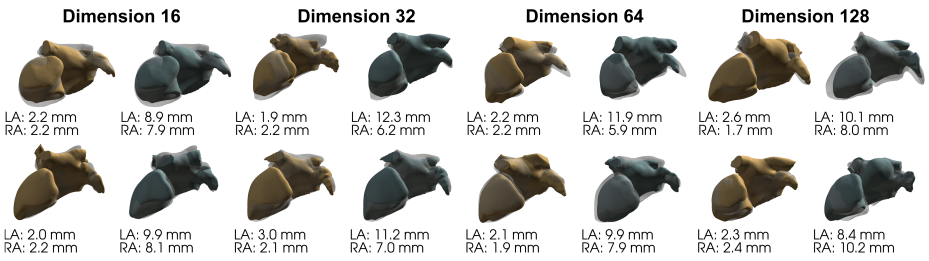

Fig. 4: Randomly generated geometries for different latent dimensions at epoch 1000 with closest neighbor according to symmetric normalized Chamfer distance (yellow color) and Hausdorff distances (green color).

characteristics or to systematic differences in volume estimation between chambers. Assuming sufficient variability in the training data, both traditional SSMs and generative DL models should, in principle, provide lower-dimensional representations of the shape independent of training set size. Compared to the existing bi-atrial SSM [16], which requires 24 eigenvectors to explain 95 % of the total shape variance, the optimal DL model required more dimensions. In this regard, our hypothesis is not supported by the results. However, this comparison is confounded by differences in training set size, as the smaller training set of the SSM by Nagel et al. may represent a simpler setting. Furthermore, since reconstruction and generation favored the same dimensionality and loss curves plateaued early, the results suggest that reconstruction-based evaluation and early stopping during training may be sufficient in practice. Although chambers generated by two separate heads can in principle intersect without an intersection penalty, in practice no intersections were observed during reconstruction or generation at dimension 64 when ablating the penalty. However, faster RA representation convergence when including the penalty suggests it benefits training efficiency.

Although in general the model captured relevant shape variability and generated anatomically plausible meshes, as indicated by low distances in the specificity evaluation, geometric accuracy of fine, complex structures such as LAA remains limited. However, visually, the level of detail aligns with that achieved in the existing bi-atrial SSM and is thus expected to be sufficient for electrophysiological simulations [16]. Furthermore, generation coverage was relatively low and smaller than training set coverage, which may reflect either failure to capture the full anatomical variability in the data, or insufficient diversity in the training data. In general, LA and RA showed similar performance but consistently larger distances were found for the LA, which can be attributed to greater topological complexity in the LA (PVs, LAA) and comparatively lower RA data fidelity in the training set. Consequently, this suggests that LA quality should be prioritized for DL model selection.

As this study focused on latent space dimensionality, no hyperparameter optimization was performed, leaving room for further improvement. Importantly, the dataset introduced additional constraints, particularly as automatic pre-

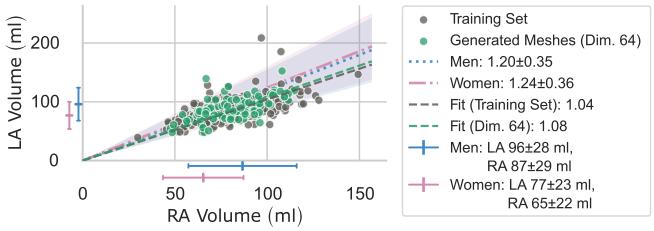

Fig. 5: LA vs. RA volume ratios for 100 generated meshes (dimension 64, epoch 1000) and 244 training set meshes. Clinically observed atrial volumes and LA/RA volume ratios for healthy male and females [19] are depicted as mean ± std after conversion from area-length to volumetric measurements [1].

processing failed for approximately half of the samples; this could, however, be overcome through semi-automatic or improved automatic preprocessing. Furthermore, RA representation is constrained by the absence of the venae cavae and coronary sinus in the segmentation. Future work should extend this framework to larger cohorts with greater anatomical variability, such as atypical PV configurations, improved RA quality, and introduce conditioning on clinical metadata [24]. Structure-specific loss weighting (e.g., for the LAA, PVs, RA/LA balance) and early stopping strategies could further improve the model and reduce computational cost. Ultimately, the model could support synthetic virtual patient cohort creation for subsequent *in silico* studies, particularly once extended to clinical metadata. Such cohorts could be made openly available, helping overcome the scarcity of large bi-atrial datasets, and providing a resource for population-based shape studies of atrial diseases.

Acknowledgments. This work was supported by the Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) – Project number 529821741.

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

References

1. Alos, B., Oresve, G., Beard, E., Beaufort, C., Bergere, V., Roxburgh, T., et al.: Normal reference intervals for left atrial volume and cardiac dimensions according to age and sex assessed by two different methods using cardiac computed tomography angiography. European Heart Journal - Imaging Methods and Practice **3**(4), qyaf125 (2025)
2. Beetz, M., Banerjee, A., Grau, V.: Generating Subpopulation-Specific Biventricular Anatomy Models Using Conditional Point Cloud Variational Autoencoders. In: International Workshop on Statistical Atlases and Computational Models of the Heart, vol. 13131, pp. 75–83. Springer International Publishing, Cham (2022)
3. Cresti, A., Camara, O.: Left Atrial Thrombus—Are All Atria and Appendages Equal? Interventional Cardiology Clinics **11**(2), 121–134 (Apr 2022)

4. Dasí, A., Nagel, C., Pope, M.T.B., Wijesurendra, R.S., Betts, T.R., Sachetto, R., et al.: *In Silico* TRials guide optimal stratification of ATrIal FIbrillation patients to Catheter Ablation and pharmacological medicaTION: the i-STRATIFICATION study. Europace **26**(6), euae150 (Jun 2024)
5. Dou, H., Virtanen, S., Ravikumar, N., Frangi, A.F.: A Generative Shape Compositional Framework to Synthesize Populations of Virtual Chimeras. IEEE Transactions on Neural Networks and Learning Systems **36**(3), 4750–4764 (Mar 2025)
6. Fedorov, A., Beichel, R., Kalpathy-Cramer, J., Finet, J., Fillion-Robin, J.C., Pujol, S., et al.: 3d slicer as an image computing platform for the quantitative imaging network. Magnetic resonance imaging **30**(9), 1323–1341 (2012)
7. Gillette, K., Gsell, M.A.F., Nagel, C., Bender, J., Winkler, B., Williams, S.E., et al.: MedalCare-XL: 16,900 healthy and pathological synthetic 12 lead ECGs from electrophysiological simulations. Scientific Data **10**(1), 531 (Aug 2023)
8. Goetz, C., Eichenlaub, M., Schmidt, K., Wiedmann, F., Rubio, E.I., Díaz, P.M., et al.: DIVAID: Consistent division of atrial geometries from multimodal imaging according to the EHRA/EACVI 15-segment bi-atrial model (Apr 2026)
9. Hansen, B., Pedersen, J., Kofoed, K.F., Camara, O., Paulsen, R.R., Sørensen, K.: A public cardiac CT dataset featuring the left atrial appendage. In: Statistical Atlases and Computational Models of the Heart. Regular and CMRxRecon Challenge Papers. pp. 171–180. Springer Nature Switzerland, Cham (2026)
10. Jia, S., Nivet, H., Harrison, J., Pennec, X., Camaioni, C., Jaïs, P., et al.: Left atrial shape is independent predictor of arrhythmia recurrence after catheter ablation for atrial fibrillation: a shape statistics study. Heart Rhythm O2 **2**(6), 622–632 (2021)
11. Kong, F., Stocker, S., Choi, P.S., Ma, M., Ennis, D.B., Marsden, A.L.: SDF4CHD: Generative modeling of cardiac anatomies with congenital heart defects. Medical Image Analysis **97**, 103293 (Oct 2024)
12. Martínez Díaz, P., Sánchez, J., Fitzen, N., Ravens, U., Dössel, O., Loewe, A.: The right atrium affects in silico arrhythmia vulnerability in both atria. Heart Rhythm **21**(6), 799–805 (Jun 2024)
13. Nagel, C., Luongo, G., Azzolin, L., Schuler, S., Dössel, O., Loewe, A.: Non-Invasive and Quantitative Estimation of Left Atrial Fibrosis Based on P Waves of the 12-Lead ECG—A Large-Scale Computational Study Covering Anatomical Variability. Journal of Clinical Medicine **10**(8), 1797 (Apr 2021)
14. Nagel, C., Sanchez, J., Azzolin, L., Zheng, T., Schuler, S., Dössel, O., et al.: A Bi-atrial Statistical Shape Model and 100 Volumetric Anatomical Models of the Atria (Jun 2021)
15. Nagel, C., Schaufelberger, M., Dössel, O., Loewe, A.: A Bi-atrial Statistical Shape Model as a Basis to Classify Left Atrial Enlargement from Simulated and Clinical 12-Lead ECGs. In: International Workshop on Statistical Atlases and Computational Models of the Heart, vol. 13131, pp. 38–47. Springer International Publishing, Cham (2022)
16. Nagel, C., Schuler, S., Dössel, O., Loewe, A.: A bi-atrial statistical shape model for large-scale in silico studies of human atria: Model development and application to ECG simulations. Medical Image Analysis **74**, 102210 (Dec 2021)
17. Niederer, S.A., Aboelkassem, Y., Cantwell, C.D., Corrado, C., Coveney, S., Cherry, E.M., et al.: Creation and application of virtual patient cohorts of heart models. Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences **378**(2173), 20190558 (Jun 2020)
18. Park, J.J., Florence, P., Straub, J., Newcombe, R., Lovegrove, S.: DeepSDF: Learning Continuous Signed Distance Functions for Shape Representation. In: Proceed-

ings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 165–174 (2019)
19. Qazi, S., Gona, P.N., Arora, G., Salton, C., O'Donnell, C.J., Manning, W.J., et al.: Relationship of Right and Left Atrial Volumes by Sex and Age Group in Healthy Adults. Circulation **140**(Suppl_1), A16196–A16196 (2019)
20. Qiao, M., McGurk, K.A., Wang, S., Matthews, P.M., O'Regan, D.P., Bai, W.: A personalized time-resolved 3D mesh generative model for unveiling normal heart dynamics. Nature Machine Intelligence **7**(5), 800–811 (May 2025)
21. Sharp, A.J., Betts, T.R., Banerjee, A.: Leveraging 3D Atrial Geometry for the Evaluation of Atrial Fibrillation: A Comprehensive Review. Journal of Clinical Medicine **13**(15), 4442 (Jul 2024)
22. Sharp, A.J., Betts, T.R., Banerjee, A.: Statistical shape modeling in cardiovascular disease: a narrative review. Journal of the Royal Society Interface **23**(235), 20250785 (Feb 2026)
23. Sørensen, K., Camara, O., Backer, O.d., Kofoed, K., Paulsen, R.: NUDF: Neural Unsigned Distance Fields for high resolution 3D medical image segmentation. In: 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI). pp. 1–5. IEEE (2022)
24. Sørensen, K., Diez, P., Margeta, J., El Youssef, Y., Pham, M., Pedersen, J.J., et al.: Spatio-Temporal Neural Distance Fields for Conditional Generative Modeling of the Heart. In: Medical Image Computing and Computer Assisted Intervention – MICCAI 2024. vol. 15003, pp. 422–432. Springer Nature Switzerland, Cham (2024)
25. Van Gelder, I.C., Rienstra, M., Bunting, K.V., Casado-Arroyo, R., Caso, V., Crijns, H.J., et al.: 2024 ESC Guidelines for the management of atrial fibrillation developed in collaboration with the European Association for Cardio-Thoracic Surgery (EACTS). European Heart Journal **45**(36), 3314–3414 (Sep 2024)
26. Verhülsdonk, J., Grandits, T., Costabal, F.S., Beiert, T., Pezzuto, S., Effland, A.: Uncertainty Quantification for Cardiac Shape Reconstruction with Deep Signed Distance Functions via MCMC methods (May 2026)
27. Zhou, Z., Zakeri, A., Dou, H., Xue, Y., MacRaild, M., Huang, J., et al.: Synthetic anatomy: Deep learning models for virtual population generation–a review. medRxiv pp. 2025–10 (2025)