

MOX-Report No. 25/2026

**Shape-informed cardiac mechanics surrogates in data-scarce regimes
via geometric encoding and generative augmentation**

Carrara, D.; Hirschvogel, M.; Bonizzoni, F.; Pagani, S.; Pezzuto, S.;
Regazzoni, F.

MOX, Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

mox-dmat@polimi.it

<https://mox.polimi.it>

Shape-informed cardiac mechanics surrogates in data-scarce regimes via geometric encoding and generative augmentation

Davide Carrara^{*1}, Marc Hirschvogel^{*1,2}, Francesca Bonizzoni¹, Stefano Pagani¹, Simone Pezzuto³, and Francesco Regazzoni¹

¹MOX Laboratory, Department of Mathematics, Politecnico di Milano, Milano, Italy

²Division 2.2 Process Simulation, Bundesanstalt für Materialforschung und -prüfung (BAM), Berlin, Germany

³Department of Mathematics, University of Trento, Trento, Italy

Abstract

High-fidelity computational models of cardiac mechanics provide mechanistic insight into the heart function but are computationally prohibitive for routine clinical use. Surrogate models can accelerate simulations, but generalization across diverse anatomies is challenging, particularly in data-scarce settings. We propose a two-step framework that decouples geometric representation from learning the physics response, to enable shape-informed surrogate modeling under data-scarce conditions. First, a shape model learns a compact latent representation of left ventricular geometries. The learned latent space effectively encodes anatomies and enables synthetic geometries generation for data augmentation. Second, a neural field-based surrogate model, conditioned on this geometric encoding, is trained to predict ventricular displacement under external loading. The proposed architecture performs positional encoding by using universal ventricular coordinates, which improves generalization across diverse anatomies. Geometric variability is encoded using two alternative strategies, which are systematically compared: a PCA-based approach suitable for working with point cloud representations of geometries, and a DeepSDF-based implicit neural representation learned directly from point clouds. Overall, our results, obtained on idealized and patient-specific datasets, show that the proposed approaches allow for accurate predictions and generalization to unseen geometries, and robustness to noisy or sparsely sampled inputs.

1 Introduction

Computational models of the heart and the cardiovascular system have emerged as powerful tools to mechanistically investigate electrochemical, mechanical, and biological processes across multiple spatial and temporal scales [1, 2]. However, high-fidelity models based on partial differential equations (PDEs) typically require computationally intensive discretization and solution techniques, often representing an obstacle to routine clinical translation and application.

To partially circumvent the computational burden without compromising accuracy, model order reduction techniques are often integrated into simulation frameworks. These reduce the problem to a lower-dimensional representation, either by exploiting snapshot information from high-fidelity models or imaging data [3, 4, 5, 6, 7], or through lumped model approaches [8, 9]. However, real-time simulation through reduced-order models (ROMs) often remains computationally prohibitive, particularly in nonlinear regimes characterized by slow decay Kolmogorov n -width and a corresponding high-number of basis functions required to accurately represent the solution [10, 11]. This has motivated the development of data-driven surrogate models, which leverage dimensionality reduction techniques better suited for nonlinear solution manifolds [12]. The resulting surrogate can be evaluated typically at a fraction of the computational cost compared to projection-based ROMs, and their evaluation does not require a computational mesh [13, 14, 15, 16, 17].

*These authors equally contributed to this work.

In this context, however, dealing with geometric variability remains an open challenge. Surrogate models trained on a single geometry do not generalize across cohorts of differently-shaped anatomies, making it impractical to train one model per patient. At the same time, training a single model on a cohort of geometries is non-trivial in the absence of point-to-point correspondence between different anatomies, motivating the need for a geometric encoding of shape variability. Universal Solution Manifold Networks (USMNs), which we employ as starting point for the present work, in their base version achieve geometric encoding by means of landmark definitions [18, 19], but do not straightforwardly generalize to complex geometries, where identifying landmarks is non-trivial. Other approaches use diffeomorphic mapping operators to map geometry variations onto a common reference space either to enhance deep learning based ROMs [20] or to extract principal components of shape variability [21]. Applied to electrophysiology models of the left ventricle, the latter leverages universal ventricular coordinates (UVCs) [22] to establish a mapping between patient-specific and reference geometries. However, both UVC fields and meshed representations are required, restricting applicability to structured data sources. An alternative paradigm for encoding geometric variability is offered by representation learning techniques, which aim to describe complex data through a set of lower-dimensional features extracted in an automatic way [23]. In the context of shape modeling, this corresponds to discovering a latent space, where each point represents a geometry in the original three-dimensional domain, together with a mapping that allows reconstructing or querying the shape from this compact representation. Deep autoencoder architectures [24, 25] implement this idea by jointly training an encoder network that maps shapes to latent codes and a decoder network that reconstructs shapes from these codes. Related approaches, known as autoencoders, instead associate each shape with a learnable latent code optimized jointly with the decoder, removing the need for an explicit encoder and simplifying the overall architecture [26]. DeepSDF [27] introduces an autoencoder architecture for the simultaneous reconstruction of multiple geometries: each shape is associated with a low-dimensional latent code that conditions the network input and is optimized jointly with the network weights. The framework is flexible with respect to input data type, operating directly on point clouds, voxels, or meshes, and an explicit mesh representation could be extracted via the Marching Cubes algorithm [28]. Smoothness and regularity of the latent space can be further enforced by controlling the Lipschitz constant of the network [29], promoting stable interpolation between geometries. The combination of DeepSDF with Lipschitz regularization has been successfully applied to reconstruct cardiac geometries in [30].

However, even when a geometric encoding method is available, learning the map from high-dimensional shape variability to the corresponding PDE solution remains a highly complex regression task. Approaches based on graph neural networks [31, 32] or attention mechanisms [33] typically require large training datasets to achieve robust generalization, and performance on unseen geometries can degrade significantly, when only limited training samples are available [34]. This is an important limitation in biomedical applications, where open-access data is often scarce and acquisition of new high-fidelity simulations is computationally expensive.

Motivated by these considerations, in this work, we present a method for building shape-informed surrogate models for cardiac mechanics from scattered point cloud data associated with a cohort of different patient anatomies. We validate this framework on two datasets: an synthetic cohort of idealized geometries and two cohorts of real left ventricular anatomies from publicly available sources [35, 36]. The proposed pipeline consists of two steps. First a shape model is fitted to the geometric data, automatically learning a compact representation of each anatomy. When working with real data, the shape model is also used to generate new realistic geometries for data augmentation. Second, a shape-informed surrogate model is trained for predicting cardiac displacement by conditioning the network on the learned geometric encoding with the augmented dataset used to improve generalization. The proposed framework, by decoupling the geometric complexity from the regression task, allows in this manner for data augmentation to compensate for the scarcity of biomedical training data.

For the shape modeling step, we propose two complementary approaches. The first, named PCA-SM (PCA-based Shape Model), is based on Principal Component Analysis (PCA). Point-to-point correspondence between training geometries is established via UVCs, enabling PCA of the shape deformations. At inference time, the latent representation of a new geometry is obtained by finding the low-dimensional code, whose projection onto the PCA modes best approximates the ground truth. The second, named SDF-SM (SDF-based Shape Model), is a neural network-based shape model that learns an implicit representation of each

geometry by approximating its signed distance function (SDF), an approach commonly used in computer vision tasks [37]. This technique is extended to multiple geometries by adopting the DeepSDF autoencoder architecture [27]: a decoder $\mathcal{NN}_{\text{sdf}}$ is conditioned on a shape code $\mathbf{z} \in \mathcal{Z}$ and a point \mathbf{x} sampled from the geometry, with optional additional geometric parameters $\boldsymbol{\mu}$; the output is the SDF value at \mathbf{x} for the geometry associated with \mathbf{z} and $\boldsymbol{\mu}$. The latent distribution \mathcal{Z} and the decoder weights are optimized jointly during training, combining a reconstruction loss, a prior on \mathcal{Z} , and a Lipschitz regularization term promoting smoothness of the latent space [29]. At inference time, the decoder weights are frozen and \mathcal{Z} acts as a posterior distribution, enabling robust latent code optimization even for noisy or sparsely sampled point clouds. To augment the dataset of real anatomies, SDF-SM is employed as a generative model: new realistic geometries are obtained by sampling the posterior distribution \mathcal{Z} and conditioning the decoder on the sampled latent codes. The second step consists in training a shape-informed surrogate model for cardiac displacement conditioned on the geometric encoding learned in the first step. The surrogate model is implemented as a neural network $\mathcal{NN}_{\text{phys}}$, optimized to predict the displacement field from the geometric encoding and the physical parameters of interest. We consider two variants depending on the source of the geometric encoding: PCA-USMNet, which uses the latent representation from PCA-SM, and SDF-USMNet, which uses the latent code learned by SDF-SM. Both variants are applied to the problem of predicting the displacement field for a finite strain mechanics problem of the left ventricle under diastolic loading conditions.

2 Results

We numerically test the proposed pipelines on two datasets: an idealized dataset and real patient anatomies. We first present the results on the shape encoding and reconstruction task (see Section 2.1) and then the performances on the surrogate modeling task (see Section 2.3). The idealized dataset comprises 512 prolate ellipsoids parametrized by the triplet (ℓ, d, w) , corresponding to long-axis length, diameter, and wall thickness, respectively, sampled on a uniform $8 \times 8 \times 8$ grid. This controlled dataset allows us to assess the shape model’s ability to learn a descriptive latent space in a setting where the true parametrization is known analytically. In Section 2.1.1, we analyze the correspondence between the learned low-dimensional representation and the ground-truth parameters, while Section 2.3.1 compares surrogate model performance when conditioned on learned latent codes versus true geometric parameters. Then, we consider a dataset of real patient anatomies consisting of 44 left ventricular geometries obtained from two publicly available cohorts: 24 heart failure patients [35] and 20 healthy patients [36]. The complete data acquisition and preprocessing pipeline is described in Section 4.1. On this dataset, we compare two shape encoding approaches, PCA- and SDF-based (respectively PCA-SM and SDF-SM), for geometric representation (Section 2.1.2) and for conditioning the surrogate model (PCA-USMNet and SDF-USMNet, respectively) (Section 2.3.2). The use of the SDF-SM for generating synthetic geometries for data augmentation is discussed in Section 2.2.

2.1 Shape reconstruction

The quality of reconstructed geometries is assessed using two classes of metrics. Chamfer Distance (CD) and Normalized Chamfer Distance (CD_{norm}) measure the discrepancy between the reconstructed surface and the ground truth, considering the original and normalized geometry sizes, respectively. Intersection over Union (IoU) and Dice Coefficient (Dice) provide a volumetric measure of the overlap between the reconstructed shape and the ground truth, both in the $[0, 1]$ range, where higher values indicate better reconstruction. For SDF-SM we also report the Mean Level Estimate (MLE) defined as the mean absolute SDF prediction on ground-truth surface points, where the correct value is zero by definition. A formal definition of all metrics is provided in the Methods (Section 4.6).

2.1.1 Idealized geometries

The reconstruction performance of SDF-SM is first evaluated on an idealized dataset. Since the parametrization is known, we also consider the case in which the network is conditioned directly on the true parameters rather than learning the latent space, and only the decoder weights are optimized. This serves as a best-case

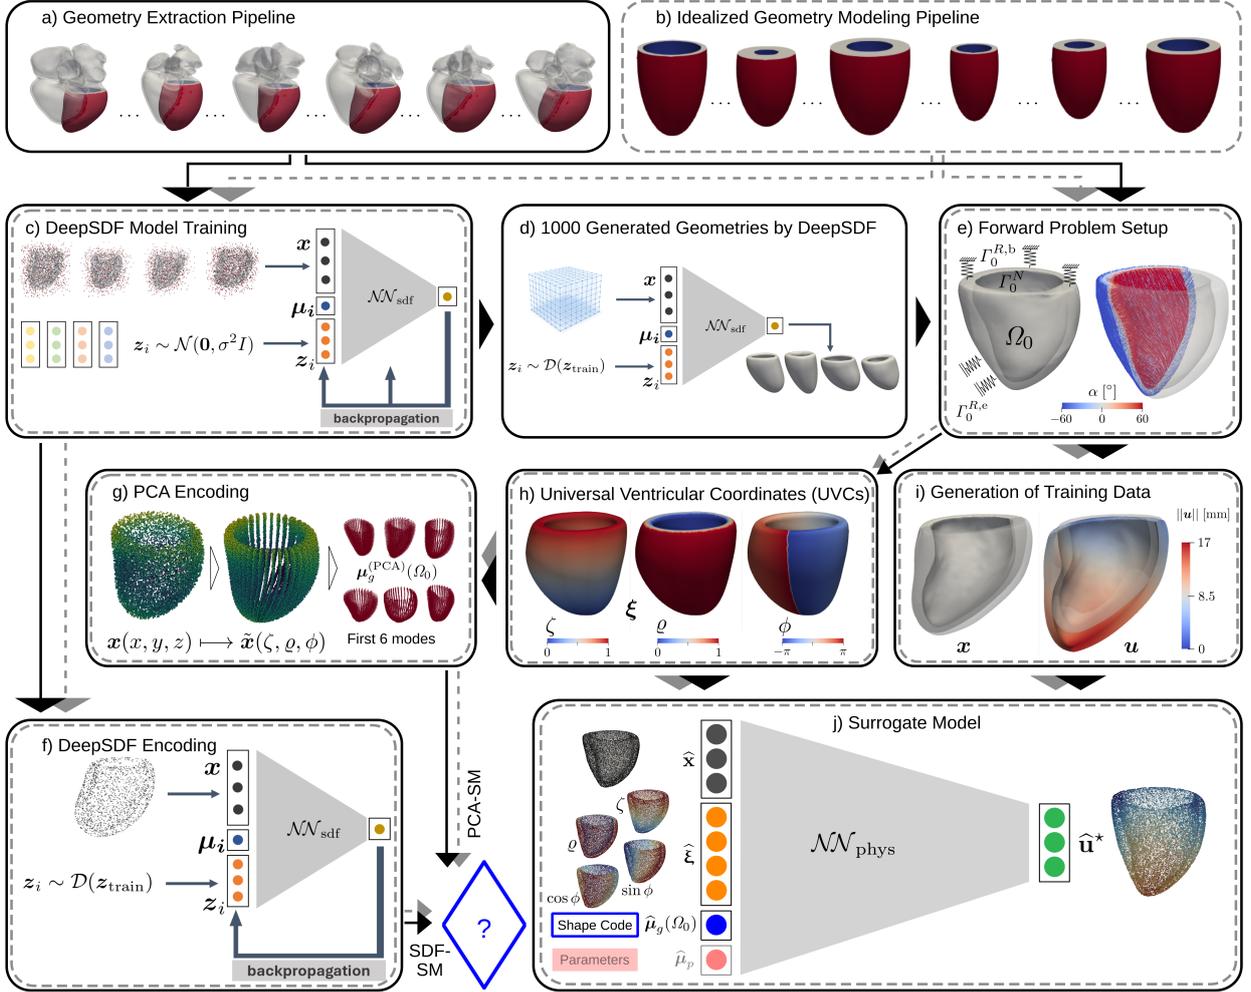


Figure 1: Data processing, modeling, and training pipelines. Solid black lines/boxes show the pipeline of patient-specific left ventricular (LV) models; dashed gray lines/boxes that of the idealized LV models. They differ in that the idealized model cohort is not augmented by synthetically generated geometries. a) LV geometry extraction from 4-chamber patient-specific heart geometries, retrieving 44 LV geometries. b) Idealized LV models generated from varying long axis, diameter, and wall thickness ℓ , d , and w , respectively, retrieving 512 models. c) SDF-SM training: learning of signed-distance function to represent the geometry boundary. d) Geometry generation by SDF-SM. e) Mechanical forward problem setup. Left: Reference geometry Ω_0 , with Robin (spring) boundaries on the base ($\Gamma_0^{R,b}$) and the epicardium ($\Gamma_0^{R,e}$) as well as a Neumann boundary at the endocardium (Γ_0^N). Right: Fiber field \mathbf{f}_0 ; color indicates the fiber angle α with respect to the circumference. f) SDF-SM encoding. g) PCA-SM encoding. h) Computation of universal ventricular coordinates (UVCs). i) Simulation of the forward model to generate displacement training data. Left: undeformed geometry. Right: deformed model with displacement field \mathbf{u} , color indicating its magnitude. j) Surrogate model, informed by cartesian coordinates \mathbf{x} , UVCs ξ , a shape code μ_g (either from PCA-SM, SDF-SM, or an analytic descriptor of shape features, e.g. features from the idealized geometries), and possibly some parameters μ_p .

comparison to decouple the representation capability of the decoder and the descriptive power of the latent space. The input to the network includes also the scaling factor μ_i associated to each shape Ω_i . Since this scaling factor relates strongly to ℓ (see Methods Section 4.3), we consider it as the third component of the latent code, without learning it.

Table 1 presents the reconstruction metrics for both approaches. The results indicate no significant difference in reconstruction quality between the learned latent codes and the original parametrization, suggesting that the network successfully captures the essential geometric information in its learned representation.

Table 1: Reconstruction metrics for SDF-SM trained with learned latent codes versus original parametric features on idealized geometries.

Code	Split	CD [mm]	CD _{norm}	IoU	Dice	MLE
Learned	Test	1.276 ± 0.282	0.0238 ± 0.0062	0.953 ± 0.027	0.976 ± 0.014	1.367 ± 0.587
	Train	1.232 ± 0.275	0.0226 ± 0.0059	0.954 ± 0.022	0.976 ± 0.012	1.257 ± 0.614
Original	Test	1.124 ± 0.290	0.0209 ± 0.0040	0.950 ± 0.036	0.974 ± 0.020	1.381 ± 0.596
	Train	1.100 ± 0.133	0.0203 ± 0.0036	0.954 ± 0.026	0.976 ± 0.014	1.261 ± 0.614

Latent space structure and interpretability. To investigate the relationship between learned latent codes and the original geometric parameters, we analyze the structure of the learned representation. Figure 2a displays the component values of both learned and original codes across all 512 geometries. The $8 \times 8 \times 8$ grid structure used to generate the original parameter triplets is clearly visible in both representations. We observe that μ_i , here indicated as the third component of the latent code, is strongly correlated with the original feature ℓ . Figure 2b shows the pairwise Pearson correlation matrices. While the learned codes capture relationships present in the original features, they are not independent. The three-dimensional distribution of the learned codes is shown in Figure 2c.

To quantify the relationship between learned codes and original features, we fit linear regression models to predict each original parameter from the learned latent codes. The models are trained on codes from the training set and evaluated on the test set. The linear regression models achieve high R^2 values both on the test set ($R_\ell^2 = 0.988$, $R_d^2 = 0.941$, $R_w^2 = 0.965$) and the train set ($R_\ell^2 = 0.995$, $R_d^2 = 0.929$, $R_w^2 = 0.957$), indicating that the original geometric parameters are recoverable from the latent representation through a linear mapping.

2.1.2 Real Geometries

Reconstruction results The model achieves a reconstruction error below 1 mm in terms of CD both on the training and test set. We also observe that there is no significant difference in CD_{norm} between the two cohorts of patients, as the size difference is effectively factored out in the preprocessing stage in Alg. 1. In terms of intersection over union, the worst performance on the test set is $\text{IoU}_{\min} = 0.797$. The complete reconstruction performances are reported in Table 2 and in Figure 3.

Table 2: Reconstruction metrics for SDF-SM by cohort and data split (mean ± std).

Cohort	Split	CD [mm]	CD _{norm}	IoU	Dice
Healthy	Test	0.994 ± 0.078	0.0205 ± 0.0026	0.826 ± 0.026	0.904 ± 0.015
	Train	0.846 ± 0.069	0.0165 ± 0.0008	0.899 ± 0.016	0.947 ± 0.008
Heart Failure	Test	0.938 ± 0.145	0.0200 ± 0.0033	0.864 ± 0.038	0.926 ± 0.022
	Train	0.763 ± 0.215	0.0170 ± 0.0033	0.912 ± 0.051	0.953 ± 0.031

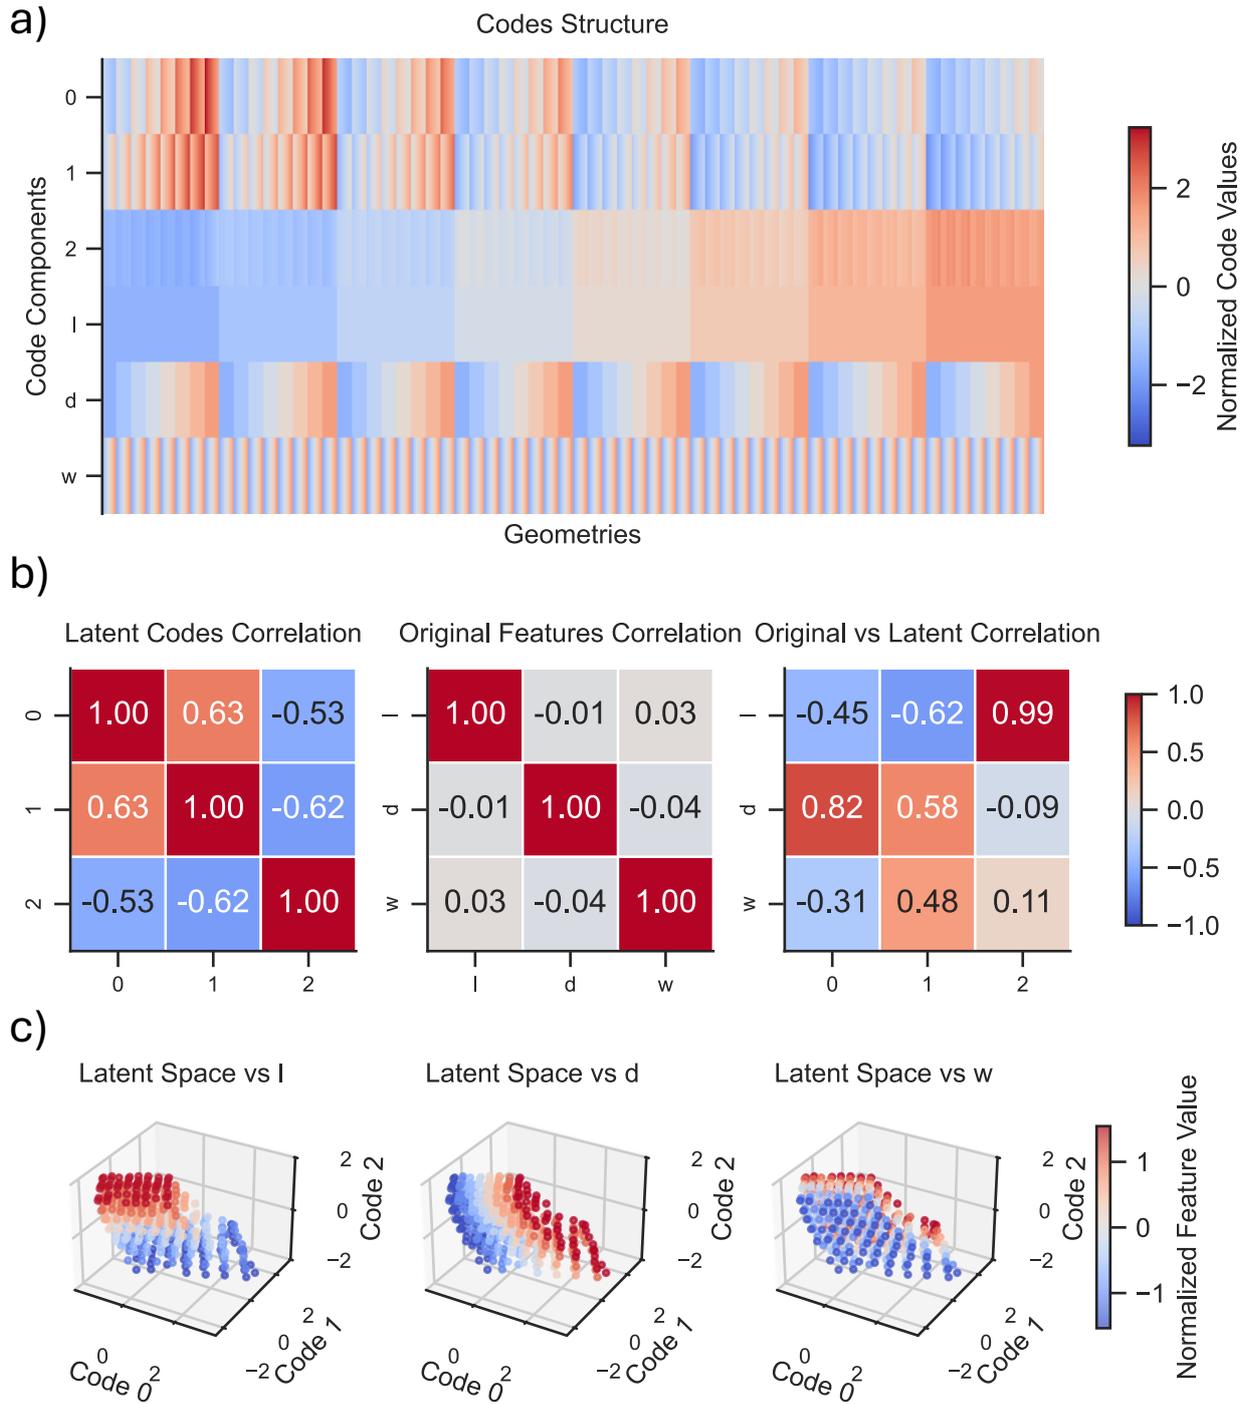


Figure 2: Visualizations of the learned latent space for the idealized geometries dataset. a) Comparison between reconstructed latent codes and original geometric features for the idealized geometries. Each column corresponds to an individual geometry, with colors representing the z-normalized value of the associated original feature to ensure homogeneous color scaling across dimensions. b) Pairwise Pearson correlation matrices between the learned latent codes and the original geometric features for the idealized geometries. c) Three-dimensional visualization of the learned latent space for the idealized geometries. Points represent individual samples in the space of the first three latent codes and are colored according to the z-normalized values of the original geometric features, illustrating how variations in physical parameters are embedded in the latent representation.

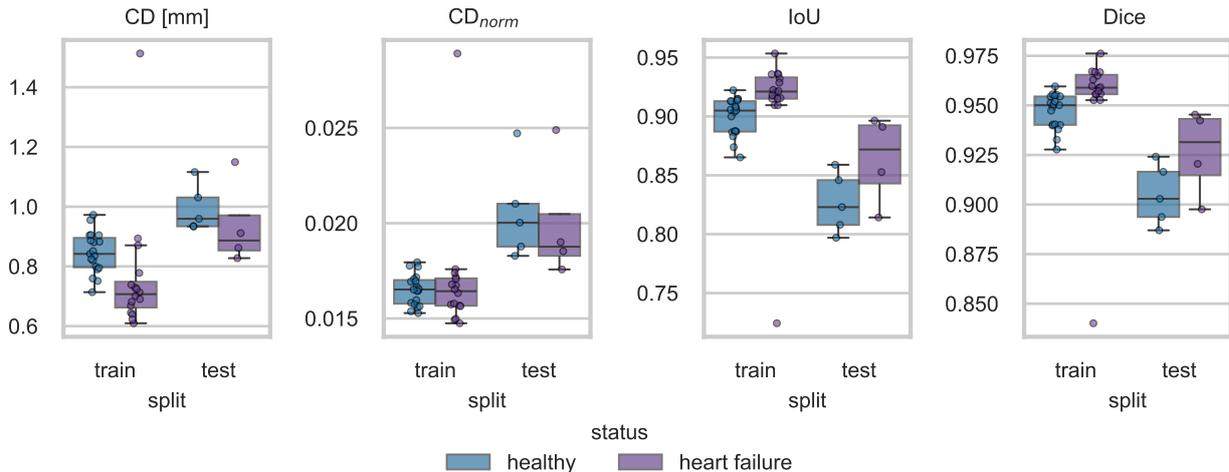


Figure 3: Performance metrics for SDF-SM reconstructed geometries.

Figure 4b compares the section of the reconstructed test geometries with the original ones, while Figure 4a provides a three-dimensional visualization.

Latent space analysis Figure 5 shows the correlation matrices and distributions among the empirical latent space components and among the latent representations of the 44 patients anatomies.

We observe that the empirical latent space shows weak pairwise correlation (range: $\rho_{\min} = -0.501$ to $\rho_{\max} = 0.3079$), indicating that the 16-dimensional latent representation captures largely independent geometric features. In contrast, more significant correlation patterns emerge among latent representations of patient anatomies, especially within cohorts. Anatomies from the same patient cohort exhibit more pronounced correlations ($\rho_{\min} = -0.699$ and $\rho_{\max} = 0.808$), which suggests that the latent space clusters geometries depending on some cohort specific anatomical characteristics.

This structure shows that the learned representation successfully captures both inter-patient and cohort-specific variability, which can be beneficial for downstream surrogate model conditioning based on anatomical features.

PCA reconstruction Figure 19 in Appendix A shows the average shape obtained by projecting each geometry onto a uniform UVC grid, together with the deviations generated by adding the first 6 modes obtained through PCA.

Table 3: Reconstruction metrics for PCA-SM by cohort and data split (mean \pm std).

Cohort	Split	CD [mm]	CD _{norm}
Healthy	Test	0.807 ± 0.063	0.0167 ± 0.0021
	Train	0.567 ± 0.129	0.0111 ± 0.0027
Heart Failure	Test	0.740 ± 0.130	0.0158 ± 0.0030
	Train	0.484 ± 0.080	0.0109 ± 0.0013

Figure 6 and Table 3 report performance metrics for PCA-SM. We observe that in this setting PCA-SM outperforms SDF-SM with respect to CD and CD_{norm}, while volume-level metrics like IoU and Dice are not computed, since PCA-SM does not produce a closed watertight surface.

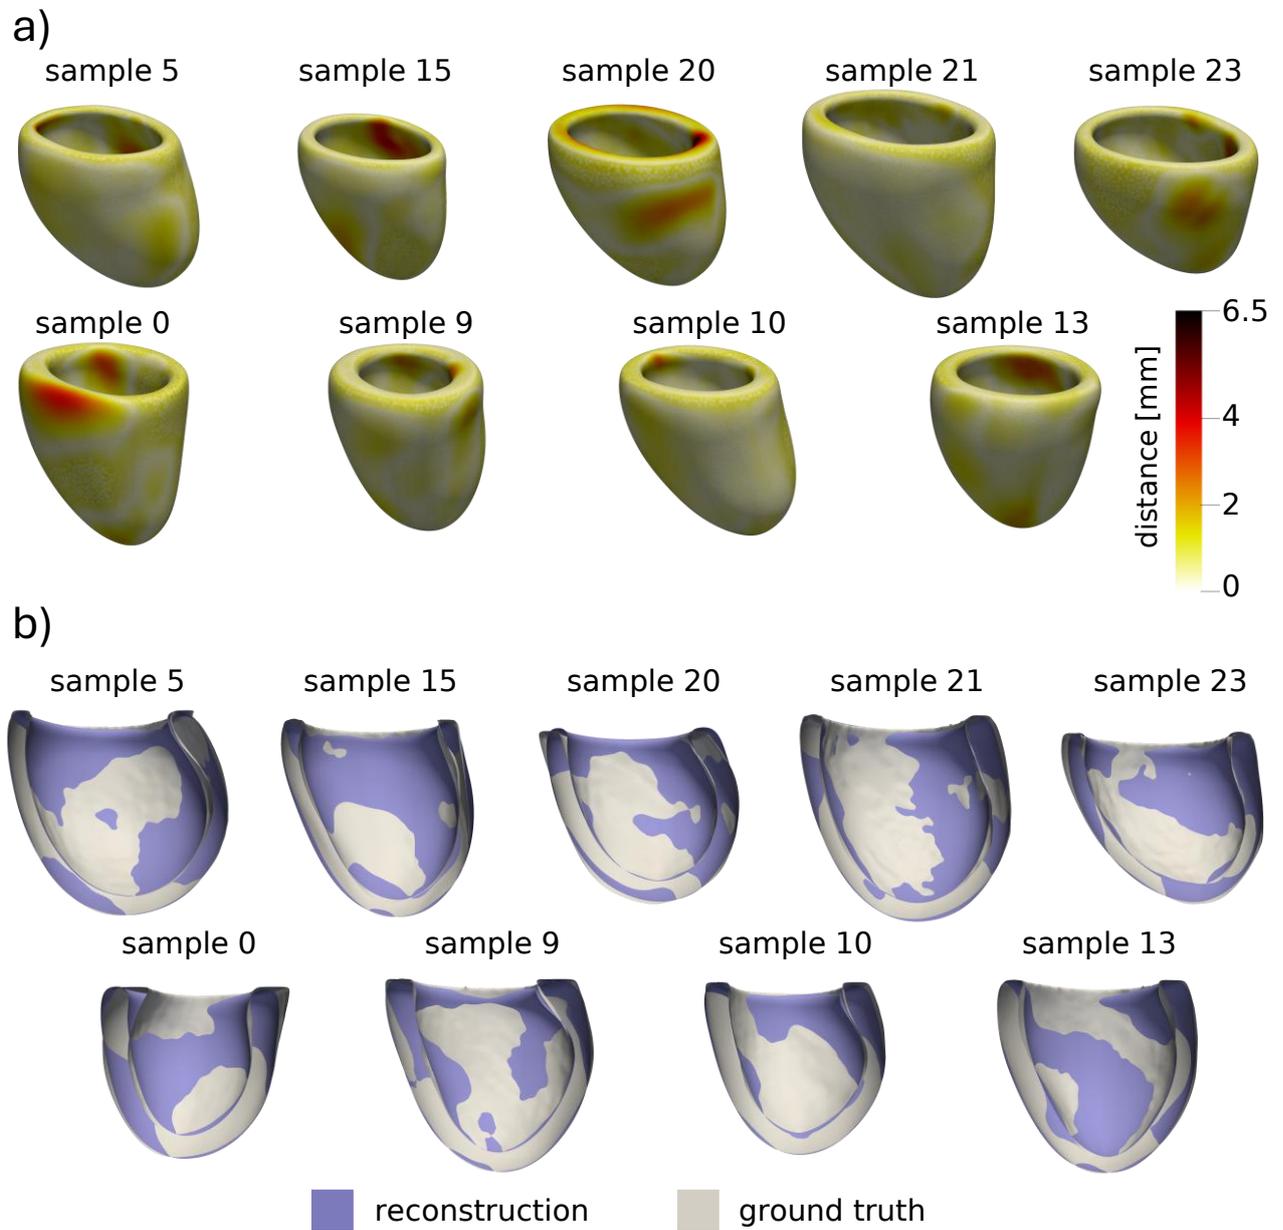


Figure 4: a) Reconstructed meshes colored according to the distance [mm] to the closest point on the target mesh. First row: heart failure patients, second row: healthy patients. b) Reconstructed meshes (in purple) and original ones (in white). First row: heart failure patients; second row: healthy patients.

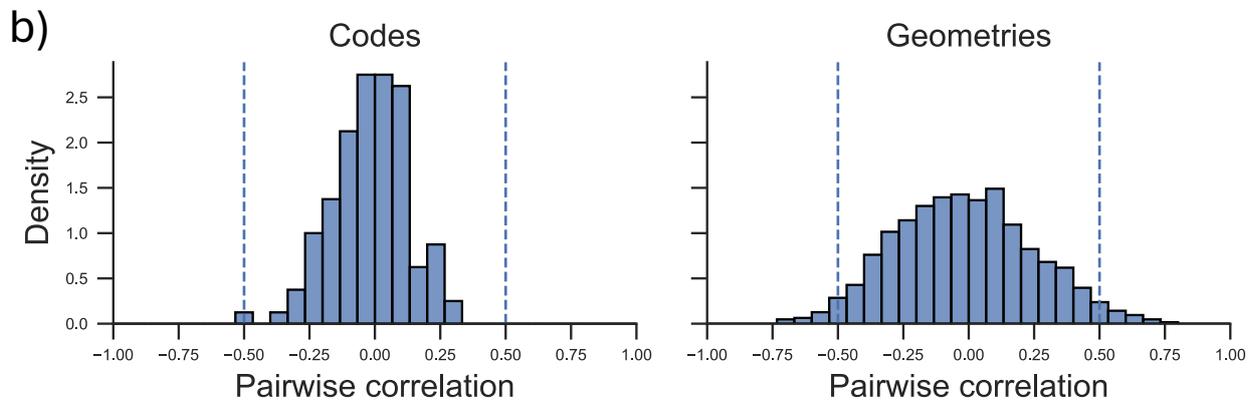
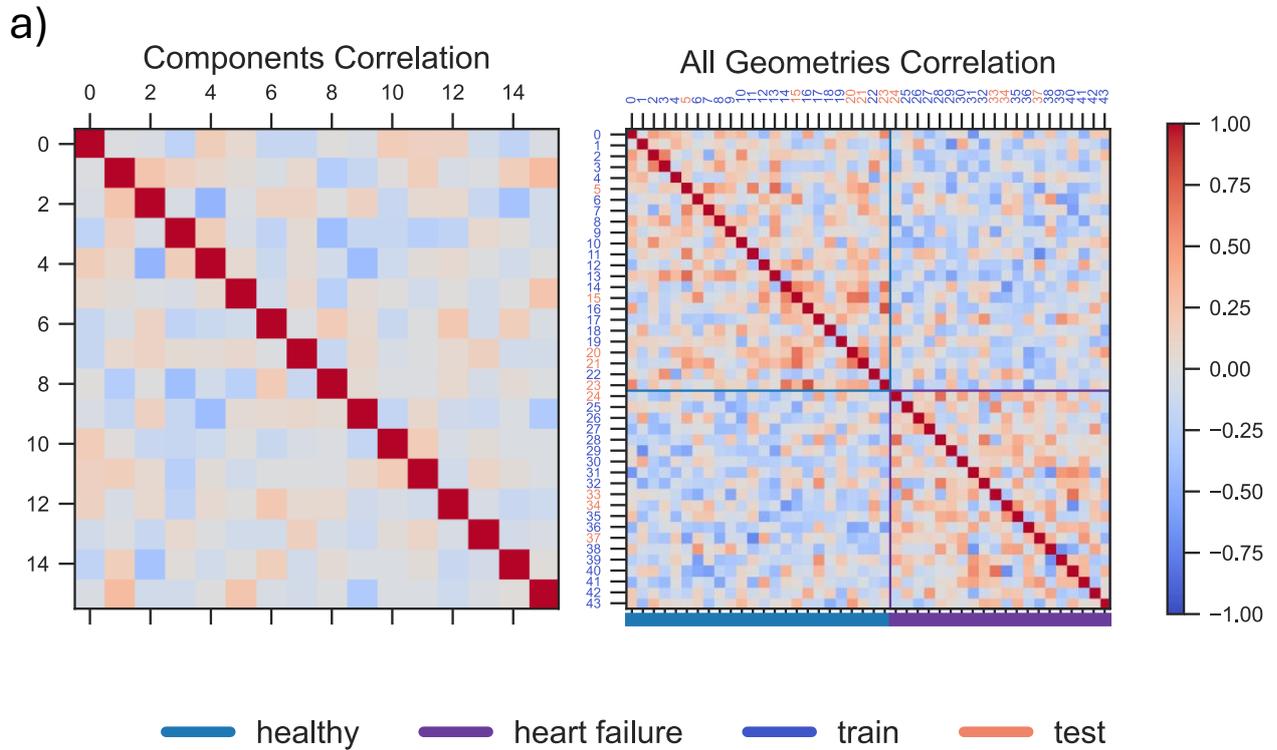


Figure 5: a) Correlation matrices computed after inference. Matrix on the left describes correlation among different components of the latent space. Matrix on the right describes correlation between codes of different geometries. Colored line at the bottom distinguishes the two cohorts of patients, while the color of the index indicates whether the geometry belongs to training or test set. b) Empirical distributions of pairwise correlations respectively between codes and geometries. Histograms show the density of the upper triangular elements of the corresponding correlation matrices (excluding the diagonal), so that each variable pair is counted once. Vertical dashed lines indicate correlation thresholds of ± 0.5 .

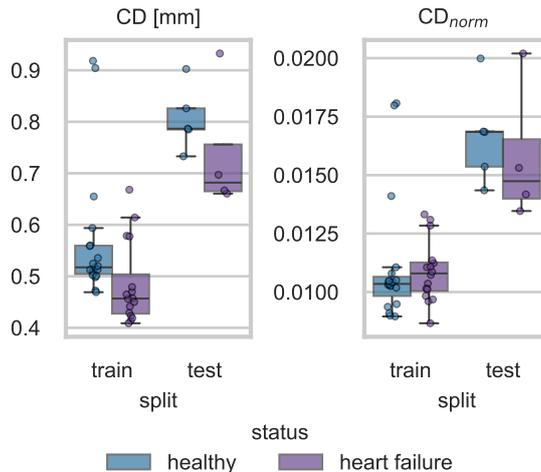


Figure 6: Performance metrics for PCA-SM on reconstructed geometries.

2.1.3 Robustness Evaluation (noise numerosity test and PCA comparison)

The reconstruction performances of the SDF-SM and PCA-SM are evaluated with varying levels of noise and sampling. The sampling value varies between $\{125, 250, 500, 1000, 2000\}$ points, while the noise standard deviation in $\{0.0, 0.0125, 0.025, 0.05, 0.1\}$ is applied after normalizing the mesh, in order to perturb all elements of the cohort with the same relative amount of noise. On average, the corresponding noise levels in the non-normalized space is $\{0.0, 0.6 \text{ mm}, 1.2 \text{ mm}, 2.4 \text{ mm}, 4.8 \text{ mm}\}$.

Figure 7 shows how the reconstruction metrics for SDF-SM vary with geometric perturbation level and sample size. As expected from the probabilistic inference formulation (Section 4.5), reconstruction quality degrades with increasing noise but improves with larger sample sizes. This behavior reflects the adaptive regularization mechanism: the Mahalanobis prior weight w_{post} decreases with sample size, allowing the optimization to rely more heavily on data when more observations are available, and more on the learned prior when data is scarce or noisy. The IoU remains above 0.5 even in the most challenging scenario (125 points with maximum noise level), demonstrating the robustness of the probabilistic latent code optimization.

Figure 8 shows that PCA-SM performance degrades with increasing noise but exhibits minimal improvement when sample size increases. This reflects the difference in the reconstruction approach, where PCA-SM coefficients are found by minimizing CD between the linearly reconstructed point cloud and noisy observation, without explicitly modeling measurement noise and sample size.

For low noise levels ($\sigma \leq 0.025$) or small sample sizes ($N \leq 500$), PCA-SM achieves lower reconstruction error, benefiting from its compact linear parameterization of the shape space learned from the training cohort. On the other hand, SDF-SM provides better performances for high noise and higher sample size regime ($N > 500, \sigma > 0.05$), where the probabilistic inference framework is able to denoise the latent code estimate by giving more importance to the learned posterior distribution.

In order to describe with a single quantity the amount of information carried by the data under different noise and sampling conditions, we define the *effective noise level* as

$$\rho = \frac{\sigma_{\text{geo}}}{\sqrt{N_{\text{geo}}}}, \quad (1)$$

where σ_{geo} is the standard deviation of the Gaussian noise applied to the point coordinates and N_{geo} is the number of sampled points. This scaling naturally arises from the statistical aggregation of independent measurements: while the variance of averaged noise decreases proportionally to $1/N$, the corresponding standard deviation scales as $1/\sqrt{N}$. The parameter ρ thus quantifies the residual noise amplitude per effective degree of information.

Evaluation Metrics Scaling on Test Geometries (Mean \pm Std)

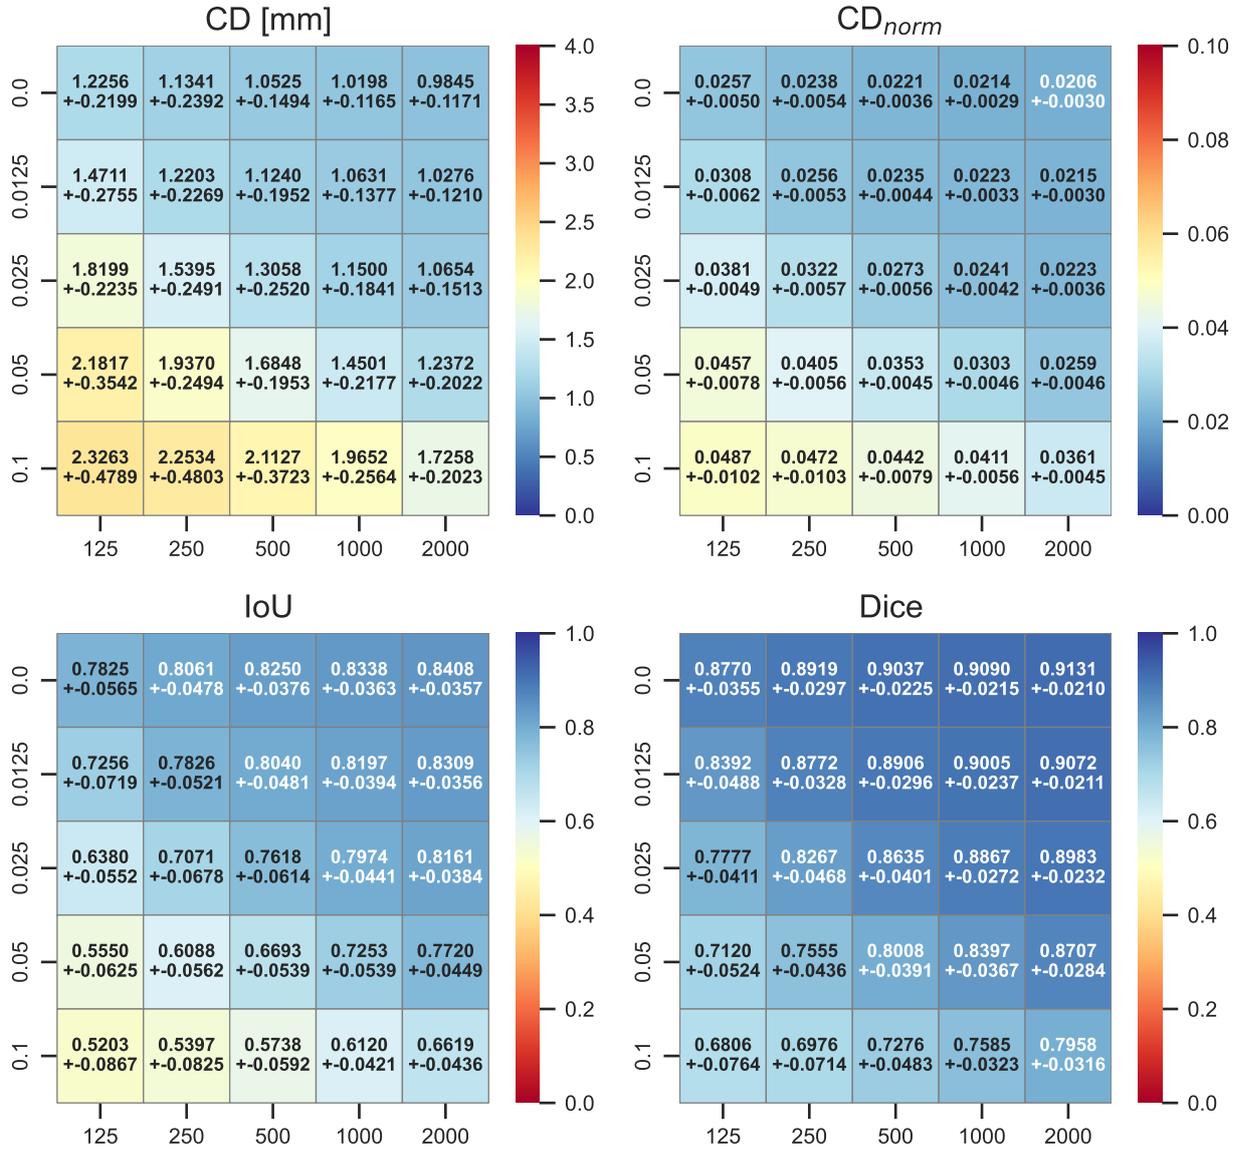


Figure 7: SDF-SM performance on test geometries as a function of sample size and geometric noise level.

Evaluation Metrics Scaling on Test Geometries (Mean \pm Std)

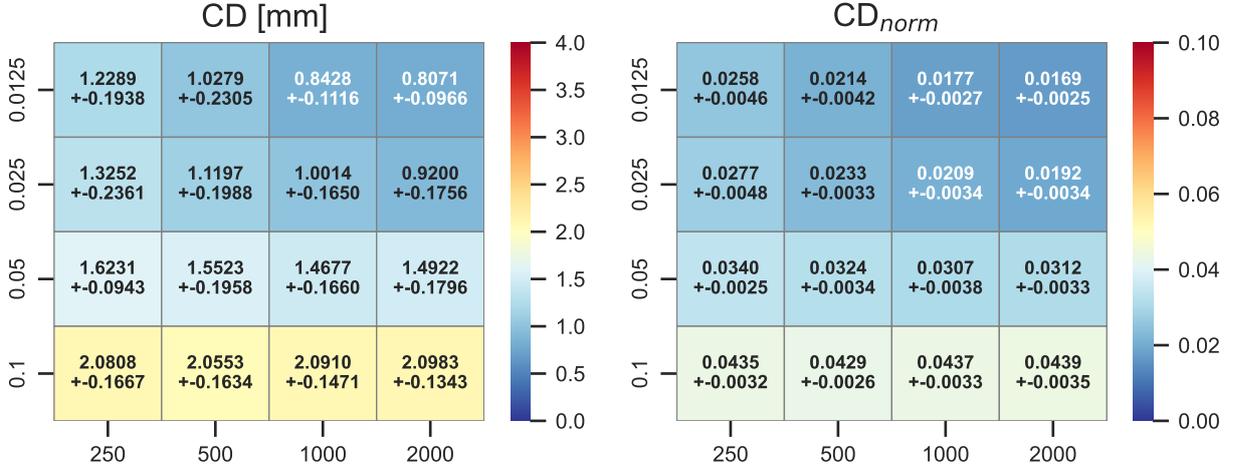


Figure 8: PCA-SM performance on test geometries as a function of sample size and geometric noise level.

The effective noise level can also be interpreted as a measure of label error. Since signed distance functions satisfy the eikonal equation $\|\nabla \text{SDF}\| = 1$, they are 1-Lipschitz continuous: perturbing a point \mathbf{x} by ϵ shifts the SDF value by at most $\|\epsilon\|_2$, as $|\text{SDF}(\mathbf{x}) - \text{SDF}(\mathbf{x} + \epsilon)| \leq \|\epsilon\|_2$. Consequently, geometric noise on the input coordinates with standard deviation σ_{geo} induces a label error of similar magnitude. When aggregated over N_{geo} independent noisy observations, the effective label uncertainty scales as $\sigma_{\text{geo}}/\sqrt{N_{\text{geo}}}$, which directly affects the quality of the latent code recovered at inference and ultimately the accuracy of the reconstructed zero-level set.

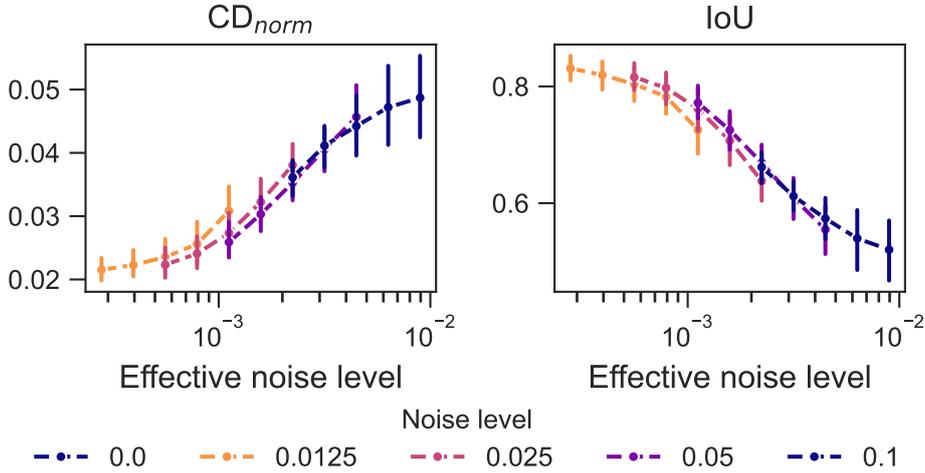


Figure 9: CD_{norm} and IoU for SDF-SM as a function of effective noise level ρ on test geometries. Each point corresponds to a different $(\sigma_{\text{geo}}, N_{\text{geo}})$ combination. Error bars represent one standard deviation.

This intuition is confirmed by Figure 9, which shows that reconstruction errors obtained for different

pairs (σ, N) collapse onto the same curve when plotted against ρ . In particular, configurations sharing the same value of ρ yield nearly identical accuracy, despite corresponding to different absolute noise levels and sampling densities. This confirms that ρ is the relevant scaling parameter governing the reconstruction performance and provides practical guidance for experimental design: increasing the number of sampled points compensates for higher noise levels according to the predictable trade-off $\rho = \sigma_{\text{geo}}/\sqrt{N_{\text{geo}}}$.

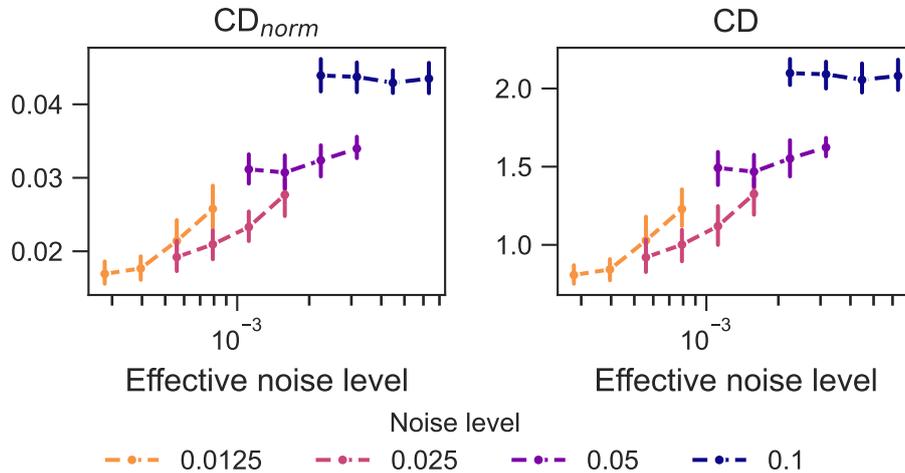


Figure 10: CD_{norm} and CD for PCA-SM as a function of effective noise level ρ on test geometries. Unlike SDF-SM (Fig. 9), points do not collapse onto a single curve, indicating that ρ does not capture the full behavior of the optimization. Error bars represent one standard deviation.

In contrast, Figure 10 shows that ρ does not characterize PCA-SM performance as effectively: points corresponding to different $(\sigma_{\text{geo}}, N_{\text{geo}})$ combinations do not collapse onto a single curve. The noise level σ_{geo} is discriminant, while the sample size N_{geo} does not impact the reconstruction quality in an equal manner.

2.2 Synthetic geometries generation

We use the latent space and decoder combination learned by SDF-SM in order to generate new realistic geometries that can be used for data augmentation in the surrogate modeling task. Following the procedure described in Section 4.7, we sample 1000 latent codes from the empirical distribution $\mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\Sigma})$ and generate corresponding synthetic geometries. Of these, 976 geometries are anatomically plausible and retained for data augmentation, while 24 samples are discarded due to topological artifacts such as holes close to the apex or degeneration in the base surface. Figure 11 show examples of successfully generated geometries, while some discarded examples are reported in Figure 18 in Appendix A.



Figure 11: Examples of five successfully generated synthetic left ventricular geometries.

2.3 Shape-dependent surrogate model

In this section, we report results on training of the cardiac mechanics surrogate $\mathcal{NN}_{\text{phys}}$. We focus on quantifying the effect of shape encoding and dataset augmentation, as well as the influence of query points and loss augmentation by strains. First, we show results on the idealized cohort of artificially generated left ventricle geometries, and then investigate the real cohort of patient-specific geometries, as standalone dataset and as dataset augmented by synthetically generated models.

2.3.1 Idealized geometries

We train the surrogate model with displacement solutions of the cohort of 512 prolate ellipsoid geometries (Figure 1b) and investigate the influence on the root mean squared error (RMSE) of: shape encoding and landmarking (Figure 12); the number of query points (Figure 13); and the presence of strain information (Figure 14).

Shape encoding and landmarking To test the influence of shape codes and geometric landmarks (e.g., UVCs), we train our neural network using $n = 15000$ randomly selected points from the geometry, and only consider displacement information in Eq. (44), hence $\lambda_s = 0$ (no strain-based regularization). For efficiency purposes, we bypass any strain calculation.

Figure 12a shows the RMSE on the training ($N_{\text{train}} = 230$), validation ($N_{\text{valid}} = 179$), and test set ($N_{\text{test}} = 103$), Figure 12b the RMSE over the number of training epochs (split into training and validation error), and Figure 12c the plotted absolute difference in ground-truth and inferred displacement for ten selected geometries from the test set. Five different shape encodings/geometric landmarkings are used:

- \mathbf{x} : Input layer only consisting of Cartesian coordinates;
- $\mathbf{x}, \boldsymbol{\xi}$: Input layer consisting of Cartesian coordinates and UVCs;
- $\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g^{(\text{SDF})}(\Omega_0)$: Input layer consisting of Cartesian coordinates, UVCs, and a DeepSDF shape code of dimension 3;
- $\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g^{(\text{PCA})}(\Omega_0)$: Input layer consisting of Cartesian coordinates, UVCs, and a PCA shape code of dimension 6;
- $\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g(\{\ell, d, w\})$: Input layer consisting of Cartesian coordinates, UVCs, and a set of length, diameter, and wall thickness of the idealized shape.

The first column block of Table 5 shows the averaged errors over the training, validation, and test sets. Compared to the naïve setting providing only \mathbf{x} to the input layer, adding UVCs, $\boldsymbol{\xi}$, alone already leads to a 40.2% (training set), 39.9% (validation set), and 40.6% (test set) improvement in RMSE, respectively. Providing PCA-SM, $\boldsymbol{\mu}_g^{(\text{PCA})}(\Omega_0)$, or SDF-SM, $\boldsymbol{\mu}_g^{(\text{SDF})}(\Omega_0)$, shape codes then improves the error further by 77.7% and 78.7% (training set), 73.3% and 79.2% (validation set), and 73.4% and 77.4% (test set), respectively. Only informing the neural network by the analytic idealized model parameterization $\boldsymbol{\mu}_g(\{\ell, d, w\})$ can achieve an even more substantiated error reduction: by 85.1% (training set), 85.8% (validation set) and 84.5% (test set) compared to using $\mathbf{x}, \boldsymbol{\xi}$ only.

Number of query points We investigate how the number of (randomly subsampled) query points per geometry used in training affects the losses evaluated on the entirety of geometric points. Figure 13 shows the RMSE across the training ($N_{\text{train}} = 230$), validation ($N_{\text{valid}} = 179$), and test set ($N_{\text{test}} = 103$), and the first column block of Tab. 4 lists the average over the respective set. The median values over all n are $4.33 \cdot 10^{-3}$ (training set), $4.30 \cdot 10^{-3}$ (validation set), and $4.35 \cdot 10^{-3}$ (test set). The relative variabilites across all n are 2.74% (training set), 2.76% (validation set), and 2.39% (test set), indicating little variability of the number of query points $n \in [5, 15000]$ on the errors.

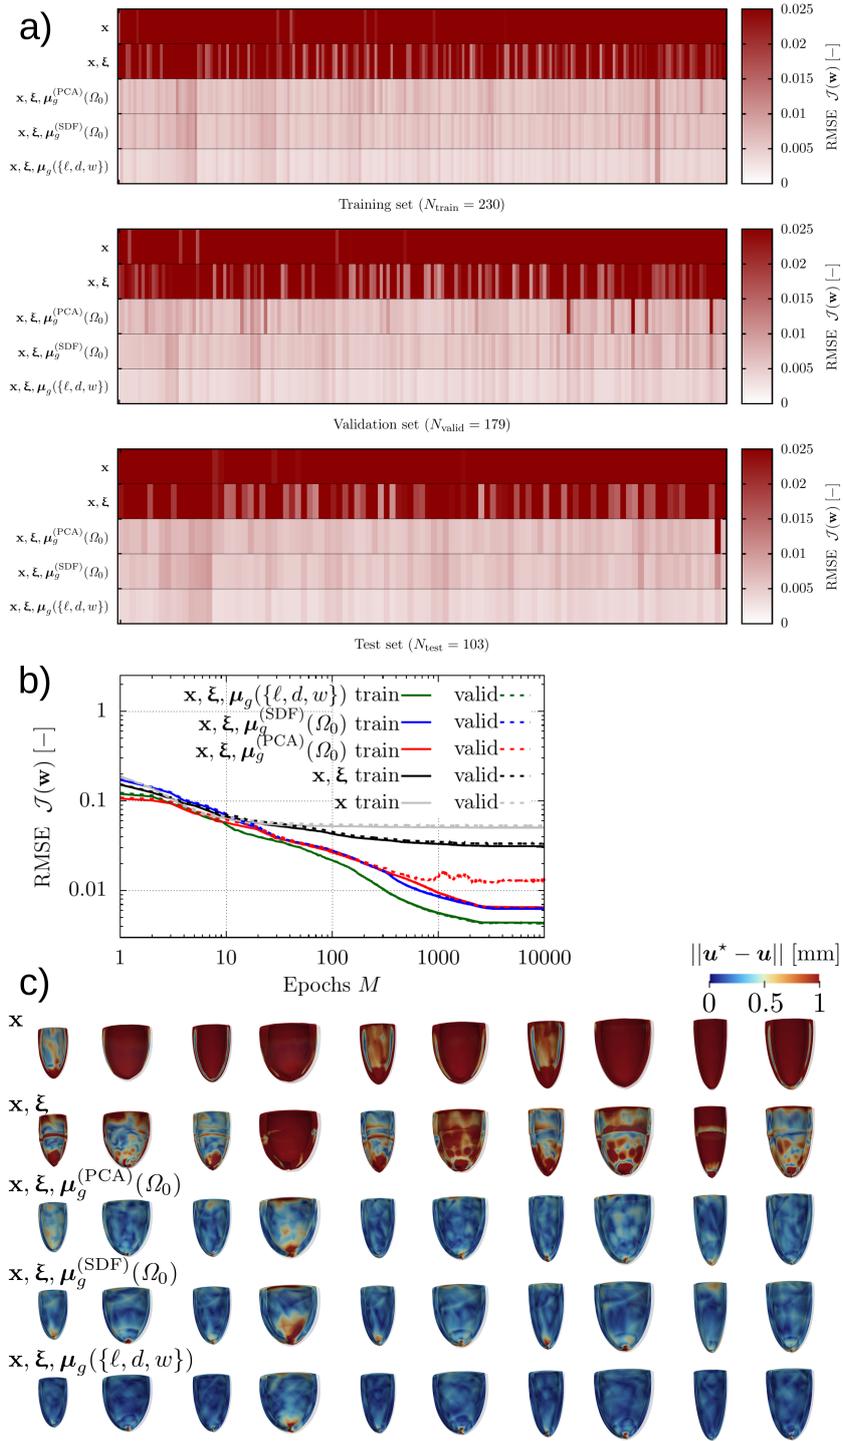


Figure 12: Idealized geometries: Comparison of different input layers and shape codes (\mathbf{x} : only Cartesian coordinates; $\mathbf{x}, \boldsymbol{\xi}$: Cartesian coordinates plus UVCs; $\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g^{(\text{PCA})}(\Omega_0)$: Cartesian coordinates, UVCs, PCA-SM shape code; $\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g^{(\text{SDF})}(\Omega_0)$: Cartesian coordinates, UVCs, SDF-SM shape code; $\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g(\{\ell, d, w\})$: Cartesian coordinates, UVCs, analytic shape code). No strain contributions to the loss function ($\lambda_s = 0$), number of query points $n = 15000$. a) RMSE on training, validation, and test set for the five different input layers. b) RMSE over training epochs, split into training and validation loss. c) Absolute difference in ground-truth and inferred displacement on ten selected test geometries.

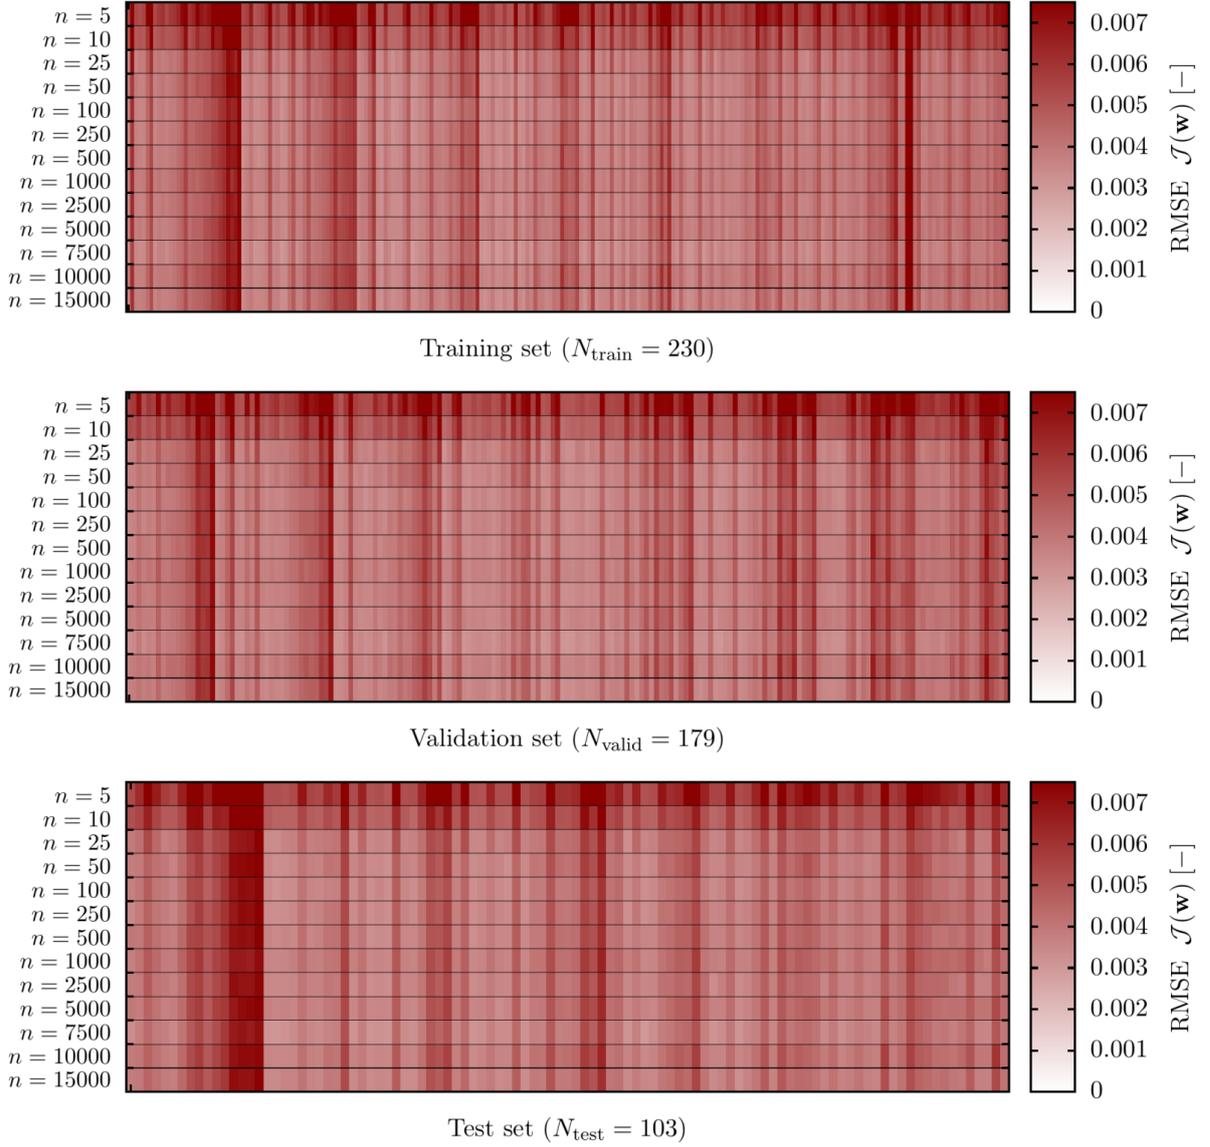


Figure 13: Idealized geometries: RMSE on training, validation, and test set for different number of query points, $n \in [5, \dots, 15000]$. No strain contributions to the loss function ($\lambda_s = 0$).

Table 4: Average root mean squared errors (RMSE) over the respective geometry cohort of idealized models. Comparison of influence of number of query points and strain augmentation of the loss function.

$\emptyset \mathcal{J}_u$	Idealized Geometries														Strain, $n = 2500$	
	Query Points $\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g(\{\ell, d, w\}), \lambda_s = 0$														$\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g(\{\ell, d, w\})$	
$\cdot 10^{-2}$	$n = 5$	$n = 10$	$n = 25$	$n = 50$	$n = 100$	$n = 250$	$n = 500$	$n = 1000$	$n = 2500$	$n = 5000$	$n = 7500$	$n = 10000$	$n = 15000$	$\lambda_g = 0$	$\lambda_g = 0.1$	
TRAIN	0.680	0.558	0.445	0.432	0.420	0.427	0.433	0.441	0.428	0.437	0.403	0.441	0.425	0.440	0.399	
VALID	0.674	0.553	0.444	0.429	0.417	0.424	0.430	0.436	0.425	0.433	0.400	0.438	0.422	0.438	0.396	
TEST	0.668	0.552	0.446	0.435	0.422	0.430	0.435	0.442	0.429	0.437	0.406	0.441	0.426	0.443	0.401	

Strain information The effects of enriching the loss function Eq. (45) by strains directly computed from the finite element space are investigated. For this, the number of query points is set to $n = 2500$, and the parameter λ_s in the combined loss Eq. (44) is chosen such that displacement and strain losses are of comparable magnitude. The input layer of the network consists of $\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g(\{\ell, d, w\})$. Figure 14a shows the *displacement* loss on the training, validation, and test set for $\lambda_s = 0$ (bypassing strain computation) and $\lambda_s = 0.1$ (adding strain contributions to the loss), and Fig. 14b shows the displacement RMSE over the epochs for both cases. The averages over the geometries are listed in the second column block of Tab. 4.

We observe mild improvements in the loss on the displacement prediction when adding strains to the overall loss function, amounting to a relative improvement of 9.32% (training set), 9.59% (validation set), and 9.48% (test set). However, we observe an increase in computational expense by a factor of ~ 5 compared to bypassing gradient computations.

2.3.2 Combined real and generated geometries

After studying different input layer settings on the idealized model cohort, we investigate the augmented dataset of real and generated geometries (44 + 976 models) and compare to training only on the real geometry cohort (44 models). For this study, we keep a fixed number of query points $n = 15000$ and bypass strain contributions to the loss ($\lambda_s = 0$). We train our network using the following three variants of input layers:

- $\mathbf{x}, \boldsymbol{\xi}$: Input layer consisting of Cartesian coordinates and UVCs;
- $\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g^{(\text{D-SDF})}(\Omega_0)$: Input layer consisting of Cartesian coordinates, UVCs, and SDF-SM shape code of dimension 16;
- $\mathbf{x}, \boldsymbol{\xi}, \boldsymbol{\mu}_g^{(\text{PCA})}(\Omega_0)$: Input layer consisting of Cartesian coordinates, UVCs, and a PCA-SM shape code of dimension 34.

Figure 15a shows the RMSE on the training ($N_{\text{train}} = 630$), validation ($N_{\text{valid}} = 305$), and test set ($N_{\text{test}} = 85$), Fig. 15b the RMSE over the number of training epochs (split into training and validation error), and Fig. 15c the plotted absolute difference in ground-truth and inferred displacement for all nine geometries from the test subset of the real cohort (HF and healthy). Results from the augmented (generated plus real) as well as the real cohort only are reported separately. The averages over the geometries are listed in the second (augmented) and third (real only) column block of Tab. 5. As for the idealized models, a substantial improvement is observable when adding shape codes retrieved from the PCA-SM and SDF-SM model compared to only considering \mathbf{x} and $\boldsymbol{\xi}$. Reductions in RMSE for the augmented dataset of generated and real geometries of 66.4% and 68.9% (training set), 58.5% and 55.9% (validation set), and 54.6% and 40.5% (test set) are achieved. Furthermore, we observe that training only on real geometries without dataset augmentation (last column block of Tab. 5) yields significantly worse results on the validation and test set, with a better training error, thus highlighting a severe overfitting when the proposed augmentation strategy is not employed.

3 Discussion

In this work, we introduced a pipeline for the development of surrogate models for shape-dependent parametrized PDEs, with specific application to cardiac mechanics. Because myocardial deformation depends nonlinearly on the domain geometry, clinically useful surrogate models must generalize to anatomies not observed during training. A central challenge in this context is the data-scarce regime typical of biomedical applications, where the cost and time required to acquire large collections of patient-specific geometries and simulations are often prohibitive. To mitigate data scarcity, we proposed decoupling geometry modeling from learning the physics response. A shape model is first trained on the available anatomies, and subsequently exploited in generative mode, to produce synthetic geometries representative of the distribution of the available shapes. These geometries are then used to augment the dataset of high-fidelity simulations employed to train the physics surrogate. Importantly, no additional real anatomies are required. The high-fidelity solver

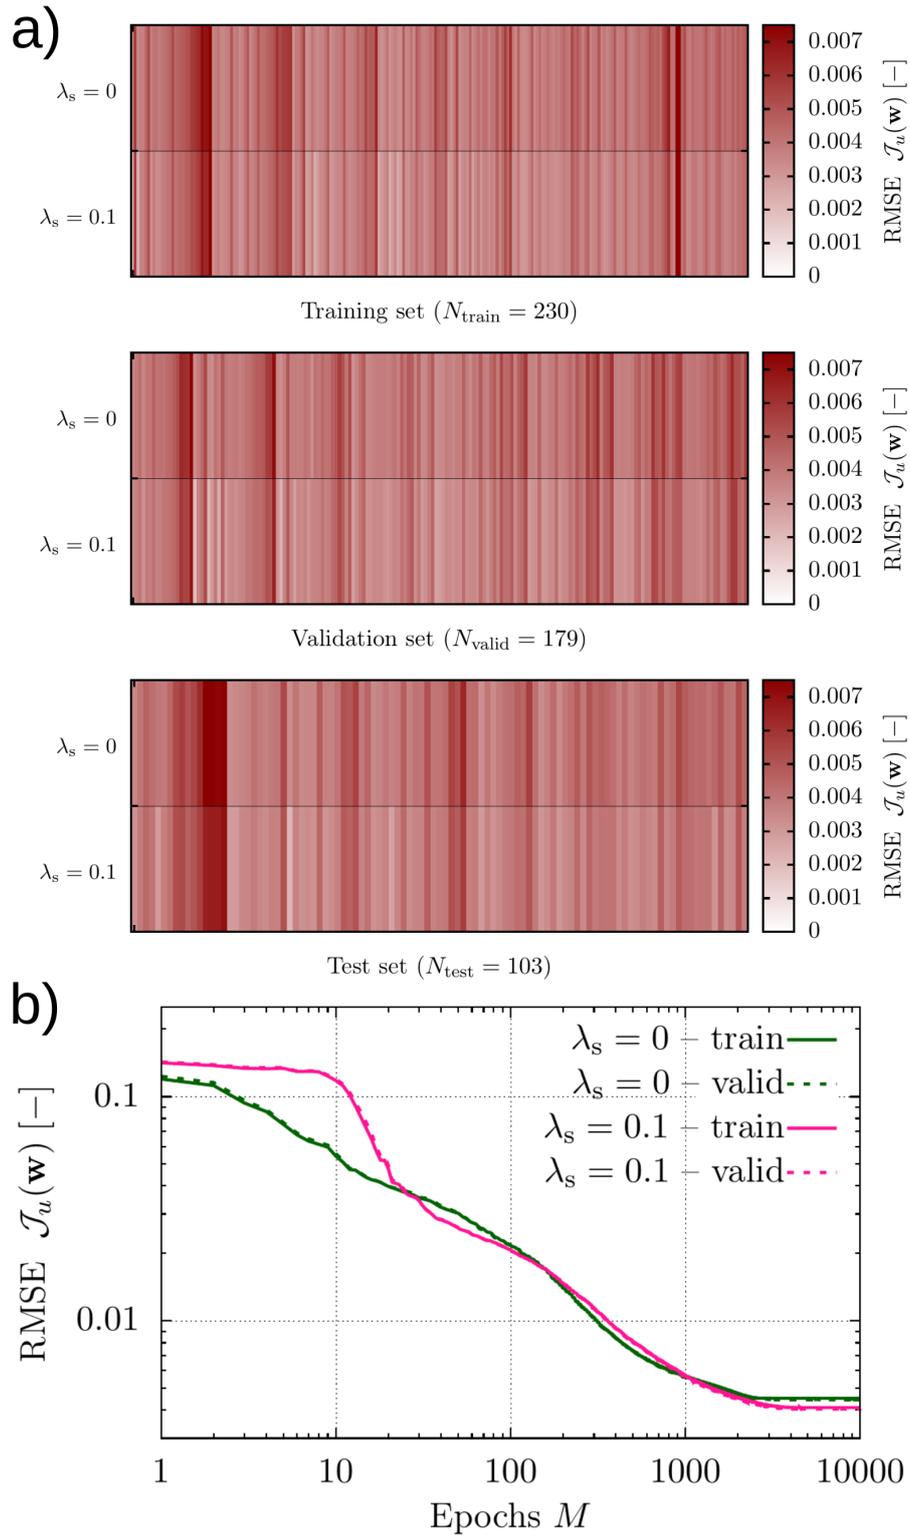


Figure 14: Idealized geometries, investigation of effects of strain contributions to the loss function. No strain ($\lambda_s = 0$) vs. strain loss, with $\lambda_s = 0.1$ chosen such that both loss contributions are of comparable magnitude; number of query points $n = 2500$. a) Displacement RMSE on training, validation, and test set. b) Displacement RMSE over number of training epochs, split into training and validation loss.

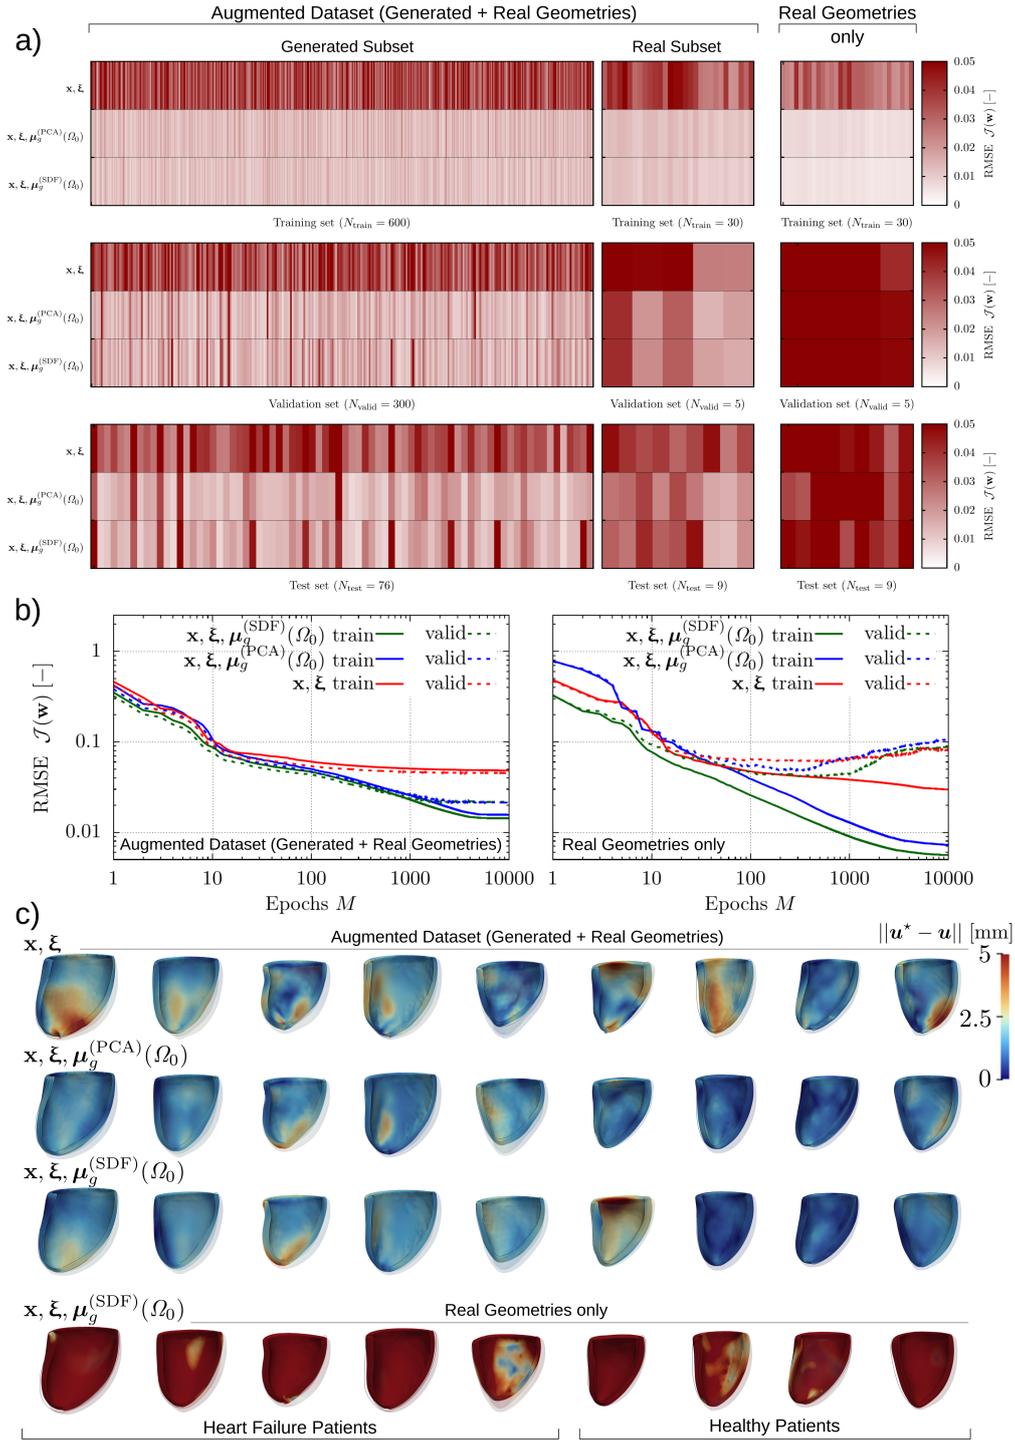


Figure 15: Generated and real geometries: Augmented (real and generated geometries) vs. non-augmented (only real geometries) datasets. Comparison of different input layers and shape codes (\mathbf{x}, ξ : only Cartesian coordinates plus UVCs; $\mathbf{x}, \xi, \mu_g^{(PCA)}(\Omega_0)$: Cartesian coordinates, UVCs, PCA-SM shape code; $\mathbf{x}, \xi, \mu_g^{(SDF)}(\Omega_0)$: Cartesian coordinates, UVCs, SDF-SM shape code). No strain contributions to the loss function ($\lambda_s = 0$), number of query points $n = 15000$. a) RMSE on training, validation, and test set for three different input layers, split into generated and real geometries. b) RMSE over training epochs, split into training and validation loss. c) Absolute difference in ground-truth and inferred displacement on all nine test geometries from the real cohort.

Table 5: Average root mean squared errors (RMSE) over the respective geometry cohort of idealized, generated, and real geometries. Comparison of influence of shape code and data augmentation.

$\emptyset \mathcal{J}_u$	Idealized Geometries					Gen. + Real Geometries			Real Geometries			
	<i>Shape</i>											
	$n = 15000, \lambda_s = 0$											
			$\mathbf{x}, \xi, \mu_g^{(\text{PCA})}(\Omega_0)$	$\mathbf{x}, \xi, \mu_g^{(\text{SDF})}(\Omega_0)$	$\mathbf{x}, \xi, \mu_g(\{\ell, d, w\})$			$\mathbf{x}, \xi, \mu_g^{(\text{PCA})}(\Omega_0)$			$\mathbf{x}, \xi, \mu_g^{(\text{PCA})}(\Omega_0)$	$\mathbf{x}, \xi, \mu_g^{(\text{SDF})}(\Omega_0)$
$\cdot 10^{-2}$	\mathbf{x}	ξ				\mathbf{x}, ξ				\mathbf{x}, ξ		
TRAIN	4.772	2.856	0.637	0.609	0.425	3.926	1.318	1.221	2.914	0.722	0.561	
VALID	4.960	2.979	0.796	0.620	0.422	3.876	1.609	1.711	7.498	8.468	8.383	
TEST	4.631	2.749	0.730	0.620	0.426	4.004	1.818	2.388	5.790	7.395	6.653	

acts indeed as a bridge between synthetic geometry generation and physics learning, effectively introducing a stronger observational bias during surrogate model training. This strategy resulted in a marked improvement in generalization to unseen patients (error reduced from 6.7% to 2.4% for SDF-SM). These results highlight that, in data-scarce settings, a well-trained shape model used in generative mode can effectively enrich the observed physical domain, and strengthen the surrogate model’s generalization capability, without requiring any new real geometries.

A key structural feature of the proposed framework is the explicit separation between a shape model and a physics surrogate, conditioned on a shape code. The shape model serves a dual role: (i) it learns the distribution of anatomies – even in data-scarce regimes – and enables sampling of novel shapes, and (ii) it provides a low-dimensional shape encoding that conditions the physics surrogate, thus allowing for generalization to unseen geometries without observable overfitting. For the surrogate model architecture, we adopted USM-Nets, a conditional neural-field supporting multiple positional encoding strategies. In this case, we used a domain-specific positional encoding, namely UVCs, which – as shown by the results – reduce the error of a naïve predictor receiving only spatial coordinates from 4.6% to 2.7% (in the idealized models), even without any shape encoding. We compared two alternative shape encoding strategies: a PCA-based approach (generalizing statistical shape models to point-cloud representations) and a DeepSDF-based implicit neural representation trained in auto-decoder mode.

Notably, the proposed decoupled framework enables a geometry-aware data augmentation strategy that is not naturally supported by several alternative approaches in the literature, such as GNN- or Transformer-based models that learn the mapping from geometry and physical parameters to the solution in a one-shot fashion. While such architectures have demonstrated strong approximation capabilities, their formulations typically rely on large training datasets to achieve stable generalization. For instance, recent graph-based and attention-based surrogate models report training regimes involving on the order of 10^3 – 10^4 simulated samples to reach competitive accuracy [38, 31, 33]. In contrast, our framework explicitly separates geometric representation from physics regression. This structural decoupling allows the shape model to learn an explicit latent distribution over anatomies, which can then be sampled to generate synthetic geometries for data augmentation. One-shot architectures that directly map geometry to solution do not explicitly model such a generative geometric distribution, and therefore do not naturally enable controlled sampling in the space of anatomies without additional modeling components.

The total error associated with our surrogate model can be decomposed into two components: the geometry encoding error and the physics prediction error. To disentangle these contributions, we introduced an idealized dataset in which an exact parameterization of the geometry is available. The results showed that the error introduced by shape encoding is, in this setting, smaller than the intrinsic physics prediction error. Specifically, normalized error increases from 0.43% (exact parameterization) to 0.73% (PCA-based encoding)

and 0.62% (DeepSDF-based encoding). The results also demonstrated robustness to point subsampling. Incorporating strain measurements alongside displacement in the loss function yielded a 9.5% improvement in accuracy, at the expense of approximately fivefold increased computational cost, highlighting a trade-off between accuracy and efficiency.

Comparing the two shape encoding techniques (PCA-based and DeepSDF-based), we showed that both are robust to noise and data subsampling when reconstructing geometries via subsequent encoding and decoding. We introduced a quantity called the effective noise level, which captures how noise magnitude and number of points interact in determining shape model accuracy. In this framework, the PCA-based shape model yields slightly more accurate results when the effective noise level is very low, but is less resilient than the DeepSDF-based model as the effective noise level increases.

When integrated into the physics surrogate, the relative performance of the two encodings depended on the dataset. In the idealized setting, DeepSDF achieved lower error (0.62% vs 0.73%), whereas for real ventricular geometries PCA performed better (1.82% vs 2.39%). Overall, the two approaches yielded comparable performance for left ventricular geometries, indicating that both are adequate to capture the dominant modes of anatomical variability in this context. However, for more complex geometries whose variability cannot be well approximated by linear combinations of modes, implicit neural representations such as DeepSDF are expected to provide greater robustness. Additionally, PCA-based encoding requires point-to-point mesh correspondence – enabled here through UVCs – which may not be available in more general applications. From this perspective, DeepSDF-based encoding provides a more flexible alternative across different application scenarios compared to PCA-based encoding.

4 Methods

4.1 Data preprocessing

Idealized data For benchmarking and testing purposes, a cohort of 512 idealized left ventricular geometries is created. Each ventricular geometry is modeled as a prolate ellipsoid, a shape commonly used to approximate the left ventricle due to its elongated geometry along the long axis. The ellipsoid is truncated along a plane orthogonal to the long axis, representing the basal opening toward the atrium. The determining geometric parameters are the ventricular long axis $\ell \in [75 \text{ mm}, 115 \text{ mm}]$, its maximum diameter $d \in [40 \text{ mm}, 80 \text{ mm}]$, and its transmural thickness $w \in [5 \text{ mm}, 15 \text{ mm}]$. From each geometric parameter range, 8 evenly spaced samples are taken, each of which is combined with every other sample, yielding a $8 \times 8 \times 8$ -sized cohort of geometries. Six exemplary models are shown in Fig. 1b.

Patient-specific geometries In order to generate high-fidelity training data on patient-specific anatomical heart models, publicly available cohorts provided by Strocchi et al. [35] and Rodero et al. [36] of volume-meshed full-heart cardiac geometries are used. Hearts in [35] are rather enlarged since they are from patients with diagnosed heart failure condition, whereas geometries in [36] stem from “healthy” patients and do not exhibit pathological changes in shape. Here, we focus on inferring physics behavior on the left ventricle (LV), hence we take steps to extract, smooth, and re-mesh the LV subset for subsequent use in a finite element forward model.

Our geometry extraction and remeshing pipeline consists of the following steps:

- Thresholding to extract all elements with LV region tag.
- Identification of 3 points on lower part of mitral valve rim for cut plane definition.
- Cut geometry at an inferior offset $\in [7.0 \text{ mm}, 22.0 \text{ mm}]$ to the defined plane (minimal distance to exclude all mitral valve rim artifacts) using a clipped threshold filter.
- Extract surface mesh on cut geometry, discarding volume mesh information (due to degenerate elements after cutting operation).
- Perform 500 smoothing iterations on the surface mesh using the Smooth Polydata Filter.

- (Re-)Generation of volume tetrahedral finite element mesh based off the surface mesh using a mesh size factor of 1.0, i.e. keeping characteristic element edge length from the original data.
- Redefinition of mesh tags, assigning surface element mesh tags to lumen, baseplane, and epicardium in order to facilitate boundary condition application in finite element model.

Six exemplary patient-specific geometries are shown in Fig. 1a.

4.2 Data alignment & normalization

To reduce the dimensionality of the anatomical space, we factor out rigid-body degrees of freedom and decouple size from the shape reconstruction problem, as size is cohort-dependent and would otherwise confound geometric comparisons. We note that the physical problems solved on the reconstructed geometries are translation- and rotation-invariant, but not scale-invariant; consequently, scaling information must be retained and incorporated into the surrogate model $\mathcal{NN}_{\text{phys}}$ to preserve correct physical behavior.

Each geometry is preprocessed as follows. First, it is centered at the LV barycenter and oriented so that the basal plane is horizontal, by performing PCA on the basal plane and retaining the first two principal components. After ensuring that the apex points downward and that the coordinates system is right-handed, we rotate the geometry around the z-axis in order to ensure that the right ventricle (RV) barycenter lays on the x-axis with positive x. Finally, the mesh is centered according to the barycenter of the LV and rescaled so that all vertices belong to the unit cube $[-1, 1]^3$. The scaling factor is stored as additional input for the shape reconstruction model $\mathcal{NN}_{\text{sdf}}$ and for the surrogate model $\mathcal{NN}_{\text{phys}}$. The detailed procedure is reported in Algorithm 1.

Algorithm 1: Orientation and Normalization of Mesh \mathcal{M}

Data: Mesh \mathcal{M} containing LV, RV, base plane segmentations

Result: Aligned and normalized mesh \mathcal{M}

Step 1: Extract anatomical vertex sets;

Basal plane V_{bas} , LV V_{lv} , RV V_{rv} , surface V_{surf} ;

Step 2: Initial centering;

Compute the barycenter of LV μ_{lv} and center all sets: $V \leftarrow V - \mu_{\text{lv}}$;

Step 3: PCA alignment;

Perform PCA on V_{bas} to obtain principal components $P \in \mathbb{R}^{3 \times 3}$;

Transform all vertex sets: $V' = VP^T$;

Recenter on basal plane by subtracting the basal plane barycenter μ_{bas} : $V' \leftarrow V' - \mu'_{\text{bas}}$;

Step 4: Orient apex-base axis;

if $\text{mean}(V'_{\text{lv}}(:,2)) > \text{mean}(V'_{\text{bas}}(:,2))$ **then**

| Flip all coordinates: $V' \leftarrow -V'$ and $P \leftarrow -P$;

end

Ensure right-handed system: **if** $\det(P) < 0$ **then** $V'^{(:,1)} \leftarrow -V'^{(:,1)}$;

Step 5: Align RV to positive x-axis;

Compute RV barycenter μ'_{rv} and rotation matrix R_{align} ;

Apply rotation: $V' \leftarrow V'R_{\text{align}}^T$ for all sets;

if $\text{mean}(V'_{\text{rv}}(:,0)) < 0$ **then**

| Apply 180° rotation around z-axis;

end

Step 6: Final normalization;

Center on LV by subtracting the new barycenter μ'_{lv} : $V'' = V' - \mu'_{\text{lv}}$;

Compute max coordinate: $s_{\text{max}} = \max_{V,i,j} |V''_{ij}|$;

Scale to unit cube: $\tilde{V} = V''/s_{\text{max}}$;

Store scaling factor s_{max} ;

return Normalized and oriented mesh $\tilde{\mathcal{M}}$ with scaling factor s_{max}

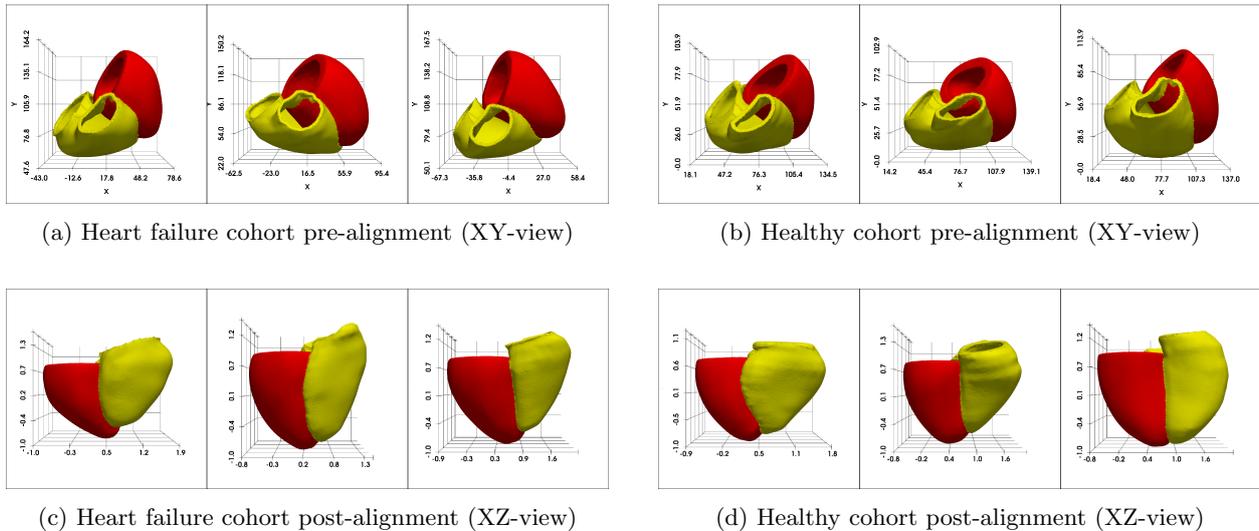


Figure 16: Views of anatomies from both cohorts pre and post alignment and orientation procedure.

4.3 Dataset generation

The dataset consists of 44 left ventricular geometries divided into training (35 geometries: 19 from [35], 16 from [36]) and test (9 geometries: 5 from [35], 4 from [36]) sets. Additionally, 512 idealized prolate ellipsoid geometries are used, divided into training ($N_{\text{train}} = 409$) and test ($N_{\text{test}} = 103$) sets.

Point cloud sampling strategy. For training SDF-SM, we construct a point cloud dataset \mathcal{X}_i for each geometry following the sampling strategy from [27]. This strategy combines three types of samples to capture different aspects of the geometry:

- 2000 points sampled on the surface and perturbed with Gaussian noise $\sigma = 0.025$
- 1500 points sampled on the surface and perturbed with Gaussian noise $\sigma = 0.25$
- 500 points sampled uniformly in the bounding box $[-1, 1]^3$

A validation set is generated by perturbing 1000 points from the training surfaces with $\sigma = 0.25$.

For each sampled point, the ground-truth SDF value is computed as the distance to its closest neighbor on the original (unperturbed) surface, obtained via KD-tree queries [39]. The sign is determined by checking whether the point is contained within a cell of the original mesh.

To preserve scale information factored out during preprocessing (Algorithm 1), each geometry i is associated with a normalized scaling parameter μ_i used as additional input to condition both \mathcal{N}_{sdf} and $\mathcal{N}_{\text{phys}}$:

$$\mu_i = \frac{s_{i_{\max}}}{\max_j s_{j_{\max}}}, \quad (2)$$

where $s_{i_{\max}}$ is the maximum absolute coordinate value of geometry i before normalization.

Inference datasets. At inference time, latent codes are optimized using point clouds sampled directly on the geometry surfaces without perturbation. For real anatomies, 4000 surface points are used; for idealized geometries, 2000 surface points are used. Target SDF values are set to zero for all surface points.

Robustness evaluation dataset. To evaluate reconstruction quality under varying noise and sampling conditions (Section 2.1.3), points are sampled on the surface and perturbed with Gaussian noise, but the ground-truth SDF value is retained as the noise-free value corresponding to the original surface point. This introduces a labeling error that simulates realistic measurement noise: the number of sampled points varies across $\{125, 250, 500, 1000, 2000\}$ and the noise standard deviation across $\{0.0, 0.0125, 0.025, 0.05, 0.1\}$.

4.4 Shape model training step

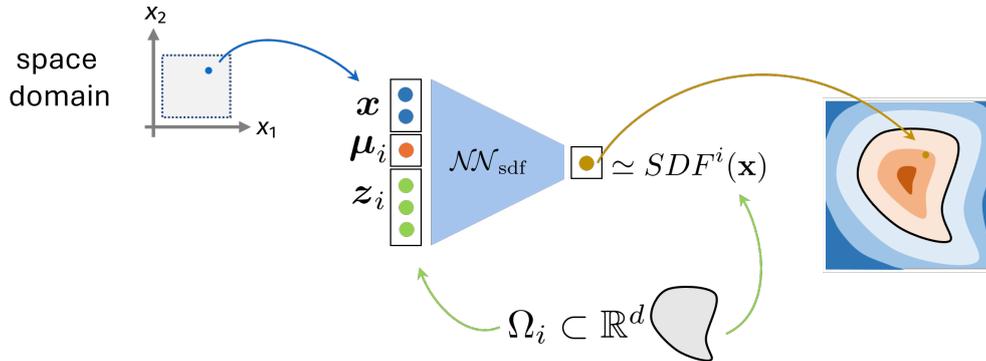


Figure 17: Auto-decoder architecture visualization. The geometrical input \mathbf{x} is concatenated with the latent code \mathbf{z}_i associated to the geometry Ω_i and with optional additional physical parameters μ_i .

In this work, we adopt an autoencoder formulation to discover a continuous latent shape space jointly with a neural implicit SDF model. We choose to represent each 3D shape Ω_i implicitly by training the autoencoder to approximate its signed distance function (SDF^{*i*}):

$$\text{SDF}^i(\mathbf{x}) = \begin{cases} -d(\mathbf{x}, \partial\Omega_i) & \text{if } \mathbf{x} \in \Omega_i \\ d(\mathbf{x}, \partial\Omega_i) & \text{if } \mathbf{x} \notin \Omega_i, \end{cases} \quad (3)$$

$$d(\mathbf{x}, \partial\Omega_i) := \inf_{\mathbf{y} \in \partial\Omega_i} d(\mathbf{x}, \mathbf{y}) \quad \forall \mathbf{x} \in X.$$

This formulation has been shown to provide smooth latent spaces, high-fidelity reconstructions, and strong generalization across shape categories, as demonstrated by the DeepSDF method [27]. The underlying mathematical formulation is based on a probabilistic perspective [27, 29].

Given a set of N shapes, each represented by the associated signed distance function $\{\text{SDF}^i\}_{i=1}^N$, we build N sets of K_i point samples, each corresponding to one shape:

$$X_i = \{(\mathbf{x}_j, \mu_i, s_j)_{j=1}^{K_i} : s_j = \text{SDF}^i(\mathbf{x}_j)\}. \quad (4)$$

In the auto-decoder setting, each geometry is paired with a shape code \mathbf{z}_i lying in the lower-dimensional latent space \mathcal{Z} . The shape code \mathbf{z}_i is learned together with the parameters \mathbf{w}_{sdf} , which parametrize the SDF likelihood.

In the Bayesian setting, the posterior distribution over the latent variable \mathbf{z}_i can be decomposed as the product of a SDF likelihood term and a prior distribution $p(\mathbf{z}_i)$ over the latent space, as in Eq. (5):

$$p_{\mathbf{w}}(\mathbf{z}_i | X_i) \propto p(\mathbf{z}_i) \prod_{(\mathbf{x}_j, \mu_i, s_j) \in X_i} p_{\mathbf{w}}(s_j | \mathbf{z}_i; \mathbf{x}_j, \mu_i), \quad (5)$$

where \mathbf{w} parametrizes the SDF likelihood. In our case, the SDF is approximated by using a neural network $\mathcal{NN}_{\text{sdf}}$, described by its set of weights \mathbf{w}_{sdf} . As a consequence, the likelihood function can be defined as:

$$p_{\mathbf{w}_{\text{sdf}}}(s_j | \mathbf{z}_i; \mathbf{x}_j, \mu_i) = \exp(-\mathcal{L}(\mathcal{NN}_{\text{sdf}}(\mathbf{z}_i, \mathbf{x}_j, \mu_i; \mathbf{w}_{\text{sdf}}), s_j)), \quad (6)$$

where \mathcal{L} is a custom loss function. The prior distribution over the latent space is modeled as a multivariate Gaussian with zero-mean and spherical covariance $\sigma_{\mathbf{z}}^2 I$ providing a simple and well-behaved initial structure for the latent manifold.

At training time, the goal is to find both the best neural network approximation of the SDF associated to each shape and the best corresponding shape code in the latent space. This can be obtained by minimizing the joint log-posterior for all shapes with respect both to the codes $\{\mathbf{z}_i\}_{i=1}^N$ and the network weights \mathbf{w}_{sdf} such that:

$$\min_{\mathbf{w}_{\text{sdf}}, \{\mathbf{z}_i\}_{i=1}^N} \sum_{i=1}^N \left(\sum_{j=1}^K \mathcal{L}(\mathcal{N}_{\text{sdf}}(\mathbf{z}_i, \mathbf{x}_j, \mu_i; \mathbf{w}_{\text{sdf}}), s_j) + \frac{1}{2\sigma_{\mathbf{z}}^2} \|\mathbf{z}_i\|_2^2 \right). \quad (7)$$

Here, $\mathcal{L}(\tilde{s}_j, s_j)$ penalizes discrepancies between the predicted SDF value $\tilde{s}_j = \mathcal{N}_{\text{sdf}}(\mathbf{z}_i, \mathbf{x}_j; \mathbf{w}_{\text{sdf}})$ and the ground-truth SDF value s_j . In this work we adopt an ℓ_1 -based loss, defined as:

$$\mathcal{L}(\tilde{s}_j, s_j) = \frac{|\tilde{s}_j - s_j|}{b}, \quad (8)$$

where b is a normalization constant.

To ensure smooth interpolation between shapes and stable latent space optimization, we control the Lipschitz constant on the decoder. This regularization bounds the sensitivity of SDF predictions to perturbations in the latent code, preventing discontinuous transitions in the manifold of learned shape codes. To achieve this we apply a weight regularization technique based on minimizing the Lipschitz constant of the network introduced in [29], and named Lipschitz Multilayer Perceptron.

Applying the proposed architecture, each layer l of \mathcal{N}_{sdf} is augmented with a Lipschitz weight normalization layer, where the bound c_l is computed from the current weights as:

$$y = \sigma(\hat{\mathbf{w}}_{\text{sdf}l} \mathbf{x} + b_l), \quad \hat{\mathbf{w}}_{\text{sdf}l} = \text{normalize}_{\infty}(\mathbf{w}_{\text{sdf}l}, \text{softplus}(c_l)), \quad \text{softplus}(c_l) = \ln(1 + e^{c_l}), \quad (9)$$

where $\mathbf{w}_{\text{sdf}l}$ is the weight matrix for layer l , and $c_l = \|\mathbf{w}_{\text{sdf}l}\|_{\infty}$ denotes the maximum absolute row sum of the weight matrix. The normalization operator $\text{normalize}_{\infty}$ scales each row of $\mathbf{w}_{\text{sdf}l}$ such that its ℓ_1 norm does not exceed $\text{softplus}(c_l)$, ensuring the layer's operator norm is bounded.

The set $\{c_l\}_{l=1}^{N_{\text{layers}}}$ is used to compute an upper bound for the Lipschitz constant of \mathcal{N}_{sdf} as:

$$C(\mathbf{w}_{\text{sdf}}) = \prod_{l=1}^{N_{\text{layers}}} \text{softplus}(c_l). \quad (10)$$

Given the Lipschitz constant of the network $C(\mathbf{w}_{\text{sdf}})$ the final training objective function becomes:

$$\min_{\mathbf{w}_{\text{sdf}}, \{\mathbf{z}_i\}_{i=1}^N} \sum_{i=1}^N \left(\sum_{j=1}^K |\mathcal{N}_{\text{sdf}}(\mathbf{z}_i, \mathbf{x}_j, \mu_i; \mathbf{w}_{\text{sdf}}) - s_j| + w_{\text{prior}} \|\mathbf{z}_i\|_2^2 + w_{\text{lip}} C(\mathbf{w}_{\text{sdf}}) \right), \quad (11)$$

where w_{lip} is a hyperparameter controlling the strength of the Lipschitz regularization and $\sum_{j=1}^K$ is the averaged sum over the set of indices $\{1, \dots, j, \dots, K\}$.

4.5 Shape model inference step

At inference time, the decoder weights \mathbf{w}_{sdf} are fixed, and we want to find the optimal latent code \mathbf{z}_i for a new shape. While in training we used a zero-mean spherical Gaussian prior, it is now possible to incorporate the empirical distribution of learned latent codes to obtain a more informative prior.

From a Bayesian perspective, observing that the training codes $\{\mathbf{z}_i\}_{i=1}^N$ form an empirical distribution with mean $\boldsymbol{\mu}$ and covariance Σ , the natural choice for the prior at inference is $p(\mathbf{z}) = \mathcal{N}(\boldsymbol{\mu}, \Sigma)$. Substituting this into the posterior distribution (analogous to (5)) and taking the negative log-posterior leads to the Mahalanobis regularization term:

$$\|\mathbf{z}_i\|_M^2 = (\mathbf{z}_i - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{z}_i - \boldsymbol{\mu}). \quad (12)$$

The Maximum a Posteriori (MAP) estimate of the latent code for a single geometry $X_i = \{(\mathbf{x}_j, \mu_i, s_j)\}_{j=1}^K$ is therefore obtained by solving

$$\mathbf{z}_i^* = \arg \min_{\mathbf{z}_i} \left[\sum_{j=1}^K |\mathcal{N}\mathcal{N}_{\text{sdf}}(\mathbf{z}_i, \mathbf{x}_j, \mu_i; \mathbf{w}_{\text{sdf}}) - s_j| + w_{\text{post}} (\mathbf{z}_i - \boldsymbol{\mu})^\top \Sigma^{-1} (\mathbf{z}_i - \boldsymbol{\mu}) \right], \quad (13)$$

where $w_{\text{post}} = \frac{b}{2K}$ balances the reconstruction term against fidelity to the learned latent distribution. The factor $\frac{1}{2}$ comes from the Gaussian log-likelihood, while b is the normalization constant from (8), and K accounts for the averaging over sample points.

We observe that $w_{\text{post}} = \frac{b}{2K}$ decreases with the number of sample points K and increases with the noise scaling factor b . As a consequence, with scarce or noisy observations, the optimization relies more heavily on the learned prior distribution $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$, regularizing the latent code toward the training manifold. On the other hand, with abundant high-quality data, the reconstruction term dominates, allowing the model to accurately fit the observed geometry.

4.6 Shape reconstruction metrics

In order to evaluate the quality of the reconstructed shapes we consider 5 metrics.

Chamfer Distance (CD). The Chamfer distance measures the dissimilarity between two point clouds \mathcal{D}_1 and \mathcal{D}_2 by computing the average nearest-neighbor distance in both directions:

$$\text{CD}(\mathcal{D}_1, \mathcal{D}_2) = \frac{1}{2} \left(\frac{1}{|\mathcal{D}_1|} \sum_{\mathbf{x} \in \mathcal{D}_1} \min_{\mathbf{y} \in \mathcal{D}_2} \|\mathbf{x} - \mathbf{y}\|_2 + \frac{1}{|\mathcal{D}_2|} \sum_{\mathbf{y} \in \mathcal{D}_2} \min_{\mathbf{x} \in \mathcal{D}_1} \|\mathbf{y} - \mathbf{x}\|_2 \right). \quad (14)$$

Normalized Chamfer Distance (CD_{norm}). Since the model is trained on shapes that are normalized in the cube $[-1, 1]^3$, in order to factor out cohort size differences, we also compute the Chamfer distance between reconstruction and normalized target. This provides a scale-invariant measure of shape similarity.

Mean Level Estimate (MLE). To assess how closely the network predicts zero on the actual surface, we compute the Mean Level Estimate as the average absolute SDF prediction across all ground-truth surface vertices and all geometries:

$$\text{MLE} = \sum_{i=1}^{N_{\text{geo}}} \left(\sum_{\mathbf{x} \in \mathbf{X}_i} |\mathcal{N}\mathcal{N}_{\text{sdf}}(\mathbf{z}_i, \mathbf{x}, \mu_i; \mathbf{w}_{\text{sdf}})| \right), \quad (15)$$

where \mathbf{X}_i denotes the set of surface vertices for the i -th geometry. Since ground-truth surface points should have SDF value zero, the MLE indicates the bias of the learned SDF representation across the dataset. Lower MLE values indicate better alignment between the predicted zero-level sets and the true geometry surfaces.

Intersection over Union (IoU). We evaluate the decoder $\mathcal{N}\mathcal{N}_{\text{sdf}}$ on a fine regular grid and classify each grid point as inside (SDF < 0) or outside (SDF ≥ 0) the reconstructed geometry. Similarly, we determine the inside/outside classification for each grid point with respect to the target mesh. The IoU is then defined as:

$$\text{IoU} = \frac{|V_{\text{pred}} \cap V_{\text{target}}|}{|V_{\text{pred}} \cup V_{\text{target}}|}, \quad (16)$$

where V_{pred} and V_{target} denote the sets of grid points classified as inside the predicted and target geometries, respectively.

Dice Coefficient (Dice). The Dice coefficient, also known as the F1 score or Sørensen–Dice coefficient, provides an alternative volumetric overlap measure:

$$\text{Dice} = \frac{2|V_{\text{pred}} \cap V_{\text{target}}|}{|V_{\text{pred}}| + |V_{\text{target}}|}. \quad (17)$$

The Dice coefficient is more sensitive to changes when the overlap is high, making it particularly useful for assessing the reconstruction quality in medical imaging applications, where accurate volume estimation is critical. We use a 3D grid with $96 \times 96 \times 96$ resolution for computation of IoU and Dice.

4.7 Shape generation

Once the decoder \mathcal{N}_{sdf} and latent codes $\{\mathbf{z}_i\}_{i=1}^N$ have been optimized, we can extract explicit mesh representations of the learned shapes.

For each training shape i , we evaluate the conditioned network $\mathcal{N}_{\text{sdf}}(\mathbf{z}_i, \cdot; \mathbf{w}_{\text{sdf}})$ on a dense 3D grid in the domain of interest. The resulting discrete SDF values are then processed using the marching cubes algorithm [28] to extract a triangular mesh corresponding to the zero-level set, which represents the shape boundary $\partial\Omega_i$.

We can exploit the learned latent space to generate new plausible geometries. We sample latent codes from the empirical distribution $\mathcal{N}(\boldsymbol{\mu}, \Sigma)$ of the training codes, where $\boldsymbol{\mu}$ and Σ are the mean and covariance calculated from $\{\mathbf{z}_i\}_{i=1}^N$. Each sampled code $\mathbf{z}_{\text{new}} \sim \mathcal{N}(\boldsymbol{\mu}, \Sigma)$ is used to condition the decoder, which is then evaluated on the same 3D grid and processed via marching cubes to obtain a mesh. We also sample a scaling factor μ_{new} from a uniform distribution obtained by taking the minimum and maximum scaling factors used during training.

Under the assumption that the latent space is sufficiently smooth and \mathcal{N}_{sdf} exhibits regularity with respect to the latent code (enforced by the Lipschitz regularization in (11)), the sampled codes produce realistic novel geometries that interpolate between training examples. These synthetic shapes can be used for data augmentation in downstream tasks, such as surrogate modeling of ventricles.

4.8 Architecture/Training/Specifics

Model architecture. The decoder \mathcal{N}_{sdf} is a fully connected network with 6 hidden layers of 64 neurons each. We use the Exponential Linear Unit (ELU) activation function instead of ReLU to ensure C^1 continuity of the learned SDF. This smoothness is necessary for downstream finite element simulations, as ReLU activations introduce sharp angles in the reconstructed zero-level set that can cause numerical instabilities.

Following [36, 40], we set the latent space dimension to $\dim(\mathcal{Z}) = 16$, when working with real patient anatomies.

Idealized geometries are generated by varying a set of 3 parameters: the ventricular long axis l , the maximum diameter d and transmural thickness w . When training the reconstruction model on them we choose to keep the same architectural structure used for the real setting and thus, knowing that the geometrical scaling value μ_i correlates strongly with l , we set $\dim(\mathcal{Z}) = 2$. In this way the scaling factor acts as third dimension of the latent space. Latent codes are initialized by sampling from $\mathcal{N}(\mathbf{0}, 10^{-2}\mathbf{I})$.

Hyperparameters. Following (11), we set the prior weight $w_{\text{prior}} = 10^{-4}$ and the Lipschitz weight $w_{\text{lip}} = 10^{-6}$ when training the model on patient anatomies and $w_{\text{lip}} = 10^{-9}$, when training on idealized geometries. The prior weight encourages latent codes to remain within a bounded region, while the Lipschitz regularization promotes smoothness in the latent space for stable interpolation. The posterior weight for the inference loss is set to $w_{\text{post}} = 6.25 \times 10^{-5}$.

Training and inference. The $\mathcal{NN}_{\text{sdf}}$ model is trained for up to 12000 epochs with ADAM optimizer and decaying learning rate starting from 5×10^{-3} . Training is automatically stopped when the root mean squared error computed on the validation points stagnates for more than 200 epochs.

At inference time the weights of the network \mathbf{w}_{sdf} are frozen and only the latent code associated with the geometry of interest is trained via the procedure indicated in section 4.5. We train each inference code for 2000 epochs with ADAM optimizer.

4.9 PCA encoding fitting

1-to-1 correspondence via universal ventricular coordinates Dimensionality reduction via principal component analysis (PCA) requires establishing a standardized coordinate system where 1-to-1 correspondence between different patient anatomies can be achieved. We employ UVCs [22], a cylindrical-like parameterization (ζ, ϱ, ϕ) , where $\zeta \in [0, 1]$ represents the apicobasal direction, $\varrho \in [0, 1]$ denotes the transmural direction, and $\phi \in [-\pi, \pi]$ corresponds to the rotational coordinate. In this space we define a uniform sampling grid with dimension $N_{\text{features}} = N_{\zeta} \times N_{\varrho} \times N_{\phi}$, creating a regular lattice of query points.

The reparametrization works as follows. The mesh vertices of each mesh, represented in Cartesian coordinates (x, y, z) , are associated with their corresponding UVCs. Spatial interpolation is performed to map the irregular patient-specific mesh onto the uniform UVC grid. Using a nearest-neighbour interpolation method, we obtain three interpolated fields $\hat{x}(\zeta, \varrho, \phi)$, $\hat{y}(\zeta, \varrho, \phi)$, and $\hat{z}(\zeta, \varrho, \phi)$ each evaluated at the grid points. At this point, we filter the grid to retain only surface points, corresponding to $z = 1$ (base) or $\varrho \in \{0, 1\}$ (endocardium and epicardium). The coordinates are concatenated and flattened to obtain a final vector representation of length N_{features} . The vector representations of all train geometries are concatenated to form the snapshot matrix $\mathbf{S}_{\text{train}}$ with shape $3N_{\text{features}} \times N_{\text{train}}$.

PCA via singular value decomposition The snapshot matrix is first centered by subtracting the mean feature vector over all train patients:

$$\tilde{\mathbf{S}}_{\text{train}} = \mathbf{S}_{\text{train}} - \bar{\mathbf{s}}\mathbf{1}^T \quad \bar{\mathbf{s}} = \frac{1}{M} \sum_{i=1}^{N_{\text{train}}} \mathbf{s}_i. \quad (18)$$

Then, the centered snapshot matrix is decomposed via truncated SVD:

$$\tilde{\mathbf{S}}_{\text{train}} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^T, \quad (19)$$

where $\mathbf{U} \in \mathbb{R}^{3N_{\text{features}} \times r}$ contains the left singular vectors (spatial modes), $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$ is the diagonal matrix of singular values σ_i , $\mathbf{V}^T \in \mathbb{R}^{r \times N_{\text{train}}}$ contains the right singular vectors, and $r = \min(3N_{\text{features}}, N_{\text{train}})$ is the rank of the decomposition.

The cumulative explained variance ratio is computed as:

$$\eta_k = \frac{\sum_{i=1}^k \sigma_i^2}{\sum_{i=1}^r \sigma_i^2}. \quad (20)$$

This metric quantifies the proportion of geometric variability captured by the first k principal modes. In our work we select $k = r - 1$ in order to retain as much geometrical information as possible, since even small variations could be relevant for downstream tasks. The final mode is excluded to avoid numerical artifacts.

The reduced basis \mathbf{U}_k is constructed by selecting the first k left singular vectors. Each patient geometry is encoded into a low-dimensional latent representation via projection onto the principal component basis:

$$\mathbf{z}_i = \mathbf{U}_k^T \tilde{\mathbf{s}}_i \in \mathbb{R}^k, \quad (21)$$

where $\tilde{\mathbf{s}}_i$ is the centered feature vector. The original geometry can be approximately reconstructed as:

$$\hat{\mathbf{s}}_i = \bar{\mathbf{s}} + \mathbf{U}_k \mathbf{c}_i, \quad (22)$$

where \mathbf{c}_i is the shape code for patient i . $\hat{\mathbf{s}}_i$ is then reshaped in the original tridimensional space.

4.10 PCA encoding inference

Since UVCs cannot be precomputed on point cloud data, shape codes need to be inferred through an optimization procedure. Given a set of observed surface points $\mathbf{X} \in \mathbb{R}^{N_{\text{obs}} \times 3}$ sampled from the test geometry, we look for the optimal shape code $\mathbf{z}^* \in \mathbb{R}^k$ that best reconstructs the observed data, by exploiting the PCA basis introduced in Section 4.9. The reconstruction model is defined as:

$$\hat{\mathbf{s}}(\mathbf{z}) = \bar{\mathbf{s}} + \mathbf{U}_k \mathbf{z}. \quad (23)$$

The flattened reconstruction $\hat{\mathbf{s}}(\mathbf{z})$ is reshaped to obtain the surface point cloud $\hat{\mathbf{s}}_{\text{surf}}(\mathbf{z}) \in \mathbb{R}^{N_{\text{features}} \times 3}$.

The objective function for the optimization procedure is the CD between the surface points and the geometry reconstructed by the PCA model:

$$\mathbf{z}^* = \arg \min_{\mathbf{z}} \text{CD}(\mathbf{X}, \hat{\mathbf{s}}_{\text{surf}}(\mathbf{z})). \quad (24)$$

The optimization is performed using the Adam optimizer with learning rate $\alpha = 5 \times 10^{-2}$ for 2000 epochs. At each step of the optimization, gradients are computed via automatic differentiation.

4.11 Forward PDE problem

Finite strain mechanics We endeavor to infer the deformation due to diastolic filling of the LV at defined cavity pressure \hat{p} . Therefore, ground truth data is generated by solving the following quasi-static finite strain solid mechanics boundary value problem on each of the (idealized and patient-specific) cardiac geometries. The strong form of the boundary value problem defined on an LV geometry Ω_0 , cf. Fig. 1e for an exemplary patient-specific model, reads as follows:

$$\begin{aligned} \mathbf{0} &= \nabla_0 \cdot (\mathbf{F}\mathbf{S}) && \text{in } \Omega_0, \\ (\mathbf{F}\mathbf{S})\mathbf{n}_0 &= -\hat{p}J\mathbf{F}^{-\text{T}}\mathbf{n}_0 && \text{on } \Gamma_0^N, \\ (\mathbf{F}\mathbf{S})\mathbf{n}_0 &= k_{\text{b}}^{\parallel}(\mathbf{I} - \mathbf{n}_0 \otimes \mathbf{n}_0)\mathbf{u} \\ &\quad + k_{\text{b}}^{\perp}(\mathbf{n}_0 \otimes \mathbf{n}_0)\mathbf{u} && \text{on } \Gamma_0^{R,\text{b}}, \\ (\mathbf{F}\mathbf{S})\mathbf{n}_0 &= k_{\text{e}}^{\parallel}(\mathbf{I} - \mathbf{n}_0 \otimes \mathbf{n}_0)\mathbf{u} \\ &\quad + k_{\text{e}}^{\perp}(\mathbf{n}_0 \otimes \mathbf{n}_0)\mathbf{u} && \text{on } \Gamma_0^{R,\text{e}}, \end{aligned} \quad (25)$$

where $\mathbf{F} = \mathbf{I} + \nabla_0 \mathbf{u}$ is the deformation gradient, $J = \det \mathbf{F}$ its determinant, and \mathbf{u} the unknown displacement field from the (stress-free) reference configuration Ω_0 to a deformed configuration Ω . Further, Γ_0^N and $\Gamma_0^{R,(\cdot)}$ are Neumann and Robin boundaries, respectively, and \mathbf{n}_0 is the outward normal with respect to the reference frame. The constitutive relation for the 2nd Piola-Kirchhoff stress tensor \mathbf{S} is of anisotropic hyperelastic type:

$$\mathbf{S} = 2 \frac{\partial \Psi(\mathbf{C})}{\partial \mathbf{C}}, \quad (26)$$

where $\mathbf{C} = \mathbf{F}^{\text{T}}\mathbf{F}$ is the right Cauchy-Green deformation tensor and Ψ the strain energy density function, here given by the anisotropic Holzapfel-Ogden model [41]:

$$\Psi = \frac{a_0}{2b_0} \left(e^{b_0(\bar{I}_C - 3)} - 1 \right) + \sum_{i \in \{f,s\}} \frac{a_i}{2b_i} \left(e^{b_i(I_{4,i} - 1)^2} - 1 \right) + \frac{a_{fs}}{2b_{fs}} \left(e^{b_{fs}I_8^2} - 1 \right) + \frac{\kappa}{2}(J - 1)^2, \quad (27)$$

where $\bar{I}_C = J^{-2/3} \text{tr} \mathbf{C}$ is the isochoric first invariant of the Cauchy-Green tensor, and $I_{4,f} = \mathbf{f}_0 \cdot \mathbf{C} \mathbf{f}_0$ and $I_{4,s} = \mathbf{s}_0 \cdot \mathbf{C} \mathbf{s}_0$ are invariants defining the squared stretches in fiber and sheet directions \mathbf{f}_0 and \mathbf{s}_0 , respectively, while $I_8 = \mathbf{f}_0 \cdot \mathbf{C} \mathbf{s}_0$ describes the shear stretch between the two (initially perpendicular) directions. Fiber and sheet directions are generated using rule-based approaches [42, 43]. Figure 1e shows the transmurally varying fiber field \mathbf{f}_0 along with its angle α with respect to the ventricle's circumference. Finally, κ

represents a bulk modulus, chosen such that deformation is nearly incompressible.

We solve problem Eq. (25) using the finite element method, seeking displacements $\mathbf{u} \in \mathcal{V}^{D,h}$ such that

$$\begin{aligned} & \int_{\Omega_0} \mathbf{S}(\mathbf{C}(\mathbf{u})) : \frac{1}{2} \delta \mathbf{C}(\delta \mathbf{u}) \, dV + \int_{\Gamma_0^N} \hat{p} J(\mathbf{u}) \mathbf{F}(\mathbf{u})^{-T} \mathbf{n}_0 \cdot \delta \mathbf{u} \, dA + \\ & \int_{\Gamma_0^{R,b}} [k_b^{\parallel} (\mathbf{I} - \mathbf{n}_0 \otimes \mathbf{n}_0) \mathbf{u} + k_b^{\perp} (\mathbf{n}_0 \otimes \mathbf{n}_0) \mathbf{u}] \cdot \delta \mathbf{u} \, dA + \int_{\Gamma_0^{R,e}} [k_e^{\parallel} (\mathbf{I} - \mathbf{n}_0 \otimes \mathbf{n}_0) \mathbf{u} + k_e^{\perp} (\mathbf{n}_0 \otimes \mathbf{n}_0) \mathbf{u}] \cdot \delta \mathbf{u} \, dA = 0 \end{aligned} \quad (28)$$

for all displacement test functions $\delta \mathbf{u} \in \mathcal{V}^h$, where $\mathcal{V}^{D,h}$ and \mathcal{V}^h are suitable trial and test spaces (identical up to potential Dirichlet conditions included in the trial space). Throughout this study, first-order piecewise polynomial spaces on tetrahedral meshes are used. Parameters of the constitutive model Eq. (27), and the boundary conditions in Eq. (25), are given in Tab. 6.

Table 6: Parameters of the constitutive model Eq. (27) as well as boundary conditions in Eq. (25).

a_0 [kPa]	b_0 [-]	a_f [kPa]	b_f [-]	a_s [kPa]	b_s [-]	a_{fs} [kPa]	b_{fs} [-]	κ [kPa]
0.059	8.023	18.472	16.026	2.481	11.120	0.216	11.436	10^2
k_b^{\perp} [kPa]	k_b^{\parallel} [kPa]	k_e^{\perp} [kPa]	k_e^{\parallel} [kPa]	\hat{p} [kPa]				
$2 \cdot 10^3$	10^2	10	1	2.67				

Universal ventricular coordinates (UVC) We compute the continuous fields of UVCs – the apicobasal, transmural, and circumferential coordinates ζ , ϱ , and ϕ , respectively – by solving three Laplace problems [22] on Ω_0 (cf. Fig. 1e for boundary labels). For the apicobasal coordinate, we solve:

$$\begin{aligned} -\nabla \cdot (\nabla \zeta) &= 0 && \text{in } \Omega_0, \\ \zeta &= 0 && \text{on } P_{\text{apex}}, \\ \zeta &= 1 && \text{on } \Gamma^{R,b}, \end{aligned} \quad (29)$$

where P_{apex} is the epicardial node of the left ventricular apex. Note that here, we do not perform geodesic distance sampling of ζ , in contrast to [22].

The transmural coordinate ϱ varies between 0 and 1 (between endo- and epicardium), hence it is a solution to:

$$\begin{aligned} -\nabla \cdot (\nabla \varrho) &= 0 && \text{in } \Omega_0, \\ \varrho &= 1 && \text{on } \Gamma_0^{R,e}, \\ \varrho &= 0 && \text{on } \Gamma_0^N. \end{aligned} \quad (30)$$

Finally, the circumferential coordinate ϕ varies between $-\pi$ and π and is solution of the following Laplace problem:

$$\begin{aligned} -\nabla \cdot (\nabla \phi) &= 0 && \text{in } \Omega_0, \\ \phi &= \pi && \text{on } \Gamma_0^c, \\ \phi &= -\pi && \text{on } \Gamma_0^{c'}. \end{aligned} \quad (31)$$

Therein, Γ_0^c is a coronary cut plane spanned by the apex, the center of the base plane, and the apex point projected onto the base plane, and $\Gamma_0^{c'}$ is a plane rotated about an infinitesimal angle $d\phi$. Problems Eq. (29)–(31) are discretized and solved using piecewise linear finite elements. The UVC fields are depicted in Fig. 1h.

4.12 Surrogate model

In order to predict the deformation field of ventricles under diastolic pressure loading conditions, a fully connected neural network (FCNN), or multi-layer perceptron (MLP), for learning the discrete field of ventricular displacements on the nodal points of a LV geometry is introduced. The FCNN is trained on the cohorts of (idealized or patient-specific) cardiac geometries, and its input layer consists of (min-max normalized versions) of the following quantities:

- $\mathbf{x} \in \mathbb{R}^{n \times 3}$: the discrete vector of spatial cartesian coordinates;
- $\boldsymbol{\xi} \in \mathbb{R}^{n \times 3}$: the discrete vector of UVCs;
- $\boldsymbol{\mu}_g(\Omega_0) \in \mathbb{R}^m$: a shape code encoding geometric variability (e.g., PCA-SM, SDF-SM);
- $\boldsymbol{\mu}_p \in \mathbb{R}^{n_p}$: a set representative of the parameterization of the physics model (here invariant).

The output layer of the FCNN is the discrete vector of displacements \mathbf{u} . The architecture of the network is depicted Fig. 1j. The discrete set of cartesian coordinates \mathbf{x} , UVCs $\boldsymbol{\xi}$, and displacements \mathbf{u} are defined according to:

$$\mathbf{x}_i = \begin{bmatrix} \mathbf{x}_1 \\ \vdots \\ \mathbf{x}_n \end{bmatrix}_i = \begin{bmatrix} x_1 & y_1 & z_1 \\ \vdots \\ x_n & y_n & z_n \end{bmatrix}_i, \quad \boldsymbol{\xi}_i = \begin{bmatrix} \boldsymbol{\xi}_1 \\ \vdots \\ \boldsymbol{\xi}_n \end{bmatrix}_i = \begin{bmatrix} \zeta_1 & \varrho_1 & \sin \phi_1 & \cos \phi_1 \\ \vdots \\ \zeta_n & \varrho_n & \sin \phi_n & \cos \phi_n \end{bmatrix}_i, \quad \mathbf{u}_i = \begin{bmatrix} \mathbf{u}_1 \\ \vdots \\ \mathbf{u}_n \end{bmatrix}_i = \begin{bmatrix} u_{x,1} & u_{y,1} & u_{z,1} \\ \vdots \\ u_{x,n} & u_{y,n} & u_{z,n} \end{bmatrix}_i, \quad (32)$$

where the index n corresponds to the number of discrete points (nodes) of the i -th geometry. Further, x, y, z denote the axes of the cartesian reference frame, and $\zeta \in [0, 1]$, $\rho \in [0, 1]$, and $\phi \in [-\pi, \pi]$ are the universal apicobasal, transmural, and circumferential ventricular coordinates. Due to the sharp discontinuity of ϕ from $-\pi$ to π , we prefer to use continuous representations $\sin \phi$ and $\cos \phi$, instead.

For scaling and performance purposes, we perform a min-max normalization of the input data as follows. Minimum and maximum values of geometry and displacement are computed as follows:

$$x_{\min} = \min \left\{ \min \begin{bmatrix} \|\mathbf{x}_1\| \\ \vdots \\ \|\mathbf{x}_n\| \end{bmatrix}_1, \dots, \min \begin{bmatrix} \|\mathbf{x}_1\| \\ \vdots \\ \|\mathbf{x}_n\| \end{bmatrix}_N \right\}, \quad x_{\max} = \max \left\{ \max \begin{bmatrix} \|\mathbf{x}_1\| \\ \vdots \\ \|\mathbf{x}_n\| \end{bmatrix}_1, \dots, \max \begin{bmatrix} \|\mathbf{x}_1\| \\ \vdots \\ \|\mathbf{x}_n\| \end{bmatrix}_N \right\}, \quad (33)$$

$$u_{\min} = \min \left\{ \min \begin{bmatrix} \|\mathbf{u}_1\| \\ \vdots \\ \|\mathbf{u}_n\| \end{bmatrix}_1, \dots, \min \begin{bmatrix} \|\mathbf{u}_1\| \\ \vdots \\ \|\mathbf{u}_n\| \end{bmatrix}_N \right\}, \quad u_{\max} = \max \left\{ \max \begin{bmatrix} \|\mathbf{u}_1\| \\ \vdots \\ \|\mathbf{u}_n\| \end{bmatrix}_1, \dots, \max \begin{bmatrix} \|\mathbf{u}_1\| \\ \vdots \\ \|\mathbf{u}_n\| \end{bmatrix}_N \right\}, \quad (34)$$

and the min-max normalized data for the i -th observation is

$$\hat{\mathbf{x}}_i = \frac{1}{x_{\max} - x_{\min}} \begin{bmatrix} \mathbf{x}_1 - x_{\min} \mathbf{1} \\ \vdots \\ \mathbf{x}_n - x_{\min} \mathbf{1} \end{bmatrix}_i = \begin{bmatrix} \hat{\mathbf{x}}_1 \\ \vdots \\ \hat{\mathbf{x}}_n \end{bmatrix}_i, \quad \hat{\mathbf{u}}_i = \frac{1}{u_{\max} - u_{\min}} \begin{bmatrix} \mathbf{u}_1 - u_{\min} \mathbf{1} \\ \vdots \\ \mathbf{u}_n - u_{\min} \mathbf{1} \end{bmatrix}_i = \begin{bmatrix} \hat{\mathbf{u}}_1 \\ \vdots \\ \hat{\mathbf{u}}_n \end{bmatrix}_i, \quad (35)$$

so that values range $\in [0, 1]$. The UVC fields are min-max normalized per component¹ as follows:

$$\hat{\boldsymbol{\xi}}_i = \frac{\boldsymbol{\xi}_i - \boldsymbol{\xi}_{\min}}{\boldsymbol{\xi}_{\max} - \boldsymbol{\xi}_{\min}}, \quad (36)$$

¹Note that $\zeta, \varrho \in [0, 1]$ already, hence these components are not affected by min-max normalization.

with the vector of minimum and maximum UVCs being

$$\boldsymbol{\xi}_{\min} = [\min \{(\boldsymbol{\xi}_1)_1, \dots, (\boldsymbol{\xi}_N)_1\}, \dots, \min \{(\boldsymbol{\xi}_1)_4, \dots, (\boldsymbol{\xi}_N)_4\}]^T, \quad (37)$$

$$\boldsymbol{\xi}_{\max} = [\max \{(\boldsymbol{\xi}_1)_1, \dots, (\boldsymbol{\xi}_N)_1\}, \dots, \max \{(\boldsymbol{\xi}_1)_4, \dots, (\boldsymbol{\xi}_N)_4\}]^T. \quad (38)$$

Analogously, the min-max normalized vectors of shape codes for the i -th geometry and parameters for the k -th parametric configuration read

$$\widehat{\boldsymbol{\mu}}_{g,i} = \frac{\boldsymbol{\mu}_{g,i} - \mu_{g,\min} \mathbf{1}}{\mu_{g,\max} - \mu_{g,\min}}, \quad \widehat{\boldsymbol{\mu}}_{p,k} = \frac{\boldsymbol{\mu}_{p,k} - \boldsymbol{\mu}_{p,\min}}{\boldsymbol{\mu}_{p,\max} - \boldsymbol{\mu}_{p,\min}}, \quad (39)$$

with the minimum and maximum shape code across geometries being $\mu_{g,\min} = \min \{ \min \boldsymbol{\mu}_{g,1}, \dots, \min \boldsymbol{\mu}_{g,N} \}$ and $\mu_{g,\max} = \max \{ \max \boldsymbol{\mu}_{g,1}, \dots, \max \boldsymbol{\mu}_{g,N} \}$. The vectors of minimum and maximum parameters read

$$\boldsymbol{\mu}_{p,\min} = \left[\min \{ (\boldsymbol{\mu}_{p,1})_1, \dots, (\boldsymbol{\mu}_{p,N_p})_1 \}, \dots, \min \{ (\boldsymbol{\mu}_{p,1})_{n_p}, \dots, (\boldsymbol{\mu}_{p,N_p})_{n_p} \} \right]^T, \quad (40)$$

$$\boldsymbol{\mu}_{p,\max} = \left[\max \{ (\boldsymbol{\mu}_{p,1})_1, \dots, (\boldsymbol{\mu}_{p,N_p})_1 \}, \dots, \max \{ (\boldsymbol{\mu}_{p,1})_{n_p}, \dots, (\boldsymbol{\mu}_{p,N_p})_{n_p} \} \right]^T, \quad (41)$$

with the number of parameters n_p and the number of parametric configurations N_p (here, we only consider $N_p = 1$).

The FCNN outputs the (min-max normalized) displacement $\widehat{\mathbf{u}}^*$ at the set of query points $\widehat{\mathbf{x}}$,

$$\mathbf{u}^*(\widehat{\mathbf{x}}; \widehat{\boldsymbol{\mu}}_p, \Omega_0) \simeq \mathcal{N}_{\text{phys}}(\widehat{\mathbf{x}}, \widehat{\boldsymbol{\xi}}, \widehat{\boldsymbol{\mu}}_g(\Omega_0), \widehat{\boldsymbol{\mu}}_p; \mathbf{w}_{\text{phys}}), \quad (42)$$

with the trainable parameters (weights and biases) \mathbf{w} . We seek an optimal set of trainable parameters

$$\mathbf{w}^* = \arg \min_{\mathbf{w}} \mathcal{J}(\mathbf{w}) \quad (43)$$

being the minimizer of the loss function

$$\mathcal{J}(\mathbf{w}) = \mathcal{J}_u(\mathbf{w}) + \lambda_s \mathcal{J}_{\nabla u}(\mathbf{w}), \quad (44)$$

which here is decomposed into the mean squared error loss relating to the displacement,

$$\mathcal{J}_u(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \left(g_i \frac{1}{n} \sum_{j=1}^n \|\widehat{\mathbf{u}}_j^* - \widehat{\mathbf{u}}_j\|^2 \right), \quad (45)$$

and one that penalizes displacement gradients, i.e. strains:

$$\mathcal{J}_{\nabla u}(\mathbf{w}) = \frac{1}{N} \sum_{i=1}^N \left(g_i \frac{1}{n_e} \sum_{j=1}^{n_e} \|\mathbf{E}(\mathbf{u}_j^*) - \mathbf{E}(\mathbf{u}_j)\|^2 \right), \quad (46)$$

where g_i are weighting factors. Here, \mathbf{E} is the vector of (elemental) Green-Lagrange strains in Voigt notation:

$$\mathbf{E} = [E_{11} \quad E_{22} \quad E_{33} \quad 2E_{12} \quad 2E_{13} \quad 2E_{23}]^T, \quad (47)$$

n_e is the number of cells per model i , and N is the total number of observations, i.e. geometries. The regularization parameter λ_s in Eq. (44) controls the contribution of gradient information that adds to the loss. Further, note that the strain is computed with the de-normalized displacement.

The displacement gradient here is computed using the linear tetrahedral finite element space, with the four nodal coordinates and displacements for one tetrahedron being

$$\mathbf{x}^{(e)} = \begin{bmatrix} x^1 & y^1 & z^1 \\ x^2 & y^2 & z^2 \\ x^3 & y^3 & z^3 \\ x^4 & y^4 & z^4 \end{bmatrix} \quad \text{and} \quad \mathbf{u}^{(e)} = \begin{bmatrix} u_x^1 & u_y^1 & u_z^1 \\ u_x^2 & u_y^2 & u_z^2 \\ u_x^3 & u_y^3 & u_z^3 \\ u_x^4 & u_y^4 & u_z^4 \end{bmatrix}, \quad (48)$$

respectively. The vector of nodal finite element shape functions and its derivatives in finite element reference space read

$$\mathbf{N} = \begin{bmatrix} N^1 \\ N^2 \\ N^3 \\ N^4 \end{bmatrix} = \begin{bmatrix} 1 - \xi_1 - \xi_2 - \xi_3 \\ \xi_1 \\ \xi_2 \\ \xi_3 \end{bmatrix} \quad \text{and} \quad \mathbf{N}_{,\xi} = \begin{bmatrix} -1 & -1 & -1 \\ 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{bmatrix}, \quad (49)$$

yielding the element-centered displacement gradient

$$\nabla_0 \mathbf{u}^{(e)} = \mathbf{u}^{(e)\text{T}} \mathbf{N}_{,x} = \mathbf{u}^{(e)\text{T}} \mathbf{N}_{,\xi} \mathbf{J}^{-1} = \mathbf{u}^{(e)\text{T}} \mathbf{N}_{,\xi} \left(\mathbf{x}^{(e)\text{T}} \mathbf{N}_{,\xi} \right)^{-1}. \quad (50)$$

Finally, the element Green-Lagrange strain is

$$\mathbf{E}^{(e)} = \frac{1}{2} \left(\mathbf{F}^{(e)\text{T}} \mathbf{F}^{(e)} - \mathbf{I} \right), \quad (51)$$

with the deformation gradient $\mathbf{F}^{(e)} = \mathbf{I} + \nabla_0 \mathbf{u}^{(e)}$.

4.13 Surrogate training

Shape-informed inference of the displacement field from end-diastolic LV inflation is performed on two different geometry cohorts. We aim at quantifying the different types of shape encoding on the quality of the inferred displacement field, test the influence on the number of query points, and assess the impact of enriching the loss function by gradient information (strains).

First, we perform tests on the idealized cohort (Sec. 4.1) of 512 geometries, cf. Fig. 1b. The set is once randomly split into 230 training, 179 validation, and 103 test geometries. The same split is kept throughout all the studies on that cohort. Second, we investigate the performance of shape encoding on the augmented cohort of 44 real geometries (24 heart failure, 20 healthy), Fig. 1c, and 976 generated models (Sec. 4.7), cf. Fig. 1d.² In order to balance loss contributions from real vs. generated models, we weigh the loss term of real geometries by $g_i = 3$ compared to generated ones ($g_i = 1$) in Eq. (45). We compare this augmented dataset to using the real-only dataset of 44 models and show the benefits of data augmentation. Patient-specific and generated geometries are randomly split into 600 and 30 training, 300 and 5 validation, as well as 76 and 9 test geometries, respectively. Again, the same split is used on all studies with these cohorts.

Hyperparameter tuning is initially carried out for the idealized models and for the augmented generated plus real cohort individually. We assume a fixed number of three layers between input and output layer. We perform 500 epochs and evaluate the error on the validation set, with 5 startup trials and 50 warm-up steps, and seek the number of neurons $n_{\ell_1} \in [10, 50]$, $n_{\ell_2} \in [5, 40]$, and $n_{\ell_3} \in [3, 20]$ in layer 1, 2, and 3, respectively. An optimal neural network architecture for the idealized cohort of $n_{\ell_1} = 50$, $n_{\ell_2} = 39$, $n_{\ell_3} = 14$ is identified. For the augmented real and generated geometries, we identify the optimal numbers of neurons as $n_{\ell_1} = 50$, $n_{\ell_2} = 29$, and $n_{\ell_3} = 20$.

We use hyperbolic tangent (tanh) activation functions for each of the neurons, and training is performed using a BFGS optimizer with a maximum number of $M = 10000$ epochs. For memory and efficiency reasons, we randomly subsample $n = 15000$ spatial points (nodes) per geometry (if not indicated otherwise) and use the same subset of points for repeated, differently configured runs to achieve optimal comparability.

4.14 Software

Data preprocessing for the shape modeling step was performed using *SciPy* [44] and *scikit-learn* [45]. Geometry creation of the idealized models is done in *FreeCAD* [46]. Geometry reconstructions are carried out using *vtk* [47]. All meshing is performed in *gmsh* [48]. The finite element solution of problem Eq. (28) is solved

²Out of 1000 generated heart models, 24 disqualified for meshing and model generation due to degenerate geometric features. See Appendix Fig. 18 for some exemplary geometries which had to be discarded.

using `lifex` [49]. The shape model $\mathcal{NN}_{\text{sdf}}$ and the mechanics surrogate model $\mathcal{NN}_{\text{phys}}$ were implemented in `JAX` [50] and `Flax` [51], and trained using `Optax` [52]. Hyperparameter tuning for $\mathcal{NN}_{\text{phys}}$ is carried out using `Optuna` [53]. Post-processing and mesh extraction were performed using `scikit-image` [54] and `PyVista` [55].

Acknowledgements

The present research has received support from the project FIS, MUR, Italy 2025-2028, Project code: FIS-2023-02228, CUP: D53C24005440001, "SYNERGIZE: Synergizing Numerical Methods and Machine Learning for a new generation of computational models". The present research has been supported by the project PRIN2022, MUR, Italy, 2023-2025, P2022N5ZNP "SIDDMs: shape-informed data-driven models for parametrized PDEs, with application to computational cardiology", funded by the European Union (Next Generation EU, Mission 4 Component 2). D.C., M.H., F.B., S.Pagani, and F.R. acknowledge the grant Dipartimento di Eccellenza 2023-2027, funded by MUR, Italy. D.C., F.B., S.Pagani, S.Pezzuto, and F.R. are members of GNCS, "Gruppo Nazionale per il Calcolo Scientifico" (National Group for Scientific Computing) of INdAM (Istituto Nazionale di Alta Matematica). F.B. acknowledges the "INdAM - GNCS Project", codice CUP E53C24001950001.

Appendices

A Additional Materials

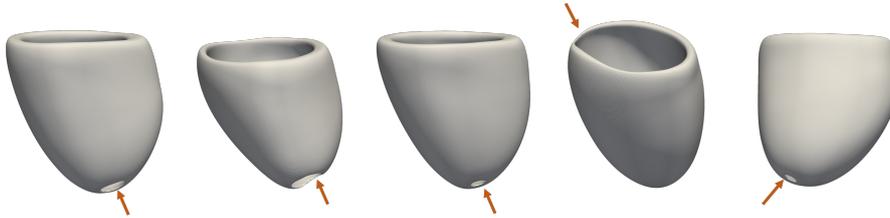


Figure 18: Examples of five discarded synthetic geometries (samples 6, 39, 184, 280, 361) exhibiting artifacts near the apex or base.

Figure 18 shows 5 examples of the discarded synthetic geometries. Out of 1000, 24 geometries were discarded mostly due to small holes close to the apex or degeneration at the cut plane level.

Figure 19 reports the mean geometry and the deformation effect of the first 6 out of 34 PCA modes. The first modes describe mainly the length of long of the long-axis, the orientation of the ventricles and the different behaviors close to the apex.

References

- [1] A. Quarteroni, ed., *Modeling the Heart and the Circulatory System*. Springer Cham Heidelberg New York Dordrecht London, MS&A–Modeling, Simulation & Applications (Vol. 14) ed., 2015.
- [2] E. L. Schwarz, L. Pegolotti, M. R. Pfaller, and A. L. Marsden, "Beyond CFD: Emerging methodologies for predictive simulation in cardiovascular health and disease," *Biophysics Reviews (Melville)*, vol. 4, no. 1, p. 011301, 2023.

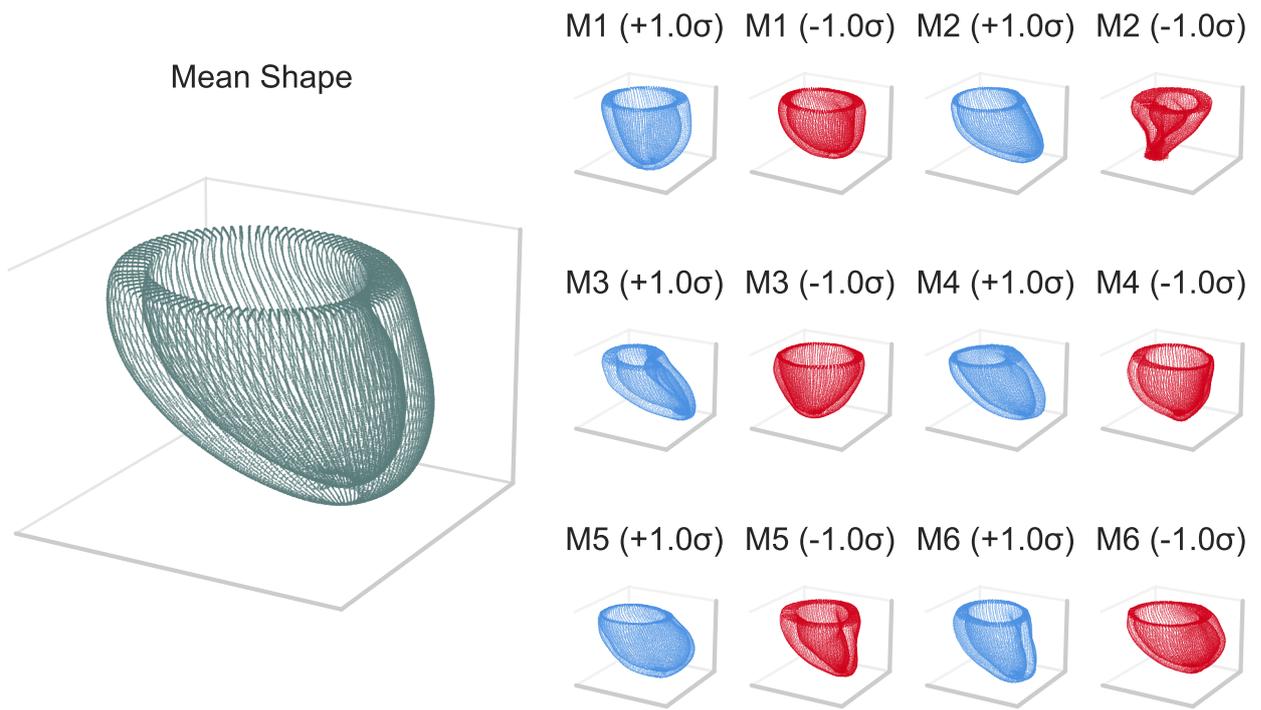


Figure 19: Mean shape and first 6 principal modes of variation obtained via PCA. The mean anatomy is shown on the left, while the first PCA modes are visualized on the right by displaying the shapes obtained by adding and subtracting σ along each mode.

- [3] M. Hirschvogel, M. Balmus, M. Bonini, and D. Nordsletten, “Fluid-reduced-solid interaction (FrSI): Physics- and projection-based model reduction for cardiovascular applications,” *Journal of Computational Physics*, vol. 506, p. 112921, 2024.
- [4] L. Cicci, S. Fresca, S. Pagani, A. Manzoni, and A. Quarteroni, “Projection-based reduced order models for parameterized nonlinear time-dependent problems arising in cardiac mechanics,” *Mathematics in Engineering*, vol. 5, no. 2, pp. 1–38, 2022.
- [5] M. R. Pfaller, M. C. Varona, C. B. J. Lang, and W. A. Wall, “Using parametric model order reduction for inverse analysis of large nonlinear cardiac simulations,” *International Journal for Numerical Methods in Biomedical Engineering*, vol. 36, no. 4, p. e3320, 2020.
- [6] C. Farhat, P. Avery, T. Chapman, and J. Cortial, “Dimensional reduction of nonlinear finite element dynamic models with finite rotations and energy-based mesh sampling and weighting for computational efficiency,” *International Journal for Numerical Methods in Engineering*, vol. 98, no. 9, pp. 625–662, 2014.
- [7] M. Boulakia, E. Schenone, and J.-F. Gerbeau, “Reduced-order modeling for cardiac electrophysiology. Application to parameter identification,” *International Journal for Numerical Methods in Biomedical Engineering*, vol. 28, no. 6–7, pp. 727–744, 2012.
- [8] F. Regazzoni, M. Salvador, P. C. Africa, M. Fedele, L. Dede’, and A. Quarteroni, “A cardiac electromechanical model coupled with a lumped-parameter model for closed-loop blood circulation,” *Journal of Computational Physics*, vol. 457, p. 111083, 2022.
- [9] A. Quarteroni, A. Veneziani, and C. Vergara, “Geometric multiscale modeling of the cardiovascular system, between theory and practice,” *Computer Methods in Applied Mechanics and Engineering*, vol. 302, pp. 193–252, 2016.
- [10] C. Greif and K. Urban, “Decay of the kolmogorov n-width for wave problems,” *Applied Mathematics Letters*, vol. 96, pp. 216–222, 2019.
- [11] S. Pagani, A. Manzoni, and A. Quarteroni, “Numerical approximation of parametrized problems in cardiac electrophysiology by a local reduced basis method,” *Computer Methods in Applied Mechanics and Engineering*, vol. 340, pp. 530–558, 2018.
- [12] K. Lee and K. T. Carlberg, “Model reduction of dynamical systems on nonlinear manifolds using deep convolutional autoencoders,” *Journal of Computational Physics*, vol. 404, p. 108973, 2020.
- [13] F. Regazzoni, S. Pagani, M. Salvador, L. Dede’, and A. Quarteroni, “Learning the intrinsic dynamics of spatio-temporal processes through Latent Dynamics Networks,” *Nature Communications*, vol. 15, no. 1834, 2024.
- [14] P. Du, M. H. Parikh, X. Fan, X.-Y. Liu, and J.-X. Wang, “Conditional neural field latent diffusion model for generating spatiotemporal turbulence,” *Nature Communications*, vol. 15, no. 1, p. 10416, 2024.
- [15] A. Dokuchaev, F. Bonizzoni, S. Pagani, F. Regazzoni, and S. Pezzuto, “Learning geometry-dependent lead-field operators for forward ecg modeling,” 2026.
- [16] N. Boullé and A. Townsend, “A mathematical guide to operator learning,” in *Handbook of Numerical Analysis*, vol. 25, pp. 83–125, Elsevier, 2024.
- [17] A. Quarteroni, P. Gervasio, and F. Regazzoni, “Combining physics-based and data-driven models: advancing the frontiers of research with scientific machine learning,” *Mathematical Models and Methods in Applied Sciences*, vol. 35, no. 04, pp. 905–1071, 2025.

- [18] F. Regazzoni, S. Pagani, and A. Quarteroni, “Universal solution manifold networks (USM-Nets): Non-intrusive mesh-free surrogate models for problems in variable domains,” *Journal of Biomechanical Engineering*, vol. 144, no. 12, p. 121004, 2022.
- [19] L. Zhang, S. Pagani, J. Zhang, and F. Regazzoni, “Shape-informed surrogate models based on signed distance function domain encoding,” *Journal of Computational Physics*, vol. 538, p. 114178, 2025.
- [20] S. Brivio, S. Fresca, and A. Manzoni, “Handling geometrical variability in nonlinear reduced order modeling through continuous geometry-aware dl-roms,” *Computer Methods in Applied Mechanics and Engineering*, vol. 442, p. 117989, 2025.
- [21] M. Yin, N. Charon, R. Brody, L. Lu, N. Trayanova, and M. Maggioni, “A scalable framework for learning the geometry-dependent solution operators of partial differential equations,” *Nature Computational Science*, vol. 4, pp. 928–940, 2024.
- [22] J. Bayer, A. J. Prassl, A. Pashaei, J. F. Gomez, A. Frontera, A. Neic, G. Plank, and E. J. Vigmond, “Universal ventricular coordinates: A generic framework for describing position within the heart and transferring data,” *Medical Image Analysis*, vol. 45, pp. 83–93, 2018.
- [23] Y. Bengio, A. Courville, and P. Vincent, “Representation learning: A review and new perspectives,” *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 8, pp. 1798–1828, 2013.
- [24] A. Dai, C. R. Qi, and M. Nießner, “Shape completion using 3D-encoder-predictor CNNs and shape synthesis,” *CoRR*, vol. abs/1612.00101, 2016.
- [25] D. P. Kingma and M. Welling, “Auto-encoding variational bayes,” in *2nd International Conference on Learning Representations, ICLR 2014, Banff, AB, Canada, April 14-16, 2014, Conference Track Proceedings* (Y. Bengio and Y. LeCun, eds.), 2014.
- [26] P. Bojanowski, A. Joulin, D. Lopez-Paz, and A. Szlam, “Optimizing the latent space of generative networks,” in *International Conference on Machine Learning*, 2017.
- [27] J. J. Park, P. Florence, J. Straub, R. Newcombe, and S. Lovegrove, “DeepSDF: Learning continuous signed distance functions for shape representation,” in *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 165–174, 2019.
- [28] W. E. Lorenson and H. E. Cline, “Marching cubes: A high resolution 3D surface construction algorithm,” *ACM SIGGRAPH Computer Graphics*, vol. 21, p. 163–169, Aug. 1987.
- [29] H.-T. D. Liu, F. Williams, A. Jacobson, S. Fidler, and O. Litany, “Learning smooth neural functions via lipschitz regularization,” 2022.
- [30] J. Verhülndonk, T. Grandits, F. S. Costabal, T. Pinetz, R. Krause, A. Auricchio, G. Haase, S. Pezzuto, and A. Effland, “Shape of my heart: Cardiac models through learned signed distance functions,” in *Proceedings of The 7nd International Conference on Medical Imaging with Deep Learning* (N. Burgos, C. Petitjean, M. Vakalopoulou, S. Christodoulidis, P. Coupe, H. Delingette, C. Lartizien, and D. Mateus, eds.), vol. 250 of *Proceedings of Machine Learning Research*, pp. 1584–1605, PMLR, 2024.
- [31] A. Sanchez-Gonzalez, J. Godwin, T. Pfaff, R. Ying, J. Leskovec, and P. W. Battaglia, “Learning to simulate complex physics with graph networks,” in *Proceedings of the 37th International Conference on Machine Learning, ICML’20, JMLR.org*, 2020.
- [32] T. Pfaff, M. Fortunato, A. Sanchez-Gonzalez, and P. W. Battaglia, “Learning mesh-based simulation with graph networks,” in *International Conference on Learning Representations*, 2021.

- [33] L. Serrano, T. X. Wang, E. Le Naour, J.-N. Vittaut, and P. Gallinari, “Aroma: Preserving spatial structure for latent pde modeling with local neural fields,” in *Advances in Neural Information Processing Systems* (A. Globerson, L. Mackey, D. Belgrave, A. Fan, U. Paquet, J. Tomczak, and C. Zhang, eds.), vol. 37, pp. 13489–13521, Curran Associates, Inc., 2024.
- [34] C. W. Lanyon, C. Rodero, A. Qayyum, T. M. Baptiste, S. A. Niederer, and R. D. Wilkinson, “Weaving the digital tapestry: Methods for emulating cohorts of cardiac digital twins using gaussian processes,” *Annals of Biomedical Engineering*, vol. 54, p. 514–530, Nov. 2025.
- [35] M. Strocchi, C. M. Augustin, M. A. F. Gsell, E. Karabelas, A. Neic, K. Gillette, O. Razeghi, A. J. Prassl, E. J. Vigmond, J. M. Behar, J. Gould, B. Sidhu, C. A. Rinaldi, M. J. Bishop, G. Plank, and S. A. Niederer, “A publicly available virtual cohort of four-chamber heart meshes for cardiac electro-mechanics simulations,” *PLoS ONE*, vol. 15, no. 6, p. e0235145, 2020.
- [36] C. Rodero, M. Strocchi, M. Marciniak, S. Longobardi, J. Whitaker, M. D. O’Neill, K. Gillette, C. Augustin, G. Plank, E. J. Vigmond, P. Lamata, and S. A. Niederer, “Linking statistical shape models and simulated function in the healthy adult human heart,” *PLOS Computational Biology*, vol. 17, pp. 1–28, 04 2021.
- [37] Y. Xie, T. Takikawa, S. Saito, O. Litany, S. Yan, N. Khan, F. Tombari, J. Tompkin, V. Sitzmann, and S. Sridhar, “Neural fields in visual computing and beyond,” *Computer Graphics Forum*, vol. 41, no. 2, pp. 641–676, 2022.
- [38] M. Taghizadeh, M. A. Nabian, and N. Alemazkoo, “Multifidelity graph neural networks for efficient and accurate mesh-based partial differential equations surrogate modeling,” *Computer-Aided Civil and Infrastructure Engineering*, vol. 40, no. 7, pp. 841–858, 2025.
- [39] S. Maneewongvatana and D. M. Mount, “Analysis of approximate nearest neighbor searching with clustered point sets,” 1999.
- [40] M. Unberath, A. Maier, D. Fleischmann, J. Hornegger, and R. Fahrig, “Open-source 4d statistical shape model of the heart for x-ray projection imaging,” in *2015 IEEE 12th International Symposium on Biomedical Imaging (ISBI)*, pp. 739–742, 2015.
- [41] G. A. Holzapfel and R. W. Ogden, “Constitutive modelling of passive myocardium: A structurally based framework for material characterization,” *Philosophical Transactions of the Royal Society A*, vol. 367, no. 1902, pp. 3445–3475, 2009.
- [42] J. D. Bayer, R. C. Blake, G. Plank, and N. A. Trayanova, “A novel rule-based algorithm for assigning myocardial fiber orientation to computational heart models,” in *Annals of Biomedical Engineering*, vol. 40, pp. 2243–2254, 2012.
- [43] R. Doste, D. Soto-Iglesias, G. Bernardino, A. Alcaine, R. Sebastian, S. Giffard-Roisin, M. Sermesant, A. Berruezo, D. Sanchez-Quintana, and O. Camara, “A rule-based method to model myocardial fiber orientation in cardiac biventricular geometries with outflow tracts,” *International Journal for Numerical Methods in Biomedical Engineering*, vol. 35, no. 4, p. e3185, 2019.
- [44] P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, J. K. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, Í. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors, “SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python,” *Nature Methods*, vol. 17, pp. 261–272, 2020.
- [45] F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, J. Vanderplas, A. Passos, D. Cournapeau, M. Brucher, M. Perrot, and

- E. Duchesnay, “Scikit-learn: Machine learning in Python,” *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [46] D. Falck and B. Collette, *FreeCAD [How-to]: Solid Modeling with the power of Python*. Packt Publishing Birmingham, 2012.
- [47] W. Schroeder, K. Martin, and B. Lorensen, *The Visualization Toolkit (4th ed.)*. Kitware, 2006.
- [48] C. Geuzaine and J.-F. Remacle, “Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities,” *International Journal for Numerical Methods in Engineering*, vol. 79, no. 11, pp. 1309–1331, 2009.
- [49] P. C. Africa, “lifex: A flexible, high performance library for the numerical solution of complex finite element problems,” *SoftwareX*, p. 101252, 2022.
- [50] J. Bradbury, R. Frostig, P. Hawkins, M. J. Johnson, C. Leary, D. Maclaurin, G. Necula, A. Paszke, J. VanderPlas, S. Wanderman-Milne, and Q. Zhang, “JAX: composable transformations of Python+NumPy programs,” 2018. Version 0.3.13.
- [51] J. Heek, A. Levskaya, A. Oliver, M. Ritter, B. Rondepierre, A. Steiner, and M. van Zee, “Flax: A neural network library and ecosystem for JAX,” 2024. Version 0.12.4.
- [52] DeepMind, I. Babuschkin, K. Baumli, A. Bell, S. Bhupatiraju, J. Bruce, P. Buchlovsky, D. Budden, T. Cai, A. Clark, I. Danihelka, A. Dedieu, C. Fantacci, J. Godwin, C. Jones, R. Hemsley, T. Hennigan, M. Hessel, S. Hou, S. Kapturowski, T. Keck, I. Kemaev, M. King, M. Kunesch, L. Martens, H. Merzic, V. Mikulik, T. Norman, G. Papamakarios, J. Quan, R. Ring, F. Ruiz, A. Sanchez, L. Sartran, R. Schneider, E. Sezener, S. Spencer, S. Srinivasan, M. Stanojević, W. Stokowiec, L. Wang, G. Zhou, and F. Viola, “The DeepMind JAX Ecosystem,” 2020.
- [53] T. Akiba, S. Sano, T. Yanase, T. T. Ohta, and M. Koyama, “Optuna: A next-generation hyperparameter optimization framework,” in *The 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pp. 2623–2631, 2019.
- [54] S. van der Walt, J. L. Schönberger, J. Nunez-Iglesias, F. Boulogne, J. D. Warner, N. Yager, E. Gouillart, T. Yu, and the scikit-image contributors, “scikit-image: image processing in Python,” *PeerJ*, vol. 2, p. e453, 6 2014.
- [55] B. Sullivan and A. Kaszynski, “PyVista: 3D plotting and mesh analysis through a streamlined interface for the Visualization Toolkit (VTK),” *Journal of Open Source Software*, vol. 4, p. 1450, May 2019.

MOX Technical Reports, last issues

Dipartimento di Matematica
Politecnico di Milano, Via Bonardi 9 - 20133 Milano (Italy)

- 24/2026** Carrara, D.; Hirschvogel, M.; Bonizzoni, F.; Pagani, S.; Pezzuto, S.; Regazzoni, F.
Shape-informed cardiac mechanics surrogates in data-scarce regimes via geometric encoding and generative augmentation
- 23/2026** Ballini, E.; Muscarnera, L.; Fumagalli, A.; Scotti, A.; Regazzoni, F.
Elimination-compensation pruning for fully-connected neural networks
- 21/2026** Bottacini, G.; Torzoni, M.; Manzoni, A.
Neural Markov chain Monte Carlo: Bayesian inversion via normalizing flows and variational autoencoders
- 17/2026** Caldera, L.; Bottacini, G.; Cavinato, L.
MAGIC-Flow: multiscale adaptive conditional flows for generation and interpretable classification
- 20/2026** Caldera, L.; Cavinato, L.; Cirone, A.; Cama, I.; Garbarino, S.; Lodi, R.; Tagliavini, F.; Nigri, A.; De Francesco, S.; Cappozzo, A.; Piana, M.; Ieva, F.;
DISARM++: Beyond scanner-free harmonization
- 19/2026** Caldera, L.; Cavinato, L.; Ieva, F.
Scanner-agnostic MRI harmonization via SSIM-guided disentanglement
- 15/2026** Zecchi, A. A.; Sanavio, C.; Cappelli, L.; Perotto, S.; Roggero, A.; Succi, S.
Block encoding of sparse matrices with a periodic diagonal structure
- 14/2026** Agasisti, T.; Cannistrà, M.; Paganoni, A.M.
Nudging communication for students at risk: experimental evidence from an Italian university
- 13/2026** Dimola, N.; Coclite, A.; Zunino, P.
Neural Preconditioning via Krylov Subspace Geometry
- 12/2026** Corbetta A; Logan K.M.; Ferro M; Zuccolo L; Perola M.; Ganna A.; Di Angelantonio E.;Ieva F.
Longitudinal patterns of statin adherence and factors associated with decline in over one million individuals in Finland and Italy