A Diffusion Model for Simulation Ready Coronary Anatomy with Morpho-skeletal Control

Karim Kadry[1], Shreya Gupta[1], Jonas Sogbadji[1], Michiel Schaap[3], Kersten Petersen[3], Takuya Mizukami[4], Carlos Collet[5], Farhad R. Nezami[2], and Elazer R. Edelman[1,2]

[1] Institute of Medical Engineering, MIT
[2] Brigham and Women's Hospital, Harvard
[3] HeartFlow
[4] Showa University
[5] OLV Aalst

kkadry@mit.edu; shreyag@mit.edu; jonassog@mit.edu; mschaap@heartflow.com; kpetersen@heartflow.com; mizukami.tky@gmail.com; carloscollet@gmail.com; frikhtegarnezami@bwh.harvard.edu; ere@mit.edu;

Abstract. Virtual interventions enable the physics-based simulation of device deployment within coronary arteries. This framework allows for counterfactual reasoning by deploying the same device in different arterial anatomies. However, current methods to create such counterfactual arteries face a trade-off between controllability and realism. In this study, we investigate how Latent Diffusion Models (LDMs) can custom synthesize coronary anatomy for virtual intervention studies based on mid-level anatomic constraints such as topological validity, local morphological shape, and global skeletal structure. We also extend diffusion model guidance strategies to the context of morpho-skeletal conditioning and propose a novel guidance method for continuous attributes that adaptively updates the negative guiding condition throughout sampling. Our framework enables the generation and editing of coronary anatomy in a controllable manner, allowing device designers to derive mechanistic insights regarding anatomic variation and simulated device deployment.

Keywords: Diffusion Models · Diffusion Guidance · Morphological Constraints · Anatomic Generation · Virtual Interventions · Digital Twins

1 Introduction

Coronary artery disease is caused by the buildup of diseased tissue in the arteries that supply blood to the heart [32]. Variation in both local and global anatomic structure influence the success of cardiological interventions, such as stent deployment or balloon angioplasty [15, 22, 34]. Local structure is characterized by the morphological attributes of each 3D component within the artery wall, such as size or shape [23, 26], while global structure is defined by the tree-like skeleton that results from the artery bifurcating into one or several branches.
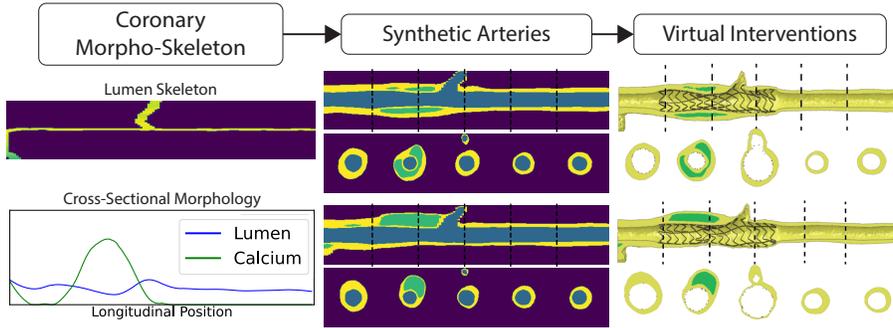
Fig. 1: We propose to control the generation of 3D multi-component diseased coronary arteries with mid-level representations such as cross-sectional morphology and tree-like skeletal structure. Synthetic arteries enable physics-based counterfactual reasoning through comparative virtual intervention studies. We synthesize semantic segmentation maps composed of luminal (blue), arterial wall (yellow) and calcified (green) tissues.

Virtual intervention platforms aim to disentangle the complex relationship between coronary anatomy, biophysics, and pathophysiology [9, 39], enabling physics-based simulations of medical interventions such as stent and balloon expansion [18, 19, 29]. However, current simulation-based platforms face a trade-off between using anatomies that are either realistic or controllable. Using simplified parametric geometry [8] such as tubular cylinders to represent coronary arteries can isolate the effect of certain morphological features on device outcomes. However, this approach cannot capture the intricate 3D micro-morphology of coronary lesions [23, 42], limiting the development of clinically relevant insights. On the other hand, reconstructing digital twins from medical images produces realistic anatomy but prevents control over anatomic attributes [9, 17, 45]. As such, digital twins are unable to assess causal relationships between anatomic features and device effectiveness, as many of these features are correlated together.

To address such limitations, we study how Latent Diffusion Models (LDMs) can be used as a controllable data source for comparative virtual intervention studies, producing customized anatomic counterfactuals that change the simulated outcome of device deployment. This would allow device designers to understand which subtle anatomic features influence interventional outcomes, guiding both the design process and clinical trial recruitment. We present three methodological adaptations that enable LDMs to synthesize and edit 3D multi-component coronary segmentation maps according to prespecified anatomic constraints. First, we study how to regularize LDM latent space to reduce the presence of topological defects (Sec. 3.2). Second, we enable conditional generation based on clinically interpretable and editable representations of morphology and skeletal structure (Sec. 3.3). The local morphological component is specified by one or several cross-sectional features along the length of the artery, while the global skeletal component specifies the presence, location, and shape of vessel bi-

furcations. Lastly, we extend current diffusion guidance strategies with morpho-skeletal regressor functions to enhance conditional fidelity (Sec. 3.4). In addition, we introduce a novel guidance algorithm tailored to morphological conditioning which utilizes lightweight and non-differentiable morphological regressor functions to adaptively update the null condition during sampling (Sec. 3.4). Our main contributions are thus as follows:

1. We introduce a method to improve the topological quality of 3D multi-component coronary anatomies produced by latent diffusion models through the use of a topological interaction loss during autoencoder training.
2. We develop a novel morpho-skeletal conditioning framework for latent diffusion models to enable the controlled synthesis of coronary arteries and extend loss-based guidance strategies to this task through the use of differentiable morpho-skeletal regressor functions.
3. We demonstrate the limitations of classifier-free guidance for morpho-skeletal generation and introduce an adaptive null guidance strategy. Our proposed method operates through the use of non-differentiable morphological regressors to determine the null guidance condition, enhancing conditional fidelity and computational efficiency.
4. We show that our framework can disentangle micro-morphology and arterial branch structure, condition on a variety of morphological features, and edit patient-specific coronary arteries for comparative simulation studies.

2   Related Work

2.1   Generative Models of Virtual Anatomies

Current generative models of coronary arteries are based on parametric approaches, utilizing geometric primitives such as cylinders to approximate coronary geometry [8, 21]. While such models enable geometric control, they do not capture the complex morphological or topological variation associated with coronary lesions. For simple topological organs such as the aorta, deep learning has been used to produce virtual anatomy by deforming a 3D template shape [1,2,10]. Template-based approaches have also been used for generative editing by applying localized elastic deformations to template shapes, inducing common vascular phenomena such as aneurysms and stenoses [27,28]. However, template-based approaches require anatomic correspondence and consistent topology between patients, which is not possible for multi-material coronary arteries. Alternatively, implicit representations [30, 31, 41, 44], enable flexible generation of topological structures but do not account for unphysiological topological defects. In contrast, we aim to generate implicit representations of coronary anatomy to enable flexible topological generation while selectively enforcing topological quality with topological losses.

2.2   Diffusion Model Guidance

Diffusion model guidance techniques primarily fall into two families. Loss-based methods include Classifier Guidance (CG) [7], which uses the gradient obtained

from pre-trained classifiers to influence the sampling process. Loss-based techniques have also been extended to solve inverse problems, in which the correcting gradient is derived from a loss that measures the degree to which a constraint is satisfied [4, 13, 38]. However, loss-based methods have not been studied in the context of anatomic constraints and require computationally expensive back-propagation. In contrast, null-conditioning based methods such as Classifier-Free Guidance (CFG) [14] use a weighted combination of conditional and unconditional model outputs to guide the sampling process. Null conditions can also be defined in a negative manner, guiding the sampling process away from certain conditions [11]. While effective for a variety of modalities [14, 25, 43], the use of null-guidance for enforcing morphological or skeletal constraints has not been explored. In this study, we adapt diffusion guidance strategies to the task of enforcing morpho-skeletal constraints and develop novel guidance algorithms suited to conditioning with continuous morphological features.

2.3   Topological regularization for Neural Networks

Cardiovascular anatomy exhibits a wide variety of topological structures, manifesting as connected components, loops, and voids. It also features multi-tissue interactions such as containment and adjacency. Various techniques have been developed to constrain neural network outputs in accordance with prespecified topological priors. Persistent homology (PH) losses [3, 5] can regularize multiclass segmentation networks, where the individual classes and interactions between different classes should adhere to specified topological priors. However, such PH-based methods are prohibitively expensive to use with 3D

Fig. 2: We regularize the latent space $\mathbf{z}$ with a topological interaction loss applied to the calcium and lumen segmentations.

data, especially as the number of topological interactions increase. Alternative techniques reduce topological interactions by penalizing critical voxels that violate topological priors [12]. We study the application of such strategies to diffusion models of anatomy, to which topological regularization has not been applied.

3   Methodology

3.1   Latent Diffusion Model

We employ a latent diffusion model (LDM), which is trained in the latent space of a variational autoencoder (VAE). The training process for the VAE is shown in

Fig. 2 in which an encoder E_{LDM} encodes the multi-tissue coronary segmentation maps $\mathbf{x} \in \mathbb{R}^{C \times H \times W \times D}$ into a latent representation $\mathbf{z} \in \mathbb{R}^{c \times h \times w \times d}$, where the latent dimensions are downsampled by a factor $f = H/h = W/w = D/d$. Here, C is the number of tissue classes and c is the number of latent channels. The latents are then passed to the decoder D_{LDM} back into voxel space, where a Dice cross entropy reconstruction loss $\mathcal{L}_R$, KL divergence loss $\mathcal{L}_{KL}$, and a topological interaction loss $\mathcal{L}_T$ (Sec. 3.2) are applied. For our application, we aim to sample from the conditional probability distribution of encoded coronary segmentation maps $p_{data}(\mathbf{z}|\mathbf{y})$, where $\mathbf{y}$ represents mid-level anatomic constraints regarding morphology $\mathbf{y}_m$ and skeletal structure $\mathbf{y}_s$ (Sec. 3.3). To train a diffusion model, we consider the joint distribution $p(\mathbf{z}_\sigma|\mathbf{y}; \sigma)$ obtained through a forward diffusion process, in which i.i.d Gaussian noise of standard deviation σ is added to the data, where at $\sigma = \sigma_{max}$ the data is indistinguishable from Gaussian noise. Equivalently, the forward process can be described by the stochastic equation $d\mathbf{z}_\sigma = \sqrt{2\sigma}d\mathbf{w}$ where $\mathbf{w}$ is the standard Wiener process. The reverse diffusion process is defined as the solution to the following stochastic differential equation

$$d\mathbf{z}_\sigma = -2\sigma \nabla_{\mathbf{z}_\sigma} \log p(\mathbf{z}_\sigma|\mathbf{y}; \sigma)\, dt + \sqrt{2\sigma}d\mathbf{w}, \tag{1}$$

Where $\nabla_{\mathbf{z}_\sigma} \log p(\mathbf{z}_\sigma|\mathbf{y}; \sigma)$ represents the score function that is conditioned on $\mathbf{y}$. To train the network, we apply the forward diffusion process to produce intermediately noised latents $\mathbf{z}_\sigma = \mathbf{z} + \mathbf{n}$ where $\mathbf{n} \sim \mathcal{N}(\mathbf{0}, \sigma^2\mathbf{I})$, parameterized by a noise level σ. The diffusion model is parameterized as a function F_θ, encapsulated within a denoiser D_θ, that takes the intermediately noised input $\mathbf{z}_\sigma$, the conditioning signal $\mathbf{y}$, and the noise level σ to predict the clean data $\mathbf{z}$.

$$D_\theta(\mathbf{z}_\sigma; \sigma, \mathbf{y}) = c_{\text{skip}}(\sigma)\, \mathbf{z}_\sigma + c_{\text{out}}(\sigma)\, F_\theta(c_{\text{in}}(\sigma)\, \mathbf{z}_\sigma; c_{\text{noise}}(\sigma), \mathbf{y}), \tag{2}$$

where c_{skip} allows F_θ to predict the noise $\mathbf{n}$ at low σ and the training data $\mathbf{z}$ at high σ. The variables c_{out} and c_{in} scale the input and output magnitudes to be within unit variance, and the constant c_{noise} maps the noise level σ to a conditioning input for the network [20]. The denoiser output is related to the conditional score function through the relation $\nabla_{\mathbf{z}_\sigma} \log p(\mathbf{z}_\sigma|\mathbf{y}; \sigma) = (D_\theta(\mathbf{z}_\sigma; \sigma, \mathbf{y}) - \mathbf{z})/\sigma^2$ and F_θ is chosen to be a 3D U-net (see Appendix E). The denoising loss $\mathcal{L}_D$ is then specified based on the agreement between the denoiser output and the original training data:

$$\mathcal{L}_D = \mathbb{E}_{\sigma, \mathbf{z}, \mathbf{n}}[\lambda(\sigma)||D_\theta(\mathbf{z}_\sigma; \sigma, \mathbf{y}) - \mathbf{z}||_2^2], \tag{3}$$

such that the loss weighting $\lambda(\sigma) = 1/c_{\text{out}}(\sigma)^2$ ensures an effective loss weight that is uniform across all noise levels, and σ is sampled from a log-normal distribution.

3.2   Topological Regularization

To be compatible for numerical simulation, diseased tissue and lumen must be encompassed by the vessel wall within every 2D cross section of the artery.

Non-compliance with these constraints can give rise to unrealistic biophysical phenomena. While post-processing can be applied to the segmentation map before meshing and simulation, such steps are time consuming and cumbersome, especially when done for each segmentation map [40]. As such, we aim to enhance the topological viability of the coronary arteries sampled by the diffusion model. We study how to regularize the latent space of the autoencoder during training with a topological interaction loss, as described in Gupta et al. [12]. We adapt this loss to a 3D multi-slice context, using it to identify and penalize critical pixels in each 2D cross section with a cross-entropy loss function using vessel wall as a ground truth label (see Appendix C). The topological interaction loss ($\mathcal{L}_T$) is then weighted and and added to the reconstruction ($\mathcal{L}_R$) and KL divergence ($\mathcal{L}_{KL}$) losses.
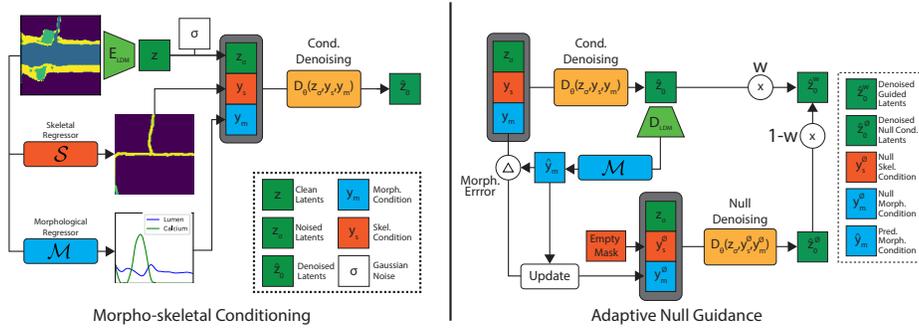
Fig. 3: Left: We condition the training process of a latent diffusion model through channel-wise concatenation of the morphological ($\mathbf{y}_m$) and skeletal ($\mathbf{y}_s$) maps to the noised latent representation ($\mathbf{z}_\sigma$). Right: Our proposed guidance algorithm operates by specifying an null skeletal condition ($\mathbf{y}_s^{\varnothing}$) and adaptively updating a null morphological condition ($\mathbf{y}_m^{\varnothing}$) based on the error between the target condition ($\mathbf{y}_m$) and the morphology exhibited by the current denoiser output ($\hat{\mathbf{y}}_m$).

3.3   Morpho-skeletal Conditioning

Morphological Conditioning The morphological conditioning process is shown in Fig. 3 (left). We use a morphological regressor $\mathcal{M}$ to compute a morphological conditioning map $\mathbf{y}_m = \mathcal{M}(\mathbf{x}) \in \mathbb{R}^{m \times h \times w \times d}$ that is concatenated to the latent representation $\mathbf{z}$ to condition the denoising process. To do this, our morphological regressor $\mathcal{M}$ spatially encodes local morphological information through cross sectional features along the depth dimension of the straightened coronary segmentation map. This process results in a set of morphological feature vectors $\mathbf{M} \in \mathbb{R}^{m \times D}$, where m represents the number of 2D morphological attributes and D signifies the number of frames along the artery centerline. The features are calculated, normalized by the 2nd and 98th percentile attribute values within

the training set, and appropriately downsampled for concatenation (see Appendix D.1).

Skeletal Conditioning The skeletal conditioning process is shown in Fig. 3 (left). We use a skeletal regressor $\mathcal{S}$ to compute a skeletal conditioning map $\mathbf{y}_s = \mathcal{S}(\mathbf{x}) \in \mathbb{R}^{1 \times h \times w \times d}$ that is concatenated to the latent representation $\mathbf{z}$ to condition the denoising process. To do this, our skeletal regressor $\mathcal{S}$ encodes global bifurcational information by finding the luminal skeleton $\mathbf{S} \in \mathbb{R}^{H \times W \times D}$, which describes the branching structure of the coronary arteries. We then compute, process, and resize the skeleton map to produce the skeletal conditioning signal (see Appendix D.2). We zero out the skeletal condition with a probability of 0.2 to jointly train an unconditional model with respect to the lumen skeleton.

3.4 Morpho-skeletal Guidance

Loss Guidance We formulate loss-based guidance in terms of the gradient derived from a morphological loss $\mathcal{L}_m$ and skeletal loss $\mathcal{L}_s$, as can be seen in Algorithm 1. Each loss term is defined as the L2 norm between the target morphoskeleton and the morpho-skeleton measured from synthetic segmentation maps by the differentiable regressor functions $\bar{\mathcal{M}}$ and $\bar{\mathcal{S}}$ (see Appendix F).

$$\mathcal{L}_m = \|\mathbf{y}_m - \bar{\mathcal{M}}(\mathbf{x})\|_2^2 \quad \text{and} \quad \mathcal{L}_s = \|\mathbf{y}_s - \bar{\mathcal{S}}(\mathbf{x})\|_2^2, \tag{4}$$

where the default choice of segmentation maps $\mathbf{x}$ is obtained by decoding the denoised latent from the diffusion model $\hat{\mathbf{z}}_0$ such as that in Chung et al. [4]. Classifier guidance [14] can be implemented by decoding the intermediately noised latents $\mathbf{z}_\sigma$ to obtain $\mathbf{x}$. Finally, we guide the conditional denoising process by using the gradient of the loss with respect to $\mathbf{z}_\sigma$ and a guidance term $(w-1)$:

$$\underbrace{D_\theta^w(\mathbf{z}_\sigma; \sigma, \mathbf{y}_m, \mathbf{y}_s)}_{\text{Guided Denoising}} = \underbrace{D_\theta(\mathbf{z}_\sigma; \sigma, \mathbf{y}_m, \mathbf{y}_s)}_{\text{Cond. Denoising}} + \underbrace{(w-1)\nabla_{\mathbf{z}_\sigma}(\mathcal{L}_m + \mathcal{L}_s)}_{\text{Gradient Guidance}}, \tag{5}$$

Adaptive Null guidance To circumvent the computational cost of backpropagating through the diffusion model, VAE decoder, and regressors, we formulate a null-conditional guidance method based on simple and non differentiable morphological measurement functions, as can be seen in Fig. 3 (right) and Algorithm 2. We take inspiration from classifier-free guidance (CFG), which guides the denoising process by a weighted combination of the denoiser output conditioned by the target conditioning signals $(\mathbf{y}_m, \mathbf{y}_s)$ and the null conditioning signals $(\mathbf{y}_m^\varnothing, \mathbf{y}_s^\varnothing)$:

$$\underbrace{D_\theta^w(\mathbf{z}_\sigma; \sigma, \mathbf{y}_m, \mathbf{y}_s)}_{\text{Guided Denoising}} = \underbrace{w D_\theta(\mathbf{z}_\sigma; \sigma, \mathbf{y}_m, \mathbf{y}_s)}_{\text{Cond. Denoising}} + \underbrace{(1-w) D_\theta(\mathbf{z}_\sigma; \sigma, \mathbf{y}_m^\varnothing, \mathbf{y}_s^\varnothing)}_{\text{Null Denoising}}, \tag{6}$$

where w is a guidance weight typically set to be larger than 1. To enable conventional classifier-free guidance, we train an unconditional diffusion model and

use it to guide sampling. We implement our proposed algorithm, adaptive null guidance (ANG), by adaptively updating the null signal for the morphological condition $\mathbf{y}_m^{\varnothing}$ based on the difference between the target morphology $\mathbf{y}_m$ and the morphology derived from the current denoised output $\hat{\mathbf{y}}_m = \mathcal{M}(\hat{\mathbf{x}}_0)$.

$$\mathbf{y}_{\mathrm{m}}^{\varnothing} = \mathcal{M}(\hat{\mathbf{x}}_0) + \underbrace{(\mathcal{M}(\hat{\mathbf{x}}_0) - \mathbf{y}_m)}_{\text{Morph. Error}}, \tag{7}$$

where $\mathcal{M}$ is a morphological regressor function that is not required to be differentiable, and the segmentation map $\hat{\mathbf{x}}_0$ is decoded from the denoiser output $\hat{\mathbf{z}}_0 = D_\theta(\mathbf{z}_\sigma, \mathbf{y}_m, \mathbf{y}_s, \sigma)$. Skeletal guidance is implemented similar to conventional CFG, where the skeletal condition $\mathbf{y}_s^{\varnothing}$ is set as an empty mask.

Morpho-skeletal Regressors for Guidance Our proposed guidance methods require morphological or skeletal regressors to guide the reverse diffusion process. For null guidance, we leverage nondifferentiable morphological regressor functions $\mathcal{M}$, as well as a hard skeletonization method $\mathcal{S}$, developed by Sato et al. [33], which produces high fidelity estimations of the coronary morphoskeleton. Loss guidance, in contrast, requires that the regressors be differentiable to calculate gradients. A soft morphological regressor $\bar{\mathcal{M}}$, consisting of differentiable image processing operations, is used when possible for morphological loss guidance. We further implement a neural morphological regressor $\bar{\mathcal{M}}_\phi$, to extend loss guidance to morphological features that cannot be easily derived in a differentiable fashion. The neural morphological regressor is trained on high fidelity morphological features from the ground truth morphological regressor $\mathcal{M}$. For skeletal loss guidance, our skeletal regressor $\bar{\mathcal{S}}$ leverages an iterative soft skeletonization algorithm [35] consisting of erosions and dilations. However, since soft skeletonization methods can exhibit limited performance compared to non-differentiable methods, we also implement a neural skeletal regressor $\bar{\mathcal{S}}_\psi$ to regress high fidelity skeletons (see Appendix F).

Algorithm 1 Loss Guidance	**Algorithm 2** Adaptive Null Guidance
Require: $D_\theta, \mathbf{y}_m, \mathbf{y}_s, \sigma, \mathbf{z}_\sigma, w, \bar{\mathcal{M}}, \bar{\mathcal{S}}$	**Require:** $D_\theta, \mathbf{y}_m, \mathbf{y}_s, \sigma, \mathbf{z}_\sigma, w, \mathcal{M}$
1: $\hat{\mathbf{z}}_0 \leftarrow D_\theta(\mathbf{z}_\sigma; \sigma, \mathbf{y}_m, \mathbf{y}_s)$	1: $\hat{\mathbf{z}}_0 \leftarrow D_\theta(\mathbf{z}_\sigma; \sigma, \mathbf{y}_m, \mathbf{y}_s)$
2: $\hat{\mathbf{x}}_0 \leftarrow D_{LDM}(\hat{\mathbf{z}}_0)$	2: $\hat{\mathbf{x}}_0 \leftarrow D_{LDM}(\hat{\mathbf{z}}_0)$
3: $\mathcal{L}_m \leftarrow \|\mathbf{y}_m - \bar{\mathcal{M}}(\hat{\mathbf{x}}_0)\|_2^2$	3: $\mathbf{y}_m^{\varnothing} \leftarrow \mathcal{M}(\hat{\mathbf{x}}_0) + (\mathcal{M}(\hat{\mathbf{x}}_0) - \mathbf{y}_m)$
4: $\mathcal{L}_s \leftarrow \|\mathbf{y}_s - \bar{\mathcal{S}}(\hat{\mathbf{x}}_0)\|_2^2$	4: $\mathbf{y}_s^{\varnothing} \leftarrow \varnothing$
5: $\hat{D}_\theta \leftarrow \hat{\mathbf{z}}_0 + (w-1)\nabla_{\mathbf{z}_\sigma}(\mathcal{L}_m + \mathcal{L}_s)$	5: $\hat{D}_\theta \leftarrow w\hat{\mathbf{z}}_0 + (1-w)D_\theta(\mathbf{z}_\sigma; \sigma, \mathbf{y}_m^{\varnothing}, \mathbf{y}_s^{\varnothing})$
6: **return** $\hat{D}_\theta$	6: **return** $\hat{D}_\theta$

4   Experiments

4.1   Dataset

We employ a dataset comprising 62 Coronary Computed Tomography (CCTA) images of patients with coronary artery disease [37], from which 222 unique 3D segmentation map patches were extracted (see Appendix B). Each segmentation map contains lumen, vessel wall, and diseased calcified tissue within the straightened coronary artery. Consequently, each 3D segmentation map has dimensions of $4 \times 128 \times 128 \times 128$. As the lateral resolution of the segmentation maps is larger than the in-plane resolution, the arterial segmentation maps in the results section are visualized with isotropic resolutions by interpolation along the depth dimension.

4.2   Evaluation Metrics

We calculate conditional fidelity, morphological quality, and topological quality for evaluation. Conditional fidelity is measured for both morphological and skeletal conditioning. Morphological conditioning fidelity is defined as the mean absolute error between the target morphological features and the features derived from the synthetic coronary segmentation map. Skeletal conditioning fidelity is defined as the mean absolute error between the number of branches derived from the target skeleton and the skeleton derived from the synthetic coronary segmentation map. To quantify morphological quality, we record 12 morphological features of each 3D artery in addition to 8 features for all 2D cross sections (see Appendix D.1). We calculate the Frechet Distance (FD), Precision, and Recall between generated and training segmentation maps for both 3D and 2D features respectively, similar to previous approaches [16]. For topological quality, we define the % of topological violations as the proportion of 2D cross sections within each coronary segmentation map that contains a critical pixel, defined by adjacency relations between the calcium, lumen, or background labels.

4.3   Enforcing Topological Constraints

We first compare the effect of applying our topological loss to the VAE latent space. We train two VAE models with and without a topological loss and use them to reconstruct 100 coronary segmentation maps in our training set. We specifically enforce that the calcium and lumen tissues be encompassed by the vessel wall. Tab. 1 shows that VAEs reconstruct arteries with similar levels of topological errors to the train set. When using a topological loss during training, the rate of topological errors is reduced by an order of magnitude. This can be seen in Fig. 4 where topological regularization reduces the presence of topological defects. We then compare the effect of topologically regularized VAEs on LDM training. We train two unconditional LDM's and sample 200 coronary artery segmentation maps from each model, using 100 sampling steps each. Tab. 1 shows topological regularization of the VAE latent space reduces topological violations in synthetic segmentation maps.

Table 1: Left: Topological interaction violations exhibited by segmentation maps that were autoencoded by various VAE models. Training a VAE with topological interaction losses improves topological quality. Right: Topological interaction violations by segmentation maps sampled from LDMs. Topological regularization of the LDM latent space improves the topological quality of generated arteries.

VAE	Topo. Viol. (%)	
Training	Lumen	Calcium
Train Set	18	16
$\mathcal{L}_R + \mathcal{L}_{KL}$	17	15
$\mathcal{L}_R + \mathcal{L}_{KL} + \mathcal{L}_T$	**0.5**	**1.1**

LDM	Topo. Viol. (%)	
Training	Lumen	Calcium
VAE w/o $\mathcal{L}_T$	1.44	9.21
VAE w $\mathcal{L}_T$	**0.10**	**2.61**
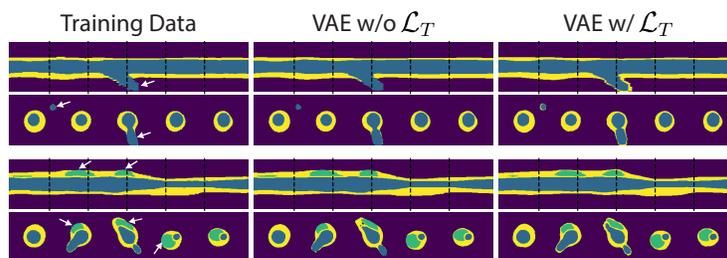
Fig. 4: Longitudinal and cross sectional slices of coronary segmentation maps. First column shows that patient-specific segmentation maps from the training set exhibit several topological defects such as lumen and calcium tissues not being fully contained within vessel wall tissue (white arrows). Second and third columns show the effect of autoencoding coronary segmentation maps with (right) or without (center) using a topological loss $\mathcal{L}_T$ during training.

4.4 Enforcing Morpho-skeletal Constraints

Conditioning Signal Ablation Study We examine the impact of morphological and skeletal conditioning on conditional fidelity and morphological quality in Tab. 2. To condition only on morphology or skeletal structure, we train the original LDM and drop the corresponding conditioning signal. For consistency, our unconditional model is trained by dropping out both conditions. Our default morphological condition consists of the cross-sectional area curves for the lumen and calcium tissues. The morphological and skeletal features were derived from the validation set and used to sample 200 segmentation maps with 100 sampling steps. Joint conditioning of both skeleton and morphology shows similar conditioning fidelity to conditioning on each constraint alone. Morphological conditioning alone improves skeletal fidelity as compared to unconditional sampling, likely due to the lumen area curve weakly indicating the presence of bifurcations.

Evaluating Morpho-skeletal Guidance Strategies We study the impact of varying the morpho-skeletal guidance weight on a morpho-skeletally conditioned
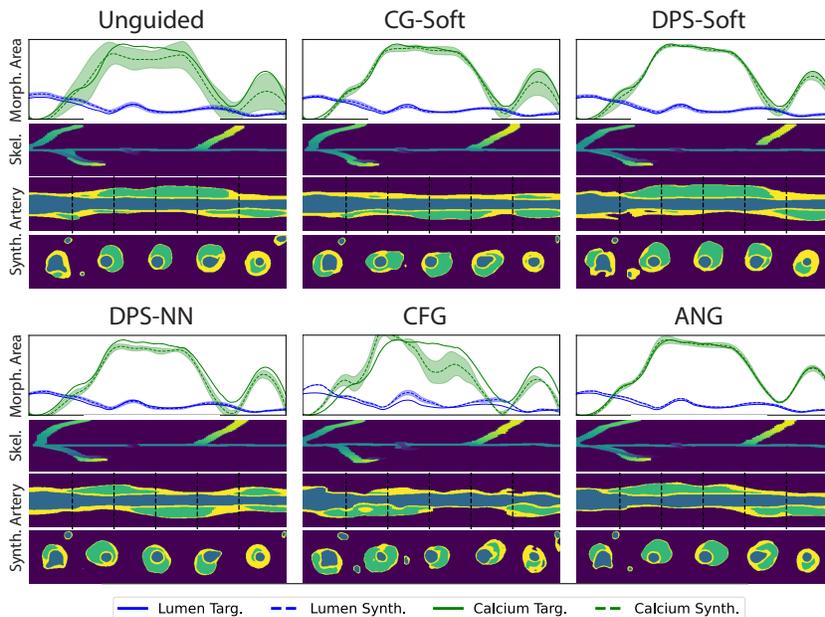
Fig. 5: Example morphological features, skeleton depth maps and synthetic segmentation map cross sections for various guidance methods. Filled in regions within morphological plots indicate standard deviation over 10 generated segmentation maps. A guidance weight of 5 is used when applicable. Our proposed guidance method (ANG) improves conditional fidelity while maintaining good visual quality.

LDM for various guidance methods. We evaluate three strategies for loss guidance and two strategies for null guidance. For loss guidance, we implement i) classifier-guidance (CG-Soft) and ii) diffusion posterior sampling with differentiable morpho-skeletal regressors $\bar{\mathcal{M}}$ and $\bar{\mathcal{S}}$ (DPS-Soft) as well as iii) a variant of DPS with neural network regressors $\bar{\mathcal{M}}_{\phi}$ and $\bar{\mathcal{S}}_{\psi}$ (DPS-NN). For null guidance, we implement i) classifier-free guidance (CFG) using a separately trained unconditional model and ii) Adaptive Null-Guidance (ANG) with a non-differentiable morphological regressor $\mathcal{M}$. We evaluated every method by sampling 200 arteries with 25 sampling steps for each combination of guidance weight and method. The input morpho-skeletal conditions were sampled from the validation set. Qualitative results can be seen in Fig. 5, where an unguided diffusion model can capture the variation in the lumen area but is unable to achieve high fidelity with respect to the calcium area. We see that morpho-skeletal guidance enhances conditional fidelity for all methods but CFG, which produces low quality segmentation maps with increasing guidance. ANG exhibits the best mix of conditioning fidelity and visual quality as compared to a method such as DPS-Soft, which has good conditional fidelity but produces un-physiological artifacts. These observations are reinforced in Fig. 6, which shows that ANG outperforms all methods in terms of both morphological and skeletal conditioning fidelity, while maintaining similar
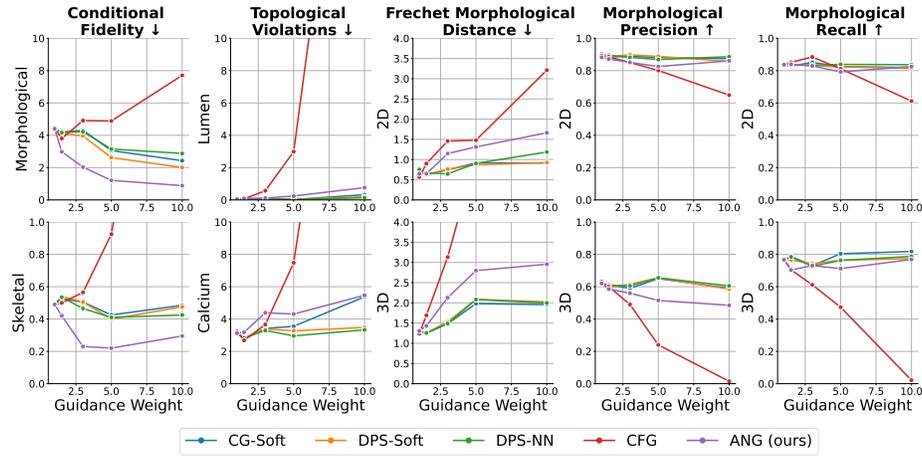
Fig. 6: Conditional fidelity and quality metrics for various guidance methods over several guidance weights. Our guidance method (ANG) enables enhanced conditioning fidelity at the cost of slightly decreased morphological and topological quality.

or slightly decreased recall values. Furthermore, ANG exhibits the lowest sampling time and memory usage for regressor based guidance strategies (Tab. 2) as it does not require the use of backpropagation. However, it can be seen that increasing the guidance weight for ANG increases the number of topological violations and causes a decrease in 3D morphological precision.
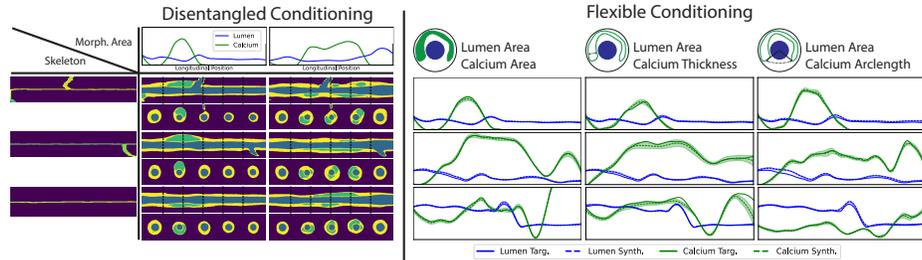
Fig. 7: Left: Synthetic coronary segmentation maps conditioned on different morphological features (columns) and lumen skeletons (rows) obtained from patient specific coronary segmentation maps. Right: Morphological features produced by conditioning on different combinations of lumen and calcium features (columns). We used adaptive null guidance with a weight of 5 to sample all coronary segmentation maps.

Table 2: Left: Ablation study for conditioning mechanisms. Joint conditioning improves conditional fidelity and morphological quality as compared to morphological conditioning alone. Right: Sampling speed and GPU memory usage comparison for a single sample over different guidance methods with 100 sampling steps. ANG is the most computationally efficient regressor-based guidance method.

Method	Conditional Fidelity		Morph. Prec.		Morph. Rec.	
	Morph.	Skel.	3D	2D	3D	2D
Uncond.	117	3.2	**0.68**	0.89	**0.85**	0.72
Morph.	**4.0**	1.2	0.65	0.88	**0.93**	0.82
Skel.	117	**0.47**	0.62	0.88	0.79	0.73
Joint.	4.5	0.48	0.64	**0.89**	0.79	**0.86**

Guidance Method	Sample Time (s)	Memory (GB)
No Regressor		
Unguided	**3**	**2.35**
CFG	9	3.07
Regressor		
CG-Soft	66	12.5
DPS-Soft	89	15.1
DPS-NN	80	8.6
ANG (ours)	**46**	**6.4**

4.5   Applications

Disentangled Morpho-skeletal Conditioning To demonstrate disentangled control over both coronary morphology and luminal skeleton, we use two sets of morphological feature vectors $\mathbf{M}$ and three skeleton heatmaps $\mathbf{S}$ derived from patient-specific arteries. We then sample 6 synthetic segmentation maps that correspond to all conditioning combinations. Fig. 7 shows that the same morphological conditioning signal can manifest differently based on the skeletal conditioning. For example, the combination of a local rise in lumen area in tandem with a skeletal branch induces a bifurcation in the generated model, but given a branch-less skeleton, the local rise in area manifests as an increase in lumen area.

Flexible Morpho-skeletal Conditioning Our framework enables morphological conditioning using a variety of features, we train our LDM on different combinations of 1) lumen area and 2) a morphological feature belonging to calcium, sampling 10 segmentation maps each. We specifically condition on calcium area, maximum calcium thickness, and maximum calcium arclength, using non differentiable morphological regressor functions. Fig. 7 demonstrates our guidance mechanism can enforce a variety of morphological constraints with high fidelity, without the need for training neural morphological regressors.

Editing with Morpho-skeletal Conditioning To virtually edit a coronary artery segmentation, we specify three components: morphological feature vectors, a skeleton heatmap, and a corresponding mask. We either specify region based masks, which mask out entire sections along the length of the artery, or tissue based masks, which mask out all voxels of pre-specified labels. We edit a coronary digital twin to create three different edited variants (Fig. 8). The first
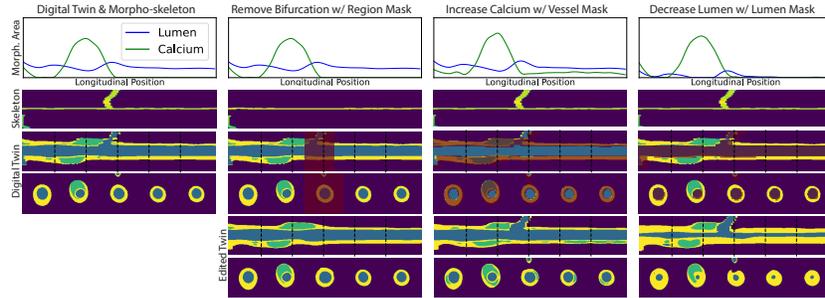
Fig. 8: Examples of morpho-skeleton based editing procedures applied to a coronary digital twin (left). The patient-specific morpho-skeleton is edited and combined with a spatial mask (red) to inpaint clinically relevant features into the artery.

edit eliminates the bifurcation within the artery, achieved through regional editing with a branch-less skeleton heatmap. The second edit increases the amount of calcium by modifying the morphological feature vectors and applying a tissue mask to maintain the boundaries of the vessel wall. The third edit results in luminal narrowing by reducing the lumen area through the morphological feature vectors and employing a dilated tissue mask over the lumen.

5   Limitations

Our method has several limitations. First, our morphological conditioning mechanism can only capture morphological features along a pre-defined direction, limiting its applicability. Second, the use of 3D voxels confines our method to low dimensional segmentation maps. Third, our guidance strategies requires the use of the decoder for every sampling step, increasing computational costs. Lastly, our topological regularization method reduces but does not eliminate topological errors, meaning that arteries sampled from our model must still undergo some amount topological post-processing.

6   Conclusions

We adapt diffusion models to conditionally generate or edit coronary atherosclerotic anatomy that can be used for comparative numerical simulation studies of coronary interventions. We introduce several methodological changes to the training process of latent diffusion models to enable precise control over topological quality, local micro-morphology, and global skeletal structure. We further compare and contrast different guidance strategies to ensure adherence to morpho-skeletal constraints, proposing a novel and lightweight guidance method that can utilize non-differentiable regressor functions to guide sampling. Our anatomic generation framework offers the flexibility and control of simplified parametric geometries while maintaining the realism and 3-dimensional complexity of patient-specific geometries derived from medical images.

Acknowledgments

References

1. Beetz, M., Banerjee, A., Grau, V.: Generating subpopulation-specific biventricular anatomy models using conditional point cloud variational autoencoders. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 75–83. Springer (2021)
2. Beetz, M., Banerjee, A., Grau, V.: Generating subpopulation-specific biventricular anatomy models using conditional point cloud variational autoencoders. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 75–83. Springer (2021)
3. Byrne, N., Clough, J.R., Valverde, I., Montana, G., King, A.P.: A persistent homology-based topological loss for cnn-based multiclass segmentation of cmr. IEEE transactions on medical imaging **42**(1), 3–14 (2022)
4. Chung, H., Kim, J., Mccann, M.T., Klasky, M.L., Ye, J.C.: Diffusion posterior sampling for general noisy inverse problems. arXiv preprint arXiv:2209.14687 (2022)
5. Clough, J.R., Byrne, N., Oksuz, I., Zimmer, V.A., Schnabel, J.A., King, A.P.: A topological loss function for deep-learning based image segmentation using persistent homology. IEEE transactions on pattern analysis and machine intelligence **44**(12), 8766–8778 (2020)
6. Dassault Systèmes: Abaqus Finite Element Analysis Software. Vélizy-Villacoublay, France, version 2023 edn. (2023), available from Dassault Systèmes, https://www.3ds.com
7. Dhariwal, P., Nichol, A.: Diffusion models beat gans on image synthesis. Advances in neural information processing systems **34**, 8780–8794 (2021)
8. Dong, P., Bezerra, H.G., Wilson, D.L., Gu, L.: Impact of calcium quantifications on stent expansions. Journal of biomechanical engineering **141**(2), 021010 (2019)
9. Doradla, P., Otsuka, K., Nadkarni, A., Villiger, M., Karanasos, A., Zandvoort, L.J.v., Dijkstra, J., Zijlstra, F., Soest, G.v., Daemen, J., et al.: Biomechanical stress profiling of coronary atherosclerosis: identifying a multifactorial metric to evaluate plaque rupture risk. Cardiovascular imaging **13**(3), 804–816 (2020)
10. Dou, H., Virtanen, S., Ravikumar, N., Frangi, A.F.: A generative shape compositional framework: Towards representative populations of virtual heart chimaeras. arXiv preprint arXiv:2210.01607 (2022)
11. Gandikota, R., Materzynska, J., Fiotto-Kaufman, J., Bau, D.: Erasing concepts from diffusion models. In: Proceedings of the IEEE/CVF International Conference on Computer Vision. pp. 2426–2436 (2023)
12. Gupta, S., Hu, X., Kaan, J., Jin, M., Mpoy, M., Chung, K., Singh, G., Saltz, M., Kurc, T., Saltz, J., et al.: Learning topological interactions for multi-class medical image segmentation. In: European Conference on Computer Vision. pp. 701–718. Springer (2022)
13. He, Y., Murata, N., Lai, C.H., Takida, Y., Uesaka, T., Kim, D., Liao, W.H., Mitsufuji, Y., Kolter, J.Z., Salakhutdinov, R., et al.: Manifold preserving guided diffusion. arXiv preprint arXiv:2311.16424 (2023)

14. Ho, J., Salimans, T.: Classifier-free diffusion guidance. arXiv preprint arXiv:2207.12598 (2022)
15. Holm, N.R., Andreasen, L.N., Neghabat, O., Laanmets, P., Kumsars, I., Bennett, J., Olsen, N.T., Odenstedt, J., Hoffmann, P., Dens, J., et al.: Oct or angiography guidance for pci in complex bifurcation lesions. New England Journal of Medicine **389**(16), 1477–1487 (2023)
16. Kadry, K., Gupta, S., Nezami, F.R., Edelman, E.R.: Probing the limits and capabilities of diffusion models for the anatomic editing of digital twins. arXiv preprint arXiv:2401.00247 (2023)
17. Kadry, K., Olender, M.L., Marlevi, D., Edelman, E.R., Nezami, F.R.: A platform for high-fidelity patient-specific structural modelling of atherosclerotic arteries: from intravascular imaging to three-dimensional stress distributions. Journal of the Royal Society Interface **18**(182), 20210436 (2021)
18. Karanasiou, G.S., Tsobou, P.I., Tachos, N.S., Antonini, L., Petrini, L., Pennati, G., Gijsen, F., Nezami, F.R., Tzafiri, R., Vaughan, T., et al.: Design and implementation of in silico clinical trial for bioresorbable vascular scaffolds. In: 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC). pp. 2675–2678. IEEE (2020)
19. Karanasiou, G.S., Tsompou, P.I., Tachos, N., Karanasiou, G.E., Sakellarios, A., Kyriakidis, S., Antonini, L., Pennati, G., Petrini, L., Gijsen, F., et al.: An in silico trials platform for the evaluation of stent design effect in post-implantation outcomes. In: 2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC). pp. 4970–4973. IEEE (2022)
20. Karras, T., Aittala, M., Aila, T., Laine, S.: Elucidating the design space of diffusion-based generative models. arXiv preprint arXiv:2206.00364 (2022)
21. Madani, A., Bakhaty, A., Kim, J., Mubarak, Y., Mofrad, M.R.: Bridging finite element and machine learning modeling: stress prediction of arterial walls in atherosclerosis. Journal of biomechanical engineering **141**(8), 084502 (2019)
22. Marlevi, D., Edelman, E.R.: Vascular lesion–specific drug delivery systems: Jacc state-of-the-art review. Journal of the American College of Cardiology **77**(19), 2413–2431 (2021)
23. Mori, H., Torii, S., Kutyna, M., Sakamoto, A., Finn, A.V., Virmani, R.: Coronary artery calcification and its progression: what does it really mean? JACC: Cardiovascular Imaging **11**(1), 127–142 (2018)
24. Myronenko, A.: 3d mri brain tumor segmentation using autoencoder regularization. In: Brainlesion: Glioma, Multiple Sclerosis, Stroke and Traumatic Brain Injuries: 4th International Workshop, BrainLes 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Revised Selected Papers, Part II 4. pp. 311–320. Springer (2019)
25. Nichol, A., Dhariwal, P., Ramesh, A., Shyam, P., Mishkin, P., McGrew, B., Sutskever, I., Chen, M.: Glide: Towards photorealistic image generation and editing with text-guided diffusion models. arXiv preprint arXiv:2112.10741 (2021)
26. Otsuka, F., Yasuda, S., Noguchi, T., Ishibashi-Ueda, H.: Pathology of coronary atherosclerosis and thrombosis. Cardiovascular diagnosis and therapy **6**(4), 396 (2016)
27. Pham, J., Kong, F., James, D.L., Marsden, A.L.: Virtual shape-editing of patient-specific vascular models using regularized kelvinlets. IEEE Transactions on Biomedical Engineering (2024)
28. Pham, J., Wyetzner, S., Pfaller, M.R., Parker, D.W., James, D.L., Marsden, A.L.: svmorph: Interactive geometry-editing tools for virtual patient-specific vascular anatomies. Journal of Biomechanical Engineering **145**(3), 031001 (2023)

29. Poletti, G., Antonini, L., Mandelli, L., Tsompou, P., Karanasiou, G.S., Papafaklis, M.I., Michalis, L.K., Fotiadis, D.I., Petrini, L., Pennati, G.: Towards a digital twin of coronary stenting: A suitable and validated image-based approach for mimicking patient-specific coronary arteries. Electronics **11**(3), 502 (2022)

30. Qiao, M., Basaran, B.D., Qiu, H., Wang, S., Guo, Y., Wang, Y., Matthews, P.M., Rueckert, D., Bai, W.: Generative modelling of the ageing heart with cross-sectional imaging and clinical data. In: International Workshop on Statistical Atlases and Computational Models of the Heart. pp. 3–12. Springer (2022)

31. Qiao, M., Wang, S., Qiu, H., de Marvao, A., O'Regan, D.P., Rueckert, D., Bai, W.: Cheart: A conditional spatio-temporal generative model for cardiac anatomy. arXiv preprint arXiv:2301.13098 (2023)

32. Ralapanawa, U., Sivakanesan, R.: Epidemiology and the magnitude of coronary artery disease and acute coronary syndrome: a narrative review. Journal of epidemiology and global health **11**(2), 169 (2021)

33. Sato, M., Bitter, I., Bender, M.A., Kaufman, A.E., Nakajima, M.: Teasar: tree-structure extraction algorithm for accurate and robust skeletons. In: Proceedings the Eighth Pacific Conference on Computer Graphics and Applications. pp. 281–449. IEEE (2000)

34. Sawaya, F.J., Lefèvre, T., Chevalier, B., Garot, P., Hovasse, T., Morice, M.C., Rab, T., Louvard, Y.: Contemporary approach to coronary bifurcation lesion treatment. JACC: Cardiovascular Interventions **9**(18), 1861–1878 (2016)

35. Shit, S., Paetzold, J.C., Sekuboyina, A., Ezhov, I., Unger, A., Zhylka, A., Pluim, J.P., Bauer, U., Menze, B.H.: cldice-a novel topology-preserving loss function for tubular structure segmentation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 16560–16569 (2021)

36. Silversmith, W., Bae, J.A., Li, P.H., Wilson, A.: Kimimaro: Skeletonize densely labeled 3d image segmentations (Sep 2021). `https://doi.org/10.5281/zenodo.5539913`, `https://doi.org/10.5281/zenodo.5539913`

37. Sonck, J., Nagumo, S., Norgaard, B.L., Otake, H., Ko, B., Zhang, J., Mizukami, T., Maeng, M., Andreini, D., Takahashi, Y., et al.: Clinical validation of a virtual planner for coronary interventions based on coronary ct angiography. Cardiovascular Imaging **15**(7), 1242–1255 (2022)

38. Song, J., Zhang, Q., Yin, H., Mardani, M., Liu, M.Y., Kautz, J., Chen, Y., Vahdat, A.: Loss-guided diffusion models for plug-and-play controllable generation (2023)

39. Stone, P.H., Saito, S., Takahashi, S., Makita, Y., Nakamura, S., Kawasaki, T., Takahashi, A., Katsuki, T., Nakamura, S., Namiki, A., et al.: Prediction of progression of coronary artery disease and clinical outcomes using vascular profiling of endothelial shear stress and arterial plaque characteristics: the prediction study. Circulation **126**(2), 172–181 (2012)

40. Straughan, R., Kadry, K., Parikh, S.A., Edelman, E.R., Nezami, F.R.: Fully automated construction of three-dimensional finite element simulations from optical coherence tomography. Computers in Biology and Medicine **165**, 107341 (2023)

41. Verhülsdonk, J., Grandits, T., Costabal, F.S., Krause, R., Auricchio, A., Haase, G., Pezzuto, S., Effland, A.: Shape of my heart: Cardiac models through learned signed distance functions. arXiv preprint arXiv:2308.16568 (2023)

42. Virmani, R., Burke, A.P., Farb, A., Kolodgie, F.D.: Pathology of the vulnerable plaque. Journal of the American College of Cardiology **47**(8S), C13–C18 (2006)

43. Wang, W., Bao, J., Zhou, W., Chen, D., Chen, D., Yuan, L., Li, H.: Semantic image synthesis via diffusion models. arXiv preprint arXiv:2207.00050 (2022)

44. Wiesner, D., Suk, J., Dummer, S., Nečasová, T., Ulman, V., Svoboda, D., Wolterink, J.M.: Generative modeling of living cells with so (3)-equivariant implicit neural representations. arXiv preprint arXiv:2304.08960 (2023)
45. Zhao, S., Wu, W., Samant, S., Khan, B., Kassab, G.S., Watanabe, Y., Murasato, Y., Sharzehee, M., Makadia, J., Zolty, D., et al.: Patient-specific computational simulation of coronary artery bifurcation stenting. Scientific reports **11**(1), 16486 (2021)

A Supplementary Materials Overview

- In Appendix B, we provide details on dataset curation and augmentation
- In Appendix C, we detail the mathematical formulation of our topological loss.
- In Appendix D, we discuss the extraction process for te morphological features and the lumenal skeleton.
- In Appendix E, we provide implementation details for our autoencoder and diffusion model
- In Appendix F, we provide implementation details for the neural networks used for morphological and skeletal regression.
- In Appendix G, we visualize the effect of adaptive null guidance with each conditioning mechanism in isolation. .
- In Appendix H, we provide details on our virtual angioplasty simulations which demonstrated that synthetic segmentation maps can be used for numerical simulation studies.

B Dataset Processing

Our study used 62 CCTA images obtained from the Precise Percutaneous Coronary Intervention Plan (P3) Study [37]. All data has been anonymized and approved for use by an IRB. The 3D lumen and vessel wall morphology were extracted using a methodology developed by HeartFlow Inc [37], and represented as 3D signed distance fields (SDFs) with a resolution of 0.25 mm. To obtain a straightened representation of the coronary artery, a virtual catheter was propagated along the main arterial branch with atherosclerotic lesions for each patient, capturing 2D cross-sections of the lumen and vessel wall at intervals of 0.2 mm. These 2D cross-sections were stacked to construct 3D segmentation maps of the lumen and vessel wall. The total number of 2D frames extracted was 28,533, where the average number of frames per artery was 460. For the segmentation of calcified tissue, thresholding was applied to the corresponding 2D frames extracted by the virtual pullback. The threshold was determined as the greater value between 450 Hounsfield units (HU) or the mean lumen density increased by one standard deviation. The anisotropic resolution of the segmentation maps was set at 60 μm in-plane and 0.2 mm out-of-plane to capture the longitudinal variation inherent in the arterial structure. To obtain the 3D segmentation maps that are fed into the VAE, we extract 3D patches consisting of 128 frames along the artery with a random starting frame. We further apply random rotations about the longitudinal axis to augment our dataset.

C Topological Loss Implementation

Our topological loss attempts to penalize topologically incorrect voxels within our autoencoded coronary segmentation maps. We specifically enforce that the vessel wall contains both lumen and calcium within any 2D frame. To enforce

such constraints, we use the topological loss as described by Gupta et al. [12] and adapt it to a compound multi-slice context, applying the loss to each 2D cross-section of the coronary segmentation map. Let $\mathbf{F} \in \mathbb{R}^{C \times H \times W \times D}$ be the multi-class segmentation map predicted by the VAE. To enforce that label-B contains label-A, we take a cross-sectional 2D label map $\mathbf{P} \in \mathbb{R}^{H \times W}$ and aim to find a critical pixel map that contains label-A pixels with a label-C neighbor and vice versa, where label-C is specifically defined as the union of all labels that are not label-A or label-B. First, each label is expanded by 3 pixels using a convolutional kernel K to obtain the neighborhood information N.

$$N_A := \mathbf{P}_A \circledast K \quad \text{and} \quad N_C := \mathbf{P}_C \circledast K, \tag{8}$$

where K is set as a 2D 4-connectivity kernel. We then use the neighborhood information to find the critical pixels of each mask ($\mathbf{V}_A$ and $\mathbf{V}_C$) that fall into the other masks neighborhood.

$$\mathbf{V}_A := \mathbf{P}_A \odot N_C \quad \text{and} \quad \mathbf{V}_C := \mathbf{P}_C \odot N_A \quad \text{and} \quad \mathbf{V} := \mathbf{V}_A \oplus \mathbf{V}_C, \tag{9}$$

where $\odot$ represents the Hadamard product and $\oplus$ represents the union operation. We set the ground truth label map $\mathbf{G} \in \mathbb{R}^{H \times W \times D}$ as the vessel wall and use the critical voxel map $\mathbf{V}_{3D}$ with a voxel-wise cross-entropy loss $\mathcal{L}_{ce}$

$$\mathcal{L}_T = \mathcal{L}_{ce}(\mathbf{F} \odot \mathbf{V}_{3D}, \mathbf{G} \odot \mathbf{V}_{3D}), \tag{10}$$

where the critical voxel map $\mathbf{V}_{3D}$ is found by stacking the critical pixel map $\mathbf{V}$ for each cross section along the depth dimension D.

D   Morpho-Skeletal Feature Extraction

D.1   Morphological Features

Morphological features were extracted from multi-channel coronary artery segmentation maps to quantify both 3D and 2D metrics and are detailed in Tab. 3. Volumes and cross-sectional areas were calculated from the mean number of positive voxels, while the stenosis ratio was calculated as the minimum lumen area divided by the average lumen volume. Vessel burden was defined as the ratio between vessel and lumen area at the site of minimum stenosis. Plaque was defined as the entire vessel cross-section, consisting of the lumen, vessel wall, and calcium. Circularity was computed as the ratio of area to the square of the perimeter. Thickness was obtained by doubling the maximum value of the Euclidean distance field for the 2D segmentation map. Arclength was calculated as the maximum angular span of the segmentation map contours per cross-sectional area. The centroid was measured as the distance between the centroid and the center of the image. When a morphological feature vector is used for conditioning, the feature curve was smoothed along the length of the artery with a Savitzky-Golay filter using a window length of 21, corresponding to approximately 16% of the feature vector length.

Table 3: Recorded morphological metrics for evaluation

2D Morph. Metrics
1) Lumen area
2) Vessel area
3) Plaque area
4) Plaque centroid
5) Plaque circularity
6) Calcium area
7) Calcium thickness
8) Calcium arclength

3D Morph. Metrics
1) Lumen volume
2) Stenosis ratio
3) Vessel volume
4) Vessel burden
5) Minimum plaque circularity
6) Mean plaque circularity
7) Calcium volume
8) Calcium length
9) Maximum calcium thickness
10) Mean calcium thickness
11) Maximum calcium arclength
12) Mean calcium arclength

Table 4: Hard Skel. Parameters

Parameter	Value
Const.	10
Scale	1.5
Prdf scale	1e5
Prdf exponent	5
Tick threshold	10

Table 5: Skel. Processing Parameters

Parameter	Value
σ_{skel}	1
Gaussian Kernel Size	3
Max pool Kernel Size	4
Max pool Stride	4

D.2   Skeletal Features

Our skeletal regressors consist of 1) a skeletonization step, which takes in a lumenal segmentation map, and 2) a pre-processing step to enable concatenation with the encoded latent representation. We leverage three types of skeletonization: non-differentiable (hard), differentiable (soft), and neural skeletonization. Our non-differntiable skeletonization process uses the Kimimaro library [36], which employs the TEASAR algorithm [33], to derive a high-quality skeleton. Our differentiable skeletonization process leverages an iterative soft skeletonization algorithm by Shit et al. [35], which is applied on lumen segmentation maps. Our neural skeletonization is done by a neural network trained to regress the processed output of the non-differentiable skeletal regressor (see Appendix F).

For hard skeletonization, the pre-processing step includes conversion into a binary voxel grid, Gaussian blurring with a standard deviation denoted by σ_{skel}, and normalization. To reduce the grid size, we employ max pooling and subsequently resize the skeleton map to fit the dimensions of the encoded latent representation. For soft skeletonization, we downsample the input lumen segmentation map to a size of $H/2 \times W/2 \times D$ before skeletonization to conserve GPU memory. The resulting skeleton is then upscaled to its original resolution and is then processed similarly to hard skeletonization. For neural skeletonization, the

processing step is equivalent to that of hard skeletonization. The configuration details for the hard skeletonization and associated processing can be found in Tabs. 4 and 5 respectively.

The outputs of our skeletonization methods along with their processed outputs are shown in Fig. 10. Our hard skeletonization method creates intact skeletons given an input lumen segmentation map. Processing the skeletons with Gaussian blurring and max pooling retains the skeletal connectivity after resizing to a coarser resolution. In contrast, the soft skeletonization method produces disconnected skeletal branches, which may or may not be repaired after processing.

E   Latent Diffusion Model Implementation

We adapt the VAE and LDM architectures specified by Kadry et al. [16] to the case of coronary segmentation map generation. The input and output channels of the VAE are set to a value of 4 to correspond to the number of tissues. The number of input channels for the LDM was set to a value of 6 such that 3 channels correspond to the encoded latent and 3 channels correspond to two morphological conditioning maps and one skeletal conditioning map. The hyperparameters and training configuration for the VAE and LDM are contained in Tab. 6 and Tab. 7 respectively.

Table 6: VAE Config.

Hyperparameter	Value
lr	1×10^{-5}
Iterations	4×10^4
Batch Size	2
Num. Channels	[64, 128, 192]
Num. Res. Blocks	2
Downscaling Factor	4
λ_{recon}	1
λ_{KL}	1×10^{-6}
λ_{topo}	2×10^{-4}

Table 7: LDM Config.

Hyperparameter	Value
Training	
lr	1×10^{-5}
Iterations	5×10^3
Batch Size	1
Num. Channels	[128, 256, 384]
Num. Res. Blocks	2
Num. Attn. Heads	1
Attn. Res.	8, 4, 2
σ_{data}	1
$p(\sigma)$ mean	1
$p(\sigma)$ std	1.2
Sampling	
σ_{min}	1×10^{-2}
σ_{max}	80
ρ	3

Neural Regressor Validation
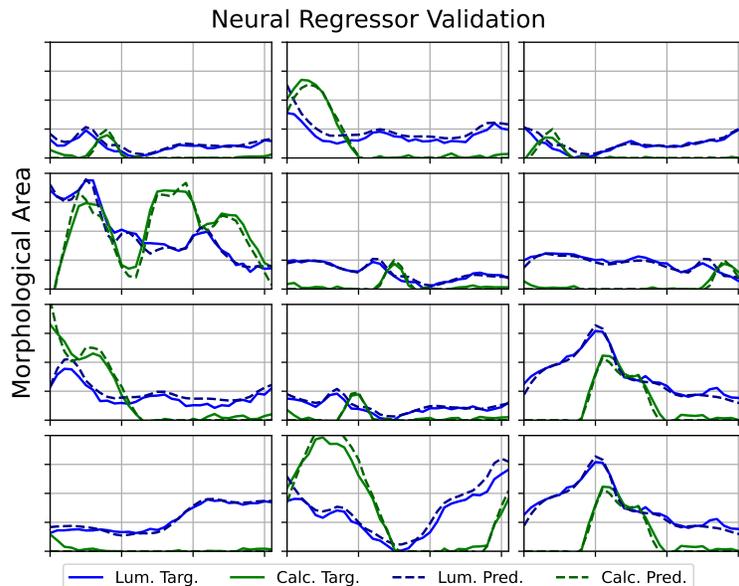
Fig. 9: Comparison of target morphological areas against predictions by our neural morphological regressor. Our regressor can predict the morphological features associated with coronary segmentation maps to a high degree of fidelity.

F   Morpho-skeletal Regressor Implementation

F.1   Morphological Regressor Training

The neural morphological regressor $\bar{\mathcal{M}}_\phi$ consists of a 3D encoder, an adaptive pooling layer, and an MLP head. To regress the morphological features, we first encode the coronary segmentation map $\mathbf{x}$ into a latent representation of size $\alpha \times h \times w \times d$ where α is the number of latent channels. This latent representation is then processed through an adaptive average pooling layer, resulting in a feature map of dimensions $\alpha \times 1 \times 1 \times d$. Subsequently, an MLP head with two fully connected layers is applied to transform this feature map into 1D morphological features of size $m \times d$, where m is the number of regressed morphological features. The output is then resized to match the output shape of our original morphological regressor $\mathcal{M}$. Tab. 8 details the hyperparameters and training details for the morphological regressor. The network is trained to minimize a morphological regression loss $\mathcal{L}_{\mathrm{morph}}$ with the Adam optimizer.

$$\mathcal{L}_{\mathrm{morph}} = \mathcal{L}_{\mathrm{MSE}}(\mathcal{M}(\mathbf{x}), \bar{\mathcal{M}}_\phi(\mathbf{x})) \tag{11}$$

The performance of our trained morphological regressor can be seen in Fig. 9, in which a regressor is applied to coronary segmentation maps drawn from the validation set. The predictions from our morphological regressor show good correlation with the ground truth morphology as measured by the ground truth morphological regressor $\mathcal{M}$.

F.2   Skeletal Regressor Training

For the skeletal regressor $\bar{\mathcal{S}}_\psi$, we employ the SegResNet architecture [24], a U-net based architecture designed to regress the skeletal heatmap for an input coronary segmentation map $\mathbf{x}$. Following the decoder stage, the resulting skeleton voxel grid undergoes blurring and maxpooling similar to the output of the hard skeletonization process. Tab. 9 details the hyperparameters and training details for the skeletal regressor. The network is trained to minimize a skeletal regression loss $\mathcal{L}_{\mathrm{skel}}$ with the Adam optimizer.

$$\mathcal{L}_{\mathrm{skel}} = \mathcal{L}_{\mathrm{MSE}}(\mathcal{S}(\mathbf{x}), \bar{\mathcal{S}}_\psi(\mathbf{x})) \tag{12}$$

The skeletal regressor predictions are visualized in the rightmost column of Fig. 10. It can be seen that while the regressor was trained on the processed output of the hard skeletonization method after pre-processing, it has learned to regress skeletons that are similar to the soft skeletonization process, but with a higher degree of continuity.
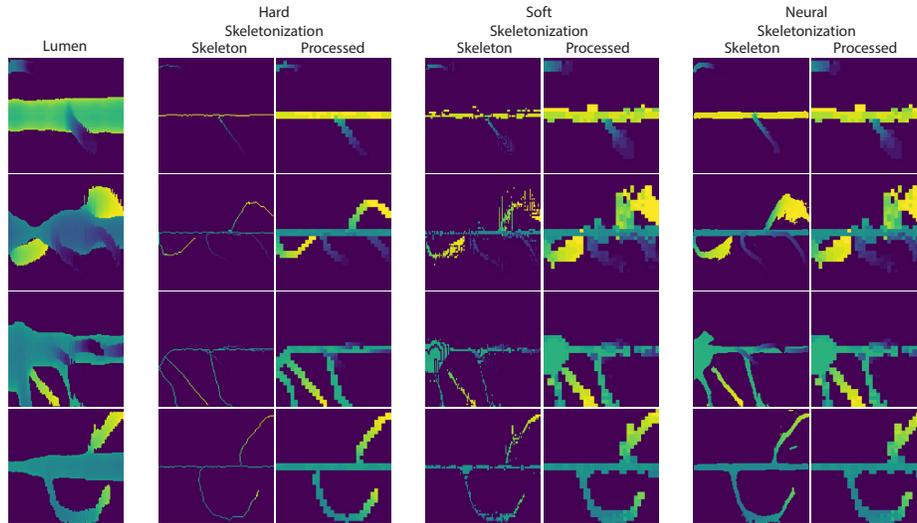
Fig. 10: Comparison of different skeletonization methods before and after processing. Left column displays lumen depth maps from validation set. Center left, center right, and rightmost columns indicate skeletonization outputs for hard, soft, and neural regressors respectively. Hard skeletonization outperforms both soft skeletonization and the neural skeletal regressor.

G   Additional Ablation Results

In this section, we further study the effect of ablating one or more conditioning signals during LDM training. We use four LDMs corresponding to 1) no-

Table 8: Morph. Regressor Config.

Parameter	Value
lr	1×10^{-4}
Iterations	4000
Batch size	1
Num. Channels	[32, 64, 96]
Num. Res. Blocks	2
Num. Latent Channels	64
MLP Channels	1024

Table 9: Skel. Regressor Config

Parameter	Value
lr	1×10^{-4}
Iterations	4000
Batch size	1
Encoder Channels	[16, 32, 32, 64]
Decoder Channels	[64, 64, 64, 64]

conditioning, 2) morphological conditioning only, 3) skeletal conditioning only, and 4) joint conditioning. For each LDM, we sample 100 segmentation maps with 25 sampling steps, sweeping over a range of guidance weights for all but the unconditional model. Unless stated otherwise, we used adaptive null guidance with a weight of 5, and the target morpho-skeletons were derived from the validation set. In addition to recording metrics for sampling quality and conditional fidelity (Fig. 18), we visualize the distribution of morphological features as a set of 1D KDE plots (Figs. 11 to 14) and visualize the segmentation maps in Fig. 17.

Unconditional Generation Fig. 11 shows that unconditional generation can recapitulate the distribution in various morphological features in the training set, with good coverage compared to the validation set. Unconditionally generated segmentation maps can be seen in Fig. 17, where unconditional sampling can synthesize a wide array of morphological and skeletal features.

Joint Conditioning Fig. 12 shows that guidance improves the similarity of the synthetic distribution to the validation distribution for most morphological features. This can also be seen in Fig. 16, where guidance increases morphological conditioning fidelity for a variety of target morphological features. Similarly, Fig. 15 displays skeleton depth maps for various adaptive null guidance values. Our joint conditioning framework can produce a wide variety of skeletal configurations without components that are disconnected from the main branch. Synthetic segmentation maps produced by joint conditioning can be seen in Fig. 17, which exhibits a wide array of morphological and skeletal features.

Morphological Conditioning Only Fig. 13 shows that morphological guidance slightly improves the similarity of between the morphological distributions of synthetic and validation segmentation maps. This can be seen in Fig. 18, which shows enhanced morphological conditioning fidelity with increasing guidance strength, similar to joint guidance. segmentation maps sampled by morphological conditioning can be seen in Fig. 17, which exhibit non-physiological bifurcation structures.

No Conditioning

Lum. Vol.   Ves. Vol.   Calc. Vol.   Ves. Burden

Stenosis Rat.   Mean Plaque Circ.   Min Plaque Circ.   Calc. Len.

Mean Calc. Thick.   Max Calc. Thick.   Mean Calc Arclen.   Max Calc. Arclen.

Training   Validation   Unconditional

Joint Conditioning

Lum. Vol.   Ves. Vol.   Calc. Vol.   Ves. Burden

Stenosis Rat.   Mean Plaque Circ.   Min Plaque Circ.   Calc. Len.

Mean Calc. Thick.   Max Calc. Thick.   Mean Calc Arclen.   Max Calc. Arclen.

Validation   Unguided   Guided

Fig. 11: KDE plots of morphological features comparing the training set, validation set, and the set of segmentation maps sampled through unconditional generation.
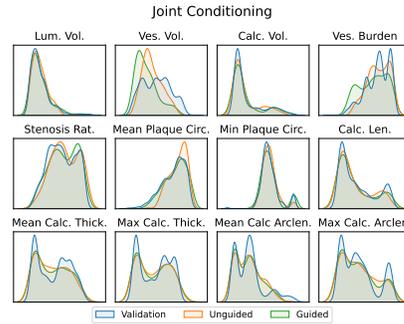
Fig. 12: KDE plots comparing the morphological feature distributions of synthetic segmentation maps produced by joint conditioning.

Skeletal Conditioning Only Fig. 14 shows that unguided skeletal conditioning does not significantly change the morphological distribution of features as compared to the validation set. However, with skeletal guidance, the morphological distributions shift towards exhibiting smaller lumen, vessel, and calcium volumes. Fig. 18 shows that skeletal guidance alone can slightly improve skeletal fidelity as compared to joint guidance. This comes at the cost of reduced recall and increased Frechet distance, likely due to the morphological bias towards smaller volumes. Segmentation maps produced by skeletal conditioning can be seen in Fig. 17, where skeletal conditioning produces a variety of realistic bifurcation structures.

H   Virtual Angioplasty Simulation

To demonstrate the compatibility of our synthetic arteries with numerical simulation, we simulate a virtual intervention one two arteries generated from a patient-specific morpho-skeleton. Our synthetic segmentation map is processed by the methodology outlines by Straughn et al. [40], in which the lumen, vessel wall, and calcium labels are morphologically interpolated and processed to produce an isotropic segmentation map which is used to produce a 3-dimensional multi-material tetrahedral mesh. We then utilize the 3D mesh for virtual angioplasty simulations using the finite element method through the commercial code ABAQUS Explicit [6]. Briefly, the simulation platform consists of the coronary artery, stent, and deployment balloon. The artery materials included vessel wall, which was defined as a relatively soft hyperelastic-plastic material, and calcium, which was defined as a stiff elastic material [40]. The stent was geometrically modelled with hexahedral elements and exhibits an elasto-plastic constitutive law [29]. The balloon was geometrically modelled with shell elements and ex-
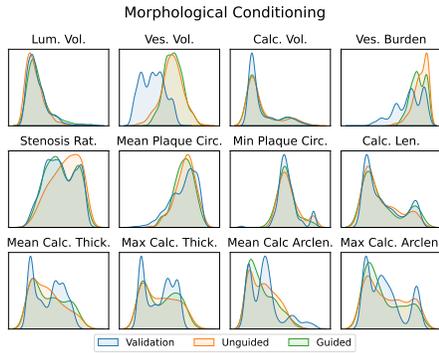
Fig. 13: KDE plots comparing the morphological feature distributions of synthetic segmentation maps produced by morphological conditioning alone.
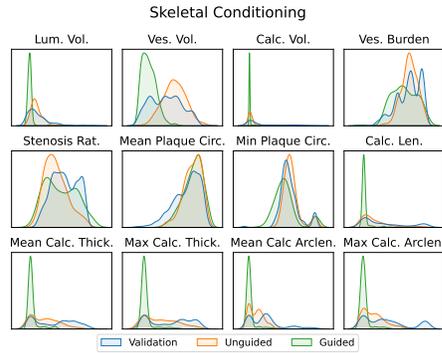
Fig. 14: KDE plots comparing the morphological feature distributions of synthetic segmentation maps produced by skeletal conditioning alone.

hibits a hyperelastic constitutive law [29]. The first stage of the simulation consists of stent crimping, in which the stent begins at a nominal diameter of 3mm and is compressed to a smaller diameter of 1mm with rigid crimping planes. In the second stage, the stent and a folded balloon are inserted into the straightened artery model, at which point the balloon is inflated to a pressure of 14 atm, expanding the stent into the arterial wall.
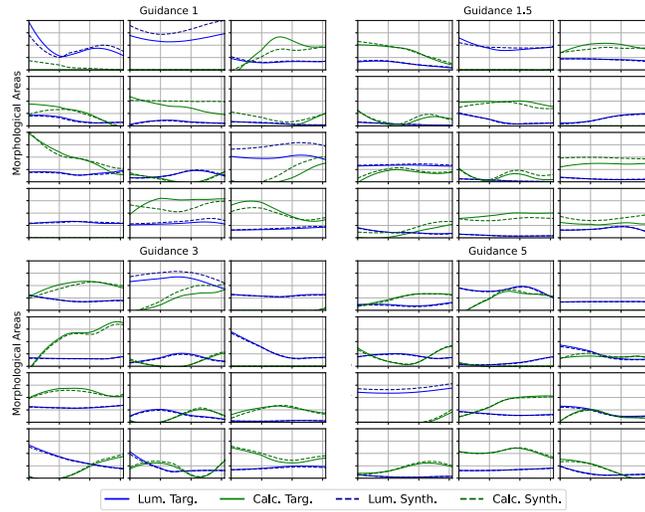
Fig. 15: Comparison of target morphological features for lumen and calcium against the equivalent morphological features derived from synthetic segmentation maps using adaptive null guidance at various guidance weights. Our guidance method enhances morphological conditional fidelity as guidance weight increases.
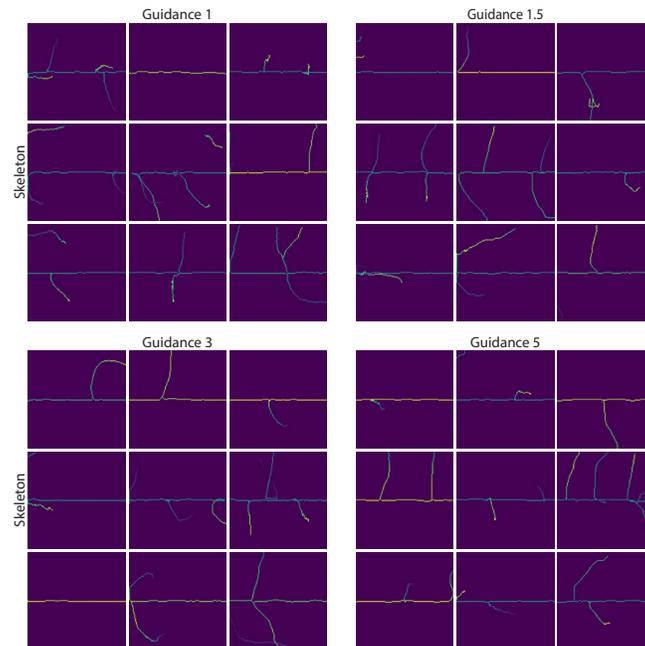
Fig. 16: Comparison of skeleton depth maps derived from synthetic segmentation maps using adaptive null guidance at various guidance weights.
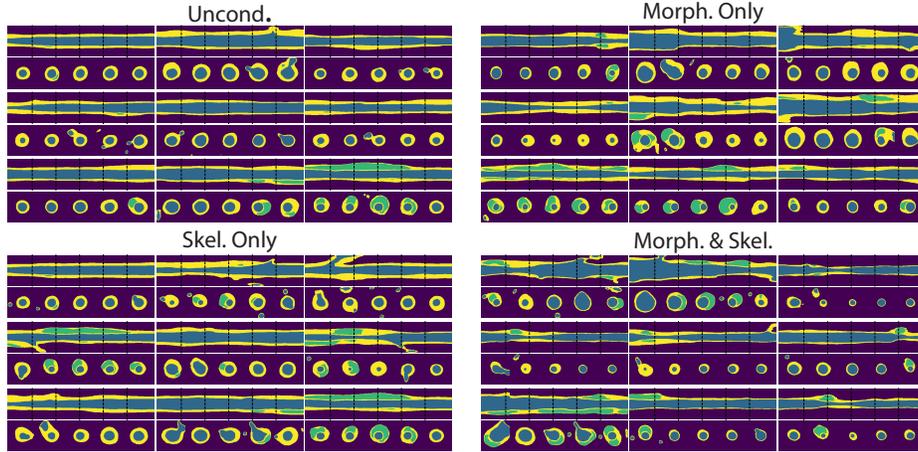
Fig. 17: Example segmentation map cross sections for unconditional sampling, morphological conditioning only, skeletal conditioning only, and joint morphological and skeletal conditioning. No guidance was used during sampling.
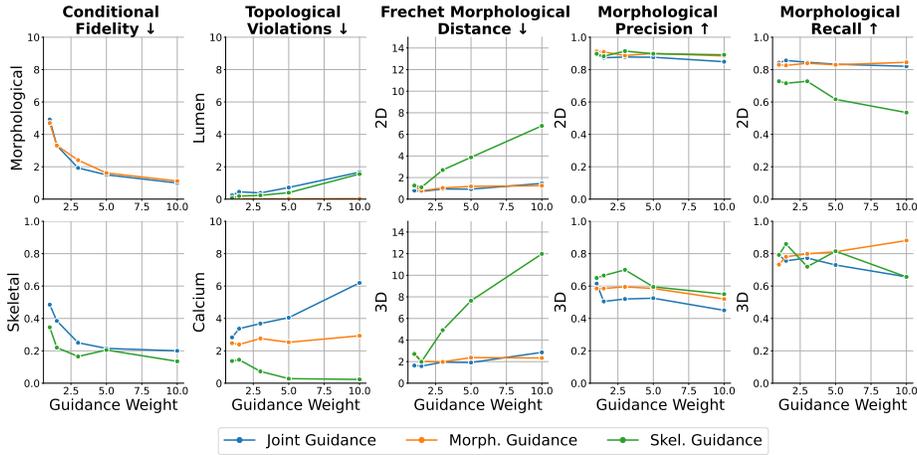
Fig. 18: Conditional fidelity and sampling quality metrics for adaptive null guidance using morphological, skeletal, and joint conditioning.