CoroLift: Fast Single-View 3D Coronary Reconstruction with Anatomical Priors

Anonymous authors
Paper under double-blind review

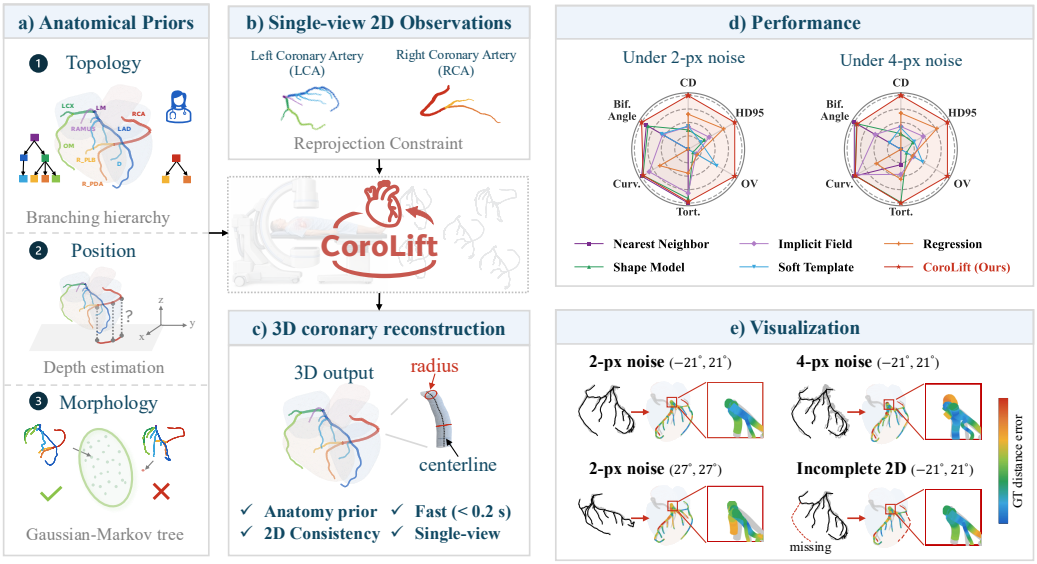


Figure 1: **Overview of CoroLift.** CoroLift combines anatomical priors with patient-specific 2D observations to resolve depth ambiguity in single-view 3D coronary reconstruction.

Abstract

3D coronary reconstruction from 2D observations is important for clinical hemodynamic assessment, yet is inherently ill-posed, particularly in the single-view setting without multi-view geometric cues. Existing methods therefore typically rely on multiple views, leaving single-view reconstruction largely unexplored. Our key idea is to leverage anatomical priors to constrain the missing 3D geometry from a single view. We propose **CoroLift**, a coarse-to-fine framework for single-view 3D coronary reconstruction that integrates anatomical priors on topology, position, and morphology into initialization and refinement. To construct these priors, we decompose coronary trees into slots between bifurcations and endpoints, with shared anatomical identities across patients. We derive a topology prior from population-level connectivity statistics, a position prior from spatial distributions, and a morphology prior that combines family-specific statistical shape models with tree-structured inter-slot dependencies. CoroLift first formulates reconstruction as depth estimation along imaging rays. Topology and position priors guide residual-based 3D endpoint localization and continuous branch-depth estimation, yielding an initial 3D coronary tree. The morphology prior then regularizes the joint vessel geometry while maintaining consistency with the 2D observations. Experiments show that CoroLift reduces geometric reconstruction error by 26.6% compared with existing baselines, while substantially improving morphological fidelity and reconstructing each 3D coronary tree in only 0.17 s.

Reviewers: please read the Reviewer Guidelines (iclr.cc/Conferences/2027/ReviewerGuidelines) and the AI Policy for Reviewers (iclr.cc/Conferences/2027/AIPolicyForReviewers). If you used AI to expand, edit, or polish your review, please provide the input text to the LLM. Better still, consider skipping the LLM and submitting your original text: we, and the authors, are much more interested in your unedited thoughts than in what an LLM has to say. AI-assisted or not, you are putting your name and reputation behind your review: LLM-generated falsehoods, hallucinations or misrepresentations are subject to disciplinary action, which may include desk-rejecting all papers you have authored.

1 INTRODUCTION

Recovering 3D coronary anatomy from 2D observations is fundamentally ill-posed (Çimen et al., 2016a), as different 3D vessel configurations can produce similar 2D projections. This ambiguity is particularly severe in the single-view setting, where multi-view geometric cues are unavailable. Existing methods therefore typically rely on multiple views and geometric correspondence (Chen & Carroll, 2000; Yang et al., 2009; Galassi et al., 2018; Banerjee et al., 2020). Learning-based approaches offer an alternative by learning 3D coronary regularities directly from paired data (Bransby et al., 2023; Iyer et al., 2023; Wang et al., 2024; Fu et al., 2024; Wang et al., 2025). However, learning a direct 2D-to-3D mapping does not eliminate the fundamental depth ambiguity of a single projection, particularly when paired 2D-3D coronary data are limited Consequently, single-view 3D coronary reconstruction remains largely unexplored.

Our key insight is that population-level anatomical regularities can constrain the solution space of this underdetermined problem. Coronary trees exhibit characteristic patterns in branch connectivity, spatial organization, and morphology. For example, as illustrated in Fig. 1(a), the left main (LM) coronary artery typically bifurcates into the left anterior descending (LAD) and left circumflex (LCX) arteries. The LAD typically courses along the anterior interventricular groove, whereas the LCX follows the left atrioventricular groove (Clemente et al., 2020). We therefore model these population-level coronary regularities explicitly as an *Anatomy Prior* to complement the geometric information missing from single-view observations during patient-specific 3D reconstruction.

We propose **CoroLift**, a coarse-to-fine framework that integrates anatomical priors on topology, position, and morphology into single-view 3D coronary reconstruction. To construct these priors, we decompose each coronary tree into *slots*, defined as vessel segments between bifurcations and endpoints, with shared anatomical identities across patients. This representation establishes fine-grained anatomical correspondence across patients, enabling population-level modeling of corresponding coronary segments. Specifically, we derive a topology prior from branch connectivity statistics, a position prior from spatial distributions of corresponding slots, and a morphology prior that combines family-specific statistical shape models with tree-structured inter-slot dependencies using a Gaussian-Markov model (Lauritzen, 1996; Willsky, 2002). CoroLift first performs 3D initialization by formulating reconstruction as depth estimation along imaging rays. The topology and position priors guide residual-based 3D endpoint localization and continuous branch-depth estimation, lifting the observed 2D centerlines into an initial 3D coronary tree. CoroLift then refines the initial tree by regularizing the joint vessel geometry with the morphology prior while enforcing consistency with the observed 2D projection.

We evaluate CoroLift on 3D coronary anatomies derived from coronary computed tomography angiography (CCTA), using calibrated angiographic projections under a strict single-view protocol. We further evaluate diverse observation conditions, including centerline noise, acquisition-view variations, and incomplete observations. CoroLift consistently improves reconstruction accuracy and morphological fidelity across these conditions, while remaining robust to degraded observations. Our main contributions are:

- We propose CoroLift, a coarse-to-fine framework for single-view 3D coronary reconstruction that formulates the problem as anatomy-guided depth estimation followed by morphology-constrained refinement.
- We introduce an Anatomy Prior built on fine-grained anatomical correspondence across patients, capturing population-level regularities in coronary topology, position, and morphology.
- We conduct extensive experiments under a strict single-view protocol and diverse observation conditions, demonstrating improved reconstruction accuracy, morphological fidelity, and robustness over existing baselines.

2 RELATED WORK

Multi-view coronary reconstruction. Classical coronary reconstruction relies on multiple angiographic views, using imaging geometry and cross-view correspondence for triangulation or reprojection optimization (Chen & Carroll, 2000; Yang et al., 2009; Galassi et al., 2018; Banerjee et al.,

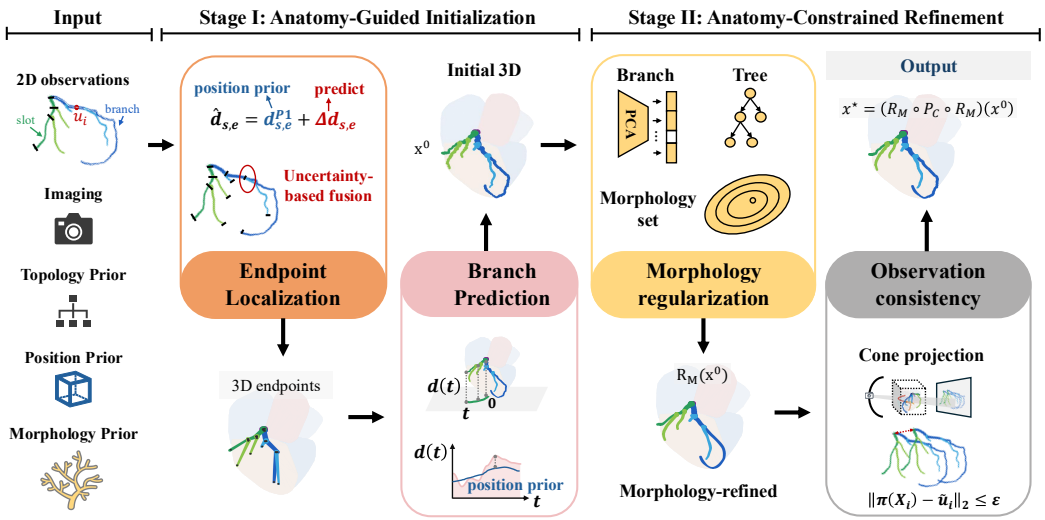


Figure 2: **CoroLift framework.** CoroLift reconstructs a 3D coronary tree from a single 2D view in two stages. **Stage I** uses topology and position priors for anatomical endpoint localization and continuous branch-depth estimation. **Stage II** uses the morphology prior and observation consistency to refine the initial 3D tree.

2020; Çimen et al., 2016a). Statistical approaches further introduce structural information through probabilistic centerline models and cardiac priors (Çimen et al., 2014; 2016b). Recent learning-based methods incorporate learned representations into sparse- and multi-view reconstruction, including graph-based, implicit, generative, and diffusion-based models (Iyer et al., 2023; Bransby et al., 2023; Wang et al., 2024; Fu et al., 2024; Wang et al., 2025; Lau et al., 2026). These methods nevertheless exploit information from multiple projections. In contrast, CoroLift targets the more underconstrained single-view setting, where cross-view geometric cues are unavailable and missing depth must instead be constrained by anatomical knowledge.

Vascular anatomy priors and generative models. Statistical shape models characterize anatomical variability through a mean shape and principal modes of variation (Cootes et al., 1995; Heimann & Meinzer, 2009), and statistical anatomical priors have also been explored for coronary reconstruction (Çimen et al., 2014; 2016b). Recent generative models extend vascular modeling to complete trees using implicit, diffusion, hierarchical, and autoregressive representations (Sinha & Hamarneh, 2024; Kuipers et al., 2024; Prabhakar et al., 2024; Feldman et al., 2025a;b; Chen et al., 2025; Guo et al., 2025). Rather than synthesizing vascular anatomy, CoroLift organizes population-level coronary regularities into topology, position, and morphology priors and uses them to constrain patient-specific reconstruction from an observed 2D tree.

Prior-guided inverse problems. Inverse problems recover latent signals by combining consistency with observed measurements and prior knowledge (Engl et al., 1996). Learned approaches extend this principle by coupling physical forward models with learned priors, including plug-and-play methods (Venkatakrishnan et al., 2013; Kamilov et al., 2023) and score- or diffusion-based methods (Song et al., 2022; Chung et al., 2022; Zhu et al., 2023). CoroLift applies the same general principle in a structured geometric setting: calibrated projection geometry defines the observation constraint, while population-level anatomical priors constrain the unobserved 3D coronary geometry.

3 METHOD

We consider 3D coronary reconstruction from traced 2D centerlines u with anatomical labels under a calibrated projection π. Since $\pi(x) \approx u$ does not uniquely determine the 3D coronary tree x, we construct an Anatomy Prior $\mathcal{A} = \{\mathcal{A}_{\text{topo}}, \mathcal{A}_{\text{pos}}, \mathcal{A}_{\text{morph}}\}$ from population regularities in topology,

position, and morphology. The topology and position priors guide 3D initialization, followed by morphology regularization and observation consistency.

3.1 Anatomy Prior Construction

The Anatomy Prior captures complementary population-level regularities in coronary topology, position, and morphology.

Topology prior. We harmonize vessel annotations in coronary computed tomography angiography (CCTA) into 11 major vessel families and exclude minor branches of limited clinical relevance. Partitioning each vessel at anatomical bifurcations yields 43 *slots*, defined as vessel segments between anatomical keypoints (endpoints or bifurcations) with consistent identities across patients. For each patient, the parent $p(s)$ of an observed slot s is determined from shared endpoints of the labeled 2D traces. This slot representation establishes anatomical correspondence across patients and provides the structural basis for position and morphology modeling.

Position prior. The position prior comprises slot-specific anatomical-keypoint statistics μ_s^{P1} and family-specific vessel trajectories $\bar{P}_f$. The latter are obtained by concatenating same-family slots, resampling by normalized arc length, and averaging corresponding 3D points over the training population.

Morphology prior. To characterize morphology independently of spatial position, we represent each slot by its endpoint-anchored residual curve, obtained by subtracting the straight chord between its endpoints. For each vessel family, we learn a shared PCA basis from the residual curves of all training slots in that family; each slot s is then represented by a $K_M = 16$ dimensional morphology coefficient c_s in its family-specific basis. To capture dependencies between anatomically corresponding slots, we model these coefficients using a Gaussian-Markov tree (Lauritzen, 1996; Willsky, 2002):

$$c_s = A_s c_{p(s)} + b_s + \varepsilon_s, \qquad \varepsilon_s \sim \mathcal{N}(0, Q_s). \tag{1}$$

Applying these dependencies to the patient-specific slot graph induces a joint morphology distribution $c \sim \mathcal{N}(\mu, \Sigma)$ over the observed coronary tree. Additional construction details are provided in Appendix A.1.

3.2 Anatomy-guided Initialization

Given an observed centerline point u, its 3D position lies along the corresponding imaging ray:

$$X(u) = o + d(u)q(u), \tag{2}$$

where o is the imaging source, $q(u)$ is the unit imaging ray, and $d(u)$ is the unknown depth. We estimate $d(u)$ from sparse anatomical endpoints to continuous vessel geometry using the topology and position priors.

Anatomical endpoint localization. For each endpoint e of slot s, the position prior provides a reference depth $d_{s,e}^{P1}$ under the current imaging view. An Endpoint Transformer predicts a patient-specific residual and uncertainty:

$$\hat{d}_{s,e} = d_{s,e}^{P1} + \Delta d_{s,e}, \qquad \sigma_{s,e} > 0. \tag{3}$$

The uncertainty is learned using a heteroscedastic Gaussian likelihood (Kendall & Gal, 2017). Predictions corresponding to the same anatomical node are fused by uncertainty-weighted averaging,

$$d_j^{\text{fuse}} = \frac{\sum_{(s,e)\in\mathcal{E}(j)} \sigma_{s,e}^{-2} \hat{d}_{s,e}}{\sum_{(s,e)\in\mathcal{E}(j)} \sigma_{s,e}^{-2}}. \tag{4}$$

The fused depths are further refined using slot-specific population geometry, yielding patient-specific 3D anatomical nodes $\{Y_j\}$.

Continuous depth prediction. To estimate depth continuously between localized anatomical nodes, we use the family-specific population trajectory as a reference and predict a low-dimensional patient-specific residual:

$$d_v(t) = \bar{d}_{f(v)}(t) + \sum_{q=0}^{Q-1} a_{v,q} T_q(2t-1), \qquad Q = 12, \tag{5}$$

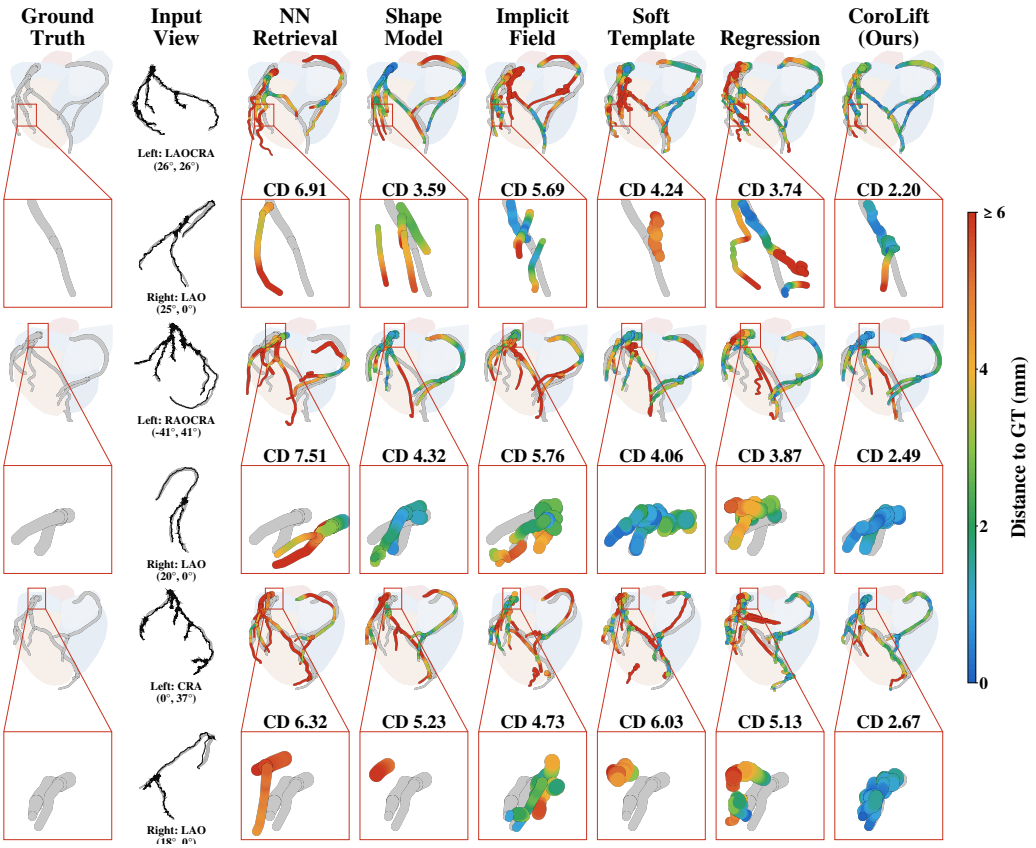


Figure 3: **Qualitative comparison on representative test cases** at 2 px trace noise. Reconstructions are colored by distance to the ground-truth centerlines (grey), with case-wise CD (mm) reported below. Enlarged regions highlight local branch and bifurcation geometry.

where $t \in [0, 1]$ is normalized arc length, $\bar{d}_{f(v)}(t)$ is the depth profile induced by the family-specific trajectory prior, T_q is the q-th Chebyshev polynomial, and $a_{v,q}$ is predicted by a tree-level Transformer. The profile is anchored to the localized anatomical nodes and lifted along the observed rays using Eq. equation 2, yielding the initial reconstruction x^0. Architecture and implementation details are provided in Appendix A.2.

3.3 Anatomy-constrained Refinement

The anatomy-guided initialization resolves the missing depth while preserving the observed imaging rays, but its joint 3D morphology is not explicitly constrained by the population morphology prior. We therefore refine x^0 by first regularizing its joint morphology and then restoring consistency with the observed 2D centerlines.

Morphology regularization. To constrain the reconstructed tree toward anatomically plausible joint configurations, we regularize its morphology using the patient-specific distribution induced by the population morphology prior. Using the joint morphology distribution $\mathcal{N}(\mu, \Sigma)$, we define the plausible morphology region

$$\mathcal{M}_t = \left\{ c : (c - \mu)^\top \Sigma_\alpha^{-1} (c - \mu) \le t^2 D \right\}, \tag{6}$$

where $c \in \mathbb{R}^D$ is the joint morphology vector, $D = K_M n$ for n observed slots, and Σ_α is the covariance regularized with α along near-null directions. We use $\alpha = 0.1$ and $t = 0.5$. The morphology operator $R_\mathcal{M}$ maps coefficients outside $\mathcal{M}_t$ to its boundary and reconstructs the corresponding slot curves. Since the residual morphology representation vanishes at both endpoints, $R_\mathcal{M}$ preserves the localized anatomical nodes and their connectivity.

Observation consistency. Anatomical plausibility alone is insufficient, as morphology regularization can perturb the projected vessel trajectories away from the observed 2D centerlines. We therefore apply an observation-consistency operator $P_{\mathcal{C}}$ that enforces

$$\|\pi(X_i) - \tilde{u}_i\|_2 \leq \varepsilon, \qquad \varepsilon = 2 \text{ pixels}, \tag{7}$$

where X_i is a reconstructed 3D point and $\tilde{u}_i$ is its corresponding point on the observed centerline, determined from the current projected curve by normalized arc length. $P_{\mathcal{C}}$ applies local transverse corrections to points that violate this tolerance. The morphology regularization and observation-consistency operators are applied in alternation, and the final reconstruction is

$$x^\star = (R_{\mathcal{M}} \circ P_{\mathcal{C}} \circ R_{\mathcal{M}})(x^0). \tag{8}$$

where the second $R_{\mathcal{M}}$ restores anatomical plausibility after $P_{\mathcal{C}}$ pulls the reconstruction back toward the observed 2D centerlines. Details of the morphology regularization and observation-consistency update are provided in Appendix A.3.

4 Experiments

4.1 Setup

Dataset. We use coronary artery trees from 836 patients in the public ImageCAS coronary CT angiography dataset (Zeng et al., 2023), with per-point vessel radius and anatomical branch labels. The patients are split into $671/83/82$ for training, validation, and test. Synthetic angiographic observations are generated from the CCTA centerlines using calibrated C-arm cone-beam geometry. Unless otherwise specified, all experiments follow a strict single-view protocol with 2 px trace noise. Population statistics used by the anatomy prior are estimated exclusively from the training split; model selection and hyperparameter choices use the validation split. Data generation, acquisition geometry, and trace-noise calibration are detailed in Appendix B.

Metrics. For centerline geometry, we report symmetric Chamfer distance (CD, mm), the 95th-percentile nearest-neighbor distance (HD95, mm), and Rotterdam overlap (OV) (Schaap et al., 2009), with CD as the primary metric. Morphological fidelity is assessed by the absolute errors in tortuosity, mean curvature, and bifurcation angle. Vessel radius is recovered analytically from the observed 2D vessel width and reconstructed depth; radius evaluation is reported in Appendix C.4. CD, HD95, and OV are computed on densely arc-length-resampled centerlines, while the morphology metrics use uniformly arc-length-resampled vessel segments. All metrics are first computed per test case and reported as medians over the 82 test cases. Metric definitions and evaluation details are provided in Appendix B.4.

4.2 Main Results

Comparison with reconstruction baselines. Table 1 compares CoroLift with five reconstruction baselines organized by which anatomy component they model. Template-based baselines without positional adaptation (NN Retrieval (Ramakrishna et al., 2012), Shape Model (Cootes et al., 1995)) give CD 7.88 and 6.43 mm at 2 px noise; Implicit Field (Wang et al., 2024) recovers geometry without an anatomy prior and reaches 6.19 mm. Adding topology and position priors (Soft Template (Hwang et al., 2021), Regression (Tai et al., 2015; Socher et al., 2011)) further reduces CD to 6.15 and 5.25 mm, but the tree-structured regressor still yields a large bifurcation-angle error (48.6°). CoroLift is the only method that jointly models topology, position, and morphology: it reaches 3.85 mm CD and 0.110 OV with the lowest bifurcation-angle error (24.2°) at competitive inference time ($\sim$ 0.17 s). CoroLift maintains its advantage under increased trace noise at 4 px. Additional results and qualitative reconstructions are provided in Appendix C.

4.3 Ablation

We ablate the three components of the anatomy prior (topology, position, morphology), the anatomy-guided initialization, the morphology model design, and the anatomy-constrained refinement. All experiments keep the remaining components fixed at their default settings.

Table 1: **Single-view 3D coronary reconstruction.** Methods are evaluated at 2 and 4 px trace noise; check marks indicate the anatomical priors used by each method. We report centerline accuracy, morphological fidelity, and inference time. CD/HD95 in mm; Bif. Angle in degrees; Time in seconds.

Method	Anatomical Prior			Noise	Centerline			Morphology			Time ↓
	Topo.	Pos.	Morph.		CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓	
NN Retrieval	✓	–	✓	2	7.884	17.556	0.017	**0.037**	**0.036**	26.71	1.27
				4	7.980	17.943	0.016	0.595	**0.036**	**26.36**	
Shape Model	✓	–	✓	2	6.427	14.634	0.027	0.055	0.053	27.55	5.43
				4	6.778	15.620	0.025	0.056	0.055	27.16	
Implicit Field	–	✓	–	2	6.187	13.783	0.033	0.077	0.130	35.19	5.716
				4	6.228	13.639	0.036	0.456	0.086	34.36	
Soft Template	✓	✓	–	2	6.145	15.318	0.074	0.253	0.611	34.65	∼0.11
				4	6.327	15.373	0.050	0.815	0.928	39.79	
Regression	✓	✓	–	2	5.251	11.172	0.035	0.157	0.261	48.64	0.102
				4	5.198	11.025	0.034	0.386	0.456	44.16	
CoroLift (Ours)	✓	✓	✓	2	**3.854**	**9.190**	**0.110**	0.042	0.054	**24.19**	∼0.17
				4	**3.847**	**9.013**	**0.089**	**0.045**	0.054	27.40	

Table 2: **Ablation of the Anatomy Prior.** Effects of the topology $\mathcal{A}_{\text{topo}}$, position $\mathcal{A}_{\text{pos}}$, and morphology $\mathcal{A}_{\text{morph}}$ priors at 2 px trace noise. Each component is evaluated at the stage where it is applied.

$\mathcal{A}_{\text{topo}}$	CD ↓	HD95 ↓	OV ↑
–	3.959	9.368	0.105
✓	**3.854**	**9.190**	**0.110**

$\mathcal{A}_{\text{morph}}$	CD ↓	OV ↑	Tort. ↓	Curv. ↓
–	3.832	0.127	0.243	0.595
✓	**3.854**	**0.110**	**0.042**	**0.054**

$\mathcal{A}_{\text{pos}}$		CD ↓	HD95 ↓	OV ↑	Tort. ↓
Endpoint	Branch				
–	–	6.271	14.769	0.059	0.038
–	✓	3.880	9.036	0.107	0.038
✓	–	4.211	10.228	0.097	0.042
✓	✓	**3.854**	**9.190**	**0.110**	**0.042**

Anatomy prior. Table 2 isolates the topology, position, and morphology components of the anatomy prior. Topology $\mathcal{A}_{\text{topo}}$ replaces nearest-neighbor correspondence with anatomical-slot correspondence, reducing CD from 3.959 to 3.854 mm and improving OV from 0.105 to 0.110. For position $\mathcal{A}_{\text{pos}}$, the slot-keypoint prior μ_s^{P1} and family-trajectory prior $\bar{P}_f$ are individually useful (4.211 and 3.880 mm CD vs. the 6.271 mm population baseline); combining them yields 3.854 mm CD with 0.110 OV, confirming their complementary role in anatomical localization and continuous vessel geometry. Morphology $\mathcal{A}_{\text{morph}}$ trades a small CD change (3.832 → 3.854 mm) for a large morphology gain, cutting tortuosity error from 0.243 to 0.042 and curvature error from 0.595 to 0.054. Complete metrics are provided in Appendix C.1.

Anatomy-guided initialization. Table 3 evaluates the two initialization components directly at x^0, before anatomy-constrained refinement. Endpoint localization provides strong patient-specific spatial anchors, reducing CD from 8.687 to 5.111 mm and increasing OV from 0.040 to 0.110 over the population initialization. Continuous branch-depth prediction alone provides a smaller improvement, whereas combining it with the localized endpoints substantially improves the reconstruction to 3.832 mm CD, 8.834 mm HD95, and 0.127 OV. These results indicate that the two components play complementary roles: endpoint localization constrains the sparse

Table 3: **Anatomy-guided initialization.** Ablation of endpoint localization and continuous branch-depth prediction, evaluated at initialization x^0 before anatomy-constrained refinement.

Initialization	CD ↓	HD95 ↓	OV ↑
Neither	8.687	21.479	0.040
Endpoint only	5.111	15.582	0.110
Branch only	8.045	19.312	0.040
Endpoint + Branch	**3.832**	**8.834**	**0.127**

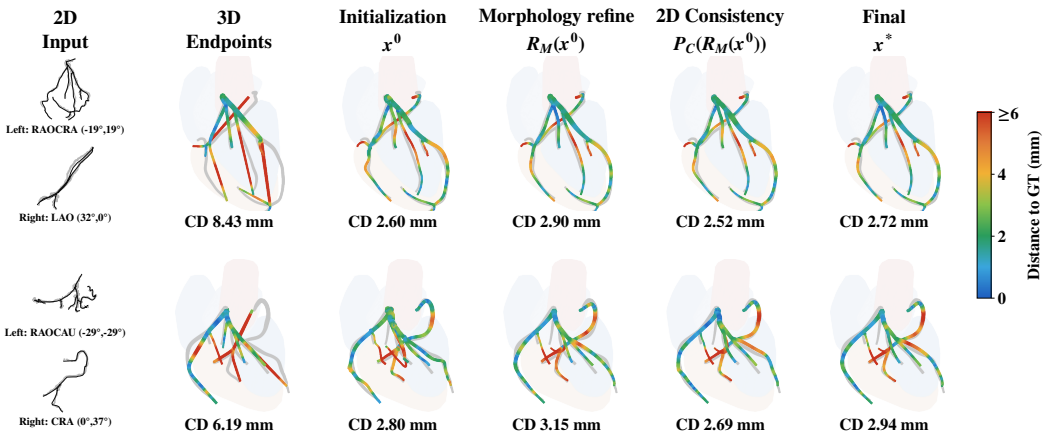


Figure 4: **Stage-wise qualitative reconstruction.** From left to right: 2D input, fused 3D endpoints, initialization x^0, $R_M(x^0)$, $P_C(R_M(x^0))$, and final reconstruction $x^\star$. Colors indicate distance to the ground truth (gray).

anatomical nodes, while continuous depth prediction recovers the intervening vessel geometry under these anchors. Additional initialization diagnostics and design choices are provided in Appendix C.2.

Morphology model design. Table 4 evaluates the two design choices of the morphology model, measured after a single application of R_M to x^0. Encoding vessel shape using family-specific PCA of endpoint-anchored residuals reduces CD from 4.792 to 3.931 mm and doubles OV ($0.050 \rightarrow 0.105$) compared with coordinate-based PCA. Coupling slots through a structured Gaussian-Markov model further cuts tortuosity from 0.086 to 0.041, curvature from 0.102 to 0.055, and bifurcation angle from $33.04°$ to $24.33°$ at essentially unchanged CD.

Table 4: **Design of the morphology model.** Effect of shape representation (top) and inter-slot dependency structure (bottom), measured after a single application of R_M to the ray-consistent initialization x^0.

Shape Representation	CD ↓	HD95 ↓	OV ↑
Coordinate PCA	4.792	10.824	0.050
Residual PCA	**3.931**	**9.395**	**0.105**

Dependency	Tort. ↓	Curv. ↓	Bif. Angle ↓
Independent	0.086	0.102	33.038
Structured GM	**0.041**	**0.055**	**24.325**

Table 5: **Ablation of anatomy-constrained refinement.** Complementary roles of morphology regularization R_M and observation consistency P_C under the strict single-view setting with 2 px trace noise.

Refinement	CD ↓	OV ↑	Tort. ↓	Curv. ↓
x^0	3.832	0.127	0.243	0.595
$P_C(x^0)$	3.832	0.127	0.243	0.595
$R_M(x^0)$	3.973	0.105	0.041	0.055
$(P_C \circ R_M)(x^0)$	3.696	0.142	0.108	0.290
$(R_M \circ P_C \circ R_M)(x^0)$	**3.854**	**0.110**	**0.042**	**0.054**

Anatomy-constrained refinement. Table 5 shows the complementary roles of morphology regularization R_M and observation consistency P_C. Applying P_C directly to x^0 leaves the reconstruction unchanged, since x^0 already satisfies the reprojection constraint. R_M alone reduces tortuosity from 0.243 to 0.041 and curvature from 0.595 to 0.055, but increases CD from 3.832 to 3.973 mm by pulling geometry off the observed 2D centerlines. Applying P_C afterward restores observation consistency and yields the best CD (3.696 mm) and OV (0.142) while preserving part of the morphology gain (tortuosity 0.108, curvature 0.290). A final R_M step then recovers the anatomical morphology (tortuosity 0.042, curvature 0.054) at a small CD cost, giving the composition $x^\star = (R_M \circ P_C \circ R_M)(x^0)$ used in the main comparison of Table 1. Complete metrics and image-space diagnostics are in Appendix C.1, C.2. Figure 4 qualitatively illustrates the stage-wise reconstruction process and the complementary roles of anatomy-guided initialization, morphology refinement, and observation consistency.

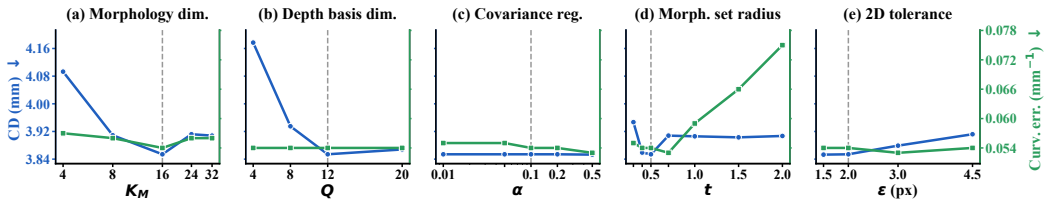


Figure 5: **Parameter analysis.** Sensitivity to the main hyperparameters. Blue: CD (mm, left axis). Green: curvature error (mm^{-1}, right axis). Dashed grey lines denote the default settings.

4.4 Parameter and Robustness Analysis

Parameter sensitivity. Figure 5 evaluates sensitivity to CoroLift's main hyperparameters. CD stays within 3.85–4.18 mm across $K_M \in [4, 32]$, $Q \in [4, 20]$, $t \in [0.3, 2.0]$, and $\varepsilon \in [1.5, 4.5]$ px, while $\alpha \in [0.01, 0.5]$ affects CD by less than 0.01 mm. The radius t and tolerance ε trade off adherence to the observed 2D centerlines against morphology regularization (small values favor CD; large values favor curvature). We fix $K_M = 16$, $Q = 12$, $\alpha = 0.1$, $t = 0.5$, $\varepsilon = 2$ px throughout the main experiments.

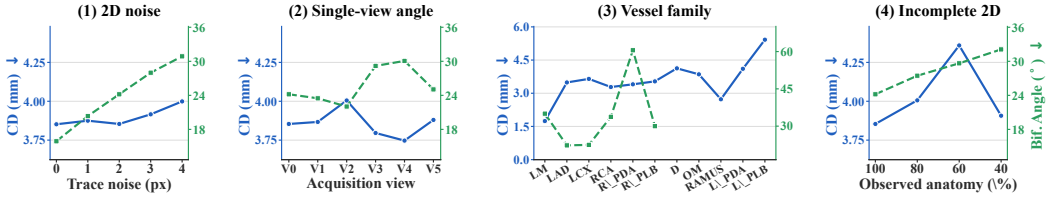


Figure 6: **Robustness analysis.** Sensitivity to trace noise, single-view acquisition geometry, vessel family, and incomplete observed anatomy. Blue: CD (mm, left axis). Green: bifurcation-angle error (°, right axis).

Robustness to observation conditions. Figure 6 evaluates robustness along four axes. *Trace noise* $0 \to 4$ px raises CD only from 3.85 to 4.00 mm, with morphology metrics degrading gracefully. *Acquisition geometry*: across six clinically representative single-view angle pairs, CD stays within 3.75–4.00 mm. *Vessel family*: main branches (LM, LAD, LCX, RCA) reach 1.7–3.7 mm CD; distal branches remain within 2.7–5.4 mm. *Incomplete observation*: we progressively remove complete distal branches while preserving a connected observed subtree, and re-evaluate on the common S_{40} vessel set to isolate context loss from evaluation-composition change; reconstruction is stable from 100% to 60% observed anatomy (3.30–3.36 mm CD) and degrades only at 40% (3.91 mm). Full protocols and results are in Appendix C.1, C.3.

5 Conclusion

We presented CoroLift for fast single-view 3D coronary reconstruction. CoroLift addresses the inherent ambiguity of single-view reconstruction with an explicit anatomy prior comprising topology, position, and morphology. The topology and position priors establish patient-specific anatomical structure and guide depth estimation along the observed imaging rays, yielding a ray-consistent 3D initialization. The morphology prior then regularizes the reconstructed tree toward anatomically plausible configurations, while an observation-consistency step constrains its projection to remain consistent with the observed 2D centerlines. Experiments demonstrate accurate single-view reconstruction while maintaining anatomical connectivity and consistency with the observed image geometry, with robustness across trace noise, acquisition angles, vessel families, and incomplete observations. These results highlight the complementary role of population-level anatomical priors in resolving the ambiguity of single-view coronary reconstruction.

ETHICS STATEMENT

All data used in this study are de-identified, and CoroLift is intended for research use only.

AI USE STATEMENT

Generative AI tools were used to assist with language editing . All scientific content, experimental results, and conclusions were reviewed and verified by the authors.

REFERENCES

Abhirup Banerjee, Francesca Galassi, Ernesto Zacur, Giovanni Luigi De Maria, Robin P. Choudhury, and Vicente Grau. Point-cloud method for automated 3D coronary tree reconstruction from multiple non-simultaneous angiographic projections. *IEEE Transactions on Medical Imaging*, 39: 1278–1290, 2020.

Kit Mills Bransby, Vincenzo Tufaro, Murat Çap, Greg Slabaugh, Christos Bourantas, and Qianni Zhang. 3D coronary vessel reconstruction from bi-plane angiography using graph convolutional networks. In *IEEE International Symposium on Biomedical Imaging (ISBI)*, 2023.

S. James Chen and John D. Carroll. 3-D reconstruction of coronary arterial tree to optimize angiographic visualization. *IEEE Transactions on Medical Imaging*, 19(4):318–336, 2000.

Siqi Chen, Guoqing Zhang, Jiahao Lai, Bingzhi Shen, Sihong Zhang, Caixia Dong, Xuejin Chen, and Yang Li. Hierarchical part-based generative model for realistic 3D blood vessel. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2025.

Hyungjin Chung, Jeongsol Kim, Michael T Mccann, Marc L Klasky, and Jong Chul Ye. Diffusion posterior sampling for general noisy inverse problems. *arXiv preprint arXiv:2209.14687*, 2022.

Serkan Çimen, Corné Hoogendoorn, Paul D. Morris, Julian Gunn, and Alejandro F. Frangi. Reconstruction of coronary trees from 3DRA using a 3D+t statistical cardiac prior. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pp. 619–626, 2014.

Serkan Çimen, Ali Gooya, Michael Grass, and Alejandro F. Frangi. Reconstruction of coronary arteries from X-ray angiography: A review. *Medical Image Analysis*, 32:46–68, 2016a.

Serkan Çimen, Ali Gooya, Nishant Ravikumar, Zeike A. Taylor, and Alejandro F. Frangi. Reconstruction of coronary artery centrelines from X-ray angiography using a mixture of Student's t-distributions. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, pp. 291–299, 2016b.

Alberto Clemente, Sara Seitun, Cesare Mantini, Giovanni Gentile, Duccio Federici, Andrea Barison, Andrea Rossi, Magdalena Cuman, Alessandra Pizzuto, Lamia Ait-Ali, et al. Cardiac ct angiography: normal and pathological anatomical features—a narrative review. *Cardiovascular diagnosis and therapy*, 10(6):1918, 2020.

T. F. Cootes, C. J. Taylor, D. H. Cooper, and J. Graham. Active shape models—their training and application. *Computer Vision and Image Understanding*, 61(1):38–59, 1995.

Heinz W. Engl, Martin Hanke, and Andreas Neubauer. *Regularization of Inverse Problems*. Kluwer Academic Publishers, Dordrecht, 1996.

Paula Feldman, Miguel Fainstein, Viviana Siless, Claudio Delrieux, and Emmanuel Iarussi. Recursive variational autoencoders for 3D blood vessel generative modeling. *Medical Image Analysis*, 105:103703, 2025a.

Paula Feldman, Martin Sinnona, Claudio Delrieux, Viviana Siless, and Emmanuel Iarussi. Vesselgpt: Autoregressive modeling of vascular geometry. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 662–672. Springer, 2025b.

Xueming Fu, Yingtai Li, Fenghe Tang, Jun Li, Mingyue Zhao, Gao-Jun Teng, and S. Kevin Zhou. 3DGR-CAR: Coronary artery reconstruction from ultra-sparse 2D X-ray views with a 3D gaussians representation. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2024.

Francesca Galassi, Mohammad Alkhalil, Regent Lee, Philip Martindale, Rajesh K Kharbanda, Keith M Channon, Vicente Grau, and Robin P Choudhury. 3d reconstruction of coronary arteries from 2d angiographic projections using non-uniform rational basis splines (nurbs) for accurate modelling of coronary stenoses. *PloS one*, 13(1):e0190650, 2018.

Zhanqiang Guo, Zimeng Tan, Jianjiang Feng, and Jie Zhou. VesselDiffusion: 3D vascular structure generation based on diffusion model. *IEEE Transactions on Medical Imaging*, 44(9), 2025.

Tobias Heimann and Hans-Peter Meinzer. Statistical shape models for 3D medical image segmentation: a review. *Medical Image Analysis*, 13(4):543–563, 2009.

Minki Hwang, Sa-Bin Hwang, Hyosang Yu, Jaehyeok Kim, Daehyun Kim, Wonjae Hong, Ah-Jin Ryu, Han Yong Cho, Jinlong Zhang, Bon Kwon Koo, and Eun Bo Shim. A simple method for automatic 3D reconstruction of coronary arteries from X-ray angiography. *Frontiers in Physiology*, 12:724216, 2021. doi: 10.3389/fphys.2021.724216.

Kritika Iyer, Brahmajee K. Nallamothu, C. Alberto Figueroa, and Raj Rao Nadakuditi. A multi-stage neural network approach for coronary 3D reconstruction from uncalibrated X-ray angiography images. *Scientific Reports*, 13:17603, 2023.

Ulugbek S Kamilov, Charles A Bouman, Gregery T Buzzard, and Brendt Wohlberg. Plug-and-play methods for integrating physical and learned models in computational imaging: Theory, algorithms, and applications. *IEEE Signal Processing Magazine*, 40(1):85–97, 2023.

Alex Kendall and Yarin Gal. What uncertainties do we need in bayesian deep learning for computer vision? In *Advances in Neural Information Processing Systems*, volume 30, 2017.

Thijs P. Kuipers, Praneeta R. Konduri, Henk Marquering, and Erik J. Bekkers. Generating cerebral vessel trees of acute ischemic stroke patients using conditional set-diffusion. In *Medical Imaging with Deep Learning (MIDL)*, 2024.

Dominik Bernard Lau, Hubert Malinowski, Jerzy Szyjut, and Tomasz Dziubich. Xra-diff: Guided diffusion as a refinement for backprojected coronary angiography x-ray images. *Applied Soft Computing*, pp. 116236, 2026.

Steffen L Lauritzen. *Graphical Models*. Oxford University Press, 1996.

Chinmay Prabhakar, Suprosanna Shit, Fabio Musio, Kaiyuan Yang, Tamaz Amiranashvili, Johannes C. Paetzold, Hongwei Bran Li, and Bjoern Menze. 3D vessel graph generation using denoising diffusion. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2024.

Varun Ramakrishna, Takeo Kanade, and Yaser Sheikh. Reconstructing 3D human pose from 2D image landmarks. In *European Conference on Computer Vision (ECCV)*, 2012.

Michiel Schaap, Coert T Metz, Theo van Walsum, Alina G van der Giessen, Annick C Weustink, Nico R Mollet, Christian Bauer, Hrvoje Bogunović, Carlos Castro, Xiang Deng, et al. Standardized evaluation methodology and reference database for evaluating coronary artery centerline extraction algorithms. *Medical Image Analysis*, 13(5):701–714, 2009.

Ashish Sinha and Ghassan Hamarneh. TrIND: Representing anatomical trees by denoising diffusion of implicit neural fields. In *Medical Image Computing and Computer Assisted Intervention (MICCAI)*, 2024.

Richard Socher, Cliff Chiung-Yu Lin, Andrew Y. Ng, and Christopher D. Manning. Parsing natural scenes and natural language with recursive neural networks. In *International Conference on Machine Learning (ICML)*, 2011.

Yang Song, Liyue Shen, Lei Xing, and Stefano Ermon. Solving inverse problems in medical imaging with score-based generative models. In *International Conference on Learning Representations (ICLR)*, 2022.

Kai Sheng Tai, Richard Socher, and Christopher D. Manning. Improved semantic representations from tree-structured long short-term memory networks. In *Annual Meeting of the Association for Computational Linguistics (ACL)*, 2015.

Singanallur V. Venkatakrishnan, Charles A. Bouman, and Brendt Wohlberg. Plug-and-play priors for model based reconstruction. In *IEEE Global Conference on Signal and Information Processing (GlobalSIP)*, pp. 945–948, 2013.

Yiying Wang, Abhirup Banerjee, and Vicente Grau. NeCA: 3D coronary artery tree reconstruction from two 2D projections via neural implicit representation. *Bioengineering*, 11(12):1227, 2024.

Yiying Wang, Abhirup Banerjee, Robin P. Choudhury, and Vicente Grau. DeepCA: Deep learning-based 3D coronary artery tree reconstruction from two 2D non-simultaneous X-ray angiography projections. In *IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 337–346, 2025.

Alan S Willsky. Multiresolution markov models for signal and image processing. *Proceedings of the IEEE*, 90(8):1396–1458, 2002.

Jian Yang, Yongtian Wang, Yue Liu, Songyuan Tang, and Wufan Chen. Novel approach for 3-D reconstruction of coronary arteries from two uncalibrated angiographic images. *IEEE Transactions on Image Processing*, 18(7):1563–1572, 2009.

An Zeng, Chunbiao Wu, Guisen Lin, Wen Xie, Jin Hong, Meiping Huang, Jian Zhuang, Shanshan Bi, Dan Pan, Najeeb Ullah, Kaleem Nawaz Khan, Tianchen Wang, Yiyu Shi, Xiaomeng Li, and Xiaowei Xu. ImageCAS: A large-scale dataset and benchmark for coronary artery segmentation based on computed tomography angiography images. *Computerized Medical Imaging and Graphics*, 109:102287, 2023.

Yuanzhi Zhu, Kai Zhang, Jingyun Liang, Jiezhang Cao, Bihan Wen, Radu Timofte, and Luc Van Gool. Denoising diffusion models for plug-and-play image restoration. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2023.

APPENDIX

CONTENTS

A Implementation Details **13**
A.1 Anatomy Prior Construction 13
A.2 Anatomy-guided Initialization 14
A.3 Anatomy-constrained Refinement 15

B Data and Evaluation Protocol **16**
B.1 Dataset 16
B.2 Single-view Acquisition Geometry 16
B.3 Observation Model 17
B.4 Evaluation Metrics 17
B.5 Baseline Protocol 18
B.6 Robustness Protocol 19

C Additional Results **19**
C.1 Full Quantitative Results 19
C.2 Additional Diagnostic Analysis 20
C.3 Incomplete-observation Analysis 21
C.4 Vessel Radius Evaluation 21
C.5 Additional Qualitative Results 22

D Mathematical Properties and Proofs **22**
D.1 Gaussian-Markov Morphology Model 22
D.2 Morphology Set and Projection 24
D.3 Morphology Projection in Curve Space 27
D.4 Single-view Observation Constraint 28
D.5 Relation to the Implemented Observation Operator 29

A IMPLEMENTATION DETAILS

This section provides implementation details omitted from the main text. We follow the same organization as the Method section, covering Anatomy Prior Construction, Anatomy-guided Initialization, and Anatomy-constrained Refinement.

A.1 ANATOMY PRIOR CONSTRUCTION

Anatomical slots and topology prior. We harmonize the coronary annotations into 11 major vessel families and partition the annotated vessels at anatomical bifurcations into 43 anatomical slots. Each slot is identified by its vessel family and proximal-to-distal segment index within that family. Minor branches of limited clinical relevance are excluded. The observed 2D traces carry these slot identities, and only slots present in the observation are instantiated for each patient. The topology prior is estimated from parent–child connectivity statistics in the training population. At inference, the observed slot connectivity defines the patient-specific coronary topology, allowing the position

and morphology models to operate on the anatomical structure present in each patient rather than imposing a fixed complete tree.

Position prior. The position prior contains two complementary population statistics: slot-specific anatomical endpoint statistics and family-specific continuous vessel trajectories. The former provides sparse spatial references for endpoint localization, whereas the latter provides a continuous reference for vessel-depth prediction. For each anatomical slot s, we estimate population statistics of its proximal and distal 3D endpoints from the training set. Slots with limited training support are regularized toward the corresponding vessel-family statistics. We denote the resulting slot-specific endpoint references by μ_s^{P1}. To obtain continuous position statistics, connected slots belonging to the same vessel family are concatenated in proximal-to-distal order. Each training vessel is normalized by 3D arc length and resampled to a common parameterization. A family-specific mean trajectory $\bar{P}_f$ is then estimated from the training population. Both the endpoint and trajectory statistics are represented in the canonical heart coordinate system and are converted to view-dependent depth references under the imaging geometry at inference time.

Morphology prior. The morphology prior models local vessel shape while preserving anatomical endpoints. For each slot s from vessel family $f(s)$, let $X_s \in \mathbb{R}^{N\times 3}$ denote its sampled 3D centerline. We construct the endpoint chord $\ell_s \in \mathbb{R}^{N\times 3}$ by linearly interpolating between its proximal and distal endpoints and define the residual morphology as

$$r_s = \mathrm{vec}(X_s - \ell_s). \tag{9}$$

Because the residual is zero at both endpoints, modifying the residual morphology does not alter the anatomical endpoints. For each vessel family f, we compute the population mean residual $\bar{r}_f$ and perform PCA on the centered training residuals. Using the first K_M principal components Φ_f, a slot is encoded by

$$c_s = \Phi_{f(s)}^{\top}\left(r_s - \bar{r}_{f(s)}\right), \tag{10}$$

and reconstructed as

$$X_s = \ell_s + \mathrm{unvec}\left(\bar{r}_{f(s)} + \Phi_{f(s)} c_s\right). \tag{11}$$

To capture dependencies between connected coronary segments, we model the PCA coefficients with a Gaussian-Markov tree. For each root slot r,

$$c_r \sim \mathcal{N}(\mu_r, \Sigma_r), \tag{12}$$

and for each non-root slot s with parent $p(s)$,

$$c_s = A_s c_{p(s)} + b_s + \varepsilon_s, \qquad \varepsilon_s \sim \mathcal{N}(0, Q_s). \tag{13}$$

All parameters are estimated exclusively from the training trees. At inference, the model is instantiated over the slots present in the patient. Propagating the linear Gaussian conditionals along the patient-specific tree yields a joint morphology distribution

$$c \sim \mathcal{N}(\mu, \Sigma), \tag{14}$$

where c concatenates the morphology coefficients of all observed slots. Near-null covariance directions are regularized using a median-scaled eigenvalue floor with $\alpha = 0.1$.

A.2 Anatomy-guided Initialization

Ray-based parameterization. Under the calibrated single-view geometry, each observed 2D point u defines an imaging ray $q(u)$ from source o. Its 3D position can therefore be parameterized by a single unknown depth:

$$X(u) = o + d(u)q(u). \tag{15}$$

This representation preserves the observed image coordinates and reduces single-view reconstruction to estimating depth along the known imaging rays. The 3D position prior is converted to depth references under the current imaging geometry. The slot-specific endpoint statistics μ_s^{P1} provide reference depths $d_{s,e}^{P1}$ for anatomical endpoints, while the family-specific trajectory $\bar{P}_f$ provides a continuous population depth profile $\bar{d}_f(t)$ along each observed vessel.

Anatomical endpoint localization. For endpoint e of slot s, the Endpoint Transformer predicts a depth residual $\Delta d_{s,e}$ relative to its population reference together with an uncertainty $\sigma_{s,e} > 0$:

$$\hat{d}_{s,e} = d^{P1}_{s,e} + \Delta d_{s,e}. \tag{16}$$

The endpoint predictions are trained with a heteroscedastic Gaussian negative log-likelihood so that $\sigma_{s,e}$ reflects the confidence of the corresponding depth estimate. Multiple slot endpoints may correspond to the same anatomical bifurcation. Their depth estimates are fused using uncertainty-based precision weighting. For anatomical node j, let $\mathcal{E}(j)$ denote its associated endpoint predictions and $w_{s,e} = \sigma^{-2}_{s,e}$. The fused depth is

$$d^{\text{fuse}}_j = \frac{\sum_{(s,e)\in\mathcal{E}(j)} w_{s,e}\hat{d}_{s,e}}{\sum_{(s,e)\in\mathcal{E}(j)} w_{s,e}}. \tag{17}$$

The fused nodes are further regularized under the patient-specific coronary topology to obtain a coherent set of anatomical 3D keypoints.

Continuous branch-depth prediction. Connected slots from the same vessel family are concatenated into continuous vessels and parameterized by normalized arc length $t \in [0, 1]$. A Transformer encoder processes the observed vessels jointly and predicts Q Chebyshev coefficients for each continuous vessel. The depth profile is represented as

$$d_v(t) = \bar{d}_{f(v)}(t) + \sum_{q=0}^{Q-1} a_{v,q} T_q(2t-1), \tag{18}$$

where $\bar{d}_{f(v)}(t)$ is the family-specific population depth profile, T_q is the q-th Chebyshev polynomial, and $a_{v,q}$ is the predicted patient-specific coefficient. The localized endpoints and internal anatomical boundaries provide depth keypoints along each continuous vessel. The predicted profile is adjusted to interpolate these keypoints and is then lifted along the original observed imaging rays using Eq. equation 15. The resulting continuous vessels are divided at their anatomical slot boundaries and resampled to a common representation. This ray-preserving reconstruction defines the initialization x^0. For all experiments, we use $Q = 12$ Chebyshev coefficients. The endpoint and branch-depth Transformers use a hidden dimension of 128; the branch-depth Transformer contains three encoder layers and four attention heads.

A.3 Anatomy-constrained Refinement

Morphology regularization. Given a reconstructed tree x, we encode all observed slots using the family-specific residual PCA models and concatenate their coefficients into c. Using the patient-specific joint Gaussian from Eq. equation 14, we define the plausible morphology region as

$$\mathcal{M}_t = \left\{ c : (c-\mu)^\top \Sigma_\alpha^{-1} (c-\mu) \le t^2 D \right\}, \tag{19}$$

where D is the dimension of the joint coefficient vector and Σ_α denotes the regularized covariance. The morphology operator $R_{\mathcal{M}}$ leaves coefficients inside $\mathcal{M}_t$ unchanged and projects coefficients outside this region onto its boundary. The projected coefficients are decoded through the corresponding family-specific PCA models. Because the endpoint chord is retained during decoding, $R_{\mathcal{M}}$ modifies the interior vessel morphology while preserving the localized anatomical endpoints.

Observation consistency. Morphology regularization can move interior vessel points away from their observed image locations. We therefore enforce consistency with the input 2D centerlines under the known projection geometry. Each reconstructed sample X_i is associated with the corresponding observed centerline point $\tilde{u}_i$ under normalized arc-length parameterization. The observation constraint is

$$\|\pi(X_i) - \tilde{u}_i\|_2 \le \varepsilon, \tag{20}$$

where $\pi(\cdot)$ denotes calibrated cone projection. The operator $P_{\mathcal{C}}$ locally corrects samples violating this tolerance using the known projection geometry, while samples already satisfying the constraint are left unchanged.

Final reconstruction. Morphology regularization and observation consistency play complementary roles: $R_{\mathcal{M}}$ constrains the joint vessel geometry toward the population morphology distribution, whereas $P_{\mathcal{C}}$ restores consistency with the observed 2D centerlines. We apply morphology regularization, observation consistency, and a final morphology projection in sequence:

$$\boxed{x^\star = (R_{\mathcal{M}} \circ P_{\mathcal{C}} \circ R_{\mathcal{M}})(x^0)}. \tag{21}$$

The final application of $R_{\mathcal{M}}$ restores anatomical plausibility after $P_{\mathcal{C}}$ enforces consistency with the observed 2D centerlines. Unless otherwise stated, we use $K_M = 16$, $Q = 12$, $\alpha = 0.1$, $t = 0.5$, and $\varepsilon = 2$ pixels in the primary experiments.

Vessel radius recovery. Vessel radius is recovered analytically rather than predicted by the reconstruction networks. Given the observed 2D vessel half-width $w_{2\mathrm{D}}$, reconstructed source-to-point depth d, and calibrated source-to-detector focal length f, the physical vessel radius is

$$R = \frac{w_{2\mathrm{D}} d}{f}. \tag{22}$$

The recovered radius is associated with the reconstructed centerline and is used for radius-aware 3D vessel representation.

B Data and Evaluation Protocol

B.1 Dataset

CCTA dataset. We use coronary artery trees from the public ImageCAS coronary computed tomography angiography (CCTA) dataset (Zeng et al., 2023). The original cohort contains 881 patients. After preprocessing the vessel annotations, we exclude 45 cases whose processed annotations do not satisfy the required coronary-tree representation. The final cohort contains 836 patients and is divided at the patient level into 671, 83, and 82 cases for training, validation, and testing, respectively. All observations generated from the same patient remain in the same split.

The heterogeneous vessel annotations are harmonized into 11 clinically relevant vessel families and partitioned at anatomical bifurcations into 43 anatomical slots. Each slot is identified by its vessel family and proximal-to-distal segment index. The observed slot labels determine the vessels present in each case, while patient-specific connectivity is inferred from shared endpoints of the observed 2D traces.

All population statistics used by CoroLift, including the position and morphology priors, are estimated exclusively from the training split. The reconstruction networks are also trained on the training split, while the validation split is used for model and hyperparameter selection.

Coordinate system. The CCTA trees are represented in a patient-centered physical coordinate system used to estimate training-set population statistics and generate synthetic observations. At inference, reconstruction is performed under the calibrated acquisition geometry of the observed projection.

Cone-beam projection. Synthetic angiographic observations are generated using a calibrated C-arm cone-beam camera. For a camera with imaging source c, source-to-detector direction $\hat{d}$, detector axes $(\hat{u}, \hat{v})$, and source-to-image distance SID, a 3D point X is projected as

$$\pi(X) = \mathrm{SID} \begin{bmatrix} (X - c) \cdot \hat{u} \\ (X - c) \cdot \hat{v} \end{bmatrix} / \left((X - c) \cdot \hat{d} \right). \tag{23}$$

The primary RAO/LAO angle and secondary cranial/caudal angle determine the camera orientation. The detector sampling corresponds to 0.434 mm per image pixel at the isocenter. Rendering provides projected centerlines and vessel-width observations from which the reconstruction inputs are generated.

B.2 Single-view Acquisition Geometry

Clinical acquisition geometries. We evaluate representative projection configurations commonly used for left and right coronary angiography. Figure 7 summarizes the single-view acquisition geometries used in our robustness experiments.

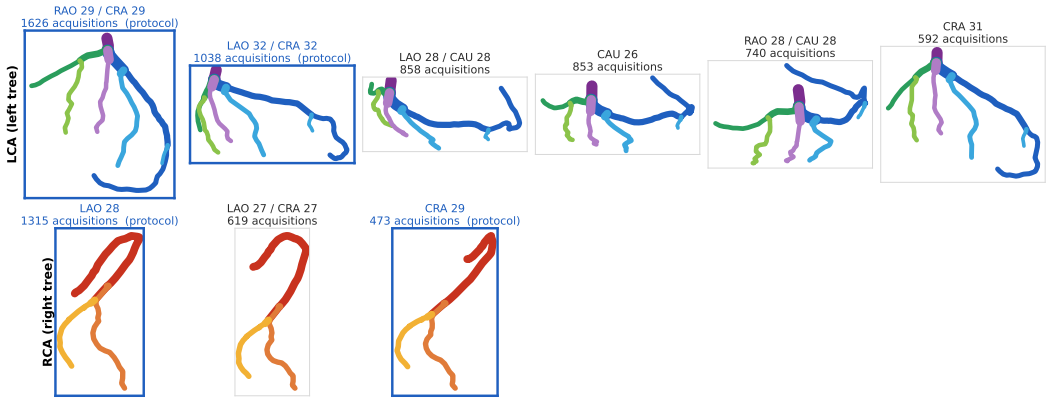


Figure 7: **Single-view acquisition geometries.** Representative clinical projection configurations used in the acquisition-angle robustness experiments.

Single-view protocol. All main experiments follow a strict single-view protocol. A single calibrated projection is retained for each coronary tree, and all methods receive the same noisy 2D traces, anatomical labels, vessel-width observations, and camera model. The acquisition-angle robustness experiment changes only the projection geometry while keeping the reconstruction and evaluation protocols fixed.

B.3 Observation Model

Centerline trace noise. We simulate errors in manually or automatically extracted angiographic centerlines using spatially correlated Gaussian perturbations along each observed slot. The perturbations are normalized to a prescribed RMS displacement and scaled to the target noise level. Shared bifurcation points receive consistent perturbations across connected slots so that the noisy observations preserve the traced topology. Unless otherwise specified, the main experiments use 2 px trace noise.

For the trace-noise robustness analysis, we evaluate noise levels of 0, 1, 2, 3, and 4 px. For each test case, the same normalized perturbation is scaled across noise levels, producing nested observation conditions. All methods evaluated under the same condition receive identical noisy observations.

Reconstruction inputs. At test time, each case is represented only by the labelled 2D centerline traces, corresponding vessel-width observations, calibrated camera geometry, and population priors estimated from the training split. The observed labels and shared trace endpoints determine the patient-specific slot set and connectivity. The underlying CCTA geometry and clean projections are used only for quantitative evaluation.

B.4 Evaluation Metrics

Centerline resampling. Before evaluation, reconstructed and ground-truth centerlines are densely resampled at a common physical spacing of 0.1 mm in arc length. Using the same sampling density prevents differences in the internal point representation of individual reconstruction methods from affecting the geometric metrics. All distances are evaluated in physical millimetres.

Centerline accuracy. Let $\hat{P}$ and P denote the concatenated reconstructed and ground-truth centerline samples over the anatomical slots included in the evaluation of a case. We report the symmetric unsquared Chamfer distance

$$\mathrm{CD}(\hat{P}, P) = \frac{1}{2}\left(\frac{1}{|\hat{P}|}\sum_{x\in\hat{P}}\min_{y\in P}\|x-y\|_2 + \frac{1}{|P|}\sum_{y\in P}\min_{x\in\hat{P}}\|x-y\|_2\right). \tag{24}$$

HD95 is the 95th percentile of the pooled bidirectional nearest-neighbor distances. Rotterdam overlap (OV) is computed using the local ground-truth vessel radius as the correspondence toler-

ance (Schaap et al., 2009). All geometric metrics are first computed independently for each test case and then summarized by the median across cases.

Radius accuracy. We report the absolute error between the recovered and ground-truth vessel radii in millimetres. Radius error is evaluated along the reconstructed centerline using the same dense geometric correspondence as the centerline evaluation and is summarized by the median across test cases.

Tortuosity and curvature. Morphological measurements are computed after resampling each reconstructed and ground-truth slot to 32 equally spaced arc-length samples. Slots with endpoint chord length below 5 mm are excluded from tortuosity and curvature evaluation.

Tortuosity is defined as the ratio between centerline arc length and endpoint chord length. We report the absolute difference between reconstructed and ground-truth tortuosity.

For curvature, let t_i denote the unit tangent at sample i. We use the discrete tangent-change measure

$$\kappa_i = \|t_{i+1} - t_i\|_2, \tag{25}$$

and average it along each slot. We report the absolute difference between the reconstructed and ground-truth mean curvature measures. Slot-wise errors are summarized by their median within each case, followed by the median across test cases.

Bifurcation angle. For each parent slot with at least two observed children, we compute the angle between the local take-off directions of anatomically matched child branches. We report the absolute difference between reconstructed and ground-truth bifurcation angles. Errors are summarized by their median within each case and then by the median across test cases.

Validation protocol. The validation split contains 83 patients and is not included in the reported test metrics. It is used for model selection and for fixing reconstruction hyperparameters. Unless otherwise specified, the same fixed configuration is used throughout the reported test experiments.

B.5 Baseline Protocol

All baselines are evaluated under the same strict single-view protocol as CoroLift. For each experimental condition, all methods receive the same noisy 2D observations and calibrated camera geometry. Population-based baselines use statistics estimated exclusively from the training split, while learning-based baselines are trained on the training split with model selection performed on the validation split.

Each baseline is evaluated using its native reconstruction output, without CoroLift's morphology regularization or observation-consistency operators. All reconstructed centerlines are converted to the common anatomical-slot representation and evaluated using the same metric pipeline described above.

Baseline definitions. We compare CoroLift with representative population-based and learning-based single-view reconstruction strategies.

NN Retrieval (Ramakrishna et al., 2012) retrieves the training-set coronary tree whose single-view projection best matches the observed labeled centerlines under the calibrated imaging geometry.

Shape Model (Cootes et al., 1995) fits a PCA-based statistical shape model to anatomically corresponding coronary branches and optimizes its shape coefficients by minimizing reprojection error with population-shape regularization.

Implicit Field (Wang et al., 2024) uses an implicit neural representation conditioned on the observed 2D centerlines to predict continuous 3D coronary geometry, followed by test-time reprojection refinement.

Soft Template uses CoroLift's topology and position priors to produce a population-anchored reconstruction without morphology refinement.

Regression uses a tree-structured recurrent network to directly regress 3D coronary coordinates from the observed 2D centerlines while propagating features along the coronary topology (Tai et al., 2015; Socher et al., 2011).

Table 6: **Anatomy-guided initialization at x^0.** Ablation of endpoint localization and continuous branch-depth prediction, evaluated at initialization x^0 before anatomy-constrained refinement.

Initialization	CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓
Neither	8.687	21.479	0.040	0.409	0.558	30.47
Endpoint only	5.111	15.582	0.110	0.351	0.516	29.80
Branch only	8.045	19.312	0.040	0.321	0.627	35.67
Endpoint + Branch	**3.832**	**8.834**	**0.127**	**0.243**	**0.595**	**33.33**

Soft Template uses CoroLift's topology and position priors to produce a population-anchored reconstruction without patient-specific morphology refinement.

Regression uses a tree-structured recurrent network to directly regress 3D coronary coordinates from the observed 2D centerlines while propagating features along the coronary topology (Tai et al., 2015; Socher et al., 2011).

B.6 Robustness Protocol

Trace noise. We evaluate sensitivity to centerline-tracing error at 0, 1, 2, 3, and 4 px noise. For each test case, the nonzero conditions are generated by scaling the same normalized spatially correlated perturbation. The reconstruction model, population statistics, and remaining hyperparameters are kept fixed across all noise levels.

Single-view acquisition angle. We evaluate sensitivity to projection direction using clinically observed single-view acquisition geometries derived from the angiographic archive. Each condition contains only one calibrated view. The reconstruction model and all fixed hyperparameters remain unchanged across acquisition geometries.

Vessel family. We report reconstruction performance across the 11 harmonized anatomical vessel families. The analysis groups the fixed test-set predictions by vessel family and does not involve retraining or family-specific parameter tuning.

Incomplete observation. We evaluate robustness to incomplete observed anatomy by progressively removing complete distal leaf slots while preserving a connected subtree containing the coronary root. This produces nested observations containing approximately 100%, 80%, 60%, and 40% of the original slots. Only retained slots are provided to the reconstruction pipeline, and population statistics and trained models remain fixed across all completeness levels. For retained-vessel evaluation, each condition is evaluated against the ground-truth slots retained at the corresponding observation level.

Common-support evaluation. To separate the effect of reduced anatomical context from changes in the composition of the evaluated vessels, we additionally use the patient-specific slot set retained at the 40% condition as a common evaluation support. Reconstructions obtained from the 100%, 80%, 60%, and 40% observations are all evaluated on this same subset. Thus, the evaluated anatomy remains fixed while the amount of surrounding anatomical context available to reconstruction is progressively reduced.

C Additional Results

C.1 Full Quantitative Results

This section provides the complete quantitative results corresponding to the experiments reported in the main paper. Unless otherwise stated, all experiments follow the strict single-view protocol with 2 px trace noise and use the same 82 test cases.

Anatomy-guided initialization at x^0. Table 6 reports the initialization ablation directly at x^0 (before anatomy-constrained refinement), complementing the main-paper Table 3, which reports the same configurations after refinement.

Table 7: **Ablation of the anatomy prior: position prior (full metrics).** Contributions of the slot-specific anatomical-keypoint (endpoint) prior μ_s^{P1} and the family-specific trajectory (branch) prior P_f, measured at initialization x^0. *Junction Gap is reported only over shared anatomical nodes actually reconstructed at initialization.

Endpoint	Branch	CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓	Junction Gap ↓
✓	✓	**3.832**	**8.834**	**0.127**	**0.243**	0.595	33.33	**0.000**
✓	–	5.111	15.582	0.110	0.351	0.516	**29.80**	11.80*
–	✓	8.045	19.312	0.040	0.321	0.627	35.67	8.06*
–	–	8.687	21.479	0.040	0.409	0.558	30.47	8.31*

Table 8: **Design of the morphology model (full metrics).** Effect of shape representation and inter-slot dependency structure on a single application of the morphology-prior operator R_M to the ray-consistent initialization x^0.

Component	Variant	CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓
Representation	Coordinate PCA	4.792	10.824	0.050	**0.045**	**0.060**	**27.347**
	Residual PCA	**3.931**	**9.395**	**0.105**	0.095	0.122	30.148
Dependency	Independent Gaussian	**3.919**	9.457	**0.110**	0.086	0.102	33.038
	Non-Markov Gaussian	3.970	9.344	0.102	0.044	0.059	27.642
	Typical-parent Gaussian-Markov	3.971	9.353	0.104	0.044	**0.054**	25.348
	Patient-specific Gaussian-Markov	3.973	**9.338**	0.105	**0.041**	0.055	**24.325**

Position prior. Table 7 reports the full-metric ablation of the position prior at initialization x^0. The anatomical keypoint and population trajectory priors provide complementary information for endpoint localization and continuous depth estimation.

Morphology model design. Table 8 compares the shape representations and dependency models used by the morphology prior under a single application of $R_{\mathcal{M}}$ to the ray-consistent initialization x^0. Residual shape modeling improves geometric accuracy over coordinate-based PCA while retaining morphology regularization. Modeling inter-vessel dependencies further changes the balance between geometric and morphological fidelity.

Anatomy-guided initialization. Table 9 reports the complete ablation of anatomy-guided initialization, including anatomical endpoint localization, endpoint fusion, and continuous branch-depth parameterization.

Anatomy-constrained refinement. Table 10 reports the successive refinement stages, including the initialization x^0, morphology regularization $R_{\mathcal{M}}(x^0)$, observation consistency $P_{\mathcal{C}}(R_{\mathcal{M}}(x^0))$, and the final reconstruction

$$x^\star = R_{\mathcal{M}}(P_{\mathcal{C}}(R_{\mathcal{M}}(x^0))).$$

The intermediate stages illustrate the complementary effects of morphology regularization and observation consistency.

Parameter analysis. Table 11 reports sensitivity to the morphology dimension K_M, Chebyshev order Q, covariance regularization α, morphology-set radius t, and observation tolerance ε. Unless otherwise stated, one parameter is varied at a time while the remaining parameters are fixed at $K_M = 16$, $Q = 12$, $\alpha = 0.1$, $t = 0.5$, and $\varepsilon = 2$ px.

Robustness analysis. Table 12 provides the complete robustness results under varying trace noise, single-view acquisition geometry, vessel family, and incomplete observed anatomy. The reconstruction model, population statistics, and fixed hyperparameters remain unchanged throughout these analyses.

C.2 Additional Diagnostic Analysis

We provide additional diagnostics of the structural and image-space properties of the reconstruction pipeline.

Table 9: **Ablation of anatomy-guided initialization.** We evaluate the sparse-to-continuous design, endpoint fusion, and continuous-depth parameterization under the strict single-view setting with 2 px trace noise.

Component	Variant	CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓
Stages	Continuous-only	**3.759**	**8.827**	**0.131**	**0.238**	0.624	**32.90**
	Two-stage (Full)	3.832	8.834	0.127	0.243	**0.595**	33.33
Endpoint Fusion	Average Fusion	3.839	**8.826**	**0.129**	**0.241**	0.597	33.74
	Precision-weighted Fusion	**3.832**	8.834	0.127	0.243	**0.595**	**33.33**
Depth parameterization	Point-wise (32 knots)	3.825	8.747	0.126	0.229	0.614	33.76
	Chebyshev ($Q = 12$)	3.832	8.834	0.127	0.243	0.595	33.33

Table 10: **Ablation of anatomy-constrained refinement.** We evaluate the complementary roles of morphology regularization R_M and observation consistency P_C under the strict single-view setting with 2 px trace noise. The final reconstruction is $x^* = R_M(P_C(R_M(x^0)))$; $P_C(R_M(x^0))$ is an intermediate stage.

Refinement Strategy	CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓
x^0	3.832	8.834	0.127	0.243	0.595	33.32
$P_C(x^0)$	3.832	8.834	0.127	0.243	0.595	33.32
$R_M(x^0)$	3.973	9.338	0.105	**0.041**	0.055	24.33
$P_C(R_M(x^0))$	**3.696**	**8.751**	**0.142**	0.108	0.290	30.50
$R_M(P_C(R_M(x^0))) = x^*$ (Full)	3.854	9.190	0.110	0.042	**0.054**	**24.19**

Initialization diagnostics. Table 13 analyzes endpoint localization, junction consistency, and ray-depth accuracy during anatomy-guided initialization. Explicit anatomical-node localization improves endpoint and junction consistency, while shared anatomical nodes eliminate discontinuities at bifurcations.

Observation-consistency diagnostics. Table 14 examines image-space consistency across the intermediate refinement stages. The ray-consistent initialization lies on the observed centerlines, whereas morphology regularization can introduce image-space deviations. Applying P_C subsequently reduces the reprojection error and restores the intermediate reconstruction to the prescribed 2 px observation tolerance.

C.3 Incomplete-observation Analysis

The retained-vessel robustness experiment evaluates each completeness level on the anatomy observed at that level. Because progressively removing distal slots also changes the composition of the evaluated vessels, we additionally perform a common-support analysis.

For each patient, we use the vessel subset retained at the 40% observation level as a fixed evaluation support. Reconstructions obtained with 100%, 80%, 60%, and 40% observed anatomy are all evaluated on this same patient-specific subset. This keeps the evaluated anatomy fixed while progressively reducing the surrounding anatomical context available to reconstruction.

The common-support analysis separates sensitivity to reduced anatomical context from changes in the composition of the evaluated vessels. Reconstruction remains relatively stable under moderate context reduction, while more severe incompleteness leads to larger degradation.

C.4 Vessel Radius Evaluation

Vessel radius is recovered analytically from the observed 2D vessel width and reconstructed depth as described in Appendix A, rather than predicted by the reconstruction networks. We report radius accuracy for both the ray-consistent initialization and the final reconstruction.

Table 11: **Parameter analysis.** Sensitivity of the final reconstruction $x^* = R_M(P_C(R_M(x^0)))$ to the main hyperparameters of CoroLift; light-blue rows denote the fixed default settings used throughout the main experiments.

(a) Morphology dimension K_M

K_M	CD ↓	HD95 ↓	Curv. ↓
4	4.093	9.147	0.057
8	3.909	9.139	0.056
16	3.854	9.190	0.054
24	3.912	9.187	0.056
32	3.908	9.233	0.056

(b) Depth basis dimension Q

Q	CD ↓	HD95 ↓	Curv. ↓
4	4.177	10.491	0.054
8	3.935	9.071	0.054
12	3.854	9.190	0.054
16	4.490	11.033	0.050
20	3.868	9.420	0.054

(c) Covariance regularization α

α	CD ↓	HD95 ↓	Curv. ↓
0.01	3.854	9.190	0.055
0.05	3.854	9.190	0.055
0.10	3.854	9.190	0.054
0.20	3.854	9.191	0.054
0.50	3.853	9.188	0.053

(d) Morphology radius t

t	CD ↓	HD95 ↓	Curv. ↓
0.3	3.947	9.220	0.055
0.4	3.859	9.187	0.054
0.5	3.854	9.190	0.054
0.7	3.908	9.231	0.053
1.0	3.906	9.307	0.059
1.5	3.903	9.372	0.066
2.0	3.907	9.399	0.075

(e) Observation tolerance ε

ε (px)	CD ↓	HD95 ↓	Curv. ↓
1.5	3.853	9.159	0.054
2.0	3.854	9.190	0.054
3.0	3.879	9.208	0.053
4.5	3.912	9.267	0.054

C.5 Additional Qualitative Results

We provide additional three-dimensional reconstruction examples from the strict single-view test set using the same reconstruction and visualization protocol as in the main paper.

D Mathematical Properties and Proofs

This section establishes the mathematical properties underlying the anatomy-constrained refinement in CoroLift. We first characterize the joint Gaussian induced by the patient-specific Gaussian-Markov morphology model. We then analyze the morphology set and show that the morphology operator $R_{\mathcal{M}}$ is an exact Euclidean projection for fixed anatomical endpoints. Finally, we characterize the geometry of the single-view observation constraint under fixed image-space correspondence and distinguish this constraint set from the finite observation-consistency operator $P_{\mathcal{C}}$ used in the implementation.

D.1 Gaussian-Markov Morphology Model

Let $\mathcal{S}$ denote the patient-specific set of observed anatomical slots, and let $p(s)$ denote the parent of each non-root slot $s \in \mathcal{S}$. For each slot, $c_s \in \mathbb{R}^{K_M}$ denotes its principal component analysis (PCA) morphology coefficients. Each root slot r follows

$$c_r \sim \mathcal{N}(\mu_r, \Sigma_r), \tag{26}$$

and each non-root slot follows

$$c_s = A_s c_{p(s)} + b_s + \epsilon_s, \qquad \epsilon_s \sim \mathcal{N}(0, Q_s), \tag{27}$$

where the innovations $\{\epsilon_s\}$ are mutually independent and independent of the root coefficients.

For notational simplicity, the proof below is written for one rooted component. If the observed coronary anatomy contains multiple rooted components, the same construction is applied independently to each component and the resulting coefficient vectors are stacked.

Table 12: **Robustness analysis.** Performance of the final reconstruction $x^* = R_M(P_C(R_M(x^0)))$ under varying trace noise, single-view acquisition geometry, vessel family, and incomplete observed anatomy.

(a) 2D trace noise

Noise (px)	CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓
0	3.852	9.473	0.136	0.039	0.050	15.91
1	3.875	9.260	0.134	0.037	0.052	20.35
2	3.854	9.190	0.110	0.042	0.054	24.19
3	3.916	9.220	0.086	0.046	0.058	27.99
4	3.999	9.253	0.076	0.048	0.061	30.88

(b) Single-view acquisition geometry

View	CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓
Default	3.854	9.190	0.110	0.042	0.054	24.19
RAOCRA / LAO	3.867	9.165	0.108	0.043	0.053	23.48
LAOCRA / LAOCRA	4.005	9.668	0.091	0.037	0.054	22.03
LAOCAU / CRA	3.796	9.131	0.113	0.040	0.052	29.19
CAU / LAO	3.747	9.228	0.107	0.042	0.056	30.07
RAOCAU / LAOCRA	3.880	9.369	0.098	0.039	0.053	25.05

(c) Vessel family

Family	N	CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓
LM	82	1.743	2.825	0.152	0.000	0.002	35.02
LAD	82	3.488	7.267	0.107	0.038	0.067	22.28
LCX	81	3.649	8.534	0.070	0.044	0.058	22.47
RCA	82	3.281	6.953	0.126	0.042	0.039	33.75
D	82	4.126	9.042	0.045	0.073	0.089	–
OM	72	3.856	7.477	0.043	0.071	0.101	–
RAMUS	41	2.728	4.975	0.051	0.046	0.057	–
R_PDA	76	3.404	5.239	0.000	0.066	0.108	60.54
R_PLB	74	3.537	6.845	0.038	0.076	0.099	30.00
L_PDA	3	4.105	5.156	0.000	–	–	–
L_PLB	23	5.420	9.102	0.008	0.087	0.162	–

(d) Incomplete observed anatomy

Observed anatomy	CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓
100%	3.854	9.190	0.110	0.042	0.054	24.19
80%	4.006	10.123	0.113	0.037	0.044	27.46
60%	4.359	10.936	0.122	0.034	0.041	29.65
40%	3.907	8.999	0.138	0.019	0.022	32.10

Proposition 1 (Joint Gaussian induced by the patient-specific tree)**.** *For any fixed patient-specific parent map, the stacked morphology coefficients follow a joint Gaussian distribution,*

$$c \sim \mathcal{N}(\mu, \Sigma). \tag{28}$$

Proof. Order the slots topologically as $s_1, \ldots, s_{|\mathcal{S}|}$, with the root preceding all of its descendants. Stack the root coefficients and all Gaussian innovations into

$$\eta = \begin{bmatrix} c_{s_1}^\top & \epsilon_{s_2}^\top & \cdots & \epsilon_{s_{|\mathcal{S}|}}^\top \end{bmatrix}^\top. \tag{29}$$

The independence assumptions give

$$\eta \sim \mathcal{N}\left(m_\eta, \Sigma_\eta\right), \tag{30}$$

Table 13: **Diagnostics of anatomy-guided initialization.** Endpoint, junction-localization, and ray-depth errors are reported at initialization. Junction gap measures the discontinuity between connected vessel slots.

Variant	Endpoint (mm)	Junction Loc. (mm)	Ray Depth (mm)	Junction Gap ↓ (mm)
Without Shared Nodes	3.613	3.481	3.389	8.9
With Shared Nodes	3.405	3.380	3.232	0.0
Average Fusion	3.413	3.416	3.235	0.0
Precision-weighted Fusion	3.405	3.380	3.232	0.0

Table 14: **Observation-consistency diagnostics.** Reprojection errors are measured from reconstructed 3D samples to their corresponding observed 2D centerlines. The final observation-consistency step restores all samples to the 2 px tolerance.

Reconstruction	Median (px)	P95 (px)	Maximum (px)	> 2 px (%)
x^0	0.000	0.000	0.0001	0.00
$P_{\mathcal{C}}(x^0)$	0.000	0.000	0.0001	0.00
$R_{\mathcal{M}}(x^0)$	0.358	2.464	19.378	7.89
$P_{\mathcal{C}}(R_{\mathcal{M}}(x^0))$	0.284	1.510	1.999	0.00

with

$$m_\eta = \begin{bmatrix} \mu_{s_1}^\top & 0 & \cdots & 0 \end{bmatrix}^\top, \tag{31}$$

and

$$\Sigma_\eta = \operatorname{blkdiag}\left(\Sigma_{s_1}, Q_{s_2}, \ldots, Q_{s_{|\mathcal{S}|}}\right). \tag{32}$$

Equation equation 27 expresses every child coefficient as an affine function of its parent coefficient and its own innovation. Recursively substituting the parent equations in topological order therefore expresses every slot coefficient as an affine function of η. Hence there exist a matrix L and vector h, determined by the patient-specific parent map and Gaussian-Markov parameters, such that

$$c = L\eta + h, \tag{33}$$

where

$$c = \begin{bmatrix} c_{s_1}^\top & \cdots & c_{s_{|\mathcal{S}|}}^\top \end{bmatrix}^\top. \tag{34}$$

An affine transformation of a jointly Gaussian vector is jointly Gaussian. Therefore,

$$c \sim \mathcal{N}(\mu, \Sigma), \tag{35}$$

with

$$\mu = Lm_\eta + h, \qquad \Sigma = L\Sigma_\eta L^\top. \tag{36}$$

□

D.2 Morphology Set and Projection

The Gaussian-Markov model defines a joint distribution over the morphology coefficients of the anatomical slots present in a patient. After covariance regularization, let

$$\Sigma_\alpha \succ 0 \tag{37}$$

denote the regularized patient-specific joint covariance. We define the morphology set as

$$M_t = \left\{c : (c-\mu)^\top \Sigma_\alpha^{-1}(c-\mu) \le t^2 D\right\}, \tag{38}$$

where

$$D = \dim(c) = K_M|\mathcal{S}|. \tag{39}$$

Table 15: **Common-support analysis of incomplete observations.** Final reconstructions $x^* = R_M(P_C(R_M(x^0)))$ obtained with different amounts of observed anatomy are evaluated on the same patient-specific S_{40} vessel set, separating reduced anatomical context from changes in evaluation composition.

Observed Anatomy	CD ↓	HD95 ↓	OV ↑	Tort. ↓	Curv. ↓	Bif. Angle ↓
100%	3.295	7.445	0.142	0.013	0.019	35.07
80%	3.360	7.865	0.132	0.015	0.022	35.36
60%	3.327	7.799	0.132	0.016	0.021	33.97
40%	3.907	8.999	0.138	0.019	0.022	32.10

The radius t controls the strength of morphology regularization and is not interpreted as a statistical confidence level.

Proposition 2 (Convexity of the morphology set)**.** *If $\Sigma_\alpha \succ 0$, then M_t is a non-empty, compact, and convex ellipsoid.*

Proof. Define

$$q(c) = (c-\mu)^\top \Sigma_\alpha^{-1}(c-\mu). \tag{40}$$

Since $\Sigma_\alpha \succ 0$, its inverse is also positive definite. The Hessian of q is

$$\nabla^2 q(c) = 2\Sigma_\alpha^{-1} \succ 0, \tag{41}$$

so q is continuous and strictly convex. Equation equation 38 defines M_t as a sublevel set of q, hence M_t is closed and convex.

Moreover,

$$q(\mu) = 0 \le t^2 D, \tag{42}$$

so $\mu \in M_t$ and the set is non-empty.

Let $\gamma_{\min} > 0$ denote the smallest eigenvalue of Σ_α^{-1}. Then

$$q(c) \ge \gamma_{\min}\|c-\mu\|_2^2. \tag{43}$$

Thus every $c \in M_t$ satisfies

$$\|c-\mu\|_2^2 \le \frac{t^2 D}{\gamma_{\min}}, \tag{44}$$

so M_t is bounded. A closed and bounded subset of finite-dimensional Euclidean space is compact. □

The coefficient-space morphology operator is defined as

$$R_{\mathcal{M}}(c_0) = \arg\min_{c \in M_t} \frac{1}{2}\|c-c_0\|_2^2. \tag{45}$$

Proposition 3 (Exact morphology projection)**.** *$R_{\mathcal{M}}$ is the exact Euclidean projection onto M_t. For $c_0 \notin M_t$,*

$$R_{\mathcal{M}}(c_0) = \mu + U\,\mathrm{diag}\left(\frac{\lambda_j}{\lambda_j+\nu}\right)U^\top(c_0-\mu), \tag{46}$$

where

$$\Sigma_\alpha = U\,\mathrm{diag}(\lambda_1,\dots,\lambda_D)U^\top, \qquad \lambda_j > 0, \tag{47}$$

and $\nu > 0$ is the unique solution of

$$\sum_{j=1}^{D} \frac{\lambda_j a_j^2}{(\lambda_j+\nu)^2} = t^2 D, \qquad a = U^\top(c_0-\mu). \tag{48}$$

Proof. If $c_0 \in M_t$, the unconstrained minimizer $c = c_0$ is feasible and has zero objective value. Hence

$$R_{\mathcal{M}}(c_0) = c_0. \tag{49}$$

Now suppose that $c_0 \notin M_t$. Since M_t is non-empty, closed, and convex by Proposition 2, the Euclidean projection exists and is unique. The optimum must lie on the boundary of M_t. Otherwise, an interior optimum would satisfy the unconstrained first-order condition $c = c_0$, contradicting $c_0 \notin M_t$.

Consider the Lagrangian

$$\mathcal{L}(c, \nu) = \frac{1}{2}\|c - c_0\|_2^2 + \frac{\nu}{2}\left[(c-\mu)^\top \Sigma_\alpha^{-1}(c-\mu) - t^2 D\right], \tag{50}$$

with $\nu \geq 0$. The stationarity condition is

$$c - c_0 + \nu\Sigma_\alpha^{-1}(c-\mu) = 0. \tag{51}$$

Since $c_0 \notin M_t$, the constraint is active at the solution and $\nu > 0$. Rearranging Eq. equation 51 gives

$$\left(I + \nu\Sigma_\alpha^{-1}\right)(c-\mu) = c_0 - \mu, \tag{52}$$

and therefore

$$c - \mu = \left(I + \nu\Sigma_\alpha^{-1}\right)^{-1}(c_0 - \mu). \tag{53}$$

Using

$$\Sigma_\alpha = U\Lambda U^\top, \qquad \Lambda = \mathrm{diag}(\lambda_1, \ldots, \lambda_D), \tag{54}$$

and

$$a = U^\top(c_0 - \mu), \tag{55}$$

Eq. equation 53 becomes

$$c - \mu = U\left(I + \nu\Lambda^{-1}\right)^{-1} a. \tag{56}$$

For each eigen-direction,

$$\left(1 + \frac{\nu}{\lambda_j}\right)^{-1} = \frac{\lambda_j}{\lambda_j + \nu}, \tag{57}$$

which gives Eq. equation 46.

It remains to determine ν. Substituting Eq. equation 46 into the active ellipsoid constraint gives Eq. equation 48. Define

$$g(\nu) = \sum_{j=1}^{D} \frac{\lambda_j a_j^2}{(\lambda_j + \nu)^2}. \tag{58}$$

Since $c_0 \notin M_t$,

$$g(0) = (c_0 - \mu)^\top \Sigma_\alpha^{-1}(c_0 - \mu) > t^2 D, \tag{59}$$

while

$$\lim_{\nu \to \infty} g(\nu) = 0. \tag{60}$$

Furthermore,

$$g'(\nu) = -2\sum_{j=1}^{D} \frac{\lambda_j a_j^2}{(\lambda_j + \nu)^3} < 0. \tag{61}$$

Thus g is continuous and strictly decreasing from a value greater than $t^2 D$ to zero. There is therefore a unique $\nu > 0$ satisfying Eq. equation 48. □

Corollary 1 (Projection properties of $R_{\mathcal{M}}$)**.** *For a fixed morphology set M_t, $R_{\mathcal{M}}$ is idempotent and firmly non-expansive.*

Proof. Proposition 3 identifies $R_{\mathcal{M}}$ with the Euclidean projector onto the non-empty closed convex set M_t. Euclidean projections onto non-empty closed convex sets are idempotent, hence

$$R_{\mathcal{M}}(R_{\mathcal{M}}(c)) = R_{\mathcal{M}}(c). \tag{62}$$

They are also firmly non-expansive, giving

$$\|R_{\mathcal{M}}(c) - R_{\mathcal{M}}(c')\|_2^2 \leq \langle R_{\mathcal{M}}(c) - R_{\mathcal{M}}(c'), c - c' \rangle. \tag{63}$$

□

D.3 Morphology Projection in Curve Space

The coefficient-space projection can be lifted to the corresponding 3D curve representation. This distinction is useful since an input curve need not lie exactly in the retained PCA subspace.

Fix the anatomical endpoints e, and let ℓ_e denote the stacked endpoint chords determined by these endpoints. Let Φ denote the block-diagonal matrix formed by the family-specific PCA bases of the observed slots. The PCA bases are orthonormal, so

$$\Phi^\top \Phi = I. \tag{64}$$

For fixed endpoints, the endpoint chords and patient-specific Gaussian parameters are fixed. The corresponding curve-space morphology set is

$$\mathcal{M}_e = \{\ell_e + \Phi c : c \in M_t(e)\}, \tag{65}$$

where $M_t(e)$ denotes the coefficient-space morphology set associated with the fixed endpoints.

For an arbitrary stacked curve u, define the curve-space realization of the morphology operator as

$$R_{\mathcal{M}}(u) = \ell_e + \Phi R_{\mathcal{M}}\left(\Phi^\top (u - \ell_e)\right), \tag{66}$$

where the inner occurrence of $R_{\mathcal{M}}$ denotes the coefficient-space projection in Eq. equation 45. We use the same symbol for the two representations when the domain is clear from context.

Proposition 4 (Exact curve-space morphology projection)**.** *For fixed endpoints e, $\mathcal{M}_e$ is a non-empty closed convex set, and*

$$R_{\mathcal{M}}(u) = P_{\mathcal{M}_e}(u) \tag{67}$$

for any curve u, including curves outside the retained PCA subspace.

Proof. Take an arbitrary curve u. Its residual from the endpoint chords admits the orthogonal decomposition

$$u - \ell_e = \Phi a + r_\perp, \tag{68}$$

where

$$a = \Phi^\top (u - \ell_e), \tag{69}$$

and

$$r_\perp = (I - \Phi\Phi^\top)(u - \ell_e). \tag{70}$$

Equation equation 64 gives

$$\Phi^\top r_\perp = 0. \tag{71}$$

For any feasible coefficient vector $c \in M_t(e)$, the corresponding curve is

$$v = \ell_e + \Phi c \in \mathcal{M}_e. \tag{72}$$

Using Eq. equation 68,

$$u - v = \Phi(a - c) + r_\perp. \tag{73}$$

The two terms are orthogonal, and therefore

$$\|u - v\|_2^2 = \|\Phi(a - c)\|_2^2 + \|r_\perp\|_2^2 \tag{74}$$

$$= \|a - c\|_2^2 + \|r_\perp\|_2^2, \tag{75}$$

where the second equality follows from $\Phi^\top \Phi = I$.

The term $\|r_\perp\|_2^2$ does not depend on c. Consequently,

$$P_{\mathcal{M}_e}(u) = \arg\min_{v \in \mathcal{M}_e} \|u - v\|_2^2 \tag{76}$$

$$= \ell_e + \Phi \arg\min_{c \in M_t(e)} \|a - c\|_2^2 \tag{77}$$

$$= \ell_e + \Phi R_{\mathcal{M}}(a). \tag{78}$$

By Eq. equation 66, this is exactly $R_{\mathcal{M}}(u)$.

Finally, $M_t(e)$ is non-empty, compact, and convex by Proposition 2. The affine map

$$c \mapsto \ell_e + \Phi c \tag{79}$$

therefore maps $M_t(e)$ to a non-empty compact convex set. Hence $\mathcal{M}_e$ is non-empty, closed, and convex. The anatomical endpoints remain fixed by construction. □

D.4 Single-view Observation Constraint

We next characterize the geometry of the observation constraint used by CoroLift. This analysis concerns a fixed image-space correspondence. The finite operator $P_{\mathcal{C}}$ used in the implementation is discussed separately below.

Consider a reconstructed 3D sample X_i with fixed corresponding point

$$\tilde{u}_i = \begin{bmatrix} \tilde{u}_i^{(1)} & \tilde{u}_i^{(2)} \end{bmatrix}^\top \tag{80}$$

on the observed 2D centerline. Let the calibrated single-view camera matrix be

$$P = \begin{bmatrix} p_1^\top \\ p_2^\top \\ p_3^\top \end{bmatrix}, \tag{81}$$

and let

$$\bar{X} = \begin{bmatrix} X^\top & 1 \end{bmatrix}^\top \tag{82}$$

denote the homogeneous coordinate of X.

For a point with positive projective depth,

$$p_3^\top \bar{X} > 0, \tag{83}$$

the calibrated perspective projection is

$$\pi(X) = \begin{bmatrix} \dfrac{p_1^\top \bar{X}}{p_3^\top \bar{X}} \\ \dfrac{p_2^\top \bar{X}}{p_3^\top \bar{X}} \end{bmatrix}. \tag{84}$$

The image-space observation constraint is

$$\|\pi(X) - \tilde{u}_i\|_2 \le \varepsilon. \tag{85}$$

Proposition 5 (Convexity of the fixed-correspondence observation set)**.** *For a fixed image-space target $\tilde{u}_i$, the closed back-projection of the ε-tolerance disk is a non-empty closed convex cone with apex at the camera center.*

Proof. For positive projective depth, Eq. equation 85 is equivalent to

$$\left\| \begin{bmatrix} \dfrac{p_1^\top \bar{X}}{p_3^\top \bar{X}} - \tilde{u}_i^{(1)} \\ \dfrac{p_2^\top \bar{X}}{p_3^\top \bar{X}} - \tilde{u}_i^{(2)} \end{bmatrix} \right\|_2 \le \varepsilon. \tag{86}$$

Since $p_3^\top \bar{X} > 0$, multiplying both sides by the positive denominator preserves the inequality:

$$\left\| \begin{bmatrix} (p_1 - \tilde{u}_i^{(1)} p_3)^\top \bar{X} \\ (p_2 - \tilde{u}_i^{(2)} p_3)^\top \bar{X} \end{bmatrix} \right\|_2 \le \varepsilon p_3^\top \bar{X}. \tag{87}$$

Define

$$A_i = \begin{bmatrix} (p_1 - \tilde{u}_i^{(1)} p_3)^\top \\ (p_2 - \tilde{u}_i^{(2)} p_3)^\top \end{bmatrix}. \tag{88}$$

The closed back-projected set can be written as

$$C_i = \left\{ X : \|A_i \bar{X}\|_2 \le \varepsilon p_3^\top \bar{X}, \quad p_3^\top \bar{X} \ge 0 \right\}. \tag{89}$$

Equation equation 89 is the inverse image of a second-order cone under an affine map of X. Second-order cones are closed and convex, and inverse images of closed convex sets under affine maps are closed and convex. Hence C_i is closed and convex.

To see the conic geometry, let o denote the camera center. For any ray direction r whose image lies in the tolerance disk and any scalar $\tau > 0$,

$$X = o + \tau r \tag{90}$$

has the same image coordinates. Including the camera center in the closure therefore makes the set closed under non-negative scaling about o. Thus C_i is a convex cone with apex at the camera center.

Finally, the ray passing through the target image point $\tilde{u}_i$ belongs to the back-projected tolerance region, so C_i is non-empty. □

For fixed correspondences over all reconstructed samples, the full observation set is the Cartesian product

$$C = \prod_i C_i. \tag{91}$$

Corollary 2 (Convexity of the fixed-correspondence observation set)**.** *For fixed image-space correspondences, the full single-view observation set C is non-empty, closed, and convex.*

Proof. Each C_i is non-empty, closed, and convex by Proposition 5. A finite Cartesian product of non-empty closed convex sets is non-empty, closed, and convex. □

D.5 Relation to the Implemented Observation Operator

The convexity result above characterizes the geometry of the observation constraint for a *fixed* image-space correspondence. The implemented observation-consistency operator $P_{\mathcal{C}}$ should not be identified with the exact Euclidean projector Π_C onto the set in Eq. equation 91.

Given the morphology-regularized reconstruction

$$y = R_{\mathcal{M}}(x^0), \tag{92}$$

CoroLift first projects each current 3D slot into the image and associates its samples with the observed 2D centerline using normalized arc-length correspondence. These correspondences are then held fixed during the local observation-consistency update.

For a sample X_i whose reprojection already satisfies

$$\|\pi(X_i) - \tilde{u}_i\|_2 \le \varepsilon, \tag{93}$$

the operator leaves the point unchanged. For a violating sample, let

$$r_i(X) = \pi(X) - \tilde{u}_i \tag{94}$$

denote its image-space residual and

$$J_i(X) = \frac{\partial \pi(X)}{\partial X} \in \mathbb{R}^{2\times 3} \tag{95}$$

the calibrated projection Jacobian. The implementation applies a minimum-norm local Gauss–Newton correction that reduces the residual toward the interior target

$$\|r_i\|_2 = 0.999\varepsilon. \tag{96}$$

Two such local updates are used for each violating sample.

The resulting $P_{\mathcal{C}}$ is therefore a finite local constraint-satisfaction operator. It is not claimed to compute the exact Euclidean projection

$$\Pi_C(y) = \arg\min_{z\in C} \|z - y\|_2^2, \tag{97}$$

and the exact-projector properties of Euclidean projection are therefore not assumed for $P_{\mathcal{C}}$. In particular, the analysis does not require $P_{\mathcal{C}}$ to be idempotent, firmly non-expansive, or globally convergent under repeated application.

The role of Proposition 5 is instead to characterize the feasible geometry induced by a fixed single-view observation. Each observed image point and tolerance defines a convex back-projected cone, while the implemented finite update provides an efficient local mechanism for returning morphology-regularized points to this observation-consistent region.

The final reconstruction applies morphology regularization before and after the observation-consistency update,

$$\boxed{x^\star = R_{\mathcal{M}}\left(P_{\mathcal{C}}\left(R_{\mathcal{M}}(x^0)\right)\right).} \tag{98}$$

The projection properties of $R_{\mathcal{M}}$ established above apply for a fixed morphology set. When the anatomical endpoints change between refinement stages, the corresponding morphology set may also change. Therefore, these properties are not used to make a convergence claim for the complete $R_{\mathcal{M}}$–$P_{\mathcal{C}}$–$R_{\mathcal{M}}$ refinement composition.