Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Evaluating Bresenham-based Binary versus Siddon-based Weighted Ray-Tracing Algorithms for Sparse-View 3D Reconstruction of Coronary Vasculature

Alireza Binesh

alireza.binesh@dokt.p.lodz.pl

Lodz University of Technology

Piotr M. Szczypínski

Lodz University of Technology

Article

Keywords:

Posted Date: June 30th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10017555/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: No competing interests reported.

Version of Record: A version of this preprint was published at Scientific Reports on November 11th, 2026. See the published version at https://doi.org/10.1038/s41598-026-79689-3.

Evaluating Bresenham-based Binary versus Siddon-based Weighted Ray-Tracing Algorithms for Sparse-View 3D Reconstruction of Coronary Vasculature

Alireza Binesh[1,*] **and Piotr M. Szczypiński**[1]

[1]Łódź University of Technology, Łódź, Poland
[*]alireza.binesh@dokt.p.lodz.pl

ABSTRACT

Interventional C-arm imaging is a standard tool for guiding minimally invasive cardiac procedures, as it reduces radiation exposure by requiring only a limited number of projections compared to computed tomography (CT). However, a major limitation of this approach is the inability to uniquely reconstruct three-dimensional data from a sparse number of projections, especially under clinical low-dose conditions characterized by high photon noise. In this work, we investigate a method for 3D reconstruction based on a limited number of C-arm images (9 views). To address the resulting ill-posed inverse problem and suppress noise, we employ Total Variation (TV) minimization and consider two alternative ray-tracing approaches within the solver: Siddon's (weighted) and Bresenham's (Binary) algorithms. The methods are evaluated using cardiovascular models with known spatial geometries. The C-arm imaging process is computationally simulated with Poisson noise, followed by the reconstruction of the 3D structures. The reconstructed results are then visually and quantitatively compared with the ground-truth data. The experiments demonstrate a significant geometric failure of Bresenham's binary approach under noisy conditions, as evidenced by a mean Dice coefficient of only 0.61 ± 0.05 and an artificial inflation of vessel volume by a factor of 1.95 ± 0.47. Siddon's method shows improved performance; combined with TV regularization, it attenuates reconstruction artifacts while providing improved preservation of morphological structures, achieving a mean Dice coefficient of 0.80 ± 0.06 and limiting average volume deviation to 5%. The results demonstrate that accurate reconstruction of vascular geometry from low-dose C-arm imaging is feasible, provided that suitable backprojection algorithms and regularization methods are employed.

Introduction

Clinical Background

Cardiovascular diseases remain the leading cause of mortality globally, driving the demand for precise interventional tools[1]. X-ray angiography using C-Arm systems is the primary modality for visualizing coronary arteries during catheterization—a minimally invasive procedure where a flexible tube is navigated through the patient's vasculature to the heart to diagnose or mechanically treat blockages. However, standard 2D projections inherently lack depth information. They capture the cumulative attenuation of X-rays as they pass through the chest, meaning the clinical input consists merely of 2D grayscale shadows of the vessel tree. When the complex, multi-branching 3D anatomy (Figure 1) is flattened into these 2D planes, vessel overlap and foreshortening occur. Foreshortening is a geometric distortion where vessels angled toward or away from the X-ray source appear artificially shorter than their true anatomical length. This loss of dimensionality creates significant ambiguity in complex scenarios, such as assessing stenosis severity in tortuous segments or guiding stent placement in bifurcations[2,17], necessitating a transition to fully volumetric 3D reconstruction.

The Problem: Sparse-View Reconstruction

The fundamental challenge in 3D reconstruction is the ill-posed inverse problem: recovering the original 3D volume based solely on its resultant 2D projections. While standard computed tomography (CT) systems acquire hundreds of projection angles to form a well-posed system of equations, interventional C-arms typically capture a limited number of projections to comply with the ALARA (As Low As Reasonably Achievable) principle for patient radiation safety[6].

Consequently, clinical protocols to minimize radiation often result in sparse-view scenarios. During these procedures, the C-arm rotates around the patient's isocenter along a circular trajectory. The number of projections acquired in low-dose C-arm imaging protocols typically consists of only a fraction of the views captured in conventional CT. In this study, an average of nine projections was assumed to reflect common clinical practice, representing a practical compromise between minimizing radiation exposure and ensuring sufficient information for reliable 3D reconstruction. Mathematically, this transforms the task into an ill-posed, underdetermined problem.

In conventional clinical CT, Filtered Back Projection (FBP) serves as the standard analytical reconstruction algorithm. However, FBP is fundamentally predicated on the Nyquist-Shannon sampling theorem, requiring hundreds of densely spaced projections. Under the sparse 9-view conditions evaluated in this study, analytical methods like FBP fail, producing severe aliasing artifacts due to angular undersampling and resulting in vascular reconstructions that are unsuitable for further clinical analysis or processing[7].

Therefore, to achieve meaningful geometric fidelity, it is necessary to abandon analytical approaches in favor of iterative reconstruction techniques. This study investigates the integration of Total Variation (TV) regularization frameworks and the choice of ray-tracing projector—specifically, the mathematical approximation of the continuous X-ray path through the discrete voxel grid. We investigate whether the combination of a high-fidelity continuous projector and gradient-domain regularization enables the recovery of distal vessel connectivity and mitigates the volumetric dilation inherent to clinical noise conditions in low-dose interventional environments.

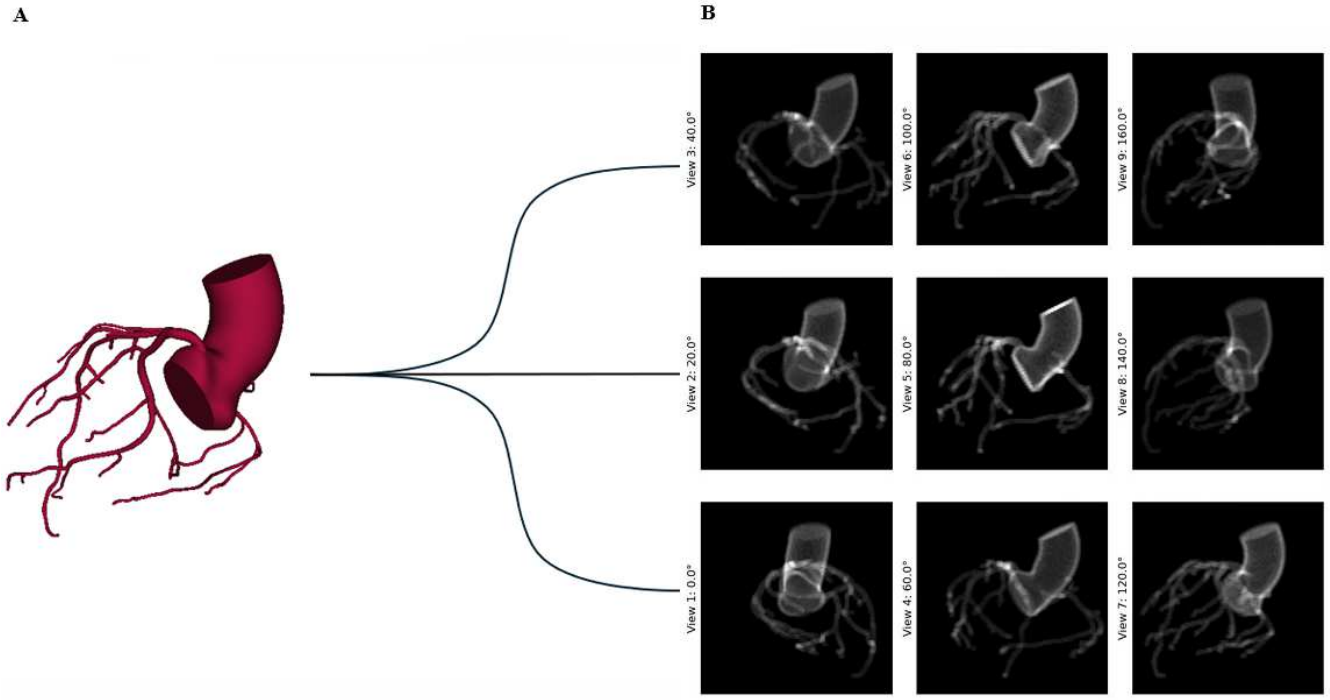


Figure 1. The Inverse Problem in Sparse-View Reconstruction. (A) The complex 3D anatomy of the patient-specific coronary tree (ground truth). For computational simulation, the mesh is treated as a solid volume uniformly filled with contrast agent. (B) Representative samples of the simulated 2D projections used as input for reconstruction. The varying grayscale intensities represent the line integral of X-ray attenuation; brighter pixels correspond to longer X-ray path lengths through the volume, highlighting thicker segments and overlapping branches. The sparse angular sampling (9 views) creates an ill-posed inverse problem, making it difficult to recover the distal branches shown in (A) without artifacts.

Related Work

Historically, 3D reconstruction in medical imaging relied on analytical direct-inversion methods, primarily Filtered Back-Projection (FBP)[7]. FBP serves as the discrete computational implementation of the mathematical inverse Radon transform[19]. These algorithms are computationally efficient and mathematically exact when provided with a well-posed dataset containing dense angular sampling (e.g., hundreds of continuous views in conventional Computed Tomography). However, the transition to

interventional C-arm environments renders these traditional methods inadequate. The physical acquisition of X-ray projections corresponds to sampling the object's Radon transform. According to the Nyquist-Shannon sampling theorem, calculating an exact, artifact-free inverse Radon transform requires a sufficient number of angular views to densely populate the Radon space[19]. In our sparse-view scenario, the large angular intervals leave large gaps in the sampled projection data. This undersampling violates the Nyquist condition, rendering the inverse problem ill-posed. Forcing a standard analytical inversion on this sparse data introduces aliasing and streak artifacts that obscure fine coronary vessel connectivity, necessitating the shift toward the regularized, iterative algebraic methods[6] utilized in this study. To address the limitations of analytical methods, the field shifted toward Iterative Reconstruction (IR) and Compressed Sensing (CS) techniques[25].

Unlike FBP, which executes an analytical direct inversion in a single computational pass, Iterative Reconstruction solves the linear system $\mathbf{Ax} = \mathbf{b}$ through a process of successive numerical approximation. Algorithms such as the Simultaneous Iterative Reconstruction Technique (SIRT)[11] execute a continuous cycle of forward projection ($\mathbf{Ax}$), residual error calculation ($\mathbf{Ax} - \mathbf{b}$), and gradient-based back-projection to incrementally update the volumetric data. While this iterative feedback loop allows the solver to mathematically compensate for the missing angular information, relying purely on unregularized SIRT introduces a risk of high-frequency noise amplification. If left unconstrained, the successive iterations will eventually begin to overfit the statistical quantum noise rather than converging on the true physical anatomy, which explicitly necessitates the integration of the regularization constraint[26] discussed in the Methods section.

While compressed sensing (CS) frameworks and total variation minimization have enabled robust sparse-view reconstruction, their success fundamentally relies on the mathematical accuracy of the underlying system matrix $\mathbf{A}$. If the forward projector utilizes a coarse, binary approximation (such as Bresenham's binary method[8]), it feeds inaccurate geometric intersection data directly into the iterative optimization loop[27]. No amount of downstream regularization can correct this fundamental geometrical error at the root of the system; the solver will successfully minimize the total variation penalty, but it will converge on a smooth yet morphologically dilated vessel tree.

In recent years, Deep Learning (DL) has emerged as a prominent area of research for solving inverse problems. Convolutional Neural Networks (CNNs) and multi-stage architectures[10,16,33] can learn complex statistical priors from large datasets, demonstrating the ability to segment and reconstruct coronary arteries by synthesizing plausible vascular structures from sparse input data. However, the efficacy of these data-driven models is fundamentally limited by the quality of their training data and their internal physics consistency. While powerful, they often act as non-linear "black boxes" that obscure the fundamental physical geometry of the reconstruction process and carry a clinical risk of hallucinating non-physical anatomical artifacts. Because the primary objective of this study is to isolate and evaluate the fundamental mathematical impact of the forward projector matrix ($\mathbf{A}$), a deterministic and interpretable framework is required. Therefore, we deliberately employ total variation (TV) minimization. Unlike "black-box" DL models, it provides a mathematically transparent, physics-based regularization environment, ensuring that any observed differences in the final reconstruction stem from the chosen ray-tracing algorithm[28].

Many DL architectures are trained on Digital Twins—synthetic vascular trees forward-projected into 2D[1]. If the forward projector utilized to generate this training corpus relies on a coarse binary approximation, the neural network learns a fundamentally flawed representation of X-ray physics. Consequently, the artificial vessel thickening (morphological dilation) inherent to binary ray-tracing is permanently encoded into the network's learned parameters, forcing the model to reproduce these geometric errors even when processing real clinical data[29,30].

State-of-the-art unrolled networks attempt to bridge this gap by directly integrating the system matrix $\mathbf{A}$ into the neural architecture as a differentiable physical layer[4]. In this coupled environment, the mathematical rigor of $\mathbf{A}$ dictates the success of the entire network. Therefore, establishing a physically accurate, anti-aliased forward projector is a mandatory prerequisite for developing the next generation of reliable, AI-driven interventional reconstruction[16].

While general-purpose algorithms and deep learning models address broad image quality, the reconstruction of coronary arteries presents unique geometric challenges. Unlike static bones, coronary arteries are narrow, tortuous structures where preserving connectivity is critical. Recent research has focused on model-based approaches. Techniques using Non-Uniform Rational B-Splines (NURBS) or centerline extraction have been proposed to mathematically force the reconstruction to resemble a tubular structure[14,32]. Similarly, Quantitative Coronary Angiography (QCA) methods attempt to infer 3D properties from 2D projections by assuming circular cross-sections[2,31].

Despite these algorithmic advancements, a critical limitation remains in how modern reconstruction frameworks interact with the physical ray-tracing physics. Many contemporary solvers—whether advanced iterative techniques or novel neural networks—continue to rely on simplified, binary forward projectors to map the 3D domain back to the 2D sensor. Consequently, any geometric inaccuracies inherent to these coarse binary approximations, such as artificial vessel thickening, are unavoidably inherited by the final reconstructed volume. Addressing this root morphological error is clinically paramount; as demonstrated

in recent studies, downstream diagnostic applications that rely on these static 3D models—such as non-invasive Fractional Flow Reserve (FFR) estimations—are exponentially sensitive to even minor deviations in reconstructed vessel radius[3,30]. Therefore, before introducing complex solvers, the fundamental geometric accuracy of the underlying system matrix **A** must first be evaluated and optimized.

If the underlying forward projector relies on binary voxel sampling, it inherently introduces morphological dilation across the entire reconstructed volume. This creates an irreducible baseline error; no amount of advanced algorithmic regularization can accurately recover fine, distal vessel branches if the fundamental physics model fails to calculate exact, sub-voxel intersection lengths. Therefore, isolating and quantifying the impact of this low-level mathematical dependency—by systematically evaluating binary versus weighted ray-tracing algorithms—constitutes the primary contribution of this work.

Methods

In a C-arm system, X-rays are emitted from a point source, pass through the patient—where they are attenuated according to tissue density—and are subsequently detected by a flat-panel detector, forming a two-dimensional projection. To model this mathematically for 3D reconstruction, the continuous physical space containing the patient is discretized into a 3D voxel grid. As illustrated in Figure 2, a sparse-view setup involves simulating geometric rays originating from discrete source positions (e.g., S_1 and S_2 along a rotation arc) that intersect this voxelized domain. The fundamental task of the forward projection model is to calculate the cumulative attenuation along each ray path as it travels from the source to a corresponding pixel **P** on the detector plane. The geometric magnification of the resulting image and the specific sequence of voxels intersected by each ray depend on the ratio of the source-to-detector distance (d_{SP}) to the source-to-object distance (d_{SO}).

The core component of any 3D reconstruction algorithm is the projector—the mathematical model that simulates how an X-ray traverses this grid. Specifically, this geometric calculation establishes the exact mathematical mapping between the spatial domains, determining precisely which discrete 3D voxels contribute to the attenuation intensity recorded at any single 2D detector pixel. To reconstruct the volume, this process is inverted via back-projection, in which the measured intensity is distributed along the corresponding ray path within the three-dimensional grid. The accuracy of this process depends on how the continuous ray path is represented in discrete voxel space.

To investigate this geometric dependency, we evaluate two fundamentally distinct forward projection models. While both algorithms inherently operate within the discrete integer coordinate space of the 3D voxel grid, they differ fundamentally in how they assign attenuation weights: the Bresenham algorithm (a binary method) and the Siddon algorithm (a weighted method).

The Binary Approach (Bresenham Algorithm)

Originally designed for rasterizing lines in computer graphics, Bresenham's algorithm approximates the continuous X-ray path using efficient integer arithmetic. Rather than calculating exact geometric intersection lengths, the algorithm steps incrementally along the primary axis of the ray path and maintains an internal error variable. This error term tracks the mathematical divergence between the true continuous line and the discrete voxel centers. At each grid interval, the algorithm uses this error term to make a binary decision, selecting the single voxel that best approximates the line's trajectory. If a voxel is selected, it is fully incorporated into the ray sum with a normalized weight of 1. Conversely, all other adjacent voxels are entirely excluded with a weight of 0.

Operationally, this approach utilizes Bresenham's line algorithm[8] to traverse the grid by incrementing along the primary axis of the ray one integer unit at a time. As illustrated in Figure 3B, this strict binary sampling yields a jagged, aliased trajectory—commonly referred to as a "staircase" effect—rather than a precise geometric line. Consequently, this binary approximation introduces a geometric artifact: the simulated ray effectively assumes a non-zero thickness dictated by the voxel dimensions. During reconstruction, this expanded spatial footprint artificially dilates the structural boundaries, leading to a significant overestimation of vessel volume.

The Weighted Approach (Siddon's Algorithm)

Conversely, the Siddon algorithm is referred to as weighted because it computes a continuous, fractional contribution for every voxel intersected by the ray path. Unlike the binary inclusion criterion of the Bresenham method, this approach calculates the precise geometric length of the line segment traversed by the ray within each individual voxel. This exact path length, normalized by the voxel dimensions, serves directly as the voxel's specific attenuation weight. Computationally, this technique[9] analytically determines the exact spatial coordinates where the continuous ray enters and exits each voxel boundary. By

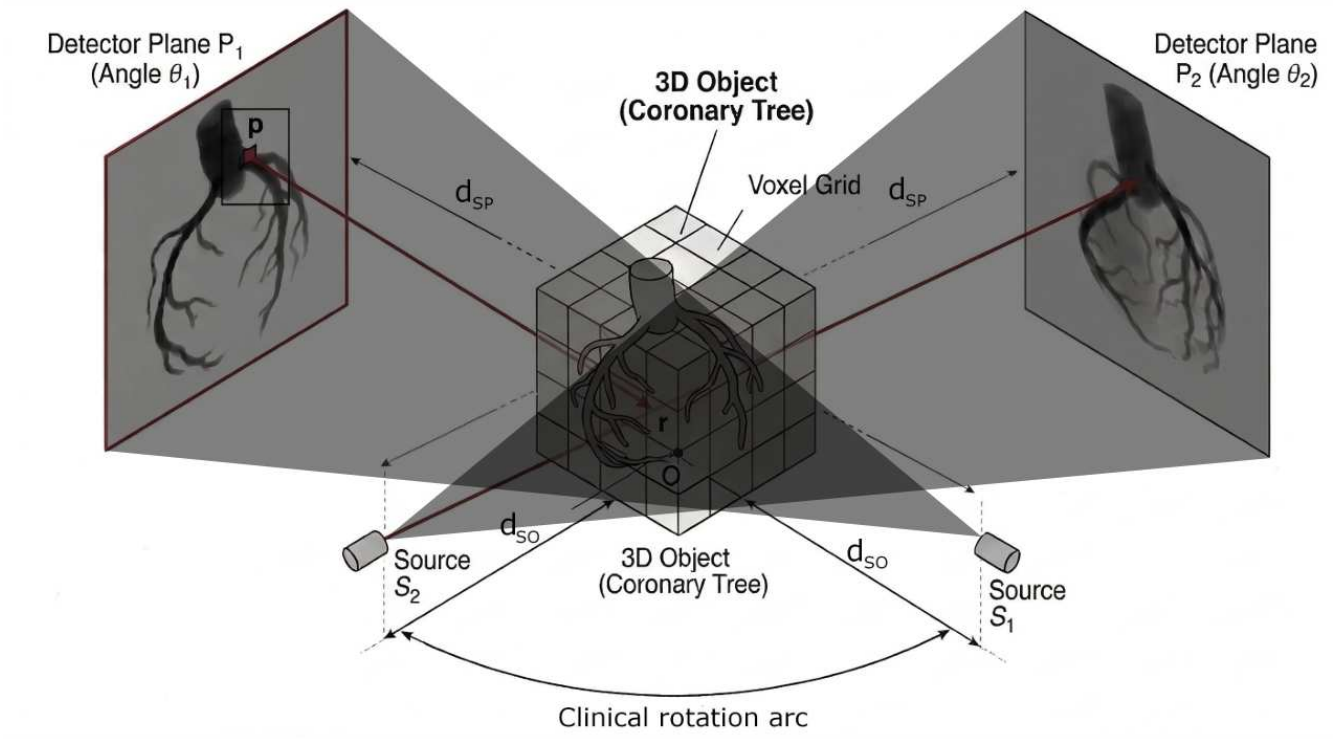


Figure 2. Geometric schematic of the cone-beam X-ray projection system for 3D coronary tree reconstruction. X-ray sources (S_1, S_2) are positioned along a standard clinical rotation arc. The 3D object (coronary vascular tree) is centered at the isocenter O and enclosed within a discrete voxel grid. For visual clarity, full expanding cone beams are omitted to emphasize the central geometric rays (red, R_1, R_2), which intersect the 3D volume and terminate at pixels on the opposing 2D detector planes.

modeling the physical principles of X-ray attenuation (where cumulative intensity reduction is proportional to the continuous path length traversed), this method accurately accounts for the partial volume effect at the boundaries of the vessels. This mathematically exact footprint is essential for preventing artificial dilation and preserving the precise morphology of fine coronary branches during reconstruction[12].

Having established how individual X-ray paths are computationally modeled, we can now formalize the overarching reconstruction challenge. The ultimate goal is to invert this forward projection process to accurately recover the 3D volume from the 2D projections. However, because interventional C-arm systems rely on extremely sparse data, this inversion is underdetermined. To understand exactly why analytical methods fail under these constraints, we must examine the mathematical nature of this specific inverse problem.

Following the physical principles of Computed Tomography (CT), the exponential X-ray attenuation (governed by the Beer-Lambert law[19]) is log-transformed to establish a linear dependence between the measured attenuation and the object's internal linear attenuation coefficients. Consequently, the 3D reconstruction process is modeled mathematically as a large-scale system of linear equations:

$$\mathbf{Ax} = \mathbf{b} \tag{1}$$

Where:

- $\mathbf{A}$ is the $M \times N$ system matrix (the Projector). Because the X-ray path geometry changes as the C-arm rotates, $\mathbf{A}$ is angle-dependent. It is composed of sub-matrices corresponding to each specific projection view, encoding the exact intersection weights of every ray at every angle.
- $\mathbf{x}$ is a column vector representing the unknown 3D volume. It contains N total elements, where N corresponds to the total number of discrete voxels in the 3D grid.

- **b** is a column vector of length M representing the measured projection data. Instead of a traditional sinogram used in slice-based CT, **b** is constructed by vectorizing the pixel intensities from all acquired 2D projection images.

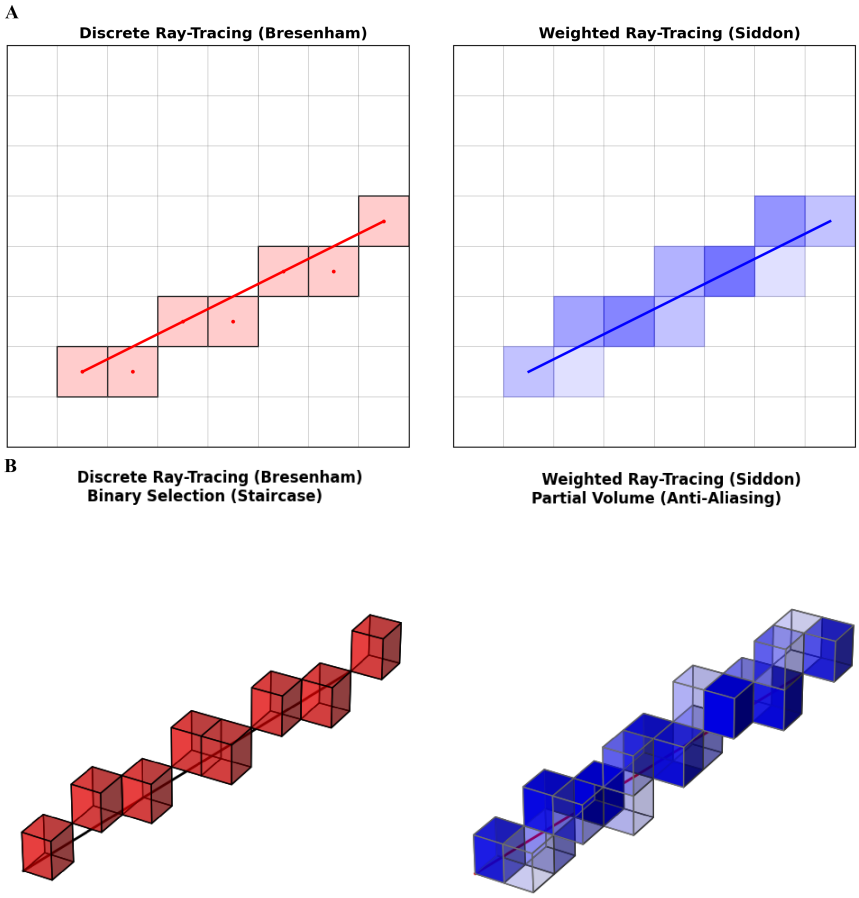


Figure 3. Geometric Ray-Tracing Models. (A) 2D Conceptualization: A simplified view showing how a ray traverses a pixel grid. The binary method (red) selects pixels based on center proximity—notably omitting pixels that the ray physically intersects if they fail this proximity condition—while the weighted method (blue) accounts for the exact intersection length. (B) 3D Implementation: The actual voxel-based operations used in this study. The binary (Bresenham) approach creates a binary "staircase" of voxels (red), causing artificial thickening. The weighted (Siddon) approach utilizes the exact chord length through each voxel (blue transparency), providing a sub-voxel accurate model of the X-ray beam.

The mathematical stability of this linear system is dictated by the ratio of measurements (M) to unknowns (N). In a standard clinical CT scan, hundreds of projection angles are acquired such that the number of measurements is equal to or greater than the number of voxels ($M \geq N$). This creates a well-posed or overdetermined system with a unique solution.

However, in our interventional C-arm scenario, we are limited to an extremely sparse dataset of only 9 views. Consequently, the number of unknown voxels vastly exceeds the number of measured pixels ($M \ll N$), creating a severely underdetermined and ill-posed inverse problem. This underdetermination manifests as a high-dimensional null space, meaning that an infinite number of distinct 3D volumetric distributions for **x** can theoretically satisfy the equations and yield the exact same 2D projection data vector **b** from a restricted set of angles. This mathematical ambiguity renders a direct analytical inversion fundamentally unfeasible. If standard techniques such as Filtered Back-Projection are applied under these sparse conditions, the violation of the Nyquist criteria generates severe streak artifacts and aliasing artifacts, because the unregularized mathematics lacks the constraints necessary to distinguish the true continuous coronary anatomy from an infinite combination of non-physical ghost solutions[7].

Because direct inversion of the underdetermined system $\mathbf{Ax} = \mathbf{b}$ is mathematically non-viable, an iterative solver must be employed. However, to prevent the algorithm from converging on noisy, non-physical "ghost solutions" within the infinite null space, the solver must be constrained. This necessity is amplified in true clinical low-dose scenarios, where the measured projection data ($\mathbf{b}$) is inherently corrupted by Poisson quantum noise (photon starvation).

To navigate these challenges and isolate the correct anatomical solution, we formulate the reconstruction as a regularized convex optimization problem[25]:

$$\hat{\mathbf{x}} = \arg\min_{\mathbf{x}\geq\mathbf{0}} \left(|\mathbf{Ax} - \mathbf{b}|^2 + \lambda \mathbf{R}(\mathbf{x}) \right) \tag{2}$$

Where:

- $\hat{\mathbf{x}}$: The final estimated 3D volume. The optimization is subjected to a non-negativity constraint ($\mathbf{x} \geq \mathbf{0}$), reflecting the physical reality that X-ray attenuation coefficients cannot be less than zero.
- $|\mathbf{Ax} - \mathbf{b}|^2$: The Data Fidelity term. It calculates the squared L_2-norm of the residual error, ensuring that the simulated forward projection of the 3D model accurately reproduces the measured 2D X-ray shadows.
- $\mathbf{R}(\mathbf{x})$: The Regularization term. To address the ill-posed nature of this sparse-view inverse problem, we specifically employ total variation (TV) minimization[6,11]. Mathematically, this is defined as the L_1-norm of the spatial gradient magnitude of the volume, which effectively suppresses noise while preserving the sharp morphological boundaries of the vessels.
- λ : The scalar weighting parameter that balances the optimization trade-off. It controls how much influence the TV regularization has versus the data fidelity.

Specifically, the isotropic total variation (TV) of the discrete volume is computed by evaluating the spatial gradients across the 3D voxel grid. Mathematically, it is formulated as:

$$\mathbf{R}(\mathbf{x}) = |\nabla \mathbf{x}|_1 = \sum_{j=1}^{N} \sqrt{(\mathbf{D_x x_j})^2 + (\mathbf{D_y x_j})^2 + (\mathbf{D_z x_j})^2} \tag{3}$$

Where:

- $|\nabla \mathbf{x}|_1$: Denotes the L_1-norm of the discrete gradient magnitude across the entire volume.
- N: Is the total number of voxels in the reconstructed volume.
- $\mathbf{x_j}$: Represents the attenuation value of the j-th voxel.
- $\mathbf{D_x}, \mathbf{D_y}, \mathbf{D_z}$: Represent the linear, first-order finite difference operators along the three spatial dimensions. Although $\mathbf{x}$ is manipulated as a 1D column vector during the $\mathbf{Ax}$ projection step, these operators explicitly map the localized spatial differences between neighboring voxels in the reconstructed 3D domain.

By minimizing this sum, the TV penalty actively smooths regions of uniform tissue density while preserving the sharp, high-contrast gradient boundaries characteristic of coronary vessels filled with a contrast agent.

Crucially, the regularization term $\mathbf{R}(\mathbf{x})$ operates solely within the 3D volumetric domain and remains identical across both the binary and weighted reconstruction implementations. This controlled mathematical design ensures that any observed geometric or morphological differences in the final reconstructions are attributable entirely to the specific mathematical formulation of the forward projector matrix $\mathbf{A}$ (Bresenham vs. Siddon).

This specific regularization choice is fundamentally grounded in Compressed Sensing (CS) theory[6,11,21]. Classical sampling theorems dictate that accurately reconstructing a volume from only 9 projection views is impossible. However, CS theory demonstrates that an undersampled system can be robustly solved provided the underlying data is sparse (meaning it consists mostly of zeros) in a specific mathematical transform domain. By enforcing this sparsity in the gradient domain, TV

minimization leverages the assumption that biological structures—such as contrast-filled coronary arteries—are piecewise-smooth, consisting of uniform tissue densities separated by sharp boundaries. Consequently, this CS framework allows the reconstruction to effectively suppress high-frequency noise and undersampling streak artifacts while preserving the morphological integrity of the vessel walls.

To transition from the raw 2D projection images to the final 3D coronary model, the computational pipeline executes the following sequential steps:

1. **Volume Initialization:** Before the solver begins, the 3D volumetric space is initialized as a null vector ($\mathbf{x}^0 = \mathbf{0}$). This establishes a completely blank, non-negative baseline grid.
2. **Iterative Optimization:** The algorithm enters the optimization loop. In each iteration (k), the current volume guess ($\mathbf{x}^{(k)}$) is forward-projected using the system matrix $\mathbf{A}$. The resulting simulated shadows are mathematically compared against the actual acquired sparse projection set (the 9 views stored in $\mathbf{b}$) to calculate the data fidelity error. The volume is then updated to $\mathbf{x}^{(k+1)}$ using the regularization gradient to minimize this error while enforcing spatial sparsity.
3. **Convergence (Stopping Criterion):** This iterative loop continues until a strict stopping condition is met. The solver terminates either when the relative change between consecutive volumetric updates falls below a predefined tolerance threshold ($\|\mathbf{x}^{(k+1)} - \mathbf{x}^{(k)}\|_2 < \varepsilon$), indicating mathematical convergence, or when a maximum allowed number of iterations (k_{max}) is reached.
4. **Volumetric Segmentation:** Upon convergence, the solver outputs a continuous, grayscale 3D attenuation map. To isolate the physical geometry of the coronary arteries, a final binarization step is applied. Using Otsu's thresholding algorithm, an optimal global threshold is automatically computed to separate the high-intensity contrast-filled vessels from the low-intensity background tissue, yielding the final discrete 3D morphological model.

To address the challenge of morphologically accurate 3D coronary reconstruction from a sparse set of projection views, this study employs a digital twin validation framework[1]. In this context, a digital twin refers to a completely synthetic, mathematically defined 3D vascular model utilized for *in silico* simulation. This controlled computational approach allows for a quantitative assessment of the reconstruction algorithm's geometric fidelity. By using a synthetic model, the exact ground truth 3D anatomy is known prior to the experiment—a baseline condition that is fundamentally impossible to achieve in real-world clinical settings, where the patient's actual 3D arterial shape is the unknown variable being sought.

The experimental pipeline follows a three-stage workflow:

1. **Input (Simulated Acquisition):** We begin with high-resolution 3D anatomical models (phantoms). From these, we mathematically generate a set of sparse 2D projections (simulating the C-Arm X-ray process) to serve as the input data $\mathbf{b}$.
2. **Process (Reconstruction):** The sparse 2D data is fed into the reconstruction engine. We employ the Simultaneous Iterative Reconstruction Technique (SIRT) solver[11], comparing two distinct system matrix $\mathbf{A}$ configurations: the binary (Bresenham) projector versus the weighted (Siddon) projector.
3. **Output Validation:** The resulting 3D volumetric reconstructions ($\hat{\mathbf{x}}$) are mathematically compared against the original ground truth phantoms. The primary objective is to quantify the geometric fidelity and the extent of morphological dilation introduced by the specific choice of forward ray-tracing projector.

To ensure the robustness of our findings across varying anatomical geometries, we utilized three distinct patient-specific coronary models (visualized in Figure 4). These models were selected to represent a range of vascular complexities, including variations in tortuosity, bifurcation angles, and vessel caliber.

The digital twin phantoms utilized in this study were derived from patient-specific anatomical geometries (Models 0066_H_CORO_H, 0271_H_CORO_H, and 0286_H_CORO_H) provided by the open-source Vascular Model Repository (Cardiovascular Biomechanics Computation Lab, Stanford University)[24]. Throughout this study, they are referred to as Model 1, Model 2, and Model 3.

Model 1 (Baseline Anatomy) is derived from a high-resolution dataset representing a comprehensive human coronary tree, including both left and right coronary arteries. The mesh data was extracted and visualized using 3D Slicer, an open-source platform for medical image analysis, from the original file structure. It serves as our baseline for a standard, healthy coronary

anatomy. Model 2 and Model 3 (Complex Pathologies) were acquired from the Vascular Model Repository (VMR), an open-source database of validated cardiovascular models. These models were specifically chosen to introduce higher anatomical variability. The VMR datasets are derived from real patient imaging and preserve the intricate, non-uniform structures of diseased or tortuous vessels. Testing on these diverse geometries ensures that our findings are not artifacts of a single simple geometry but are applicable to the complex, unpredictable shapes encountered in interventional cardiology.

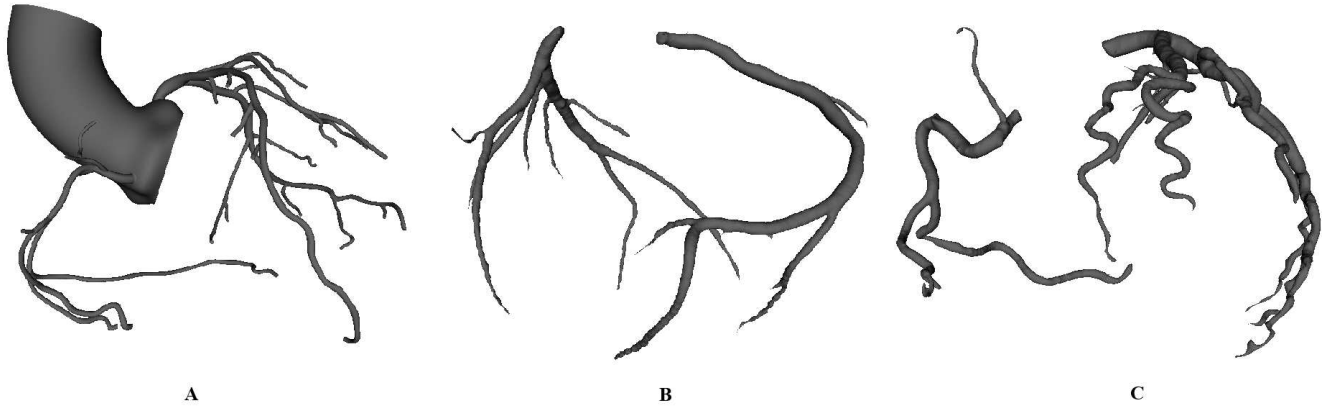


Figure 4. Patient-Specific Validation Dataset. (A) Model 1: A baseline coronary tree showing standard anatomy. (B) Model 2 and (C) Model 3: Pathological models acquired from the Vascular Model Repository, selected for their high tortuosity and complex bifurcation structures to stress-test the reconstruction algorithms.

To perform the reconstruction, we utilized the ASTRA Toolbox (All Scale Tomographic Reconstruction Antwerp)[13]. While often treated as a "black box" in general imaging tasks, we selected this framework specifically because of its modular architecture. It allows us to explicitly define the system matrix—the mathematical model of the X-ray geometry—independently of the iterative solver, while offloading the massive ray-tracing calculations to the GPU for computational efficiency.

The core of this experiment is the comparison of two distinct methods for calculating the elements of the system matrix $\mathbf{A}$. These weights determine the geometric relationship between a specific X-ray path and a specific voxel.

The baseline forward projector implements Bresenham's Line Algorithm[8], an approach originally developed for drawing linear trajectories on discrete raster displays. To model the X-ray trajectory, this algorithm forces the continuous physical ray to snap to discrete integer grid coordinates using binary, nearest-neighbor logic. Consequently, the resulting system matrix is binary: if the geometric center of a ray intersects a voxel, it is assigned a uniform weight of 1; otherwise, it is assigned a 0. This binary classification creates a jagged, "staircase" approximation of the true continuous ray. By mathematically treating a thin, linear X-ray beam as a blocky chain of uniform cubes, this method artificially widens the intersection profile, directly introducing the morphological dilation and artificial vessel thickening observed in the final reconstructed volumes.

To construct the weighted system matrix, the algorithm calculates the exact physical intersection length of the X-ray traversing each discrete cubic voxel, normalized by the physical voxel dimensions. Consequently, the resulting matrix is populated with continuous floating-point values dependent on the 3D angle of incidence. For example, a ray passing diagonally through the center of a voxel assigns a proportionally higher geometric weight than an orthogonal ray, while a ray merely grazing an outer corner yields a weight approaching zero. This continuous geometric weighting acts as an intrinsic spatial anti-aliasing filter. By mathematically accounting for these partial-volume intersections, the forward projector represents sub-voxel boundaries, thereby preventing the morphological dilation and artificial vessel thickening produced by binary ray-tracing algorithms.

To solve the inverse problem, we utilized the Simultaneous Iterative Reconstruction Technique (SIRT)[11,20]. Because direct analytical inversion of the system matrix $\mathbf{A}$ is computationally intractable for high dimensional data, SIRT approaches the solution through successive approximation. Specifically, the algorithm's update rule employs the transpose matrix $\mathbf{A}^T$—which physically corresponds to the back-projection operator—to map the computed residual error from the 2D projection domain back into the 3D image volume. Unlike sequential algebraic methods that update the volume ray-by-ray, SIRT computes these iterative updates simultaneously based on the data from all projection angles. This global update strategy is advantageous for sparse-view data, as it inherently averages out quantum noise and suppresses the convergence oscillations common in sequential algorithms.

We formulate this reconstruction task as a regularized least-squares minimization problem. To address the ill-posedness of our 9-view scenario, we apply a strict non-negativity constraint ($\mathbf{x} \geq \mathbf{0}$). This enforces physical realism—since X-ray attenuation density cannot be negative—and significantly restricts the allowable mathematical solution space to prevent artifact divergence. Ultimately, the SIRT algorithm is deployed to iteratively compute the volumetric solution ($\hat{\mathbf{x}}$) that minimizes the objective function defined previously in Equation 2.

The update rule for the k-th iteration is defined as:

$$\mathbf{x}^{(k+1)} = \mathbf{x}^{(k)} + \mathbf{C} \cdot \underbrace{\mathbf{A}^T}_{\text{back projection}} \left(\mathbf{b} - \underbrace{\mathbf{A}\mathbf{x}^{(k)}}_{\text{forward projection}} \right) \tag{4}$$

In this iterative update rule, $\mathbf{x}^{(k)}$ represents the current volumetric estimate at iteration k, and $\mathbf{x}^{(k+1)}$ represents the updated volume for the subsequent step. The vector $\mathbf{b} \in \mathbb{R}^M$ denotes the projection data acquired from the 9 sparse-view angles. The system matrix $\mathbf{A}$ serves as the forward projector, computing the simulated projections ($\mathbf{A}\mathbf{x}^{(k)}$) to calculate the residual error. In our experimental design, this matrix acts as the primary independent variable, deliberately swapped between binary ($\mathbf{A}_{\text{binary}}$) and weighted ($\mathbf{A}_{\text{weighted}}$) ray-tracing models to isolate their geometric effects. Because a direct inverse ($\mathbf{A}^{-1}$) does not exist for this rectangular system, the update relies on the transpose matrix $\mathbf{A}^T$, which functions physically as the adjoint back-projection operator to map the residual error back into the 3D voxel space. Finally, $\mathbf{C}$ represents a diagonal pre-conditioning relaxation matrix that scales the step size based on intersection sums. This strict weighting prevents numerical divergence in the center of the imaging field, ensuring uniform convergence across the entire volume.

It is critical to note that the specific geometric model chosen for the system matrix ($\mathbf{A}$) governs both the forward-projection and back-projection steps defined in Equation 4. If the binary model ($\mathbf{A}_{\text{binary}}$) is utilized, the residual error is forcefully back-projected along the same jagged, binary coordinate path. Conversely, if the weighted model ($\mathbf{A}_{\text{weighted}}$) is employed, the back-projected error is distributed proportionally across the grid based on exact partial-volume intersections, effectively mitigating spatial aliasing without falsely implying a continuous spatial domain.

To ensure a fair comparative analysis, all initialization and solver hyperparameters were kept identical across both experiments. The simulated projection data ($\mathbf{b}$) was generated by forward-projecting the digital twin phantoms from 9 equispaced angles spanning a 180° arc. This specific degree of sparsity was chosen to emulate a dose- and contrast-limited interventional scenario, capturing the minimum data required to span the angular domain while freezing cardiac motion. To avoid the well-documented inverse crime[31] of testing an algorithm using the same mathematical model used to generate the data, all ground-truth projections were generated using the high-fidelity weighted model. Furthermore, to accurately model clinical photon starvation and a spatially varying, low signal-to-noise ratio (SNR), Poisson noise was subsequently introduced to the forward-projected data by simulating an unattenuated incident X-ray intensity of $I_0 = 5 \times 10^3$ photons per detector pixel.

For the iterative reconstruction pipeline, the 3D volume was initialized as a null vector ($\mathbf{x}^0 = \mathbf{0}$). During the SIRT optimization, a hard non-negativity constraint was enforced to prevent the formation of physically impossible negative attenuation artifacts. The solver utilized a fixed stopping criterion of 150 iterations—where one iteration encompasses a full forward and backward cycle across all 9 angles. This threshold was selected empirically, as the relative volumetric change successfully plateaued by this point, guaranteeing mathematical convergence for both projection models. Finally, the Total Variation (TV) regularization hyperparameter (λ) was tuned via a systematic parameter sweep, utilizing the Root Mean Square Error (RMSE) between the reconstructed volume and the digital twin ground truth as the quantitative selection criterion. Because an excessive λ over-smooths delicate distal coronary branches, while an insufficient λ allows Poisson noise to dominate, this mathematically optimized baseline value was established to balance noise suppression against geometric fidelity, and was held rigidly constant across all comparative evaluations.

Quantitative Performance Metrics

To evaluate the impact of the projector model on reconstruction quality, we employed four complementary metrics. These were selected to assess distinct characteristics of the output: spatial overlap (Dice, Jaccard), intensity fidelity (RMSE), and geometric volume preservation (Volume Ratio).

In the formulations below, let $\mathbf{V}_{\text{gt}}$ denote the ground truth volume and $\mathbf{V}_{\text{rec}}$ denote the reconstructed volume. For the overlap-based metrics (DSC and Jaccard), these volumes are binarized using a global threshold to isolate the vessel structure.

The Dice Similarity Coefficient (DSC) is a statistical metric utilized to quantify the volumetric spatial overlap between the

reconstructed vessel tree and the baseline ground truth phantom. Because DSC operates on Boolean logic, the continuous attenuation map generated by the SIRT solver must first be binarized into a discrete vessel mask. Consequently, the sets evaluated in this metric represent the specific collections of 3D spatial coordinates (voxels) classified exclusively as positive vascular tissue. The coefficient is defined mathematically as:

$$\mathrm{DSC} = \frac{2|\mathbf{V}_{\mathrm{gt}} \cap \mathbf{V}_{\mathrm{rec}}|}{|\mathbf{V}_{\mathrm{gt}}| + |\mathbf{V}_{\mathrm{rec}}|} \tag{5}$$

where $\mathbf{V}_{\mathrm{gt}}$ represents the set of true vessel voxels in the digital twin phantom, $\mathbf{V}_{\mathrm{rec}}$ represents the set of segmented vessel voxels in the reconstructed volume, and $|\cdot|$ denotes the mathematical cardinality (the absolute count of non-zero voxels) of the respective sets.

Mathematically, the DSC yields a normalized scalar value ranging from 0 (indicating zero spatial overlap) to 1 (representing perfect voxel-for-voxel alignment). Within medical image segmentation, it is widely regarded as the gold-standard metric for evaluating global geometric fidelity. For the purposes of this study, the DSC serves as the primary quantitative indicator of morphological dilation. Because the binary forward projector inherently causes artificial vessel thickening, the resulting reconstructed volume ($|\mathbf{V}_{\mathrm{rec}}|$) erroneously expands beyond the strict physical boundaries of the baseline digital twin ($|\mathbf{V}_{\mathrm{gt}}|$). This volumetric overestimation inherently penalizes the spatial overlap calculation, allowing the Dice score to serve as a direct measurement of the geometric error introduced by the ray-tracing algorithm.

While similar to the DSC, the Jaccard Index (or Intersection over Union, IoU) imposes a mathematically stricter penalty on spatial mismatches. Defined as:

$$\mathrm{Jaccard} = \frac{|\mathbf{V}_{\mathrm{gt}} \cap \mathbf{V}_{\mathrm{rec}}|}{|\mathbf{V}_{\mathrm{gt}} \cup \mathbf{V}_{\mathrm{rec}}|} \tag{6}$$

The Jaccard Index yields a value between 0 and 1. We utilize it to penalize morphological dilation. Because the index divides the intersection by the total combined spatial union, any artificial thickening produced by the binary algorithm drastically inflates the union denominator. This geometric overestimation degrades the Jaccard score much more aggressively than the DSC, allowing for a clear quantitative distinction between a volume that is merely misaligned and one that is artificially dilated.

While the Dice similarity coefficient evaluates the geometric accuracy of the binarized vessel masks, it cannot capture subtle intensity variations, background noise, or attenuation drop-offs. Therefore, to evaluate the fidelity of the raw, continuous physical density maps produced by the solver, we calculate the root mean square error (RMSE). This metric is computed directly on the continuous grayscale volumetric output ($\mathbf{V}_{\mathrm{rec}}$) against the continuous ground truth phantom ($\mathbf{V}_{\mathrm{gt}}$), prior to any thresholding or binarization steps. The metric is defined as:

$$\mathrm{RMSE} = \sqrt{\frac{1}{N}\sum_{j=1}^{N}(\mathbf{V}_{\mathrm{gt}}(j) - \mathbf{V}_{\mathrm{rec}}(j))^2} \tag{7}$$

where N represents the total number of voxels in the 3D volume, and j denotes the linear voxel index. By squaring the intensity residuals before averaging, RMSE inherently applies a heavier mathematical penalty to large, localized reconstruction artifacts—such as the density streaks commonly caused by binary ray-tracing algorithms.

Mathematically, a lower RMSE value indicates greater physical consistency between the reconstructed attenuation map and the true contrast density. Because the weighted forward projector models the exact partial-volume attenuation path, it is hypothesized to recover the true internal density values of the contrast medium. Conversely, a high RMSE in the binary method indicates that approximation errors inherently introduce geometric distortions into the voxel intensities, corrupting the internal physical density even if the global macroscopic shape appears relatively intact.

To prepare the data for metric compatibility and ensure an objective comparison between the continuous grayscale domain evaluated by the RMSE and the discrete boolean domain required by the DSC, an automated binarization step was implemented. Following the SIRT reconstruction, the continuous volumes were segmented into binary vessel masks utilizing Otsu's thresholding method. This approach algorithmically minimizes intra-class intensity variance, providing an objective, parameter-free binarization baseline for all discrete morphological evaluations.

The volume ratio (VR) serves as the primary quantitative metric for testing the core hypothesis regarding geometric dilation and algorithmic stability. Defined simply as:

$$\mathrm{VR} = \frac{|\mathbf{V}_{\mathrm{rec}}|}{|\mathbf{V}_{\mathrm{gt}}|} \tag{8}$$

This metric measures the global volumetric output. A VR of 1.0 indicates perfect volume preservation, whereas a VR > 1.0 indicates artificial morphological dilation, and a VR < 1.0 indicates vessel narrowing. Because the binary algorithm inherently approximates continuous trajectories as block-like structures, it is mathematically expected to add false spatial volume to the reconstructed vessels. Consequently, the VR provides a quantitative measure of the morphological errors introduced by binary ray-tracing matrices.

Results

Quantitative metrics were calculated, and 3D surface models for visual assessment were generated using the optimal isosurface thresholds derived from the sensitivity analysis.

Table 1. Quantitative Evaluation of 3D Coronary Reconstruction (9 Views). This table compares the geometric and photometric accuracy of the binary (Bresenham) forward projector against the weighted (Siddon) projector, both evaluated within the identical regularized SIRT+TV solver framework. Under low-dose noise conditions, the binary method exhibits volume dilation (1.95 ± 0.47) and poor geometric overlap (Jaccard = 0.44 ± 0.05). In contrast, the weighted method successfully suppresses Poisson noise, maintaining high geometric fidelity (Jaccard = 0.67 ± 0.09) and reducing the Root Mean Square Error (RMSE) by over 50% (0.03 ± 0.03 vs 0.08 ± 0.04).

Phantom	Method	Best Dice (DSC)	Jaccard (IoU)	Volume Ratio (VR)	RMSE
Model 1	Binary (SIRT+TV)	0.56	0.39	2.50	0.13
	Weighted (SIRT+TV)	0.87	0.77	1.16	0.06
Model 2	Binary (SIRT+TV)	0.62	0.45	1.66	0.05
	Weighted (SIRT+TV)	0.76	0.62	0.94	0.02
Model 3	Binary (SIRT+TV)	0.65	0.48	1.70	0.07
	Weighted (SIRT+TV)	0.76	0.61	1.06	0.02
Mean $\pm$ SD	**Binary (SIRT+TV)**	**0.61 ± 0.05**	**0.44 ± 0.05**	**1.95 ± 0.47**	**0.08 ± 0.04**
	Weighted (SIRT+TV)	**0.80 ± 0.06**	**0.67 ± 0.09**	**1.05 ± 0.11**	**0.03 ± 0.03**

Under simulated low-dose conditions—characterized by a low signal-to-noise ratio (SNR) and Poisson noise—the binary (Bresenham) method exhibited significant morphological dilation, resulting in a mean Jaccard Index of 0.44 ± 0.05 and a mean volume ratio (VR) of 1.95 ± 0.47. Because the binary ray-tracer relies on rigid, binary voxel selection, it lacked the sub-voxel precision necessary to differentiate between true structural boundaries and adjacent background noise.



Figure 5. Visual Comparison of 3D Coronary Reconstructions under Poisson Noise. (A) Ground truth volumes. (B) The binary model (SIRT+TV) exhibits significant morphological dilation, incorporating background Poisson noise to expand the vessel volumes by an average factor of 1.95. (C) The weighted model (SIRT+TV) suppresses the Poisson noise while preserving sub-voxel geometric boundaries, recovering the hollow tubular structures and distal branches.

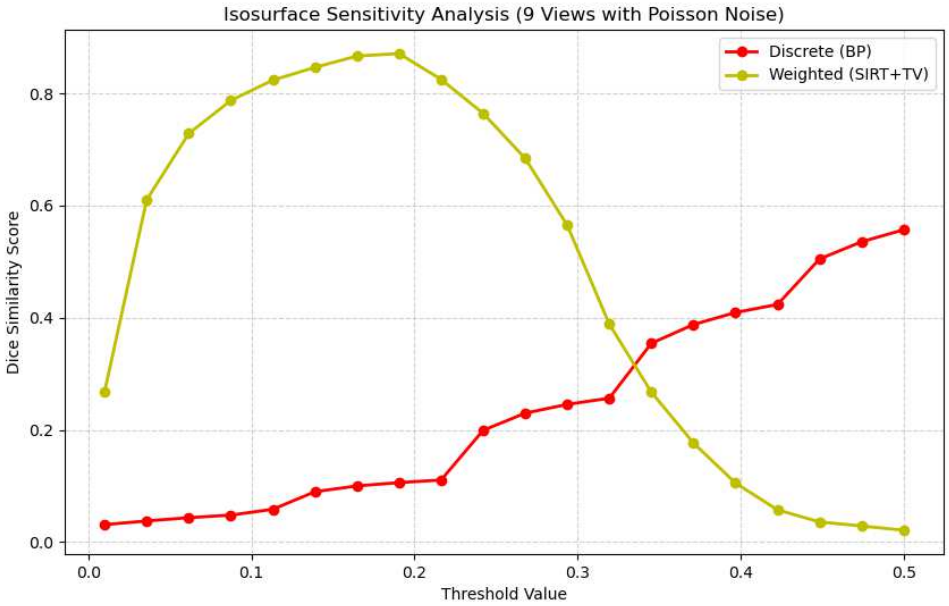


Figure 6. Isosurface Sensitivity Analysis under Poisson Noise. The graph plots the Dice Similarity Coefficient (DSC) against varying segmentation thresholds for a representative anatomical phantom (Model 1).

Quantitative evaluation of the root mean square error (RMSE) across all models confirmed the spatial observations. The binary model yielded an average RMSE of 0.08 ± 0.04, indicating a deviation from the true physical attenuation values due to noise

amplification. Conversely, the weighted model reduced this average error by more than 50% (RMSE: 0.03 ± 0.03). This demonstrates that the weighted projector accurately recovers the physical density of the contrast agent, even in clinical low-dose, sparse-view environments.

Discussion

Quantitative evaluation of the ray-tracing algorithms under simulated clinical noise conditions revealed substantial geometric limitations of the binary approach. The binary (Bresenham) method achieved a mean Jaccard Index of 0.44 ± 0.05 and produced a mean volume ratio of 1.95 ± 0.47 (Table 1), indicating considerable overestimation of vessel volume. Binary ray-tracing constrains each ray to integer voxel coordinates, introducing spatial discretization errors. In noise-free conditions, these errors manifest as morphological distortions, including artificial vessel thickening and structural fragmentation, depending on vessel diameter and orientation.

The performance of the weighted reconstruction framework can be assessed within the broader context of sparse-view tomographic reconstruction. While recent developments increasingly rely on geometric heuristics or data-driven methods that incorporate learned statistical priors, the weighted SIRT+TV framework achieved a mean DSC of 0.80 ± 0.06 and a stable mean volume ratio of 1.05 ± 0.11. These results demonstrate that the combination of a weighted ray-tracing projector and total variation (TV) regularization provides accurate geometric reconstruction under sparse-view and noisy acquisition conditions. Furthermore, unlike deep-learning-based approaches, the proposed framework remains fully interpretable, as reconstructed voxel intensities are directly related to the underlying X-ray attenuation measurements.

The isosurface sensitivity analysis of the representative phantom (Figure 6) further highlights the instability of the binary model. Under noisy conditions, the binary reconstruction failed to produce a distinct high-accuracy region across the evaluated threshold range, indicating a strong dependence on discretization artifacts and noise. In contrast, the weighted Siddon projector combined with TV regularization exhibited a broad stability plateau, indicating that reconstruction accuracy remained stable over a wide range of threshold values. Such robustness is desirable in clinical applications, where reconstruction quality should remain stable without extensive parameter tuning.

Clinical X-ray reconstructions generate continuous grayscale volumes representing attenuation coefficients rather than binary structures. Consequently, the use of binary evaluation metrics such as DSC and Jaccard Index requires threshold-based segmentation and therefore simplifies the partial-volume effects present in clinical data. Nevertheless, this evaluation remains relevant because many downstream applications, including fractional flow reserve estimation and automated vessel tracking, ultimately require a binary representation of the vessel lumen resulting from image segmentation.

To complement the binary evaluation, reconstruction accuracy was also assessed using the root mean square error (RMSE) metric presented in Table 1. The weighted method reduced the mean density error by more than 50% compared with the binary baseline (0.03 ± 0.03 versus 0.08 ± 0.04). This improvement is consistent with the visual reconstruction quality observed in Figure 5. TV Regularization suppresses noise by promoting sparsity in the image-gradient domain, whereas the weighted projector preserves smooth intensity transitions at vessel boundaries through sub-voxel ray-path modeling. Together, these mechanisms reduce quantization artifacts and morphological distortions associated with binary ray-tracing.

A limitation of TV regularization observed in this study is the partial smoothing of high-frequency anatomical details. Because distal coronary branches are extremely narrow structures, the regularization process may suppress these vessels together with high-frequency noise components. Although the weighted method substantially improves vessel continuity compared with the binary baseline, minor thinning or discontinuities may still occur in the most distal vessel segments.

This study employed digital phantoms to provide a controlled ground-truth benchmark. Although the simulations incorporated realistic Poisson noise to emulate low-dose clinical acquisitions, they did not account for cardiac motion or respiratory motion, both of which are present in clinical catheterization procedures. In addition, the evaluation was intentionally limited to an extreme sparse-view configuration consisting of nine projections. This experimental design enabled isolation of the geometric limitations of the projection models and assessment of the regularization requirements of the reconstruction algorithms. Future work will investigate the robustness of the weighted regularized framework in the presence of motion artifacts and uncalibrated *in vivo* C-arm rotational angiography data.

Conclusions

The primary finding of this study is the severe geometric inaccuracy of Bresenham's binary ray-tracing. In sparse-view (9-projection) reconstruction scenarios, the binary projector lacks the spatial precision required to accurately reconstruct

morphological boundaries, particularly for thin vascular structures. As shown in Figure 6, under simulated clinical low-dose conditions, the binary projector exhibited a substantial reduction in accuracy across all threshold values, never exceeding a DSC of 0.55, a level generally insufficient for clinical applications. Consequently, the algorithm proved unstable, artificially enlarging larger structures while failing to preserve finer anatomical details. This fundamental inability to accurately reconstruct vessel geometry resulted in significant geometric dilation, merging of adjacent parallel vessel branches, and fragmentation of thin vessel segments, yielding an average volume ratio of 1.95 ± 0.47 (Figure 5B).

In contrast, Siddon's weighted iterative reconstruction framework combined with total variation (TV) regularization demonstrated substantially greater robustness to clinical noise conditions. The weighted method achieved a mean Jaccard Index of 0.67 ± 0.09 and a mean Dice coefficient of 0.80 ± 0.06, maintaining a broad accuracy plateau that supports its suitability for automated threshold selection even in highly degraded projection data. TV regularization effectively suppressed background noise, while the sub-voxel accuracy of the weighted projector preserved vessel boundaries and structural continuity (Figure 5C). This combination enabled preservation of the tubular morphology and distal branch connectivity that were largely lost in the binary reconstructions. Nevertheless, the weighted method did not completely resolve the reconstruction of the smallest vascular structures. Some thin vessels remained partially disconnected, while closely spaced small vessels were occasionally reconstructed as a single enlarged structure.

The weighted reconstruction framework produced a mean volume ratio of 1.05 ± 0.11, corresponding to a modest geometric overestimation. This effect is consistent with the boundary-smoothing properties of TV regularization and remained substantially smaller than that observed for the binary approach. Overall, the results demonstrate that accurate sparse-view coronary reconstruction under low-dose conditions requires both weighted ray-path modeling and appropriate regularization. Among the evaluated methods, the combination of a weighted projector and TV regularization provided the highest geometric fidelity and the greatest robustness to noise.

Future work will focus on extending the proposed framework to uncalibrated *in vivo* C-arm rotational angiography data, evaluating its performance in the presence of cardiac and respiratory motion, and systematically investigating the relationship between reconstruction accuracy, noise level, and angular sampling density. In addition, further research will be directed toward the development of novel regularization strategies specifically designed to preserve and enhance thin tubular structures while maintaining effective noise suppression.

Data availability

The 3D cardiovascular geometries (Models 0066_H_CORO_H, 0271_H_CORO_H, and 0286_H_CORO_H) used to support the findings of this study are publicly distributed by the Vascular Model Repository (https://www.vascularmodel.com/), curated by the Cardiovascular Biomechanics Computation Lab at Stanford University, and were utilized in accordance with their respective open-access research licenses.

Acknowledgements

The foundational concept for this research was developed during a three-month research mobility at the University of Stavanger, under the supervision of Prof. Mahdieh Khanmohammadi. During this period, initial investigations focused on optimizing 3D reconstruction methodologies using clinical datasets. The authors gratefully acknowledge Prof. Khanmohammadi for her guidance, valuable discussions, and scientific support throughout this work.

Funding

The authors declare that no financial support was received for the research, authorship, and/or publication of this article.

References

1. Lashgari, S. M., et al. (2024). Patient-specific in silico 3D coronary model in cardiac catheterisation laboratories. *Frontiers in Cardiovascular Medicine*, 11, 1398290. https://doi.org/10.3389/fcvm.2024.1398290
2. Ding, D., et al. (2020). Accuracy of 3-dimensional and 2-dimensional quantitative coronary angiography for predicting physiological significance of coronary stenosis: a FAVOR II substudy. *Cardiovascular Diagnosis and Therapy*, 12(9), 5032. https://doi.org/10.21037/cdt.2019.09.07

3. Solanki, R., et al. (2021). The importance of three dimensional coronary artery reconstruction accuracy when computing virtual fractional flow reserve from invasive angiography. *Scientific Reports*, 11(1), 19694. https://doi.org/10.1038/s41598-021-99065-7
4. Arefinia, f., et al. (2024). Non-invasive fractional flow reserve estimation using deep learning on intermediate left anterior descending coronary artery lesion angiography images. *Scientific Reports*, 14(1), 1818. https://doi.org/10.1038/s41598-024-52360-5
5. Chabi, N., et al. (2022). Self-calibration of C-arm imaging system using interventional instruments during an intracranial biplane angiography. *International Journal of Computer Assisted Radiology and Surgery*, 17(7), 1355-1366. https://doi.org/10.1007/s11548-022-02580-9
6. Rashed, E. A., et al. (2015). Sparsity-constrained three-dimensional image reconstruction for C-arm angiography. *Computers in Biology and Medicine*, 60, 141-151. https://doi.org/10.1016/j.compbiomed.2015.04.014
7. Pan, X., Sidky, E. Y., & Vannier, M. (2009). Why do commercial CT scanners use traditional filtered back-projection? *Inverse Problems*, 10(1), 12-36. https://doi.org/10.1088/0266-5611/25/12/123009
8. Bresenham, J. E. (1965). Algorithm for computer control of a digital plotter. *IBM Systems Journal*, 4(1), 25-30. https://doi.org/10.1147/sj.41.0025
9. Siddon, R. L. (1985). Fast calculation of the exact radiological path for a three-dimensional CT array. *Medical Physics*, 12(2), 252-255. https://doi.org/10.1118/1.595715
10. Iyer, K., et=al. (2021). AngioNet: a convolutional neural network for vessel segmentation in X-ray angiography. *Scientific Reports*, 11(1), 18066. https://doi.org/10.1038/s41598-021-97355-8
11. Biguri, A., et al. (2016). TIGRE: a MATLAB-GPU toolbox for CBCT image reconstruction. *Biomedical Physics & Engineering Express*, 2(5), 055010. https://doi.org/10.1088/2057-1976/2/5/055010
12. Jacobs, f., et al. (1998). A fast algorithm to calculate the exact radiological path through a pixel or voxel space. *Journal of Computing and Information Technology*, 6(1), 89-94.
13. van Aarle, W., et al. (2016). Fast and flexible X-ray tomography using the ASTRA toolbox. *Optics Express*, 24(22), 25129-25147. https://doi.org/10.1364/OE.24.025129
14. Galassi, F., et al. (2018). 3D reconstruction of coronary arteries from 2D angiographic projections using non-uniform rational basis splines (NURBS). *PLOS ONE*, 13(1), e0190650. https://doi.org/10.1371/journal.pone.0190650
15. Feuillâtre, H., et al. (2013). Automatic determination of optimal view for the visualization of coronary lesions by rotational X-ray angiography. *IRBM*, 34(1), 29-35. https://doi.org/10.1016/j.irbm.2013.08.004
16. Iyer, K., et al. (2023). A multi-stage neural network approach for coronary 3D reconstruction from uncalibrated X-ray angiography images. *Scientific Reports*, 13(1), 14210. https://doi.org/10.1038/s41598-023-44633-2?urlappend=%3Futm_source%3Dresearchgate.net%26utm_medium%3Darticle
17. Duong, L., et al. (2017). Online C-arm calibration using a marked guide wire for 3D reconstruction of pulmonary arteries. *Proc. SPIE Medical Imaging*, 10135, 101351D. https://doi.org/10.1117/12.2254175
18. Navab, N., et al. (1996). 3D Reconstruction from Projection Matrices in a C-Arm Based 3D-Angiography System. *IEEE Transactions on Medical Imaging*, 15(3), 275-285. https://doi.org/10.1007/BFb0056194
19. Kak, A. C., & Slaney, M. (1988). *Principles of Computerized Tomographic Imaging*. IEEE Press. https://doi.org/10.1137/1.9780898719277
20. Gilbert, P. F. C. (1972). Iterative methods for the three-dimensional reconstruction of an object from projections. *Journal of Theoretical Biology*, 36(1), 105-117. https://doi.org/10.1016/0022-5193(72)90180-4
21. Candès, E. J., Romberg, J., & Tao, T. (2006). Robust uncertainty principles: Exact signal reconstruction from highly incomplete frequency information. *IEEE Transactions on Information Theory*, 52(2), 489-509. https://doi.org/10.1109/TIT.2005.862083
22. Updegrove, A., Wilson, N. M., Merkow, J., Lan, H., Marsden, A. L., & Shadden, S. C. (2017). SimVascular: an open source pipeline for cardiovascular simulation. *Annals of Biomedical Engineering*, 45(3), 525-541. https://doi.org/10.1007/s10439-016-1762-8
23. Sidky, E. Y., & Pan, X. (2008). Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization. *Physics in Medicine & Biology*, 53(17), 4777-4807. https://doi.org/10.1088/0031-9155/53/17/021

24. Rudin, L. I., Osher, S., & Fatemi, E. (1992). Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4), 259-268. https://doi.org/10.1016/0167-2789(92)90242-F

25. De Man, B., & Basu, S. (2004). Distance-driven projection and backprojection in three dimensions. *Physics in Medicine & Biology*, 49(11), 2463-2481. https://doi.org/10.1088/0031-9155/49/11/024

26. Strong, D., & Chan, T. (2003). Edge-preserving and scale-dependent properties of total variation regularization. *Inverse Problems*, 19(6), S165-S187. https://doi.org/10.1088/0266-5611/19/6/059

27. Arridge, S., Maass, P., Öktem, O., & Schönlieb, C. B. (2019). Solving inverse problems using data-driven models. *Acta Numerica*, 28, 1-174. https://doi.org/10.1017/S0962492919000059

28. Maier, A. K., Syben, C., Lasser, T., & Riess, C. (2019). A gentle introduction to deep learning in medical image processing. https://doi.org/10.1016/j.zemedi.2018.12.003

29. Wirgin, A. (2004). The inverse crime. *arXiv preprint math-ph/0401050*. https://doi.org/10.48550/arXiv.math-ph/0401050

30. Szczypiński, P. M., Klepaczko, A., & Olbrycht, R. (2025). VesselKnife–software for the analysis of tubular structures in biomedical images. *Computer Methods and Programs in Biomedicine*, 268, 108817. https://doi.org/10.1016/j.cmpb.2025.108817

31. Szczypiński, P. M. (2023, June). Radius Estimation in Angiograms Using Multiscale Vesselness Function. In *International Conference on Computational Science* (pp. 230-244). Cham: Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-36027-5_17

32. Materka, A., & Jurek, J. (2024). Using deep learning and b-splines to model blood vessel lumen from 3d images. *Sensors*, 24(3), 846. https://doi.org/10.3390/s24030846

33. Materka, A., Jurek, J., Kocinski, M., & Klepaczko, A. (2023, June). CNN-Based Quantification of Blood Vessels Lumen in 3D Images. In *International Conference on Computational Science* (pp. 653-661). Cham: Springer Nature Switzerland. https://doi.org/10.1007/978-3-031-36021-3_62