




Article

NeCA: 3D Coronary Artery Tree Reconstruction from Two 2D Projections via Neural Implicit Representation

Yiying Wang ¹, Abhirup Banerjee ^{1,2,*} and Vicente Grau ^{1,†}

¹ Institute of Biomedical Engineering, Department of Engineering Science, University of Oxford, Oxford OX3 7DQ, UK; yiying.wang@eng.ox.ac.uk (Y.W.); vicente.grau@eng.ox.ac.uk (V.G.)

² Division of Cardiovascular Medicine, Radcliffe Department of Medicine, University of Oxford, Oxford OX3 9DU, UK

* Correspondence: abhirup.banerjee@eng.ox.ac.uk

† These authors contributed equally to this work.

Abstract: Cardiovascular diseases (CVDs) are the most common health threats worldwide. 2D X-ray invasive coronary angiography (ICA) remains the most widely adopted imaging modality for CVD assessment during real-time cardiac interventions. However, it is often difficult for the cardiologists to interpret the 3D geometry of coronary vessels based on 2D planes. Moreover, due to the radiation limit, often only two angiographic projections are acquired, providing limited information of the vessel geometry and necessitating 3D coronary tree reconstruction based only on two ICA projections. In this paper, we propose a self-supervised deep learning method called NeCA, which is based on neural implicit representation using the multiresolution hash encoder and differentiable cone-beam forward projector layer, in order to achieve 3D coronary artery tree reconstruction from two 2D projections. We validate our method using six different metrics on a dataset generated from coronary computed tomography angiography of right coronary artery and left anterior descending artery. The evaluation results demonstrate that our NeCA method, without requiring 3D ground truth for supervision or large datasets for training, achieves promising performance in both vessel topology and branch-connectivity preservation compared to the supervised deep learning model.

Keywords: 3D coronary artery tree reconstruction; invasive coronary angiography; limited-projection reconstruction; neural implicit representation; self-supervised learning; deep learning



Citation: Wang, Y.; Banerjee, A.; Grau, V. NeCA: 3D Coronary Artery Tree Reconstruction from Two 2D Projections via Neural Implicit Representation. *Bioengineering* **2024**, *11*, 1227. <https://doi.org/10.3390/bioengineering11121227>

Academic Editors: Giuseppe Baselli and Riemer H.J.A. Slart

Received: 10 November 2024

Revised: 24 November 2024

Accepted: 25 November 2024

Published: 4 December 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Cardiovascular diseases (CVDs) are the most common cause of death worldwide [1]. X-ray invasive coronary angiography (ICA) remains the most widely adopted imaging modality for CVD assessment during real-time cardiac interventions [2]. ICA acquires 2D projections of the coronary tree, which makes it difficult for cardiologists in clinical practice to understand the global vascular anatomical structure due to vessel overlap and foreshortening. Moreover, potential adverse effects of the higher amount of radiographic contrast agent and higher radiation required for long-time exposure to X-rays restrict the number of angiographic projections acquired; typically, 2–5 projections are acquired, providing limited information of the vessel structures. Therefore, it is of great significance to perform 3D coronary tree reconstruction from only two 2D projections to provide spatial vascular information, which can significantly reduce the risks of subjective interpretation of the 3D coronary vasculature from 2D views and decrease the complexity of interventional surgeries.

Several conventional mathematical methods have been proposed for 3D coronary tree reconstruction from ICA projections [3–6], but they usually depend on traditional stereo-vision algorithms, requiring substantial manual interactions. The emergence and prosperity of deep neural networks have enabled 3D automated reconstruction from limited views in medical images [7–9]. Most of them need large training datasets and work in a supervised

learning manner, but the acquisition of paired data has always been a challenge in real clinics. Recently, Neural Radiance Fields (NeRF) [10] have made a significant contribution to the field of computer vision, allowing for neural implicit representation and novel view synthesis. In neural implicit representation learning, a bounded scene is parameterized by a neural network as a continuous function that maps spatial coordinates to metrics such as occupancy and color. The optimization of NeRF only relies on several images from different viewpoints. Based on NeRF, Neural Attenuation Fields [11] (NAFs) are proposed to tackle the problem of sparse-view cone-beam computed tomography (CT) reconstruction, which require at least 50 projections. Shen et al. [12] proposed a neural implicit representation learning methodology to reconstruct CT images, which performs on 10, 20, and 30 projections.

Few studies have explored deep learning for 3D vessel reconstruction from limited projections. Reconstructing the 3D cerebral vessels using deep learning has received some attention in recent years. A self-supervised learning model [13] was proposed for the 3D reconstruction of cerebral vessels based on ultra-sparse X-ray projections. Zuo [14] implemented an adversarial network for 3D neurovascular reconstruction based on biplane angiograms, but the results are limited, with flaws occurring near crossed vessels. Some deep learning-based studies also attempted 3D coronary tree reconstruction from limited projections. Wang et al. [15] used coronary computed tomography angiography (CCTA) data to simulate projections and trained a weakly supervised adversarial learning model for 3D reconstruction from two projections. However, their model requires large training datasets (8800 data in the experiments), with the 3D ground truth used in the discriminator. Wang et al. [16] also used a large CCTA dataset to simulate projections for training. Ibrahim and GEDİK [17], Uluhan and Gedik [18], Iyer et al. [19] generated 3D synthetic coronary tree data and simulated corresponding 2D projections to train supervised learning models; their models require more than two projections for training. Bransby et al. [20] used biplanar ICA data to reconstruct a single coronary tree branch in a supervised learning setup. Maas et al. [21] proposed a NeRF-based model to achieve 3D coronary tree reconstruction from limited projections without involving 3D ground truth in training. However, they tested the performance only on two 3D studies, and the number of required projections is at least four. Despite the improvement in deep neural networks, 3D coronary tree reconstruction from two projections without involving corresponding 3D ground truth and large training datasets remains challenging.

In this paper, we propose a self-supervised deep learning method named NeCA, which is based on neural implicit representation to achieve 3D coronary artery tree reconstruction from only two projections. Our method neither requires 3D ground truth for supervision nor large training datasets. It iteratively optimizes the reconstruction results in a self-supervised fashion with only the projection data of one subject as input. Our proposed method utilizes the advantages of the multiresolution hash encoder [22] to encode point coordinates, residual multilayer perceptrons (MLP) to predict point occupancy, and a differentiable cone-beam forward projector layer [23] to simulate projections. The simulated projections are then learned from the input projections by minimizing the projection error in a self-supervised manner. Our method aims to learn and optimize the neural representation for the entire image and can directly reconstruct the target image by incorporating the forward model of the imaging system. We use a public CCTA dataset [24] to validate our model's feasibility on the task based on six metrics. The evaluation results indicate that our proposed NeCA model, without 3D ground truth for supervision or large datasets for training, achieves promising performance in both vessel topology preservation and maintaining branch connectivity compared to an equivalent supervised learning model. The code of our work is available at <https://github.com/WangStephen/NeCA>, 10 November 2024. The main contributions of this work are:

1. **3D coronary tree reconstruction using self-supervised learning from only two projections:** Our proposed deep learning method achieves 3D coronary artery tree

reconstruction from two projections where neither 3D ground truth for supervision nor large training datasets are required.

2. **Neural implicit representation learning:** We leverage the advantages of MLP neural networks as a continuous function to represent the coronary tree in 3D space in order to enable mapping from encoded coordinates to corresponding occupancies.
3. **The applications of multiresolution hash encoder and differentiable cone-beam forward projector layer:** We combine a learnable hash encoder and a differentiable projector layer in our model to allow for efficient feature encoding and self-supervised learning from 2D input projections.
4. **Evaluations:** We perform thorough evaluation of our model on the right coronary artery and left anterior descending artery in terms of six quantitative metrics.

2. Materials and Methods

2.1. Dataset

We use a public CCTA dataset [24] containing binary segmented coronary trees for our study, splitting the coronary trees into the right coronary artery (RCA) and left anterior descending (LAD) artery. Since our model is an optimization-based method for each individual data point, we do not need training/validation split. We use 67 RCA data and 79 LAD data points as the test set. We perform cone-beam forward projections on the CCTA data to generate the input projections with simulated attenuated X-ray intensities based on the Operator Discretization Library (ODL) [23]. For each CCTA data point, we generate only two projections to use in our model for 3D coronary tree reconstruction. The projection geometries for RCA and LAD are illustrated in Table 1, which mimic the ones generally used in clinics. Figure 1 illustrates an example of two projections generated from both RCA and LAD.

Table 1. The projection geometry to simulate cone-beam forward projections for both RCA and LAD. DSD: distance for source to detector; DSO: distance for source to origin.

Data	Geometry	First Projection Plane	Second Projection Plane
RCA and LAD	Detector spacing	$0.2769 \times 0.2769 \text{ mm}^2$ to $0.2789 \times 0.2789 \text{ mm}^2$	
	Detector size	512 × 512	
	Volume spacing	$90 \times 90 \times 90 \text{ mm}^3$ to $105 \times 105 \times 105 \text{ mm}^3$	
	Volume size	128 × 128 × 128	
RCA	DSD	970 mm to 1010 mm	1050 mm to 1070 mm
	DSO	745 mm to 785 mm	±3 mm to the 1st projection
	Primary angle	18° to 42°	−8° to 8°
	Secondary angle	−8° to 8°	18° to 42°
LAD	DSD	1030 mm to 1090 mm	+70 mm to the 1st projection
	DSO	740 mm to 760 mm	+3 mm to the 1st projection
	Primary angle	−8° to 8°	−47° to −23°
	Secondary angle	18° to 42°	21° to 45°




Figure 1. An example of two projections generated from RCA and LAD data.

2.2. Proposed Model

Our proposed model NeCA consists of five stages and allows for end-to-end learning. First, we normalize the coordinate index in the image spatial field according to resolution. Then, for each voxel point, we use a multiresolution hash encoder [22] to encode their normalized coordinates to obtain the corresponding multiresolution spatial feature vectors. These feature vectors are next sent to the residual MLP to predict the occupancy at the position of that point. The occupancy predictions of all the points form the 3D coronary tree reconstruction results. After that, we simulate the X-ray forward projections from the 3D predicted reconstruction based on the projection geometry of the input. Finally, these simulated projections are learned iteratively against the input projections in a self-supervised way. Stages 2 to 5 of our proposed model are illustrated in Figure 2.

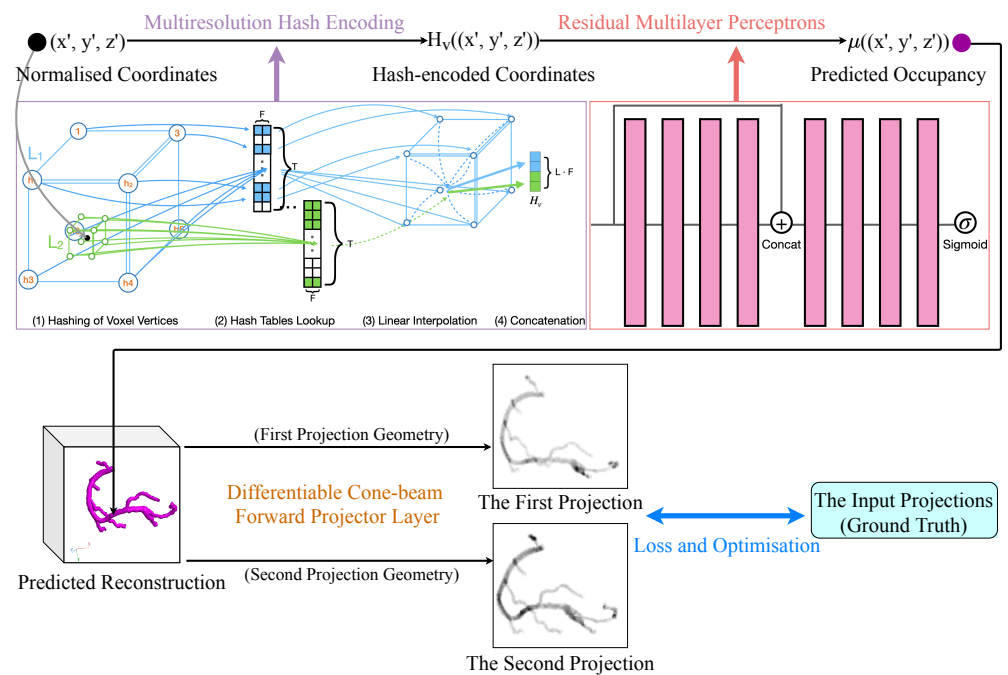


Figure 2. The proposed NeCA model (stages 2–5). The multiresolution hash encoder illustrates an example of 2 resolution levels (coloured in green and blue) from fine to coarse resolution for one sampled point (in black).

2.2.1. Coordinate Normalization

The input to the model is a set of integer coordinates $\mathbf{x} = (x, y, z)$ based on the number of voxels $n_{vx} \times n_{vy} \times n_{vz}$ in 3D volume ranging in $(1 \text{ to } n_{vx}, 1 \text{ to } n_{vy}, 1 \text{ to } n_{vz})$. We normalize the coordinates from these voxels according to the voxel spacing $s_{vx,vy,vz}$ along each axis, as calculated in Equation (1). These normalized coordinates $\mathbf{x}' = (x', y', z')$ are then sent to a multiresolution hash encoder at the next stage to efficiently obtain the corresponding spatial feature vectors.

$$n'_{x,y,z} = \frac{n_{vx,vy,vz} \times s_{vx,vy,vz} - s_{vx,vy,vz}}{2}, \tag{1}$$

$$\mathbf{x}' = \text{Norm}((x, y, z)) = (-n'_x + (x - 1) \times s_{vx}, -n'_y + (y - 1) \times s_{vy}, -n'_z + (z - 1) \times s_{vz}).$$

2.2.2. Multiresolution Hash Encoding

We use the multiresolution hash encoder [22] $H_v = \text{enc}(\mathbf{x}'; \Theta)$ to encode the normalized positions of sampled points, which enables fast encoding without sacrificing performance. With the multiresolution structure, it allows the encoder to disambiguate hash collisions. The multiple resolutions are arranged into L levels with different T -dimensional learnable hash tables at each level containing feature vectors with size F . The hyperparameters of our multiresolution hash encoder are shown in Table 2, and the structure of the encoder is illustrated in Figure 2.

Table 2. The hyperparameters for the multiresolution hash encoder used in our work.

Parameter	Symbol	Value
Number of levels	L	16
Maximum entries per level (hash table size)	T	2^{19}
Number of feature dimensions per entry	F	2
Coarsest resolution	N_{min}	16
Resolution growth factor	b	2
Input dimension	d	3

For each voxel, we apply L resolution levels, which are independent of each other. The resolution size N is chosen based on an exponential increment between the coarsest and finest resolutions $[N_{min}, N_{max}]$, where N_{max} is selected to match the finest detail in the training data. It is defined as:

$$N_l := \lfloor N_{min} * b^l \rfloor, \tag{2}$$

where $l \in \{0, 1, \dots, L - 1\}$, and $b = 2$ is the growth factor. For a single level N_l , the input point with normalized coordinates $\mathbf{x}' = (x', y', z') \in \mathbb{R}^3$ is geometrically scaled to a grid cube containing 2^3 vertices according to the grid resolution at this level. To implement this functionality, the original 3D volume is evenly split into a number of grid cubes according to the resolution N_l^3 , and the grid cube containing the desired sampled point is assigned to this point as the spanned grid cube. The multiresolution property in the hash encoder covers the full range from the coarsest resolution N_{min} to the finest resolution N_{max} , which ensures that all scales are contained, in spite of sparsity. The four parts of the multiresolution hash encoder are discussed in detail below.

Hashing of Voxel Vertices

For all normalized voxels after scaling at resolution level N_l , we have $(N_l + 1)^d$ vertices in total. For coarse levels when $(N_l + 1)^d \leq T$, we have one-to-one mapping from all the vertices at this resolution level N_l to hash table entries, so there is no collision. Regarding finer levels when $(N_l + 1)^d > T$, we use a hash function h to index into the feature vector array, effectively treating it as a hash table. In this case, we do not explicitly tackle hash collisions, but instead we rely on gradient-based optimization in the backpropagation of the subsequent residual MLP to automatically handle them. For instance, if two voxels have the same hash value on one or more vertices, the voxel closer to the desired object which our model is more focused on tends to have larger gradients during optimization, so this voxel takes the domination to update the collided feature vector entry. In this way, the collision issue is handled implicitly.

We assign indices to these vertices by hashing their coordinates. The spatial hash function [25] h is defined in the following form:

$$h(\mathbf{x}') = (\oplus_{i=1,2,3} \mathbf{x}'_i \pi_i) \bmod T \quad (3)$$

where \mathbf{x}' is the input point, $\mathbf{x}'_{i=1,2,3}$ are the corresponding spatial normalized coordinate values, \oplus denotes the bit-wise XOR operation, π_i are unique large primary numbers, and T is the hash table size.

Hash Tables Lookup

We now have the hash value for each vertex at each resolution level of each point. We then maintain an individual learnable hash table, which contains T numbers of F -dimensional feature vectors for each resolution level. For the hash values on all the vertices of each resolution level, we look up the corresponding entries in the level's respective feature vector array, i.e., the hash table. Next, the previously assigned indices on the vertices are replaced by the corresponding lookup feature vectors, so each resolution level conceptually stores feature vectors at the vertices of a grid cube. The hash tables at different resolution levels are the only trainable parameters Θ in the multiresolution hash encoder, and the size of these parameters is $L \times T \times F$.

Linear Interpolation

For each resolution level, we linearly interpolate the feature vectors on the vertices according to their relative positions to the sampled point within this resolution level cube. Interpolating the queried hash table entries guarantees the encoded feature vectors with the later residual MLP are continuous during network training. After interpolation, the final feature vectors with the dimension F for the sampled voxel at this resolution level are produced.

Concatenation

We concatenate the interpolated feature vectors for each resolution level to generate the final multiresolution hash encoding feature vectors $H_v \in \mathbb{R}^{L \times F}$ for the sampled point, which can then be utilized to predict the occupancy of coronary tree for this point position by the residual MLP at the next stage. The dimension $L \cdot F$ for the final encoded feature vectors of each voxel is regarded as the channel dimension for later residual MLP training.

2.2.3. Residual MLP

We exploit residual [26] MLP $m(H_v; \Phi)$ to predict the occupancy value μ from the position-encoded feature vectors H_v of each point, where Φ is the trainable weight parameters of the residual MLP. The residual MLP network serves as a continuous function to implicitly parameterize a bounded scene, i.e., the 3D coronary tree in our case, which maps spatial coordinate features to the predicted occupancy values. This, in fact, encodes the internal information of an entire 3D coronary tree into the network parameters.

The residual MLP contains eight fully connected layers, as depicted in Figure 2. We apply residual learning in the middle layer to preserve the original feature information. The residual MLP receives the feature vectors as input with $L \cdot F$ -dimensional channels and produces predicted occupancy values with a 1-dimensional channel. The feature dimensions for all the hidden layers are 256-wide. Except for the last layer followed by a sigmoid activation, all the layer outputs are followed by LeakyReLU activation [27].

2.2.4. Differentiable Forward Projector Layer

At this stage, we have all the predicted occupancy values for all the voxels, which construct the 3D coronary tree reconstruction results. After that, we simulate the X-ray cone-beam forward projections from the 3D reconstruction results based on the same projection geometry as the input projections to generate two predicted projections. The forward projection simulation is based on the theory that the intensity of an X-ray beam is

reduced by the exponential integration of attenuation coefficients along the ray path. We use ODL [23] to implement this differentiable X-ray forward projector layer that enables self-supervised loss optimization at the final stage.

2.2.5. Loss

We use Mean Square Error (MSE) loss to calculate the differences between the input projections and simulated forward projections. The loss function \mathcal{L} is defined as follows:

$$\mathcal{L}(\Theta, \Phi) = \frac{1}{N \times I} \sum_n^N \sum_i^I (P_{ni} - G_{ni})^2 \quad (4)$$

where N ($= 2$ in our work) is the number of projections, I ($= 512$ in our work) is the number of pixels in one projection, P is the simulated projection, and G is the corresponding input projection.

The loss function is used to learn the multiresolution hash tables Θ and the residual MLP Φ during training. With this, the 3D occupancy predictions are improved iteratively based on the optimization of 2D projection errors. After training, the final 3D coronary tree can be rendered with the predicted occupancy values, after binarisation with 0.5, by querying all the voxels with their coordinates from the model.

2.3. Training Setup

We implement our proposed model using PyTorch [28] and choose the Adam optimizer [29] with a learning rate of 10^{-4} . The number of epochs for optimization is 5000. The learning was performed on an HPC cluster utilizing Nvidia Tesla v100 GPUs. The package versions we used for NeCA are Python 3.8.17, PyTorch 1.9.0, and ODL 1.0.0.dev0.

2.4. Baseline Model

We use the supervised learning model 3D U-Net [30] as our baseline model. We follow the original 3D U-Net architecture with three sampling levels and a bottleneck layer using the same number of convolutional filters. The channel size for both the input and output to 3D U-Net model in our work is 1. The input to 3D U-Net is an ill-posed volume reconstructed from two clinical-angle projections of the 3D coronary tree by a conventional back-projection method, and the output is the 3D coronary tree reconstruction result. We train two 3D U-Net models based on the CCTA dataset [24] using 669 RCA data points and 788 LAD data points, respectively, where we split them into 75% training, 15% validation, and 10% test data. The test datasets here are the same datasets used for testing our proposed model.

We implement the 3D U-Net baseline model using PyTorch [28] and choose the Adam optimizer [29] with an initial learning rate of 10^{-4} . A learning rate decay policy is used, where the learning rate is decayed by 0.1 if no improvement is observed after 10 epochs. We use an early stopping strategy to avoid overfitting when there is no more improvement after 15 epochs. The training was performed with a batch size of 3 on an HPC cluster utilizing Nvidia Tesla v100 GPUs. The models are trained with MSE loss.

2.5. Evaluation Metrics

We employ six metrics for evaluation between the 3D coronary tree reconstruction results and the original CCTA data (ground truth): centerline Dice score (termed as *clDice*) [31], Dice score (termed as *Dice*), intersection over union (termed as *IoU*), reconstruction error (termed as *reError*) [32], Chamfer ℓ_2 distance (termed as CD_{ℓ_2}), and reconstruction MSE (termed as *reMSE*). *clDice* $\in [0, 1]$ where a larger value suggests a better performance in vessel topology preservation. *Dice* ($\in [0, 1]$) and *IoU* ($\in [0, 1]$) also suggest a better performance if measurement values are bigger. In terms of *reError*, CD_{ℓ_2} , and *reMSE*, a smaller value represents a better reconstruction result. Before evaluation, we apply connected

component analysis [33] on our reconstructed coronary tree to remove sparse disconnected objects with less than 25 voxels.

3. Results

We perform both quantitative and qualitative evaluations on both RCA and LAD datasets. Apart from the clinical-angle projections simulated according to Table 1, we additionally test 3D reconstructions based on two orthogonal views using our NeCA model for comparison (termed as NeCA (90°)).

3.1. Quantitative Results

We quantitatively evaluate our NeCA model, NeCA (90°), and supervised 3D U-Net model on 67 RCA test data points and 79 LAD test data points.

3.1.1. RCA Dataset

Performance over Six Metrics

We evaluate NeCA, NeCA (90°), and the 3D supervised U-Net model in terms of six metrics, namely *clDice*, *Dice*, *IoU*, *reError*, CD_{ℓ_2} , and *reMSE*. The quantitative results are presented in Table 3.

Table 3. The quantitative evaluation results of NeCA, NeCA (90°), and supervised 3D U-Net model on 67 RCA test data in terms of six metrics. The best results of each metric are in **bold**.

Model	<i>clDice</i> (%)	<i>Dice</i> (%)	<i>IoU</i> (%)	<i>reError</i>	CD_{ℓ_2} (mm)	<i>reMSE</i> (1×10^{-4})
NeCA	87.01 ± 9.93	90.43 ± 7.46	83.29 ± 11.42	0.139 ± 0.101	0.27 ± 0.37	2.74 ± 2.14
NeCA (90°)	89.07 ± 8.33	91.03 ± 6.93	84.17 ± 10.25	0.111 ± 0.087	0.22 ± 0.26	2.73 ± 2.60
3D U-Net	95.34 ± 4.16	85.18 ± 4.22	74.42 ± 6.24	0.188 ± 0.054	0.31 ± 0.16	4.63 ± 2.91

All values represent mean ± standard deviation.

From the results presented in Table 3, we can observe that our NeCA model performs better than 3D U-Net model, with relative improvements of 6.16%, 11.92%, 26.06%, 12.90%, and 40.82% in terms of *Dice*, *IoU*, *reError*, CD_{ℓ_2} , and *reMSE* metrics, respectively. 3D U-Net model is better than our NeCA model based on the *clDice* metric, with a respective improvement of 9.57%. 3D reconstruction from two orthogonal projections by our NeCA model produces the best performance in all metrics compared to our NeCA model using two clinical-angle projections. 3D U-Net model maintains the smallest standard deviations among all metrics except for *reMSE*, where our NeCA model performs the best.

Statistical Analysis

The choice of the statistical test is very important, as different tests can have different conclusions for the same evaluation. For this reason and the nature of deep learning in our work, we use the Almost Stochastic Order (ASO) test [34,35] as implemented by [36] specifically for deep learning models to compare score distributions from different models, with a significance level α . ASO returns a confidence score ϵ_{\min} , which indicates (an upper bound to) the amount of violation of stochastic order. In terms of analysis between model A and B using ASO, if $\epsilon_{\min} < \tau$ (where the rejection threshold τ is 0.5 or less), model A is said to be stochastically dominant over model B in more cases, and model A is considered superior. The lower ϵ_{\min} is, the more confidently we can conclude that model A outperforms model B. The tests from [36] show that $\tau = 0.2$ is the most effective threshold value that has a satisfactory tradeoff between Type I and Type II errors across different scenarios. Please note for metrics such as errors where a smaller value expresses a better performance, the final confidence score ϵ_{\min} should be 1 minus the returned ϵ_{\min} from ASO.

With regard to statistical significance test in our work using ASO, we choose a significance level $\alpha = 0.05$ and $\tau = 0.2$. The confidence scores for all six metrics between our NeCA model and 3D U-Net model using the ASO testing on the RCA test dataset is demonstrated in Table 4.

Table 4. The confidence scores ϵ_{\min} for six metrics between our NeCA model and 3D U-Net model using the ASO testing with a significance level $\alpha = 0.05$ on the RCA test dataset. The confidence scores where our NeCA model is found to be stochastically dominant over 3D U-Net are in **bold**, i.e., $\epsilon_{\min} < \tau = 0.2$.

	<i>clDice</i>	<i>Dice</i>	<i>IoU</i>	<i>reError</i>	CD_{ℓ_2}	<i>reMSE</i>
ϵ_{\min}	0.982350	0.198873	0.127973	0.0	0.287172	0

From Table 4, we can find that the score distributions of our NeCA model in terms of *Dice*, *IoU*, *reError*, and *reMSE* are stochastically dominant over the 3D U-Net model. Regarding the metric CD_{ℓ_2} , according to threshold $\tau = 0.2$, our NeCA model is better but not stochastically dominant over 3D U-Net. For the *clDice*, the 3D U-Net model is found to be stochastically dominant over our NeCA model.

Optimizing the Performance of our NeCA Model over Iterations

Our NeCA model is optimized for each individual data point, and we record the quantitative evaluation results of different metrics every 100 iterations. Here, we use two RCA example data points to show how the performance improves iteratively using our NeCA model with clinical-angle projections, as illustrated in Figure 3.




Figure 3. The results of all six metrics every 100 iterations for two RCA example data points (R_1 and R_2) using our NeCA model with two clinical-angle projections.

We can see from Figure 3 that the performance starts to improve after 2000 iterations. We can also find that it usually takes less than 2000 iterations to reach good results after the improvement starts.

3.1.2. LAD Dataset

Performance over Six Metrics

We perform the quantitative evaluations on the LAD test dataset the same as for the RCA data, as described in Section 3.1.1. The results are presented in Table 5.

Table 5. The quantitative evaluation results of NeCA, NeCA (90°), and 3D U-Net model on 79 LAD test data points in terms of 6 metrics. The best results of each metric are in **bold**.

Model	<i>clDice</i> (%)	<i>Dice</i> (%)	<i>IoU</i> (%)	<i>reError</i>	CD_{ℓ_2} (mm)	<i>reMSE</i> (1×10^{-4})
NeCA	76.08 ± 10.42	77.48 ± 9.93	64.28 ± 13.00	0.322 ± 0.129	0.75 ± 0.49	7.28 ± 3.61
NeCA (90°)	91.69 ± 5.62	94.27 ± 3.91	89.41 ± 6.70	0.077 ± 0.051	0.17 ± 0.18	2.26 ± 1.89
3D U-Net	83.36 ± 7.50	68.54 ± 6.87	52.54 ± 7.91	0.415 ± 0.081	0.99 ± 0.51	10.38 ± 4.22

All values represent mean ± standard deviation.

In Table 5, in contrast to the 3D U-Net model, our NeCA model shows improvements of 13.04%, 22.34%, 22.41%, 24.24%, and 29.87% in terms of *Dice*, *IoU*, *reError*, CD_{ℓ_2} , and *reMSE*, respectively. The 3D U-Net model is 9.57% better than our NeCA model with respect to *clDice*. Our NeCA model with two orthogonal projections as input maintains the best performance among all six metrics compared to both our NeCA model with clinical-angle projections and the 3D U-Net model. Furthermore, our NeCA model with two orthogonal projections as input has the smallest standard deviations among all six metrics compared to both the 3D U-Net model and NeCA with clinical-angle projections.

Statistical Analysis

For the statistical significance analysis on the LAD test dataset, we use the ASO test, as described in Section 3.1.1, where we choose a significance level of $\alpha = 0.05$ and $\tau = 0.2$. The confidence scores in terms of all six metrics between our NeCA model and the 3D U-Net model are presented in Table 6.

Table 6. The confidence scores ϵ_{\min} for six metrics between our NeCA model and the 3D U-Net model on the LAD test dataset using ASO testing with a significance level of $\alpha = 0.05$. The confidence scores where our NeCA model is tested to be stochastically dominant over 3D U-Net are in **bold**, i.e., $\epsilon_{\min} < \tau = 0.2$.

	<i>clDice</i>	<i>Dice</i>	<i>IoU</i>	<i>reError</i>	CD_{ℓ_2}	<i>reMSE</i>
ϵ_{\min}	0.992092	0.010340	0.005389	0	0	0

Table 6 demonstrates that our NeCA model evidently outperforms the 3D U-Net model in terms of five metrics, namely *Dice*, *IoU*, *reError*, CD_{ℓ_2} , and *reMSE*. In terms of the *clDice* metric, the 3D U-Net model is stochastically dominant over the NeCA model.

Optimizing the Performance of our NeCA Model Over Iterations

We record the quantitative evaluation results of different metrics every 100 iterations for each individual data point our NeCA model optimizes for. Here, we report two LAD example data points to demonstrate how the NeCA model’s performance improves iteratively, as illustrated in Figure 4.

From Figure 4, we can see that the performance does not start to improve until at least 2000 iterations, and it often takes about 2000 iterations to reach satisfactory performance after the improvement starts. The same phenomenon is also observed for the RCA dataset in Section 3.1.1.

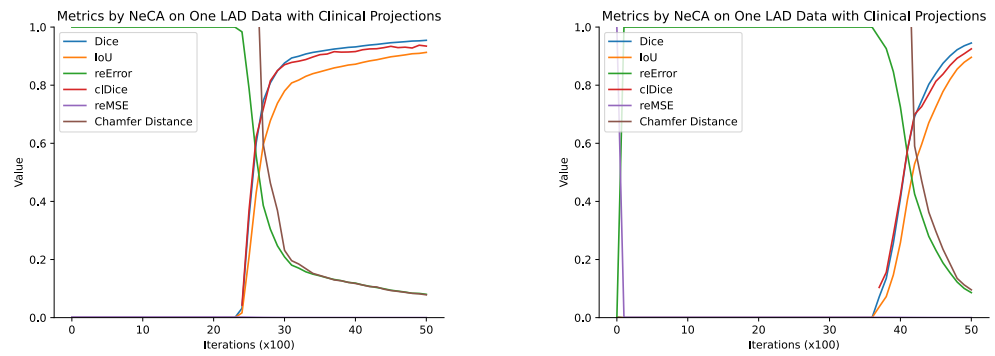


Figure 4. The quantitative results of our NeCA model over two LAD example data points (L_2 and L_3) every 100 iterations with respect to all 6 metrics and evaluated with 2 clinical-angle projections.

3.2. Qualitative Results

We present the qualitative results of 3D coronary artery tree reconstruction based on our NeCA model, NeCA (90°), and the 3D U-Net model on both the RCA and LAD test datasets. Here, we use five example data points for each dataset.

3.2.1. RCA Dataset

3D Reconstruction Results

Figure 5 illustrates five RCA examples of 3D coronary tree reconstruction using our NeCA model, NeCA (90°), and 3D U-Net model, along with the corresponding ground truth for each case. The results show that all three models can successfully perform satisfactory 3D RCA reconstruction.

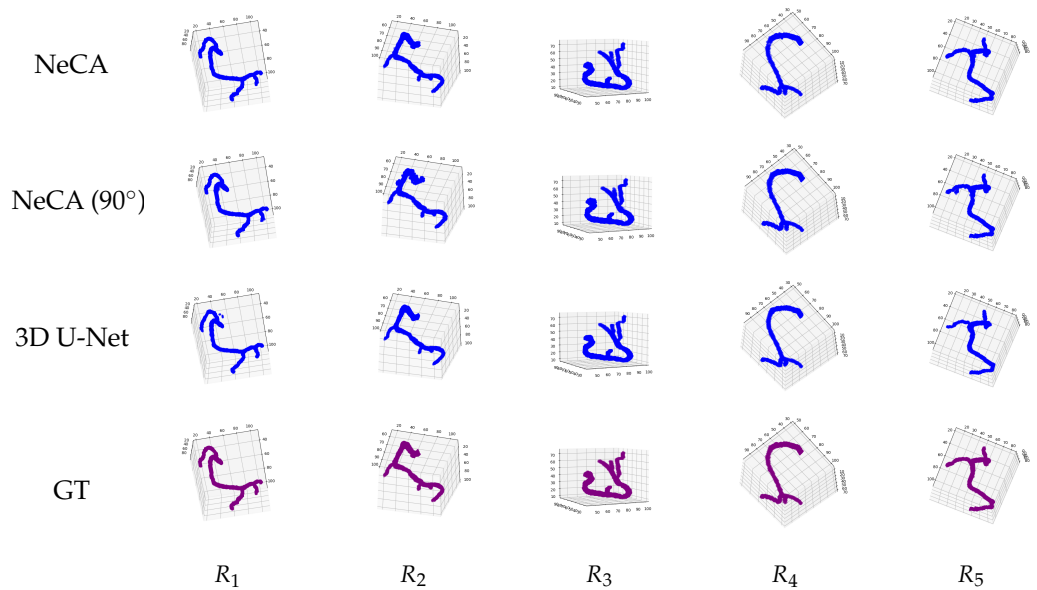


Figure 5. Five qualitative results of 3D RCA reconstruction. From left to right: five RCA data points $R_{1,2,3,4,5}$. From top to bottom: the reconstruction results from our NeCA model, NeCA (90°), 3D U-Net model, and the corresponding ground truth (GT).

Comparison Between 3D Reconstruction and Ground Truth

We additionally compare the 3D RCA reconstruction results using the NeCA, NeCA (90°), and 3D U-Net model with the corresponding ground truth in the same 3D space, as illustrated in Figure 6. These figures show that our NeCA model demonstrates better reconstruction overlap than the 3D U-Net model.

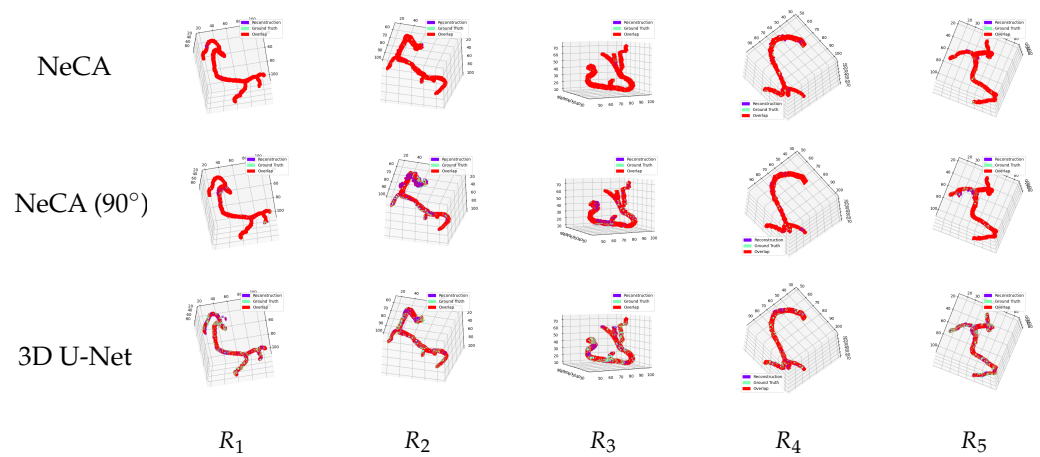


Figure 6. Five 3D RCA reconstruction results compared with the corresponding ground truth in the same 3D space. From left to right: five RCA data points $R_{1,2,3,4,5}$. From top to bottom: the comparison results from our NeCA model, NeCA (90°), and 3D U-Net model. The purple color represents the reconstruction results; green represents the ground truth; and red shows the overlap between them.

3.2.2. LAD Dataset

3D Reconstruction Results

We show in Figure 7 five 3D LAD reconstruction results using our NeCA model, NeCA (90°), and the 3D U-Net model, with the corresponding ground truth. From the results, we can observe that our NeCA model successfully reconstructs the vasculature of LAD in all five cases. On the other hand, the 3D U-Net model fails to reconstruct some branches in $L_{2,4,5}$ and loses vessel connectivity, as presented in red boxes.

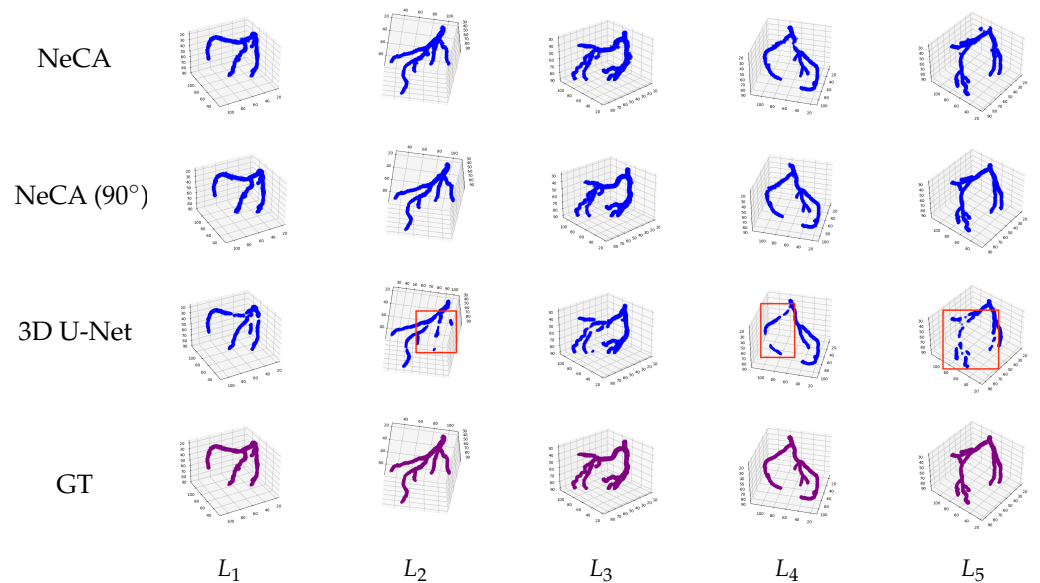


Figure 7. Five qualitative 3D LAD reconstruction results. From left to right: five LAD data points $L_{1,2,3,4,5}$. From top to bottom: the reconstruction results using our NeCA model, NeCA (90°), and 3D U-Net model, along with the corresponding ground truth.

Comparison Between 3D Reconstruction and Ground Truth

We also compare in Figure 8 the five 3D LAD reconstruction results using NeCA, NeCA (90°), and the 3D U-Net models with the corresponding ground truth in the same 3D space. The results show a similar performance to the RCA dataset; our NeCA model demonstrates better reconstruction overlap than the 3D U-Net model.

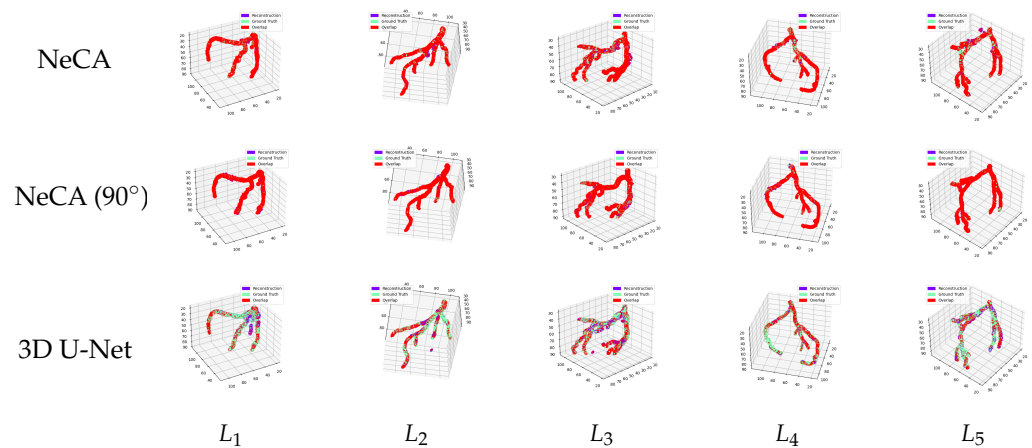


Figure 8. Five 3D LAD reconstruction results compared with the corresponding ground truth in the same 3D space. From left to right: five LAD data points $L_{1,2,3,4,5}$. From top to bottom: the comparison results from our NeCA model, NeCA (90°), and 3D U-Net model. The purple color represents the reconstruction results; green represents the ground truth; and red shows the overlap between them.

4. Discussions and Conclusions

Our evaluation on both the RCA and LAD datasets demonstrates that the NeCA model performs better than the supervised 3D U-Net model in terms of five metrics: *Dice*, *IoU*, *reError*, CD_{ℓ_2} , and *reMSE*. The NeCA model performs statistically significantly better than 3D U-Net model in four metrics for the RCA dataset and five metrics for the LAD dataset out of a total of six metrics. This indicates that our self-supervised learning model, where neither 3D ground truth for supervision nor large training datasets are required, is better than the supervised 3D U-Net model in 3D coronary tree reconstruction from only two projections. It is also demonstrated qualitatively in Section 3.2 that our NeCA model presents good vasculature reconstruction. In addition, due to the intrinsic properties of our model, we do not need to train two models for RCA and LAD separately, and as a result, it has significant potential to generalize to other tasks.

Our model optimized with two orthogonal projections (NeCA (90°)) shows consistently better performance than our model with two clinical-angle projections (Table 5), since two orthogonal projections usually contain more feature coverage and less overlapped redundant information (Figure 1). However, in real clinics such as cardiac catheterization laboratories, projections are generally not acquired at orthogonal views, thus necessitating this feature of our NeCA model.

Our NeCA model contains two trainable components: the hash tables with feature vectors Θ from the multiresolution hash encoder and network parameters Φ from the residual MLP. The residual MLP is the backbone of the neural implicit representation, so it cannot be replaced. For the multiresolution hash encoder to encode the coordinates, there are alternative encoders available, such as a frequency encoder, which is not learnable. We have tested the coordinate encoder where we have replaced our multiresolution encoder with a frequency encoder and used the same projection geometry for validation. According to our experiments, the model could not reconstruct any vessels for every case of the RCA and LAD datasets under 5000 iterations.

The supervised 3D U-Net model, once trained, can perform real-time 3D coronary tree reconstruction, while our model takes around one hour to optimize the results with a volume size of $128 \times 128 \times 128$ for 5000 iterations. We have also tested our model to optimize a coronary tree of size $64 \times 64 \times 64$, which takes on average of 11 minutes for reconstruction. Therefore, there is a tradeoff between lower reconstruction time and better reconstruction resolution for our NeCA model. The 3D U-Net model applies a pre-trained model during evaluation, so when reconstructing out-of-distribution data, it may fail to generalize, which is a serious threat during clinical applications, whereas our model is optimized for each individual data points and can generalize well. Hence, there is also a

tradeoff between real-time reconstruction and stable performance between the 3D U-Net and our NeCA model.

The input cone-beam projections to our NeCA model are based on simulation of X-ray intensity attenuation through the object, i.e., the 3D coronary tree. In our experiments using 3D segmented CCTA data, the attenuation coefficients for the coronary tree are assumed to be uniform as a value of 1. However, in real scenarios, the actual coefficients vary, usually within a certain range due to different vessel conditions. Moreover, blood and contrast injected in the vessel contribute to the X-ray attenuation as well as the other tissues and organs in the background. Though the background removal could be solved with automated coronary vessels segmentation [37,38], the 3D coronary tree reconstruction based on real X-ray projections with contrast injected and different vessel conditions needs to be explored further.

In summary, we have proposed a self-supervised deep learning method—NeCA—using neural implicit representation to achieve 3D coronary artery tree reconstruction from only two projections. Our method neither requires 3D ground truth for supervision nor large training datasets and optimizes the reconstruction results in an iterative self-supervised fashion with only the projection data of one patient as input. We leverage the advantages of a learnable multiresolution hash encoder [22] to allow for efficient feature encoding, residual MLP neural networks as a continuous function to represent the coronary tree in 3D space, and a differentiable projector layer [23] to enable self-supervised learning from 2D input projections. We use a public CCTA dataset [24] containing both RCA and LAD data to validate our model's feasibility on the task based on six quantitative metrics, and we perform a thorough evaluation. The results demonstrate that our proposed NeCA model achieves promising performance in both vessel topology preservation and branch-connectivity maintenance compared to the supervised 3D U-Net model. Our proposed model also has a high possibility to generalize to other clinical tasks where the ground truth is usually unavailable and hard to acquire.

Author Contributions: Conceptualization: Y.W., A.B., and V.G.; methodology, software, validation, formal analysis, investigation, data curation, visualization: Y.W.; resources: Y.W., A.B., and V.G.; writing—original draft preparation: Y.W.; writing—review and editing: Y.W., A.B., and V.G.; supervision: A.B. and V.G.; project administration: A.B. and V.G. All authors have read and agreed to the published version of the manuscript.

Funding: Abhirup Banerjee is supported by the Royal Society University Research Fellowship (Grant No. URF\R1\221314).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The CCTA dataset used in this study is available at <https://github.com/XiaoweiXu/ImageCAS-A-Large-Scale-Dataset-and-Benchmark-for-Coronary-Artery-Segmentation-based-on-CT> (accessed date: 3 December 2024).

Conflicts of Interest: The authors declare no conflicts of interest.

References

1. World Health Organization. *Cardiovascular Diseases (CVDs)*; WHO: Geneva, Switzerland, 2021.
2. Lashgari, M.; Choudhury, R.P.; Banerjee, A. Patient-Specific *In Silico* 3D Coronary Model in Cardiac Catheterisation Laboratories. *Front. Cardiovasc. Med.* **2024**, *11*, 1398290. [[CrossRef](#)] [[PubMed](#)]
3. Çimen, S.; Gooya, A.; Grass, M.; Frangi, A.F. Reconstruction of coronary arteries from X-ray angiography: A review. *Med Image Anal.* **2016**, *32*, 46–68. [[CrossRef](#)] [[PubMed](#)]
4. Banerjee, A.; Galassi, F.; Zacur, E.; De Maria, G.L.; Choudhury, R.P.; Grau, V. Point-Cloud Method for Automated 3D Coronary Tree Reconstruction From Multiple Non-Simultaneous Angiographic Projections. *IEEE Trans. Med Imaging* **2020**, *39*, 1278–1290. [[CrossRef](#)] [[PubMed](#)]
5. Banerjee, A.; Kharbanda, R.K.; Choudhury, R.P.; Grau, V. Automated Motion Correction and 3D Vessel Centerlines Reconstruction from Non-simultaneous Angiographic Projections. In Proceedings of the Statistical Atlases and Computational Models of the Heart. Atrial Segmentation and LV Quantification Challenges, Cham, Granada, Spain, 14 February 2019; pp. 12–20.

6. Banerjee, A.; Choudhury, R.P.; Grau, V. Optimized Rigid Motion Correction from Multiple Non-simultaneous X-ray Angiographic Projections. In Proceedings of the Pattern Recognition and Machine Intelligence, Cham, Tezpur, India, 17–20 December 2019; pp. 61–69.
7. Wang, T.; Xia, W.; Lu, J.; Zhang, Y. A review of deep learning CT reconstruction from incomplete projection data. *IEEE Trans. Radiat. Plasma Med. Sci.* **2024**, *8*, 138–152. [[CrossRef](#)]
8. Ratul, M.A.R.; Yuan, K.; Lee, W. CCX-rayNet: A Class Conditioned Convolutional Neural Network For Biplanar X-rays to CT Volume. In Proceedings of the 2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI), Nice, France, 13–16 April 2021; IEEE: Piscataway, NJ, USA, 2021; pp. 1655–1659.
9. Wang, Y.; Yang, T.; Huang, W. Limited-Angle Computed Tomography Reconstruction using Combined FDK-Based Neural Network and U-Net. In Proceedings of the 2020 42nd Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC), Montreal, QC, Canada, 20–24 July 2020; pp. 1572–1575.
10. Gao, K.; Gao, Y.; He, H.; Lu, D.; Xu, L.; Li, J. NeRF: Neural radiance field in 3D vision, a comprehensive review. *arXiv* **2022**, arXiv:2210.00379.
11. Zha, R.; Zhang, Y.; Li, H. NAF: Neural attenuation fields for sparse-view CBCT reconstruction. In Proceedings of the Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, 18–22 September 2022; Proceedings, Part VI; pp. 442–452.
12. Shen, L.; Pauly, J.; Xing, L. NeRP: Implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE Trans. Neural Netw. Learn. Syst.* **2024**, *35*, 770–782. [[CrossRef](#)] [[PubMed](#)]
13. Zhao, H.; Zhou, Z.; Wu, F.; Xiang, D.; Zhao, H.; Zhang, W.; Li, L.; Li, Z.; Huang, J.; Hu, H.; et al. Self-supervised learning enables 3D digital subtraction angiography reconstruction from ultra-sparse 2D projection views: A multicenter study. *Cell Rep. Med.* **2022**, *3*, 100775. [[CrossRef](#)] [[PubMed](#)]
14. Zuo, J. 2D to 3D Neurovascular Reconstruction from Biplane View via Deep Learning. In Proceedings of the 2021 2nd International Conference on Computing and Data Science (CDS), Stanford, CA, USA, 28–29 January 2021; IEEE: Piscataway, NJ, USA, 2021; pp. 383–387.
15. Wang, L.; Liang, D.X.; Yin, X.L.; Qiu, J.; Yang, Z.Y.; Xing, J.H.; Dong, J.Z.; Ma, Z.Y. Weakly-supervised 3D coronary artery reconstruction from two-view angiographic images. *arXiv* **2020**, arXiv:2003.11846.
16. Wang, Y.; Banerjee, A.; Choudhury, R.P.; Grau, V. Deep Learning-based 3D Coronary Tree Reconstruction from Two 2D Non-simultaneous X-ray Angiography Projections. *arXiv* **2024**, arXiv:2407.14616.
17. İbrahim, A.; GEDİK, O.S. 3D reconstruction of coronary arteries using deep networks from synthetic X-ray angiogram data. *Commun. Fac. Sci. Univ. Ank. Ser. A2-A3 Phys. Sci. Eng.* **2022**, *64*, 1–20.
18. Uluhan, G.Y.; Gedik, Ö.Ü.O.S. 3D Reconstruction of Coronary Artery Vessels from 2D X-ray Angiograms and Their Pose’s Details. In Proceedings of the 2022 30th Signal Processing and Communications Applications Conference (SIU), Safranbolu, Turkey, 15–18 May 2022; IEEE: Piscataway, NJ, USA, 2022; pp. 1–4.
19. Iyer, K.; Nallamothe, B.K.; Figueroa, C.A.; Nadakuditi, R.R. A Multi-stage Neural Network Approach for Coronary 3D Reconstruction from Uncalibrated X-ray Angiography Images. *Sci. Rep.* **2023**, *13*, 17603. [[CrossRef](#)] [[PubMed](#)]
20. Bransby, K.M.; Tufaro, V.; Cap, M.; Slabaugh, G.; Bourantas, C.; Zhang, Q. 3D Coronary Vessel Reconstruction from Bi-Plane Angiography using Graph Convolutional Networks. *arXiv* **2023**, arXiv:2302.14795.
21. Maas, K.W.H.; Pezzotti, N.; Vermeer, A.J.E.; Ruijters, D.; Vilanova, A. NeRF for 3D Reconstruction from X-ray Angiography: Possibilities and Limitations. In Proceedings of the Eurographics Workshop on Visual Computing for Biology and Medicine, Norrköping, Sweden, 20–22 September 2023; Hansen, C., Procter, J., Raidou, R.G., Jönsson, D., Höllt, T., Eds.; The Eurographics Association: Eindhoven, The Netherlands, 2023. [[CrossRef](#)]
22. Müller, T.; Evans, A.; Schied, C.; Keller, A. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Trans. Graph. (ToG)* **2022**, *41*, 1–15. [[CrossRef](#)]
23. Adler, J.; Kohr, H.; Öktem, O. *Operator Discretization Library (ODL)*; Zenodo: Genève, Switzerland, 2017. [[CrossRef](#)]
24. Zeng, A.; Wu, C.; Lin, G.; Xie, W.; Hong, J.; Huang, M.; Zhuang, J.; Bi, S.; Pan, D.; Ullah, N.; et al. ImageCAS: A large-scale dataset and benchmark for coronary artery segmentation based on computed tomography angiography images. *Comput. Med Imaging Graph.* **2023**, *109*, 102287. [[CrossRef](#)] [[PubMed](#)]
25. Teschner, M.; Heidelberger, B.; Müller, M.; Pomerantes, D.; Gross, M.H. Optimized spatial hashing for collision detection of deformable objects. In Proceedings of the 8th Workshop on Vision, Modeling, and Visualization (VMV), München, Germany, 19–21 November 2003; Volume 3, pp. 47–54.
26. He, K.; Zhang, X.; Ren, S.; Sun, J. Deep residual learning for image recognition. In Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, Las Vegas, NV, USA, 26 June–1 July 2016; pp. 770–778.
27. Maas, A.; Hannun, A.; Ng, A. Rectifier Nonlinearities Improve Neural Network Acoustic Models. In Proceedings of the International Conference on Machine Learning, Atlanta, GA, USA, 16–21 June 2013.
28. Paszke, A.; Gross, S.; Massa, F.; Lerer, A.; Bradbury, J.; Chanan, G.; Killeen, T.; Lin, Z.; Gimelshein, N.; Antiga, L.; et al. PyTorch: An imperative style, high-performance deep learning library. *Adv. Neural Inf. Process. Syst.* **2019**, *32*.
29. Kingma, D.P.; Ba, J. Adam: A method for stochastic optimization. *arXiv* **2014**, arXiv:1412.6980.

30. Çiçek, Ö.; Abdulkadir, A.; Lienkamp, S.S.; Brox, T.; Ronneberger, O. 3D U-Net: Learning dense volumetric segmentation from sparse annotation. In Proceedings of the International Conference on Medical Image Computing and Computer-assisted Intervention, Athens, Greece, 17–21 October 2016; Springer: Berlin/Heidelberg, Germany, 2016; pp. 424–432.
31. Shit, S.; Paetzold, J.C.; Sekuboyina, A.; Ezhov, I.; Unger, A.; Zhylka, A.; Pluim, J.P.; Bauer, U.; Menze, B.H. cIDice—a novel topology-preserving loss function for tubular structure segmentation. In Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, Nashville, TN, USA, 20–25 June 2021; pp. 16560–16569.
32. Bousse, A.; Zhou, J.; Yang, G.; Bellanger, J.J.; Toumoulin, C. Motion compensated tomography reconstruction of coronary arteries in rotational angiography. *IEEE Trans. Biomed. Eng.* **2008**, *56*, 1254–1257. [[CrossRef](#)] [[PubMed](#)]
33. Silversmith, W. *cc3d: Connected Components on Multilabel 3D & 2D Images*; Zenodo: Genève, Switzerland, 2021. [[CrossRef](#)]
34. Del Barrio, E.; Cuesta-Albertos, J.A.; Matrán, C. An optimal transportation approach for assessing almost stochastic order. In *The Mathematics of the Uncertain*; Springer: Berlin/Heidelberg, Germany, 2018; pp. 33–44.
35. Dror, R.; Shlomov, S.; Reichart, R. Deep Dominance—How to Properly Compare Deep Neural Models. In Proceedings of the 57th Conference of the Association for Computational Linguistics, ACL 2019, Florence, Italy, 28 July–2 August 2019; Volume 1: Long Papers; Korhonen, A., Traum, D.R., Màrquez, L., Eds.; Association for Computational Linguistics: Stroudsburg, PA, USA, 2019; pp. 2773–2785. [[CrossRef](#)]
36. Ulmer, D.; Hardmeier, C.; Frellsen, J. deep-significance: Easy and Meaningful Significance Testing in the Age of Neural Networks. In Proceedings of the ML Evaluation Standards Workshop at the Tenth International Conference on Learning Representations, Virtual, 25–29 April 2022.
37. He, H.; Banerjee, A.; Beetz, M.; Choudhury, R.P.; Grau, V. Semi-Supervised Coronary Vessels Segmentation from Invasive Coronary Angiography with Connectivity-Preserving Loss Function. In Proceedings of the 2022 IEEE 19th International Symposium on Biomedical Imaging (ISBI), Kolkata, India, 28–31 March 2022; IEEE: Piscataway, NJ, USA, 2022; pp. 1–5.
38. He, H.; Banerjee, A.; Choudhury, R.P.; Grau, V. Automated Coronary Vessels Segmentation in X-ray Angiography Using Graph Attention Network. In Proceedings of the Statistical Atlases and Computational Models of the Heart. Regular and CMRxRecon Challenge Papers, Vancouver, BC, Canada, 2 February 2024; pp. 209–219.

Disclaimer/Publisher’s Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.