NeRF for 3D Reconstruction from X-ray Angiography: Possibilities and Limitations

K.W.H. Maas[1] , N. Pezzotti[1,2] , A.J.E. Vermeer[1,3] , D. Ruijters[1,2] , and A. Vilanova[1]

[1]Eindhoven University of Technology, The Netherlands
[2]Philips Healthcare, The Netherlands
[3]Catharina Ziekenhuis Eindhoven, The Netherlands

Abstract

Neural Radiance Field (NeRF) is a promising deep learning technique based on neural rendering for three-dimensional (3D) reconstruction. This technique has overcome several limitations of 3D reconstruction techniques, such as removing the need for 3D ground truth or two-dimensional (2D) segmentations. In the medical context, the 3D reconstruction of vessels from 2D X-ray angiography is a relevant problem. For example, the treatment of coronary arteries could still benefit from 3D reconstruction solutions, as common solutions do not suffice. Challenging areas in the 3D reconstruction from X-ray angiography are the vessel morphology characteristics, such as sparsity, overlap, and the distinction between foreground and background. Moreover, sparse view and limited angle X-ray projections restrict the information available for the 3D reconstructions. Many traditional and machine learning methods have been proposed, but they rely on demanding user interactions or require large amounts of training data. NeRF could solve these limitations, given that promising results have been shown for medical (X-ray) applications. However, to the best of our knowledge, no results have been shown with X-ray angiography projections or consider the vessel morphology characteristics. This paper explores the possibilities and limitations of using NeRF for 3D reconstruction from X-ray angiography. An extensive experimental analysis is conducted to quantitatively and qualitatively evaluate the effects of the X-ray angiographic challenges on the reconstruction quality. We demonstrate that NeRF has the potential for 3D X-ray angiography reconstruction (e.g., reconstruction with sparse and limited angle X-ray projections) but also identify explicit limitations (e.g., the overlap of background structures) that must be addressed in future works.

CCS Concepts
• Computing methodologies → Reconstruction; • Applied computing → Life and medical sciences;

1. Introduction

Neural Radiance Field (NeRF) [MST*21] is a promising deep learning technique for various tasks, such as view-synthesis and 3D reconstruction. It leverages fully-connected neural networks to represent a scene as a continuous 3D function. The network is optimized with volume rendering techniques, avoiding the need for 3D ground truth information. Contrary to other deep learning methods, NeRF does not rely on learning a prior function but rather learns an individual scene given a set of images at hand. NeRFs have been created and successfully applied for photo-realistic scene reconstruction [MBRS*21].

On the other hand, the three-dimensional (3D) reconstruction of blood vessels from two-dimensional (2D) X-ray angiography is a clinically relevant problem. These angiograms are fundamentally limited by providing a 2D representation of a 3D anatomy, leading to a lack of depth perception or foreshortening. Clinical settings, such as the treatment of coronary arteries, could benefit from a 3D reconstruction from a few X-ray angiography images. During in-

terventions enhancing the 3D perception of the vessels could be achieved by a real-time geometrical reconstruction overlayed over the X-ray [PKA*18]. The main challenges in 3D X-ray angiography reconstruction are related to the X-ray system and the vessel characteristics [ÇGGF16]. The X-ray system challenges include sparse view or limited angle X-ray projections, incorrect X-ray system calibration, and intra-scan motion. Concerning the vessel characteristics, their sparsity, overlap, and poor visibility, due to contrast medium inhomogeneity or overlap of background structures in the images, may misguide the perception of geometric structures and complicate the distinction of foreground and background.

Several 3D reconstruction methods for X-ray angiography have been proposed, each with advantages and disadvantages. These methods include traditional and machine-learning approaches [ÇGGF16, JZ21]. Traditional reconstruction approaches have overcome challenges such as vessel sparsity and visibility. However, they require time-consuming user interactions, like 2D segmentations or keypoint selections, which do not fit clinical workflows. Although machine learning approaches do not re-

quire these user interactions, they require large amounts of annotated training data, which are not publicly available in such large sizes [ZZW*22, WLY*20].

NeRF does not require user interaction or depend on large amounts of annotated training data, making it an interesting methodology for 3D vessel reconstruction. Moreover, successful adaptations of NeRFs have been shown for medical imaging [ZIL*21, CFFBT*22], including X-ray applications [ZZL22, RWL*22, FML*22]. NeRF X-ray techniques have demonstrated the possibility of reconstructing with sparse view and limited angle X-ray projections. However, NeRFs are yet to be applied to X-ray angiography. It is unclear whether existing results apply to complex morphological vessel objects with challenges such as thin objects, overlap, and occlusion. Overall, despite the potential, the vessel morphology characteristics, combined with sparse and limited X-ray projections, are far from the common scenes successfully reconstructed with NeRFs.

In this paper, we define, execute and present an extensive experimental analysis of the possibilities and limitations of applying NeRF for 3D X-ray angiography reconstruction. Overall, we demonstrate that NeRF can reconstruct segmented blood vessels with surprisingly few views. Tuned to angiography applications, we analyze the effects of sparse and limited angle X-ray projections, sparse vessel structures, and vessel visibility quantitatively and qualitatively. We also consider pixel-based sampling techniques to improve reconstruction quality and rendering time. Next to quantitative analysis, we present interactive web-based visualizations that facilitate qualitative analysis of the results. These visualizations provide valuable insights into the potential use of the reconstructions in the context of angiography.

2. Background

The following sections briefly introduce the original NeRF method and the clinical background for 3D reconstruction from X-ray angiography for the basis of our experiments.

2.1. Neural Radiance Field

Neural Radiance Field (NeRF) by Mildenhall et al. [MST*21] represents a single scene as a continuous function F_Θ, which is modeled with a multilayer perceptron (MLP), a fully-connected neural network with layers with randomly-initialized weights Θ. The function $F_\Theta(\mathbf{x}, \mathbf{d}) = (\mathbf{c}, \sigma)$ takes as input a 3D coordinate $\mathbf{x} = (x, y, z)$ and 2D camera viewing direction $\mathbf{d} = (\theta, \phi)$ and outputs color $\mathbf{c} = (r, g, b)$ and volume density σ. This function models an implicit representation of the 3D volumetric space, from which a 3D representation can be obtained via (grid-wise) sampling. Using volume rendering [Max95], the color and density of 3D coordinates along a ray $\mathbf{r}$ can be accumulated into a predicted 2D image pixel color $\hat{C}(\mathbf{r})$, where the ray $\mathbf{r}$ is defined as $\mathbf{r} = \mathbf{o} + t\mathbf{d}$ with $\mathbf{o}$ and $\mathbf{d}$ as the origin and direction. The network is optimized by comparing this final predicted color $\hat{C}(\mathbf{r})$ to the ground truth pixel color $C(\mathbf{r})$. Therefore, optimization is solely based on 2D images and their respective camera poses, alleviating the need for a 3D ground truth. The camera poses are sampled on a sphere with rotation angles θ and ϕ, where θ is the azimuth angle, and ϕ is the elevation angle.

To reconstruct high-frequency details, NeRFs leverage positional encoding functions, which map the input $\mathbf{x}$ and $\mathbf{d}$ into a higher dimensional space allowing the MLPs to approximate a higher frequency function [TSM*20]. To optimize training, NeRF also utilizes a form of importance sampling within rays with two networks. For more detail on the implementation of the positional encoding and the two network settings, we refer to the work of Mildenhall et al. [MST*21].

2.2. Clinical Background

Despite the advantages of 3D imaging modalities, X-ray angiography is still the gold standard for most clinical decision-making and interventions. While 2D angiography has high spatial and temporal resolution than 3D methods, poor depth perception due to visual distortion, also referred to as *foreshortening*, limits its applicability [ÇGGF16]. However, to reconstruct a 3D volume, a significant increase number of X-ray projections is required in relation to the imaging required for diagnosis and treatment [JSGR09] This leads to a larger dosage of radiation and contrast material, which is harmful to the patient. The fact that the orientation of such projections is typically manually chosen by the angiographer further limits the ability to reconstruct the 3D structures properly [GMC*09].

The reconstruction of angiographic images enables multiple medical applications based on the expected quality of the reconstructed vessels [CS09]. The *roadmap overlay* application, aimed at providing a 3D visualization as real-time guidance for navigation during cardiac interventions, requires highly accurate topological requirements in seconds, whereas the accuracy of the vessel diameters is not as relevant [RHM*11]. On the other hand, for assessing coronary artery disease, where the quantification of the narrowing in the coronary arteries is critical, highly precise reconstructions, especially for vessel diameter quantification, are needed [CS09]. Although this application is less time-critical than roadmap overlay, but should still be obtained within minutes to fit clinical workflows. In our work, we focus on the navigation application, like roadmap overlay, as a starting point for using 3D reconstructions from NeRFs.

3. Challenges

The task of 3D reconstruction from X-ray angiography images has multiple challenges that we address in our experiments. We refer to these challenges as (**C#**), where (**C1-C2**) refer to X-ray system limitations and (**C3-C5**) to the vessel morphology characteristics.

Viewing angle limitations for X-ray imaging are [ZIL*21]: (**C1**) **Sparse view tomography** that addresses the problem of a few projections across the angular range; and (**C2**) **Limited angle tomography** that addresses the challenge that the projection angles are within a restricted angular range. Both limitations can result in the loss of crucial information regarding the 3D morphology. In general, the acquisition camera geometry of the X-ray system is fundamental to the 3D reconstruction. For this paper, we will assume a perfect system calibration.

On the other hand, the vessel shape structure itself poses other challenges, such as vessel sparsity, overlap, and visibility [ÇGGF16]. These challenges mainly hinder the completeness

of the vascular morphologies, as they may cause missing structures in the 3D reconstructions. The challenges are: **(C3)** The **vessel sparsity** refers to the thin nature of the vessels, requiring a dense sampling approach for reconstruction; **(C4)** The **vessel overlap** addresses the occlusion of relevant parts of the vessels in the 2D projection; **(C5)** The **vessel visibility** causes parts of the structures to be invisible through contrast inhomogeneity or the overlap of background structures. The typical solution to these challenges is 2D segmentations, which require time-consuming and error-prone manual input. For certain vascular morphologies, such as cerebral arteries, digital subtraction angiography is utilized to address background structures. However, these solutions for background and foreground distinction do not apply to cases like coronary angiography. In our work, we focus on the challenges that emerge from the 3D reconstruction of coronary angiography images and the potential of using NeRFs in this setting. The challenges that emerge from cardiac and respiratory motion in coronary angiography are considered out of the scope.

4. Related work

Our work is related to NeRFs in general and 3D reconstruction techniques for X-ray angiography, whose state-of-the-art is summarized in the following sections in relation to the challenges discussed in Section 3.

4.1. Neural Radiance Fields

NeRF has been widely adopted for reconstructing natural scenes with only a few applications in the medical domain due to its limitations. In this section, we introduce such issues and present some of the work that aims at overcoming such limitations. A few NeRF methods have been proposed to reconstruct medical scenes [RWL*22,FML*22,ZZL22,ZIL*21]. However, these methods do not consider **(C3-C5)** the morphological characteristics of vascular structures, which are unique to our work.

The challenge of reconstruction with NeRF with **(C1)** sparse views has also been explored for synt4.1hetic and medical scenes [TTM*22, RKY*23]. NeRFs require a significant amount of views to obtain accurate reconstructions. Learning priors [YYTK21, WWG*21] or regularization [NBM*22, KSH22] techniques have been proposed to address the sparse view scenario. Our work investigates the base limitations of the original NeRF for our application, so we refer to these techniques for possible future work. In the medical context, NeRF research has shown that **(C2)** the limited angle projections cause significant decreases in performance [ZIL*21, RWL*22]. They also demonstrate that a significantly larger amount of **(C1)** sparse view projections is required for general medical scenes [RWL*22, FML*22, ZZL22, ZIL*21]. These works define a much larger number of views and angular ranges than what is possible in the angiography application domain. In addition, sparse and limited angle projections are only addressed in isolation, i.e., sparse views in the whole angular range or small angular range with many views, while it is common in our clinical applications to have them intertwined. **(C3-C5)** The morphological characteristics may enhance the difficulty of reconstruction with few views due to overlapping, occluding, and sparse

structures. Therefore, we analyze the possibilities of reconstructing vessel structures in this context.

Although NeRFs produce highly-detailed reconstructions, these come at the cost of high optimization times [RPLG21]. Several approaches have been proposed to improve the training time of NeRFs. Examples are smaller network sizes, volumetric representations, and 2D pixel sampling. While the reduction in network size is trivial, other works accelerate training with efficient volumetric representations, such as sparse grids, octrees, and hashes [TTM*22], which have also been adopted for medical scenes [RWL*22, ZZL22, FML*22]. However, these volumetric representations are tailored to the task at hand, limiting our experimental analysis. Volume rendering techniques, such as empty space skipping and early termination of rays [RPLG21, MESK22, LTK22], have also been proposed. Although these techniques are less effective than volumetric representation approaches, they are generalizable across several NeRF methods [LTK22]. We decide to leverage these techniques to pinpoint the constraints of the original NeRF method while significantly improving training time. However, more efficient approaches could be applied once the optimal parameters are deduced from our analysis. Another approach to accelerate training is reducing the number of sampled rays per training image. The sampling of rays could be more focused on reconstructing the foreground objects or objects of interest. Some works have stressed the significance of 2D pixel importance sampling to accelerate the training of NeRF through point or region sampling [YCFB*21] or pre-defined bounding boxes [XPMBB21]. They also stress that degenerate results may occur due to local minima, as information is not weighted equally. Both approaches utilize synthetic scenes, where the foreground and background are relatively easy to distinguish. In X-ray angiography, such **(C5)** background distinction is problematic, so we analyze the possibilities of adapting these methods in such a context.

4.2. X-ray angiography reconstruction techniques

Several techniques for 3D vessel reconstruction from X-ray angiography have been proposed [ÇGGF16, JZ21]. 3D reconstruction solutions for rotational angiography, which utilize many projections, have been very successful. However, this work will focus on single-plane X-ray angiography, which can still benefit from 3D reconstruction solutions. Existing techniques can be categorized into traditional approaches, tomographic and model-based, and machine-learning approaches.

Tomographic reconstruction methods produce a volume of attenuation coefficients directly from the images. These methods do not require manual interventions and can reconstruct highly-accurate vessels with **(C5)** overlapping background structures. However, they require many **(C1)** sparse projections (> 100) and **(C2)** large angular ranges ($> 120°$) for the X-ray projections for high-precision applications. Also, they rely on the need for **(C5)** background removal, limiting their applicability in the clinical workflow [ÇGGF16]. Model-based methods extract the vessels from the 2D images to reconstruct them as binary trees. Unlike tomographic reconstruction methods, these methods can reconstruct from **(C1)** sparse and **(C2)** limited angle projections fitting the clinical data. They, however, typically require manual user input for **(C5)** the

distinction of the vessel from the background, mainly through precise 2D segmentations [CS09] or keypoint selections [YCC*14], which is error-prone and time-consuming. To overcome the reliance on manual user input, approaches employing tubularity responses (e.g., the Frangi filter [FNVV98]) or automatic segmentations [THK*11] have been proposed. Nevertheless, these approaches require (**C1**) many sparse projections for accurate reconstructions to mitigate the error-prone segmentations yielded by these approaches. Some machine learning works have also been proposed [WLY*20, ZZW*22]. Although they indicate the possibility of reconstructing from (**C1**) a few sparse projections in (**C2**) limited angular ranges in projections, they rely on large amounts of annotated training data, which are unavailable.

NeRFs perform per-scene optimization and have demonstrated impressive results without requiring large amounts of training data or manual input. Therefore, NeRFs have the potential to overcome the existing limitations of state-of-the-art methods, making our investigation relevant.

5. Experimental design

The NeRF model has to be adapted for X-ray angiography since we are not reconstructing general scenes. This section describes our adaptation to the NeRF model and explains the parameters that will be analyzed. We also describe the two datasets used and introduce our evaluation approach.

### 5.1. Neural radiance fields for X-ray angiography

X-ray can be physically represented by just the absorption-only physical model, derived from the Beer-Lambert law [Max95]. Similar to other works addressing NeRF for X-ray reconstructions [RWL*22, FML*22, ZZL22], we approximate the scene with the function $F_\Theta(\mathbf{x}) = \sigma$. The volume density σ, or attenuation coefficient, expresses the density of the material that the ray passes through. We refer to the work of Max et al. [Max95] for the definition of the continuous absorption equation. Scattering effects are ignored by assuming a constant background intensity. We model the scene independently from the viewing direction, unlike the original NeRF architecture. The discretization of the absorption equation can then be expressed as

$$I(\mathbf{r}) = \prod_{i=1}^{N} \exp(-\sigma_i \delta_i),$$

where $\delta_i = t_{i+1} - t_i$ is the distance between the adjacent sample points and N is the number of sampling points between a predefined scene's bounding box. The loss is computed for predicted pixel intensity $\hat{I}(\mathbf{r})$ and ground truth pixel intensity $I(\mathbf{r})$ for a random batch of rays $\mathcal{R}_r$ as the mean squared error (MSE)

$$\mathcal{L} = \frac{1}{|\mathcal{R}_r|} \sum_{\mathbf{r} \in \mathcal{R}_r} \left(\hat{I}(\mathbf{r}) - I(\mathbf{r}) \right)^2. \quad (1)$$

Our work has chosen not to utilize positional encoding due to its limited impact in our context. Contrary to natural scenes, the individual human structures imaged in X-ray are uniform in attenuation and smooth in shape. Moreover, the contrast between the
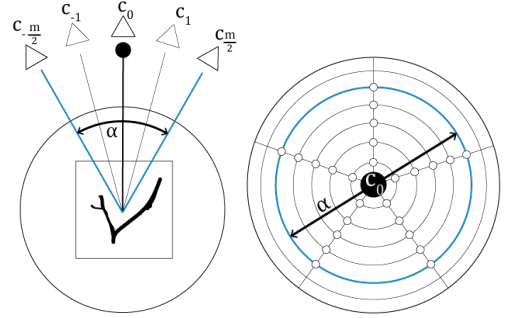
(a) *Rotation over one axis*   **(b)** *Rotation over two axes*

Figure 1: *Sparse view and limited angle camera positions, where α and m denote the limited angular range and number of sparse views, respectively. (a) Rotations over one angle α. (b) Rotations over two angles.*

background and blood vessels in our imaging is high. Both can be modeled easily by low-dimensional features. The background of X-ray angiography may contain high-frequency details, but our goal is not to reconstruct them.

Instead of the importance sampling used by the original NeRF implementation for acceleration, we adopt *NeRFAcc for NeRF acceleration* [LTK22]. For this work, we analyze varying scenes and basic bottleneck limitations, for which a more general optimization approach, like NeRFAcc, is suitable. NeRFAcc implements efficient volume rendering techniques, specifically skipping empty and occluded spaces, to significantly improve the training time of NeRFs. While this method was mainly developed for acceleration, it also improves reconstruction quality, as the network utilizes its capacity mainly for the non-empty areas [TTM*22]. Empty spaces are skipped based on an occupancy grid, where emptiness is determined based on the density $\sigma_i < h_e$, where h_e is a user-specified threshold. Occluded regions are skipped by terminating the ray marching early based on a set threshold h_o for the transparency $T_i < h_o$. The accumulated transparency T_i is calculated with the absorption equation from the ray's start to the point i.

5.2. Parameters

This section describes the analysis's parameter space and the choices made for the systematic experimental design, including projections, vessel sparsity, background, and sampling. Such parameters have been motivated by the challenges described in section 3.

Projections   The projections parameter addresses the challenge of (**C1**) sparse projections and (**C2**) limited angle induced by the physical limitations of the X-ray systems. The X-ray's source and detector rotation can be described as a spherical rotation around a center point with an x-rotation angle θ and y-rotation angle ϕ [ÇGGF16], equal to the camera rotations defined for NeRFs. Figure 1 visualizes the setting of limited angle and sparse projections with the possible camera positions $c_i = (\theta, \phi)$ for rotations over one axis (a) and two axes (b). For the one-axis rotation, limited angle projections are sampled around a pre-defined center point from a limited angular range α, highlighted in blue. The black circle indicates the center point in camera c_0. Due to the X-ray projection

transparency, the vessel geometry object can be assumed symmetric over the 180° rotation, ignoring the non-symmetries caused by perspective projection. Hence, we constrained the angular range to 180° since angles beyond this limit do not contribute any additional information. The *m* sparse projections are obtained by uniform sampling within the limited angular range, where the triangles represent each camera position. For two rotation axes, the space of possible camera rotations is visualized with a 2D circular heatmap, as shown in Figure 1b. Each of the six sides of the sphere can be represented by one 2D circular heatmap. The first axis of the heatmap is the distance from the center point, and the second axis is the angle from a point to the positive x-axis. The number of sparse projections sampled per axis is *m*. Each circular heatmap region within the angular bound represents one of these m^2 sparse projections, represented with the circles (○). Specific projections of the vessel structure may bias the results if the same center point is used for sampling. Therefore, three center points at each unique sphere side are defined to cover all possible camera rotations of the sphere. This setup enables us to examine the impact of morphology during our experimental analysis.

Vessel sparsity   Vascular trees can vary largely in the sparsity of their vessels. This parameter addresses the role of (**C3**) sparse vessels in reconstruction quality. We analyze the effect of sparsity between two vascular structures with varying amounts of branching and sizes of the vessels. Specifically, the morphologies are low and high in sparsity, respectively. Reconstruction of sparse vessels may pose a problem for NeRFs, as in rendering methods, sparse objects may require many dense samples in the 3D space. NeRFs heavily rely on sampling in 3D space in terms of quality and rendering speed. The combination of sparse and limited angle projections and sparsity is also interesting, as (**C4**) overlap may cause sparse vessels to be hidden in these projections. We also conducted experiments to determine whether different model architectures play a role in the vessel sparsity parameter. We found that larger models are relevant for capturing sparser vessels. The details can be found in the supplementary material.

Background   Visibility of the vessels is one of the main limitations in the reconstruction of X-ray angiography for traditional vessel 3D reconstruction methods, as described in Section 4.2. By distinguishing binary and background cases, this parameter addresses the challenge of (**C5**) (overlapping) background structures. The binary case simulates the reconstruction of a binary representation of the vessel, where the background has been removed based on a known high-quality 2D segmentation. The background case focuses on the tomographic reconstruction of the attenuation volume. The intensity of the vessels is assumed to be homogeneous, so contrast inhomogeneity is not considered explicitly. By comparing the reconstruction of the binary structure and the background, we can identify whether 2D segmentations are essential for the reconstruction of the vessels with sparse and limited angle projections. In summary, with this parameter, we can explore the trade-off between the availability of segmentations and amount of projections, which also plays a role in the related work.

Sampling   This parameter explores the effects of 2D pixel importance sampling on the reconstruction quality and time. The random sampling of rays of all ground truth pixels during the training of
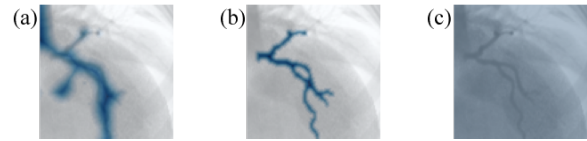

Figure 2: *(a) Frangi, (b) segmentation, and (c) random 2D pixel importance sampling approaches, where the intensity of the blue areas indicates higher importance of sampling.*

NeRF may be inefficient and unnecessary (see Section 4.1). Importance areas can be defined to determine the frequency of sampling batches of pixel rays belonging to these areas, as stated as $\mathcal{R}$ in Equation 1. We consider three settings: Frangi weighted sampling, segmentation weighted sampling, and random sampling, illustrated in Figure 2 with the importance areas highlighted in blue. We refer to the Frangi weighted sampling and segmentation weighted sampling as segmentation-based sampling approaches. The Frangi weighted sampling rays $\mathcal{R}_f$ are computed based on the distance transform of the image's Frangi response [FNVV98]. This sampling approach could speed up the reconstruction of the vessel area without user input. However, as a weak segmentation, the Frangi filter is also sensitive to non-vessel areas, so it may lead to noisy results. Segmentation weighted sampling rays $\mathcal{R}_s$ are computed based on known high-quality 2D segmentations and the resulting distance transform. Similar to the Frangi approach, segmentation sampling may speed up the reconstruction of the vessel area. Although it avoids the noise of the Frangi filter, this comes at the cost of manual input, limiting the application in an interventional setting. Lastly, the original NeRF method's random sampling approach is also considered. This approach is more computationally expensive but may be beneficial for background reconstruction. Generally, a trade-off between reconstruction quality, time, and manual user input will be analyzed.

5.3. Datasets

Two synthetic volumes are utilized for the experimental analysis. X-ray angiography imaging naturally does not involve 3D ground truth. Moreover, available datasets contain too few images per patient for the reconstruction task. Therefore, we chose synthetic volumes to fully control the sparse and limited angle projections. Both volumes are of the left coronary artery (LCA), as shown in Figure 3. The images used to train NeRF are obtained through backprojection of the volumes.

The first volume is obtained from an anonymized coronary computed tomography angiography (CCTA). Coronary X-ray angiograms (CAG) were simulated with this CCTA by applying a transfer function on manually segmented structures, as shown in Figure 3b. Specifically, a CAG's three main visual human structures
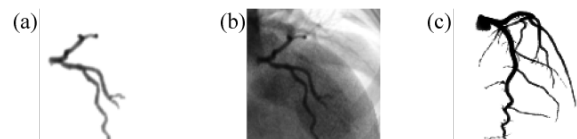

Figure 3: *The low-sparsity dataset, (a) binary and (b) background, and (c) high-sparsity dataset.*

were segmented: the heart, ribs, and LCA. The contrast medium injection of an LCA on a coronary angiogram is mimicked with higher intensity values for this structure with the transfer function. Due to the low resolution of the CT, the LCA consists of smoothened major vessels [CS09]. We, therefore, utilize this volume to simulate the low-sparsity parameter, as well as the background parameter. To generate the binary case, the intensities of the heart and ribs are set to zero, displayed in Figure 3a. To analyze the reconstruction of sparser vessels, a second volume of the LCA is used, which is generated from a silicone phantom model. This model is a highly-detailed surface, as seen in Figure 3c. The synthetic model is only utilized as a binary case, as no background information is available for this surface.

5.4. Evaluation

The reconstruction results are quantitatively and qualitatively evaluated across the parameters described in Section 5.2. A consensus regarding assessing medical imaging quality has not yet been reached [KZPD23]. As a result, common natural image quality metrics, like peak-signal-to-noise ratio (PSNR), structural similarity (SSIM), and, more recently, Learned Perceptual Image Patch Similarity (LPIPS), are mostly applied to evaluate medical imaging computer vision tasks [CP16]. However, even though these metrics capture general visual quality, they are not designed to evaluate the reconstruction's perceived topological accuracy. For example, bifurcation points, hallucinations, or missing segments may be perceived as more important. To the best of our knowledge, no metric has been proposed to evaluate this reconstruction task. Centerlines and diameters can also be utilized to evaluate the vessel morphology, but these focus on precision rather than perceived topological accuracy [RHM*11]. Therefore, we utilize a combination of existing image quality metrics and qualitative analysis for our work.

For the quantitative metrics, we consider four: PSNR, SSIM, LPIPS, and the Dice score. We compute all metric scores such that the higher the score, the better the quality. The results are evaluated in the form of 2D images and 3D volumes, where the 3D volumes are obtained by sampling the learned continuous function F_Θ in a grid. PSNR measures the amount of noise in an image utilizing MSE pixel-wise computations, which corresponds with the loss function of NeRF. Therefore, it may be sensitive to smaller patches of noise specifically and the blurriness of the imaging. SSIM and LPIPS both measure the perceptual similarity between two images. Rather than PSNR, these measures are known to be closer to human perception. They are more robust to noise but less sensitive to blurriness. Whereas SSIM takes a statistical approach to extract structural information from the images, LPIPS utilizes activations of neural networks to represent images. We use both metrics, as they each have shown to be fruitful for medical (X-ray) image evaluation [KZPD23, CP16]. Lastly, the similarity of binarized images is evaluated with the Dice score. To evaluate the vessel morphology in background images, we apply thresholding for binarization. Besides quality metrics, the computation time will also be reported.

We also evaluate the effects of the experimental parameters on the visual quality of the 2D projections guided by image quality metrics. We focus on qualitative analysis of our results in 2D, as X-ray angiography data naturally only consists of ground-truth 2D views. An interactive tool was developed to evaluate quality based on the camera positions with interactive circular heatmaps. Our openly available tool is available at `https://nerfforangiography.netlify.app/`. The explanation of the circular heatmap is provided in Section 5.2. Camera positions are sampled for every $10°$ of the sphere, independently from the number of camera positions used for reconstruction. The regions of the heatmap are colored based on the user-defined quality metric score of the respective 2D predicted projections. This visual inspection overcomes the limitations of the individual metric scores, as users can identify the visual differences relevant to them. Hovering over the regions displays the ground-truth, predicted, and difference images for the camera position (θ, ϕ). Overall, the tool allows for exploring rotation-specific trends, such as the dependency of image quality on the training camera positions or noise patches in the 3D space. More information on the tool can be found in the supplementary materials.

6. Experiments

An overview of the experiments and respective parameters is provided in Table 1. Each experiment (row) is defined and conducted based on a selection of parameters (column). The parameters are schematically displayed based on their role in the experiment. Generally, it is not possible to cover all combinations of parameters, so the experiment space is designed such that parameters can be fixed for follow-up experiments. So, the table is ordered according to the sequence of experiments. For each experiment, we vary the parameters to be explored (●) to identify their influence on the parameters of interest (●), where the remaining parameters are fixed (x). In the first experiment, the parameters have a shared role (◉), where they will be explored and of interest. As certain experiments share parameters, assumptions are made based on previous experiments (●[a]) to limit the exploration space.

The models were implemented in PyTorch and were run on a RTX A5000 GPU. The experiments are conducted with a default MLP model architecture of 4 layers and 128 hidden units, commonly used across different applications for NeRFs [MST*21]. We utilize ReLU activations between the hidden layers and a sigmoid activation for the attenuation coefficient output σ. Each model was trained for 2×10^5 iterations with ray batch size 5625. A learning rate 1×10^{-4} linearly decaying to 1×10^{-5} was used. The emptiness threshold $t_e = 1 \times 10^{-4}$ and early stop threshold $t_o = 1 \times 10^{-2}$ were set for efficient volume rendering.

6.1. Projections

This experiment aims to analyze the effect of (**C1**) sparse and (**C2**) limited angle projections on the reconstruction quality in the simplest binary setting using Frangi sampling. With this experiment, we determine a baseline for the minimum requirements for these projections. The experiments are repeated for three center points to avoid view-specific biases. No major variance was found across these center points, so we display our results as the average.

First, we analyze the effects of sparse and limited angle projections quantitatively. Figure 4(a, b) shows the effects of the limited angle projections with the PSNR and SSIM quality metrics for a

range of 4 to 49 sparse projections and 5° to 180° limited angle projections. These 2D scores are obtained by averaging the score for all projections for every 10° of the sphere. The SSIM, LPIPS, and Dice score portray similar trends for binary structures. Moreover, the scores in 2D and 3D also show similar trends. Therefore, we only portray the varying 2D metrics, the PSNR and SSIM, in the remaining binary structure experiments. The figures for other metrics can be found in the supplementary material. The figures generally show high performance across the varying sparse and limited angle projection settings. They also indicate that a minimum of 30° for limited angle projections is needed to achieve reasonable quality, matching the current demands for 3D reconstruction techniques [THK*11]. The effects of the sparse projections on the reconstruction quality are reported in Figure 4(c,d) in PSNR and SSIM. Unlike the limited angle projections, the sparse projections or the number of projections do not significantly influence the performance quality, visible due to the flat lines. This outcome is considered advantageous because a reduction in the required number of X-ray projections translates to a decreased level of radiation exposure for the patient. However, experimentation has shown that when the number of sparse projections is below 4, it leads to de-

generate results independent of the angle of the projections. These results were similar for all three center points, so only one center point will be considered for the follow-up experiments.

We also report the circular heatmaps for all possible combinations of the limited angle and sparse projections for the SSIM metric, shown in Table 2. The scenario with limited 180° and 4 sparse projections is excluded because the 4 sparse projections result in only 2 distinct projections due to the geometric configuration shown in Figure 1. The heatmaps display the same patterns for the limited angle and sparse projections as the line graphs. Specifically, they indicate stable performance from 30° projections and no significant effects on the number of sparse projections. It also highlights that outliers occur within this pattern, which will be discussed in more detail later. Moreover, it shows that perfect reconstructions can be generated if many projections (sparse > 25) and enough limited angular range (limited > 45°) are provided, as shown by the circular heatmaps colored fully in yellow, indicating high values.

Two cases are discussed to highlight the patterns visible in the heatmaps, as shown in Figure 5. The limited angle projections below 30° show a significant decrease in performance outside of the angular range. It indicates that the limited angular range does not provide enough information on the object, even if many projections are sampled. An example of this pattern is shown in Figure 5a for limited 5° and 49 sparse projections, where moving outside of the limited training range leads to a blurry vessel. However, the geometric accuracy is still reasonable enough to apply roadmap overlay due to the correct topological representation. Some outliers are also visible, such as the case of 90° limited and 16 sparse projections, where patches of camera positions show a significant decrease in performance. As seen in Figure 5b, thanks to our interactive exploration tool, a patch of structural noise outside of the vessel geometry causes this decrease in the metrics. Overlap of this patch and the geometry may have led to a decrease in the SSIM value. However, the location of this patch is away from the vessel geometry, so it does not disturb the application. Apart from this, the geometry of the reconstruction is correct and, therefore, acceptable for practical applications. Overall, we have shown that good quality reconstructions can be obtained with few sparse projections (> 4) and limited angle projections (> 30°) are available. More importantly, the geometry of the vessels is correctly represented.
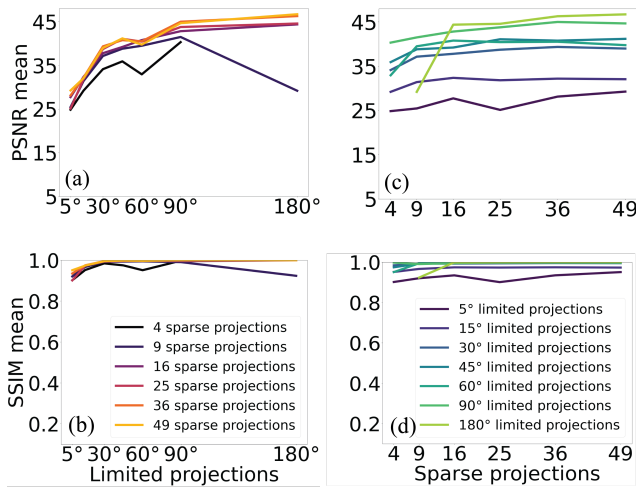

Figure 4: *Relation between limited angle and sparse projections and specified metric scores for the Projections experiment. Limited angle projections in (a) PSNR and (b) SSIM, with legend in (b). Sparse projections in (c) PSNR and (d) SSIM, with legend in (d).*

	Projections		**(C3) Vessel sparsity**		**(C5) Background**		**Sampling**		
	(C1) Sparse	**(C2) Limited**	Low	High	Binary	Background	Frangi	Segment.	Random
1. Projections	◎	◎	x		x		x		
2. Vessel sparsity	●[1]	●[1]	●	●	x		x		
3. Background	●[1]	●[1]	x		●	●	x		
4. Sampling							●	●	●
4.1 Vessel sparsity	●[2]	●[2]	●	●	x				
4.2 Background	●[3]	●[3]	x		●	●			

● = parameters of interest ●[a] = parameters to be explored based on assumptions of experiment *a*

x = fixed parameters ◎ = parameters of interest and to be explored

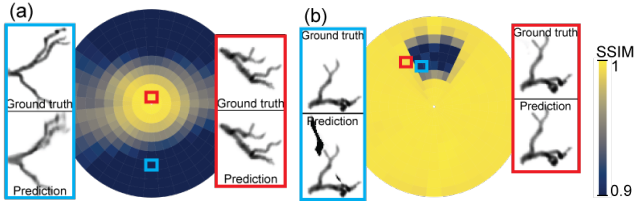**Table 1:** *Overview of experiments and their corresponding parameters.*

Figure 5: *Heatmaps for the Projections experiment. (a) Limited 5° and sparse 49. (b) Limited 90° and sparse 16.*

6.2. Vessel sparsity

This experiment explores the role of (**C3**) high vessel sparsity and (**C1**) sparse and (**C2**) limited angle projections on the reconstruction quality with Frangi sampling. We analyze the results as the mean scores for three different center points to avoid view-specific biases for the high-sparsity dataset.

The relation between limited angle and sparse projections for the quality metrics are shown in Figure 6. Comparing the PSNR scores for the low-sparsity dataset (Figure 4a) from the previous experiment and the high-sparsity dataset (Figure 6a), we observe a significant drop in performance. The requirement for the angular range of limited projections is around 30°, as shown in Figure 6(a, b), similar to the low-sparsity case. In contrast, the number of sparse projections required is significantly different. As shown in Figure 6(c, d), around 16 sparse projections is required to obtain reasonable performance. We also visualized the table of sparse and limited angle projection heatmaps for this dataset, which supports this trend. This table can be found in the supplementary material. The increase of the minimum number of required sparse projections can be explained by the sparse vessels in this dataset, which are often hidden due to (**C4**) overlap in the 2D projections.

Two qualitative cases are highlighted to support this argument, displayed in Figure 7. Figure 7a shows the case for limited 5° and 49 sparse projections, where moving beyond the limited range leads to added noise in overlapping areas of the training images. On the other hand, Figure 7b shows that few, i.e., 4 sparse projections, in limited angle projections 60° may lead to disconnected geometries, highlighted with circles. From the circular heatmap, four distinct regions with higher performance can be identified, which belong to the camera positions of the training images. The vessel structure
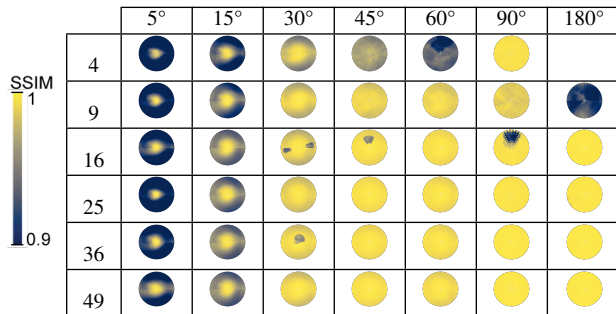
may have overlapped in the specific training images, explaining the performance difference between these regions and the rest of the camera positions. Overall, the sparsity of the vessel morphology affects the reconstruction quality negatively. Although the limited projection angles remain comparable to the previous experiment (> 30°), more sparse projections are required (> 16). This limits the applicability of these reconstructions in clinical practice due to radiation exposure and time constraints.

6.3. Background

The (**C5**) background is a major factor in the reconstruction of X-ray angiography, as precise 2D segmentations are often not available for X-ray angiograms to remove the background. Frangi sampling technique for importance sampling is utilized to minimize the effect of the background on the reconstruction. This experiment explores (**C1**) the sparse and (**C2**) limited angle projections for tomographic reconstruction with a background. Based on the previous experiment in Section 6.1, it is assumed that the minimum angular range of the projections required is 30°. Moreover, the sampling of projections is only performed around one centerpoint.

The background effects are shown in Figure 8 presented in SSIM, PSNR, and Dice for the 2D projections. The figures indicate a significant drop in performance compared to the binary structures without background, shown in Figure 4(a, b). The weak segmentation of the Frangi sampling approach leads to a higher frequency of sampling the important vessel areas. Therefore, the background reconstruction is expected to perform poorly, so the SSIM and PSNR scores, evaluating the whole image quality, are expected to be lower. The Dice score is more reliable for this evaluation, as only the thresholded binary structures are evaluated. Based on this Dice score, it may be concluded that a significantly larger angular range for limited projections is required (> 90°), as well as a larger number of sparse projections (> 16) to reconstruct the attenuation volume, given the drop of quality for fewer projections.
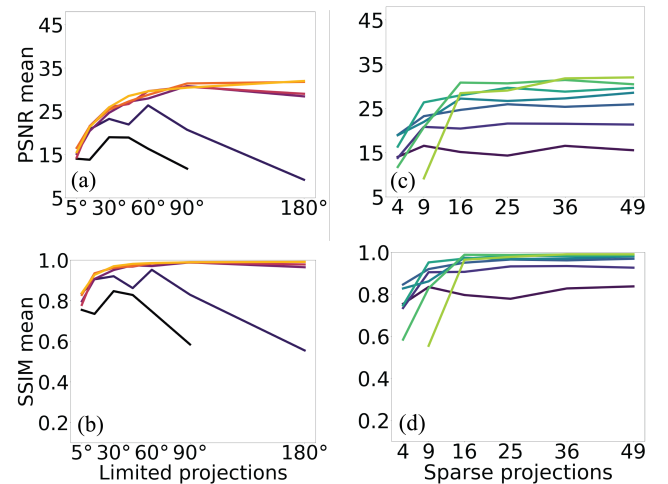

Table 2: *Heatmaps for the low vessel sparsity dataset. The rows of the tables represent the sparse projections and the columns the limited angle projections.*

Figure 6: *Relation between limited angle and sparse projections and specified metric scores for the vessel sparsity experiment. Limited angle projections in (a) PSNR and (b) SSIM. Sparse projections in (c) PSNR and (d) SSIM. The legends are in Figure 4(b,d).*
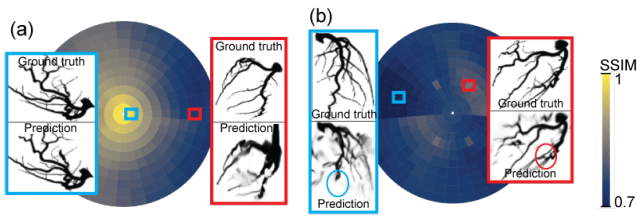
Figure 7: *Heatmaps for the Sparsity experiment. (a) Limited $5°$ and sparse 49. (b) Limited $60°$ and sparse 4.*

Figure 9: *Heatmaps for the Background experiment. (a) Limited $30°$ and sparse 49 (b) Limited $180°$ and sparse 9.*

These results significantly outperform the traditional tomographic reconstruction methods in the number of projections required for topologically accurate applications.

Two cases are qualitatively evaluated to motivate this statement shown in Figure 9. Both cases indicate that the truncation of background structures leads to artifacts in the reconstruction. For example, in the $30°$ case (Figure 9a), the model reached a bad local minimum, where noise occludes the structure from most angles. On the other hand, the sparse 9 case (Figure 9b) suffers from disconnected structures due to overlap in the sparse training views. Generally, it can be concluded that reconstruction of the background is a difficult problem with sparse and limited angle projections. Neither geometric nor precision accuracy satisfies the needs for clinical applications unless many projections are utilized for training. The problem may lie in the Frangi sampling method, which enforces the truncation of the background, as the background is less frequently considered during the reconstruction. The next section investigates whether different importance sampling methods can mitigate the bad local minima with few sparse projections.

6.4. Sampling

This experiment explores the role of 2D segmentations in the form of 2D pixel importance sampling on the reconstruction quality. Two sub-experiments are defined for the vessel sparsity and background parameters, as previous findings have shown that these parameters have distinct impacts on reconstruction quality.

Vessel sparsity We compare the effects of Frangi, segmentation, and random sampling to reconstruct sparse vessels. Due to the noisiness of the Frangi filter, the importance area obtained with this filter may be more prone to highlight larger vessels than smaller vessels. As a result, the sampling approaches may negatively impact the sparse vessels' reconstruction. Therefore, we compare the model's performance across (**C1**) the sparse and (**C2**) limited angle projections and (**C3**) low and high vessel sparsity datasets. In the
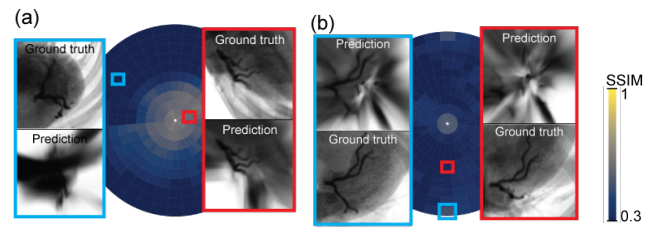
first experiment (see Section 6.1), we have shown that the reconstruction quality of the low-sparsity dataset is high, mostly independent of the projection parameter. As a result, varying the sampling approaches leads to small differences in quality. We refer to the supplementary material for figures detailing the differences in quality and time performance for the low-sparsity dataset.

The sampling methods do lead to significant differences in quality for the high-sparsity dataset. As shown in Figure 10, the random sampling approach especially outperforms the others when fewer sparse projections are given. For example, the random sampling approach provides high-quality scores with only 4 projections in a limited angular range of $15°$, whereas the other approaches lead to noisy results. The qualitative examples can be found in the supplementary material. We expect that the random sampling approach may avoid local minima by sampling more in the empty space. Therefore, more empty space can be determined, which minimizes the possibilities for non-empty vessel spaces. This is especially crucial in the case of (**C4**) overlap in vessels, which explains the effectiveness of importance sampling for the high-sparsity structure with more overlapping areas rather than the low-sparsity structure. In contrast to the results from the earlier vessel sparsity experiments (see Section 6.2), the reconstruction of sparse vessels can be performed with few projections. As such, reasonable quality 3D reconstruction can be obtained with less radiation exposure with fewer X-ray projections. Experimentation has shown that the sampling approaches do not significantly influence the training time (all approximately 50 minutes). Specifics on the training time can be found in the supplementary material. We conclude that random sampling is more effective than the segmentation-based sampling approaches for reconstructing sparse vessel structures with minor differences in running times. More importantly, we demonstrated that sparse structures can be reconstructed with a few sparse projections (> 4), which could be utilized for navigation applications.

Background This experiment explores whether the different sampling techniques can mitigate the problem of local minima when re-
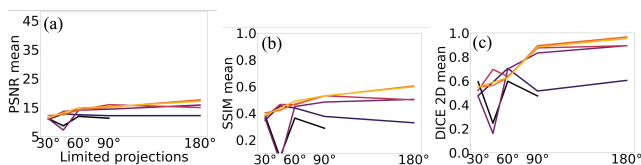

Figure 8: *Relation between limited angle projections and (a) PSNR, (b) SSIM and (c) Dice metrics for the Background experiment. The legend can be found in Figure 4b.*
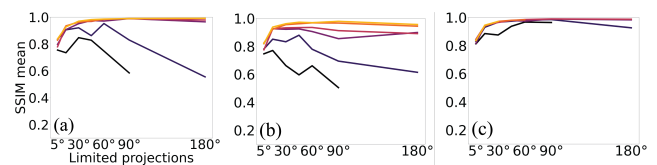

Figure 10: *Relation between limited angle projections and SSIM for (a) Frangi, (b) segmentation, and (c) random sampling approaches for the high-sparsity dataset. The legend is in Figure 4b.*

constructing background images (see Section 6.3). Figure 11 shows the relationship between limited angle projections and the Dice score for the three sampling techniques. Independent of the sampling approach, a minimum angular range of $90°$ is required to reach reasonable performance. However, the random sampling approach is significantly more robust against a few sparse projections (> 9) than the other sampling approaches. We argue that local minima in the background may be avoided by reconstructing the (**C4**) overlapping background structures, which reduces the number of possibilities for the predicted images. To highlight this point qualitatively, Figure 12 shows the Frangi and random sampling results for limited angle projections $90°$ and sparse projections 4. The binarized images show that the reconstruction's topological accuracy of the random sampling method is reasonably accurate, whereas the Frangi sampling approach generates significantly more noise around the vessel structure. The random sampling approach minimizes the accumulation of background structures, reducing the noise around the vessel structure. The training time for reconstructing the attenuation volumes is 240, 258, and 288 minutes using the Frangi, segmentation, and random sampling approaches, respectively. As the attenuation volumes are mainly non-empty, not many spaces can be skipped with the NeRFAcc approach. Therefore, the average reconstruction time of the binary structures (approximately 50 minutes) and background structures (approximately 262 minutes) is significantly different. Altogether, the reconstruction of the attenuation volume of the vessels remains challenging. While segmentation-based sampling approaches generate noisy results, random sampling leads to reasonably topologically accurate structures. However, the approaches' running times and quality remain limited and must be addressed in future work.

7. Discussion and Conclusion

This paper extensively analyzed the potential and limitations of using NeRF for X-ray angiography reconstruction. We designed an experimental framework that focuses on the main challenges of NeRFs in this context, namely (**C1**) sparse and (**C2**) limited angle X-ray projections, (**C3**) vessel sparsity, (**C4**) vessel overlap and (**C5**) background structures. We also incorporated foreseeable solutions to these challenges, such as segmentation-based 2D sampling techniques. We present the results in an interactive visualization environment (`https://nerfforangiography.netlify.app/`) to explore the reconstruction performance of NeRF models quantitatively and qualitatively within the various parameter settings. We demonstrated that NeRF can successfully reconstruct the geometries of binary sparse vessel structures, even with surprisingly few projections (> 4) acquired from a relatively limited angle
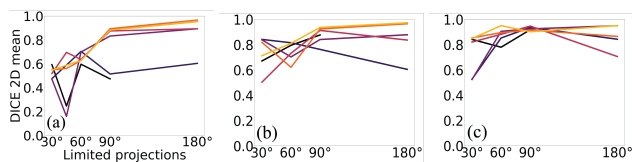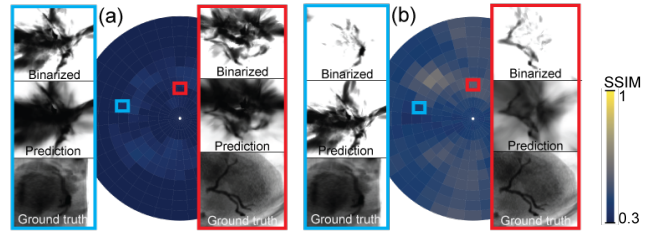

Figure 12: *Heatmaps for the background for sparse 4 and limited $90°$ for the (a) Frangi and (b) random sampling techniques.*

range ($> 30°$). Moreover, we established that with a larger amount of projections (> 9), tomographic reconstructions with reasonable geometric accuracy could also be obtained. Therefore, we established a trade-off between the need for the number of sparse projections and high-quality segmentations. Specifically, the quality of reconstructions relies mainly on the number of sparse projections available, although fewer projections are needed when high-quality segmentations are available. We also show that 2D segmentation-based sampling techniques do not benefit the reconstruction quality of binary or tomographic reconstructions. With a random sampling approach, fewer sparse projections are required than with the segmentation-based sampling approaches, especially for high-sparsity vessel structures (> 4 rather than > 16) and tomographic reconstructions (> 9 rather than > 16). The random sampling approach minimizes local minima that occur due to overlapping vessels or background structures in X-ray angiography images. Overall, our experiments with the basic NerF method show the potential for NerFs to be successfully implemented for navigation applications: geometrical 3D reconstructions from a few images enable vessel navigation.

However, reconstruction from a few sparse projections and training time must be addressed in future work to adapt NeRFs in the clinical setting, both for high-precision applications and tomographic reconstructions. Based on our results, NeRF methods that integrate regularization [KSH22, NBM*22] and learning priors [YYTK21, WWG*21] to reconstruct from fewer projections are worthwhile to investigate. They could potentially resolve the limitation of sparse projections, but it has to be explored whether (high-quality) segmentations may still be needed. Nonetheless, since these techniques depend on depth information or extensive training data, substantial modifications would be necessary to apply them to X-ray scenarios where such information is largely unavailable.

Moreover, the general training time of NeRF in the explored setting is too long for practical clinical use. It needs to be improved for binary structures, approximately 50 minutes, and tomographic reconstructions, approximately 260 minutes, to less than 10 minutes to allow adaptation in clinical practice. The investigation of methods that propose efficient volumetric representation (see Section 4.1) adapted to the X-ray angiography setting is an interesting direction. NeRF methods for X-ray applications [RWL*22, ZZL22, FML*22] should also be investigated for possible improvements in reconstruction quality and training time. If adapted NeRF methods are being proposed for X-ray angiography, our experimental analysis should be performed across these methods to evaluate the improvements and limitations. Moreover,

Figure 11: *Relation between limited angle projections and Dice for (a) Frangi, (b) segmentation, and (c) random sampling approaches for the background. The legend can be found in Figure 4b.*

state-of-the-art methods should also be included in this analysis in order to evaluate the limitations across the task-specific parameters.

Furthermore, our definition of sparse and limited angle projections could be tuned further to the application. Rather than uniformly sampling projection angles, using standard clinical angles and identifying optimal viewing angles would be interesting. As indicated in Section 5.4, we also acknowledge the need for more task-specific quality metrics for our evaluation. However, in the literature, there are no metrics that directly suit this purpose. Finally, although not addressed in this work, challenges such as calibration, scattering noise, inhomogeneous vessel contrast, and motion of coronary arteries are inherently part of X-ray angiography data and should be addressed in future works.

In this work, we presented an experimental framework that we have used to explore the potential of applying NeRF for 3D X-ray angiography reconstruction. We have identified opportunities like good results with sparse and limited angle projections and explicit limitations, such as background hindrance, high-precision reconstructions, and slow training time.

Acknowledgements

References

[CFFBT*22] CORONA-FIGUEROA A., FRAWLEY J., BOND-TAYLOR S., BETHAPUDI S., SHUM H. P., WILLCOCKS C. G.: Mednerf: Medical neural radiance fields for reconstructing 3d-aware ct-projections from a single x-ray. In *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)* (2022), IEEE, pp. 3843–3848. 2

[ÇGGF16] ÇIMEN S., GOOYA A., GRASS M., FRANGI A. F.: Reconstruction of coronary arteries from x-ray angiography: A review. *Medical image analysis 32* (2016), 46–68. 1, 2, 3, 4

[CP16] CHOW L. S., PARAMESRAN R.: Review of medical image quality assessment. *Biomedical signal processing and control 27* (2016), 145–154. 6

[CS09] CHEN S. J., SCHÄFER D.: Three-dimensional coronary visualization, part 1: modeling. *Cardiology clinics 27*, 3 (2009), 433–452. 2, 4, 6

[FML*22] FANG Y., MEI L., LI C., LIU Y., WANG W., CUI Z., SHEN D.: Snaf: Sparse-view cbct reconstruction with neural attenuation fields. *arXiv preprint arXiv:2211.17048* (2022). 2, 3, 4, 10

[FNVV98] FRANGI A. F., NIESSEN W. J., VINCKEN K. L., VIERGEVER M. A.: Multiscale vessel enhancement filtering. In *Medical Image Computing and Computer-Assisted Intervention—MICCAI'98: First International Conference Cambridge, MA, USA, October 11–13, 1998 Proceedings 1* (1998), Springer, pp. 130–137. 4, 5

[GMC*09] GARCIA J. A., MOVASSAGHI B., CASSERLY I. P., KLEIN A. J., JAMES CHEN S.-Y., MESSENGER J. C., HANSGEN A., WINK O., GROVES B. M., CARROLL J. D.: Determination of optimal viewing regions for x-ray coronary angiography based on a quantitative analysis of 3d reconstructed models. *The international journal of cardiovascular imaging 25* (2009), 455–462. 2

[JSGR09] JANDT U., SCHÄFER D., GRASS M., RASCHE V.: Automatic generation of 3d coronary artery centerlines using rotational x-ray angiography. *Medical image analysis 13*, 6 (2009), 846–858. 2

[JZ21] JIA D., ZHUANG X.: Learning-based algorithms for vessel tracking: A review. *Computerized Medical Imaging and Graphics 89* (2021), 101840. 1, 3

[KSH22] KIM M., SEO S., HAN B.: Infonerf: Ray entropy minimization for few-shot neural volume rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 12912–12921. 3, 10

[KZPD23] KASTRYULIN S., ZAKIROV J., PEZZOTTI N., DYLOV D. V.: Image quality assessment for magnetic resonance imaging. *IEEE Access 11* (2023), 14154–14168. 6

[LTK22] LI R., TANCIK M., KANAZAWA A.: Nerfacc: A general nerf accleration toolbox. *arXiv preprint arXiv:2210.04847* (2022). 3, 4

[Max95] MAX N.: Optical models for direct volume rendering. *IEEE Transactions on Visualization and Computer Graphics 1*, 2 (1995), 99–108. 2, 4

[MBRS*21] MARTIN-BRUALLA R., RADWAN N., SAJJADI M. S., BARRON J. T., DOSOVITSKIY A., DUCKWORTH D.: Nerf in the wild: Neural radiance fields for unconstrained photo collections. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 7210–7219. 1

[MESK22] MÜLLER T., EVANS A., SCHIED C., KELLER A.: Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (ToG) 41*, 4 (2022), 1–15. 3

[MST*21] MILDENHALL B., SRINIVASAN P. P., TANCIK M., BARRON J. T., RAMAMOORTHI R., NG R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM 65*, 1 (2021), 99–106. 1, 2, 6

[NBM*22] NIEMEYER M., BARRON J. T., MILDENHALL B., SAJJADI M. S., GEIGER A., RADWAN N.: Regnerf: Regularizing neural radiance fields for view synthesis from sparse inputs. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2022), pp. 5480–5490. 3, 10

[PKA*18] PIAYDA K., KLEINEBRECHT L., AFZAL S., BULLENS R., TER HORST I., POLZIN A., VEULEMANS V., DANNENBERG L., WIMMER A. C., JUNG C., ET AL.: Dynamic coronary roadmapping during percutaneous coronary intervention: a feasibility study. *European journal of medical research 23* (2018), 1–7. 1

[RHM*11] RUIJTERS D., HOMAN R., MIELEKAMP P., VAN DE HAAR P., BABIC D.: Validation of 3d multimodality roadmapping in interventional neuroradiology. *Physics in Medicine & Biology 56*, 16 (2011), 5335–5354. 2, 6

[RKY*23] REMONDINO F., KARAMI A., YAN Z., MAZZACCA G., RIGON S., QIN R.: A critical analysis of nerf-based 3d reconstruction. *Remote Sensing 15*, 14 (2023), 3585. 3

[RPLG21] REISER C., PENG S., LIAO Y., GEIGER A.: Kilonerf: Speeding up neural radiance fields with thousands of tiny mlps. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021), pp. 14335–14345. 3

[RWL*22] RÜCKERT D., WANG Y., LI R., IDOUGHI R., HEIDRICH W.: Neat: Neural adaptive tomography. *ACM Transactions on Graphics (TOG) 41*, 4 (2022), 1–13. 2, 3, 4, 10

[THK*11] TU S., HOLM N. R., KONING G., MAENG M., REIBER J. H.: The impact of acquisition angle differences on three-dimensional quantitative coronary angiography. *Catheterization and Cardiovascular Interventions 78*, 2 (2011), 214–222. 4, 7

[TSM*20] TANCIK M., SRINIVASAN P., MILDENHALL B., FRIDOVICH-KEIL S., RAGHAVAN N., SINGHAL U., RAMAMOORTHI R., BARRON J., NG R.: Fourier features let networks learn high frequency functions in low dimensional domains. *Advances in Neural Information Processing Systems 33* (2020), 7537–7547. 2

[TTM*22] TEWARI A., THIES J., MILDENHALL B., SRINIVASAN P., TRETSCHK E., YIFAN W., LASSNER C., SITZMANN V., MARTIN-BRUALLA R., LOMBARDI S., ET AL.: Advances in neural rendering. In *Computer Graphics Forum* (2022), vol. 41, Wiley Online Library, pp. 703–735. 3, 4

[WLY*20] 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 preprint arXiv:2003.11846* (2020). 2, 4

[WWG*21] WANG Q., WANG Z., GENOVA K., SRINIVASAN P. P., ZHOU H., BARRON J. T., MARTIN-BRUALLA R., SNAVELY N., FUNKHOUSER T.: Ibrnet: Learning multi-view image-based rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 4690–4699. 3, 10

[XPMBB21] XIE C., PARK K., MARTIN-BRUALLA R., BROWN M.: Fig-nerf: Figure-ground neural radiance fields for 3d object category modelling. In *2021 International Conference on 3D Vision (3DV)* (2021), IEEE, pp. 962–971. 3

[YCC*14] YANG J., CONG W., CHEN Y., FAN J., LIU Y., WANG Y.: External force back-projective composition and globally deformable optimization for 3-d coronary artery reconstruction. *Physics in Medicine & Biology 59*, 4 (2014), 975. 4

[YCFB*21] YEN-CHEN L., FLORENCE P., BARRON J. T., RODRIGUEZ A., ISOLA P., LIN T.-Y.: inerf: Inverting neural radiance fields for pose estimation. In *2021 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)* (2021), IEEE, pp. 1323–1330. 3

[YYTK21] YU A., YE V., TANCIK M., KANAZAWA A.: pixelnerf: Neural radiance fields from one or few images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition* (2021), pp. 4578–4587. 3, 10

[ZIL*21] ZANG G., IDOUGHI R., LI R., WONKA P., HEIDRICH W.: Intratomo: self-supervised learning-based tomography via sinogram synthesis and prediction. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (2021), pp. 1960–1970. 2, 3

[ZZL22] ZHA R., ZHANG Y., LI H.: Naf: Neural attenuation fields for sparse-view cbct reconstruction. In *Medical Image Computing and Computer Assisted Intervention–MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part VI* (2022), Springer, pp. 442–452. 2, 3, 4, 10

[ZZW*22] 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 ultrasparse 2d projection views: A multicenter study. *Cell Reports Medicine 3*, 10 (2022), 100775. 2, 4