

Research Article

ρ -NeRF: Leveraging Attenuation Priors in Neural Radiance Field for 3D Computed Tomography Reconstruction

Li Zhou¹, Hengyong Yu¹¹. University of Massachusetts Lowell, United States

This paper introduces ρ -NeRF, a self-supervised approach that sets a new standard in novel view synthesis (NVS) and computed tomography (CT) reconstruction by modeling a continuous volumetric radiance field enriched with physics-based attenuation priors. The ρ -NeRF represents a three-dimensional (3D) volume through a fully-connected neural network that takes a single continuous four-dimensional (4D) coordinate—spatial location (x, y, z) and an initialized attenuation value (ρ)—and outputs the attenuation coefficient at that position. By querying these 4D coordinates along X-ray paths, the classic forward projection technique is applied to integrate attenuation data across the 3D space. By matching and refining pre-initialized attenuation values derived from traditional reconstruction algorithms like Feldkamp-Davis-Kress algorithm (FDK) or conjugate gradient least squares (CGLS), the enriched schema delivers superior fidelity in both projection synthesis and image reconstruction, with negligible extra computational overhead. The paper details the optimization of ρ -NeRF for accurate NVS and high-quality CT reconstruction from a limited number of projections, setting a new standard for sparse-view CT applications.

1. Introduction

X-ray imaging, a form of penetrative imaging to capture detailed internal structures, is widely applied across fields, such as medicine, industrial inspection, materials science, *etc.* ^{[1][2]}. While high-resolution projections can produce high-quality CT images, they require extended exposure time, and excessive radiation pose health risks, particularly in medical applications. Consequently, one of the recent research topics focused on sparse-view and low-dose techniques to reduce radiation exposure and/or increase temporal resolution while maintaining diagnostic utility ^{[3][4][5]}, particularly for the situations where there are no slip-ring in the CT scanner. However, in cases like 3D cone beam CT (CBCT) reconstruction, sparse-view data often limits resolution and introduces artifacts compared with full-view data, which highlights the need for advanced computational methods to accurately synthesize views and generate anatomical structure from a limited number of projections.

CT reconstruction techniques can be categorized into analytical, iterative, and hybrid data-driven approaches. Analytical methods, such as filtered backprojection (FBP) ^[6] and its cone-beam variants ^[7], are effective with dense

projection data but fall short under sparse-view conditions. Iterative approaches, including methods like simultaneous algebraic reconstruction technique (SART) family^{[8][9]} and total variation (TV)-minimization techniques^[10], address some of these limitations by optimizing the reconstruction through iterative updates. However, these methods require considerable computational resources and are not always practical for real-time or large-scale applications. Recently, hybrid data-driven approaches that incorporate deep learning have enhanced the traditional techniques by using neural networks to predict, fill gaps, and improve denoising for sparse data^{[11][12]}. While these models achieve impressive results, they often require extensive labeled datasets for training, which limits their adaptability and generalizability to different datasets without further fine-tuning. Given these limitations, self-supervised methods offer a promising alternative by reducing reliance on labeled data and improving adaptability across various reconstruction tasks. Inspired by advancements in 3D reconstruction in computer vision, recent works^{[13][14][15][16]} have explored the applications of self-supervised methods for CT reconstruction, where internal structures are captured through X-ray transmission imaging rather than reflective imaging.

In the computer vision field, 3D reconstruction typically represents shapes as discrete point clouds or meshes. Implicit neural representations (INRs) have become popular as they map discrete points to continuous functions, enabling smoother and more accurate modeling of complex geometries. NeRF^[17], a leading model in this space, leverages INRs for reflective imaging with camera rays, mapping spatial position and viewing direction in 5D coordinates to RGB color and volume density, allowing photorealistic novel view synthesis (NVS) through volumetric rendering. However, NeRF and similar models^{[18][19][20]} are primarily designed for reflective imaging and can not be directly applied to X-ray imaging to capture internal structures. Recent INR-based approaches for CT reconstruction, such as neural attenuation field (NAF)^[15] and SAX-NeRF^[16], adapt NeRF for X-rays by mapping 3D spatial coordinates to attenuation coefficients. While these adaptations improve reconstruction quality, they only consider spatial location, and further enhancements to accuracy come with advanced mapping functions/models which significantly increased computational costs.

To address the limitations of current methods, we present ρ -NeRF framework (see Figure 1), a self-supervised approach designed for X-ray-based 3D tomographic reconstruction. We adopt the CBCT geometry as defined in TIGRE^[21], making ρ -NeRF adaptable across a range of CT imaging settings. Our key contributions are as follows:

- **Redefining Input with Attenuation Priors:** By examining the inherent properties of CT imaging, we redefine the input for implicit neural representations (INRs) as a 4D coordinate system consisting of spatial location (x, y, z) and an initialized attenuation value (ρ) . This enables ρ -NeRF to model a continuous volumetric radiance field, capturing the 3D volume with a fully-connected neural network. By querying these 4D coordinates along X-ray paths and applying forward projection technique, ρ -NeRF is compatible with other NeRF-based CT reconstruction frameworks and achieves enhanced fidelity in sparse-view applications.

- **Efficient Attenuation Initialization:** We introduce interpolation methods—mean, nearest neighbor, and trilinear—to initialize attenuation values (ρ) at a sparse spatial location from its 8 neighboring points on the grids. We use these techniques to match and initialize attenuation values, which are pre-initialized by the traditional algorithms such as FDK and CGLS. Also, a lightweight, learnable linear transformation serves as feature encoding for further refinement. Together, these techniques enhance projection synthesis and image reconstruction with minimal computational cost.
- **Benchmark Validation on X3D Dataset:** Extensive evaluations on the X3D dataset, a large-scale benchmark for X-ray 3D reconstruction, confirm ρ -NeRF’s state-of-the-art performance in both NVS and CT reconstruction across different X-ray applications. This 4D scene representation captures high-resolution geometry and enables high-fidelity anatomical reconstructions, establishing a new benchmark for higher accuracy and efficiency in sparse-view CT applications.

In summary, ρ -NeRF leverages attenuation priors to address the limitations of existing INR-based CT reconstruction methods, achieving notable improvements in performance and fidelity with negligible extra computational demands. This work sets a new benchmark for low-dose, sparse-view CT imaging, advancing the field toward efficient, high-quality 3D X-ray reconstructions.

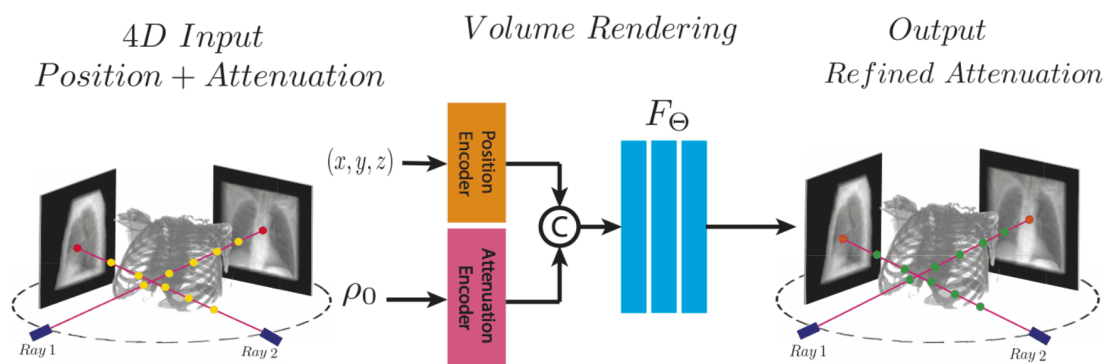


Figure 1. Overview of the ρ -NeRF framework for CT reconstruction. Discrete X-ray sample points integrate attenuation priors ρ_0 with spatial coordinates (x, y, z) . Attenuation values are updated along the X-ray paths by applying forward projection technique with encoders and the mapping function F_Θ .

2. Related Work

2.1. 3D CT Reconstruction

CT reconstruction methods are broadly categorized into analytical, iterative, and hybrid data-driven approaches. Foundational analytical algorithms, such as filtered backprojection (FBP) [6] and its cone-beam extension, Feldkamp-Davis-Kress (FDK) [7], reconstruct attenuation coefficients from projections by solving the Radon

transform. Variants with different filters and Parker weights improve image quality under dense data but produce artifacts under sparse-view conditions. Iterative algorithms address these limitations by framing reconstruction as a maximum a posteriori (MAP) problem, solved through iterative optimization. Gradient-based approaches like the algebraic reconstruction technique (ART) family (SART [9], SIRT [8], and OS-SART [10]) vary in update strategy: SART and SIRT update the full projection set at once, yielding high quality images with high computational cost, and OS-SART strikes a balance with efficient and subset-based updates. Total variation (TV) minimization based methods (e.g., ASD-POCS [22], OS-ASD-POCS [10], and AwASD-POCS [23]) iteratively refine reconstructions to suppress noise, with variants improving computational efficiency or edge preservation. Krylov subspace algorithms (e.g., CGLS [24], LSQR [25]) achieve faster convergence by focusing on the eigenvectors of the residual in descending order, which allows for increased convergence rates compared to the SART family. While the SART-based methods are effective with high-quality projection data, Krylov subspace algorithms, like CGLS, are well-suited for handling large datasets or low-quality data, offering both speed and robustness, especially in CT denoising.

Hybrid data-driven approaches combine the traditional methods with deep learning. These methods enhance reconstruction by using neural networks to 1) predict and fill gaps in limit-view or sparse-view projections [26][27][28], 2) directly infer attenuation coefficients as denoising tasks from training data [29][30][5], and 3) optimize differentiable processes to accelerate computations [31][32][33]. While these models achieve impressive results and require minimal data during inference, they depend on extensive, domain-specific datasets for training, which restricts their adaptability to different applications without further fine-tuning. Consequently, hybrid data-driven models excel in specific contexts but face challenges in generalizability across different CT reconstruction scenarios. Given these limitations, studying self-supervised methods could provide a promising alternative by reducing reliance on labeled data and improving adaptability across different reconstruction tasks.

2.2. Implicit Neural Representation and Rendering

Learning implicit neural representations (INRs) has become a popular approach in 3D scene reconstruction, offering a way to transform discrete data points into continuous functions for learning-based 3D geometry modeling [34][35][36]. Neural rendering leverages INRs to map discrete data to coordinate-based continuous representations, typically using implicit functions parameterized by neural networks [37][38][17]. One prominent method is the neural radiance field (NeRF) [17], a model that has set a standard for high-quality novel view synthesis.

For reflective imaging, NeRF models a scene by mapping spatial position (x, y, z) and viewing direction (θ, ϕ) along camera rays to RGB color (r, g, b) and volume density (σ) . This allows NeRF to produce high-quality novel views by integrating color and densities along rays through volumetric rendering [39]. Numerous works have aimed to improve NeRF's efficiency [40][41][42][43], and extend its application scope, such as generative modeling,

unbounded scenes, and RGB-D synthesis. More recently, INR-based methods have been explored for CT image reconstruction^{[15][44][16]}, transforming discrete samples into continuous representations of internal structures.

For X-ray imaging, adapting NeRF is necessary to address fundamental differences from reflective imaging. Unlike visible light, which reveals surface color and reflectance, X-ray penetrates objects to capture internal structures through attenuation. Neural attenuation fields (NAF)^[15] modify the NeRF framework for 3D CBCT reconstruction by mapping spatial positions (x, y, z) directly to attenuation coefficients (ρ) . Follow-up works^{[45][46][47][48]} refined this method by introducing advanced modeling functions or complex encoding techniques, such as Transformer-based network in SAX-NeRF^[16]. However, these adaptations often increase computational demands. To address this issue, we introduce attenuation priors to INR-based CT reconstruction algorithms, achieving performance improvement with minimal computational overhead.

3. Method

3.1. X-ray Sampling

With appropriate preprocessing steps, the measurement of X-ray attenuation can be approximately modeled by a linear integral, represented as:

$$I(\mathbf{r}) = \int_{t_n}^{t_f} \rho(\mathbf{r}(t))dt, \quad (1)$$

where $\rho(\mathbf{r}(t))$ represents the attenuation coefficient at each point $\mathbf{r}(t)$ parametrized by t along the path, and t_n and t_f are the near and far limits to account for the material's effective attenuation coefficients along the ray path. To align with the voxel grid of view, the ray is divided into evenly spaced bins within the near and far limits, with one point uniformly sampled within each bin. Eq. (1) can be discretized as:

$$I(\mathbf{r}) = \sum_{j=1}^M \rho_j \Delta t, \quad (2)$$

where M is number of points sampled along the ray, ρ_j denotes the attenuation coefficient at the each sampled bin parameterized by t_j , and Δt is the sampling interval length. The objective of tomographic reconstruction is to estimate the distribution of ρ and produce it as a discrete volume, using X-ray projections captured from N different angles.

As shown in Figure 2, we model the CBCT geometry by following the conventional definitions and incorporating a local coordinate system. An X-ray source S rotates around the object along a circular trajectory defined by the rotation angle α . Unlike reflective imaging setups, the object is centered at the origin O of a world coordinate system, and the projections are captured on a flat panel detector positioned opposite to the source. The imaging object is represented as a cube and discretized into voxel units for detailed analysis. This detector records X-ray projections in its own local coordinate system, *i.e.* an image coordinate system.




Figure 2. Geometric configuration of CBCT and X-ray sampling.

Let $\mathbb{I} = \{I_i\}_{i=1, \dots, N}$ be the set of X-ray projections obtained from N different rotation angles. Each projection I_i records the attenuation of X-rays as they traverse the object. We model an X-ray path as $\mathbf{r}(t) = \mathbf{o}_o + t\mathbf{d}_o \in \mathbb{R}^3$, bounded by the near and far limits t_n and t_f . The direction of each ray is encoded by a 3D unit Cartesian vector \mathbf{d}_o and X-ray source position \mathbf{o}_o derived from the geometric configuration between the X-ray source and the detector. Specifically, \mathbf{d}_o and \mathbf{o}_o are calculated based on the rotation angle α , the distance from source to detector (D_{SD}), and the distance from source to object (D_{SO}). At the position of source point \mathbf{o}_o

$$\mathbf{o}_o = (D_{SO} \cdot \cos(\alpha), D_{SO} \cdot \sin(\alpha), 0), \quad (3)$$

we define a local coordinate system (x_s, y_s, z_s) along three unit vectors $(-\cos(\alpha), -\sin(\alpha), 0)$, $(-\sin(\alpha), \cos(\alpha), 0)$ and $(0, 0, -1)$. We assume that the detector is perpendicular to x_s -axis. Considering a local detector coordinator system (u, v) with u -axis parallel to the $x_o y_o$ -plane and y_s -axis and v -axis parallel to z_o -axis and z_s , $(0, 0)$ is the projection position of the X-ray source. The transformation of the ray direction within the global coordinate system can be expressed through an affine transformation:

$$\mathbf{d}_I = \frac{1}{\sqrt{u^2 + v^2 + (D_{SD})^2}}(-D_{SD}, u, -v), \quad \mathbf{d}_o = R \cdot \mathbf{d}_I, \quad (4)$$

where \mathbf{d}_I is ray's direction for $\alpha = 0$, and R is the rotation matrix associated with the source rotation angle α

$$R = \begin{pmatrix} \cos(\alpha) & \sin(\alpha) & 0 \\ -\sin(\alpha) & \cos(\alpha) & 0 \\ 0 & 0 & 1 \end{pmatrix}. \quad (5)$$

3.2. Attenuation Priors Scene Representation

We follow the conventional idea of NeRF^[17], adapting it for X-ray imaging by using attenuation values instead of color and density. The following pipeline, illustrated in Figure 1, details our approach to integrate attenuation priors into a neural radiance field.

Attenuation Modeling

We model tomographic images as a continuous 4D attenuation field, represented as a function:

$$F_{\Theta} : (x, y, z, \rho_0) \rightarrow \rho, \quad (6)$$

where (x, y, z) is a 3D spatial coordinate, ρ_0 represents an initialized attenuation value, and ρ denotes the learned attenuation coefficient. The neural network, parameterized by Θ , is optimized to map each 4D input to its corresponding attenuation value, yielding a continuous representation of the attenuation field.

Attenuation Initialization

The initialization of attenuation priors begins with the traditional reconstruction algorithms, such as FDK and CGLS, to produce initial attenuation estimates, ρ_0 , for the 3D attenuation map in the global coordinate system. While these voxel-based initial values provide a structured foundation, further refinement is needed to achieve a continuous, high-fidelity reconstruction. To this end, we employ three interpolation methods—mean, nearest neighbor, and trilinear—to smooth and adapt attenuation values at points where X-ray paths intersect the voxel grid, as shown in Figure 2. The attenuation value at a sampled point is determined using 8 neighboring voxel vertices via their mean, nearest neighbor, or trilinear interpolation. This initialization pipeline offers robust priors that support accurate projection synthesis and high-quality reconstruction from any view.

Feature Encoding

Previous studies^{[49][50]} have shown that deep networks tend to favor low-frequency representations, limiting their ability to capture fine details in color and geometry. By mapping low-dimensional inputs into a higher-dimensional space with high-frequency encodings, neural networks better capture these high-frequency variations. Given that X-ray imaging naturally features homogeneous tissue regions with sharp boundaries at anatomical transitions, resulting low variation of the image, we adopt a learning-based encoding mechanism to leverage these properties. For spatial positions, we use a *hash encoder*^[51]—a sparse, learning-based encoding method—to efficiently represent position details. For attenuation values, we apply a lightweight *linear transformation* to refine the input for enhanced accuracy in reconstruction.

Model Optimization

The optimization process of ρ -NeRF involves minimizing the difference between the predicted and actual X-ray projections across the dataset. The network is trained to predict attenuation coefficients for each 4D input (spatial position and initial attenuation) by using an objective function based on the mean squared error (MSE) between the synthesized projections and the ground truth:

$$Loss(\Phi, \Psi, A) = \sum_{\mathbf{r} \in \mathbf{R}} \|I(\mathbf{r}) - \hat{I}(\mathbf{r})\|^2, \quad (7)$$

where $I(\mathbf{r})$ and $\hat{I}(\mathbf{r})$ are real and synthesized projections from ray \mathbf{r} respectively, and \mathbf{R} is a batch of rays. Φ is the learned hash encoder for spatial location (x, y, z) , Ψ is the learned linear encoder for initial attenuation ρ_0 , and A is the attenuation coefficients network. $\hat{I}(\mathbf{r})$ is computed by Eq. (2) where ρ_j is estimated by the network A and the related encoders Φ and Ψ . This is implemented by the classic forward projection technique that corresponds to the rendering model in the computer vision field. The CT reconstruction process is achieved by feeding the voxel grid coordinates, sized to the desired dimensions, into the trained neural function F_Θ to predict the corresponding attenuation coefficients. In our work, we use the MLP model architecture from [17] as the foundation for our study. Additionally, we explore a transformer-based model that incorporates X-ray structural awareness [16] to further validate the applicability of our proposed method. The final output is represented as a discrete 3D matrix.

4. Experiments and Results

4.1. Experimental Setup

Following novel Nerf-based works [15][16], we evaluate our method using publicly available human organ CT datasets, including LIDC-IDRI [52] and scientific visualization dataset [53]. Using the open-source tomographic toolbox TIGRE [21], we simulate 100 projections per case with 3% noise over a 0° – 180° range. For the novel view synthesis (NVS) task, we split these projections evenly, with 50 for training and 50 for testing. For sparse-view CT reconstruction, the CT volumes are reserved as ground truths for evaluation.

The ρ -NeRF framework is implemented in PyTorch and trained on an NVIDIA RTX 3090 GPU. We optimize the model using the Adam optimizer ($\beta_1 = 0.9, \beta_2 = 0.999$) for at most 3,000 iterations, with an initial learning rate of 0.001, reduced to 0.0001 halfway through training. For efficient processing, we use a batch size of 1,024 rays, with each ray sampled at 192 points to adequately represent the attenuation field across the CT volume. After training, we quantitatively evaluate the reconstruction by computing both peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) metrics. PSNR (in dB) provides a statistical measure of artifact suppression performance, while SSIM assesses perceptual differences between two images. Higher PSNR and SSIM values indicate more accurate reconstructions.

4.2. Main Results

We include comparisons with the conventional CT reconstruction algorithms, including FDK (physics-based), AS-POCS (iterative), and CGLS (Krylov subspace), as well as the state-of-the-art NeRF-based methods, named NAF and SAX-NeRF. We deploy the ρ -NeRF framework to NeRF-based models by utilizing their neural networks as the mapping function F_Θ , naming the enhanced models ρ -NAF and ρ -SAX-NeRF. These new models integrate attenuation priors, the nearest neighbor interpolation method, and the two proposed feature encoders, demonstrating the effectiveness and adaptability of our approach. Unless otherwise noted, this configuration is

used throughout the experiments. A detailed analysis of these design choices is provided in 4.3, with additional quantitative and qualitative results available in the supplementary materials.

3D CT Reconstruction

Table 1 and Figure 3 provide quantitative and visual comparisons of 3D CT reconstruction across baseline and ρ -NeRF-enhanced models. The ρ -based models consistently outperform their counterparts across all cases. Specifically, SSIM values improve from 0.93 to 0.95 in foot reconstructions when comparing NAF and ρ -NAF, also it surpasses SAX-NeRF-based models while maintaining a negligible computational cost, which will be discussed in Sec. 4.3. In chest scans, ρ -NAF achieves a PSNR of 35.07 dB, compared to 34.46 dB with SAX-NeRF, underscoring the positive impact of leveraging attenuation priors. These results validate that ρ -NeRF effectively reduces artifacts and yields more precise reconstructions, especially in sparse-view conditions, emphasizing the benefits of refined attenuation initialization.

| | FDK | ASD-POCS | CGLS | NAF | ρ -NAF | SAX-NeRF | ρ -SAX-NeRF |
|-----------------|--------------|--------------|--------------|--------------|---------------------|--------------|---------------------|
| Jaw | 28.57/0.7816 | 33.23/0.9322 | 28.80/0.8536 | 34.08/0.9345 | 35.23/0.9512 | 35.30/0.9503 | 35.25/0.9522 |
| Foot | 24.53/0.5999 | 29.98/0.9208 | 25.39/0.7901 | 31.76/0.9341 | 32.84/0.9506 | 32.25/0.9403 | 32.72/0.9481 |
| Chest | 22.89/0.7861 | 31.13/0.9422 | 24.44/0.7721 | 32.65/0.9609 | 35.07/0.9733 | 34.46/0.9725 | 34.81/0.9747 |
| Aneurism | 28.07/0.7295 | 34.71/0.9858 | 30.08/0.9272 | 38.02/0.9883 | 39.43/0.9928 | 41.46/0.9956 | 41.41/0.9966 |
| Bonsai | 24.53/0.7275 | 32.71/0.9529 | 26.27/0.8200 | 34.07/0.9606 | 34.79/0.9705 | 36.11/0.9751 | 36.85/0.9798 |

Table 1. Quantitative comparison of CT reconstruction from different methods in terms of PSNR/SSIM, where **bold** fonts indicate the best results.




Figure 3. Visualization of some representative CT reconstruction results for Jaw, Foot, Chest, Aneurism, and Bonsai (from top to bottom). Zoomed region-of-interests are also provided.

Novel View Synthesis

In Table 2, quantitative evaluation results of novel view synthesis (NVS) show that ρ -NAF and ρ -SAX-NeRF provide substantial improvements over approaches adopting only spatial location as inputs, and promising the effectiveness of attenuation priors mechanism. For instance, ρ -SAX-NeRF achieves a PSNR of 47.66 dB and an SSIM of 0.9996 on chest views, compared to SAX-NeRF's 47.41 dB and 0.9994, respectively. Figure 4 visually demonstrates the sharper anatomical details rendered by ρ -based models, effectively capturing fine structures. This performance, combined with supplementary results, highlights the capability of ρ -NeRF in delivering visually accurate and consistent NVS results across various anatomical regions.

| | NAF | ρ -NAF | SAX-NeRF | ρ -SAX-NeRF |
|----------|--------------|--------------|--------------|---------------------|
| Jaw | 40.10/0.9988 | 43.18/0.9991 | 42.89/0.9991 | 43.37/0.9991 |
| Foot | 38.13/0.9924 | 46.47/0.9994 | 46.64/0.9994 | 47.56/0.9994 |
| Chest | 42.26/0.9993 | 47.62/0.9996 | 47.41/0.9994 | 47.66/0.9996 |
| Aneurism | 41.03/0.9991 | 46.84/0.9996 | 52.90/0.9998 | 53.78/0.9998 |
| Bonsai | 49.12/0.9990 | 51.64/0.9992 | 52.29/0.9995 | 53.44/0.9995 |

Table 2. Quantitative comparison of novel view synthesis from different methods in terms of PSNR/SSIM, where **bold** fonts indicate the best results.

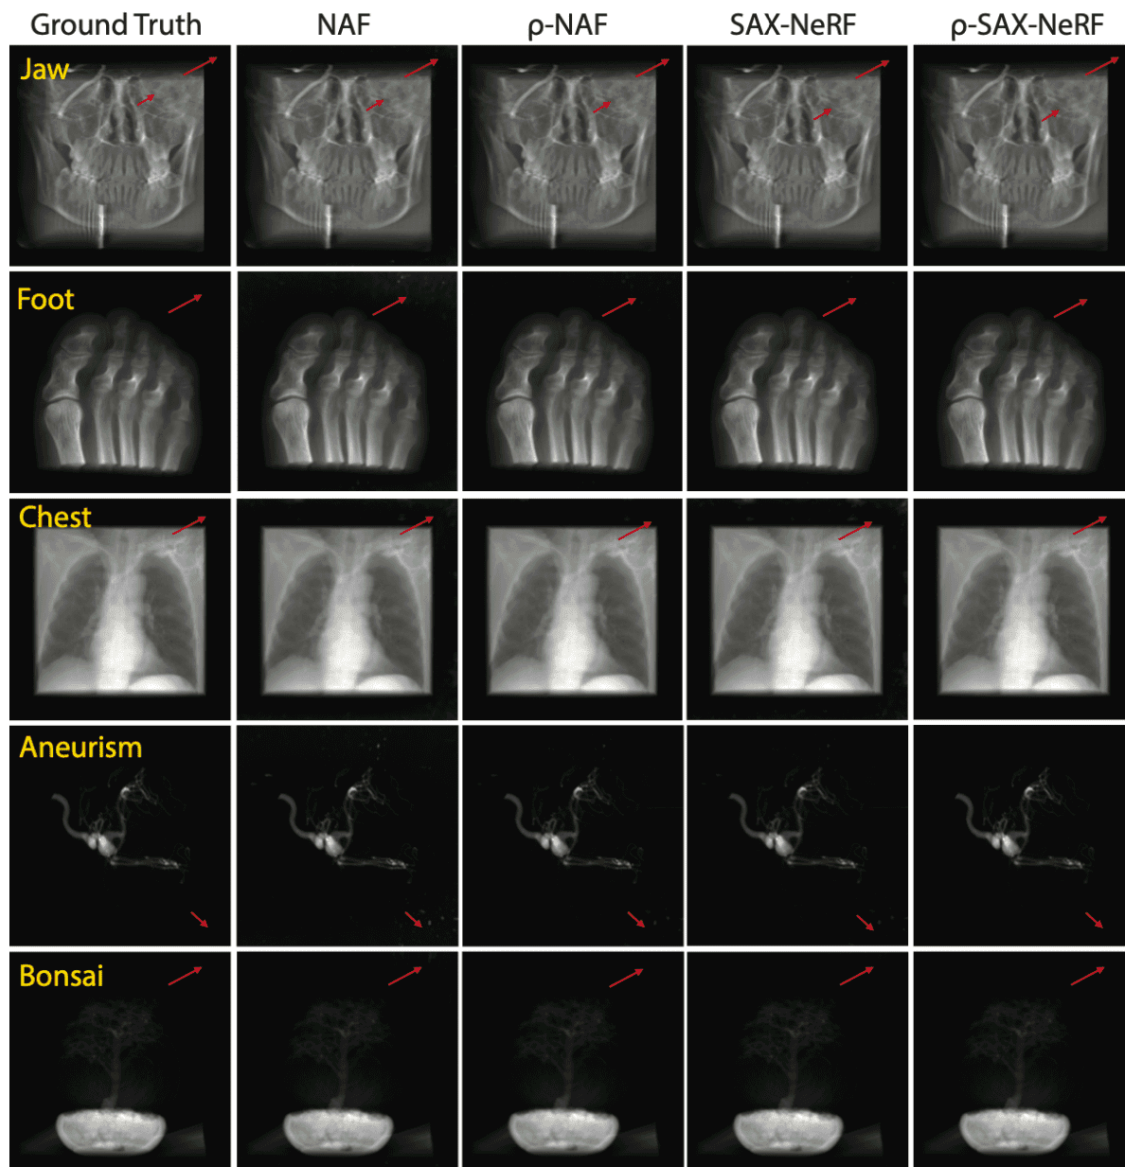


Figure 4. Visualization of some representative novel view results for Jaw, Foot, Chest, Aneurism, and Bonsai (from top to bottom).

4.3. Ablation Study

Attenuation Initialization Analysis: FDK is our primary choice for initialization in the pipeline due to its ability to retain crucial attenuation information, providing a solid baseline for the ρ -based models. We evaluate three traditional initialization algorithms—FDK, ASD-POCS, and CGLS—to estimate initial attenuation coefficients for the ρ -NeRF models. As shown in Table 3, FDK consistently achieves high SSIM and PSNR scores, confirming its effectiveness for initializing attenuation values. FDK, being a physics-based algorithm, preserves more attenuation details, whereas iterative methods with TV-regularization, such as ASD-POCS and CGLS, may result in overly

smooth reconstructions, potentially losing physical details. For these reasons, we primarily use FDK for initialization in our main experiments.

| ρ -NAF | FDK | ASD-POCS | CGLS |
|-------------|---------------------|---------------------|--------------|
| Jaw | 35.23/0.9512 | 34.48/0.9420 | 34.87/0.9466 |
| Foot | 32.84/0.9506 | 32.84/0.9507 | 32.20/0.9408 |
| Chest | 35.07/0.9733 | 33.39/0.9689 | 33.13/0.9654 |
| Aneurism | 41.46/0.9956 | 41.52/0.9973 | 39.43/0.9928 |
| Bonsai | 36.11/0.9751 | 35.31/0.9656 | 34.23/0.9627 |

Table 3. Ablation study results of attenuation initialization algorithms in terms of PSNR/SSIM, where bold fonts indicate the best results.

Model Complexity Analysis

Table 4 compares the computational cost and parameter size for baseline models (NAF and SAX-NeRF) against their ρ -based counterparts (ρ -NAF and ρ -SAX-NeRF). The ρ -based models introduce minimal increases in computational cost, with only 0.19 giga multiply-add operations per second (GMACs) increase. Parameter counts remain nearly identical across versions, with a negligible increase of less than 0.1%. Combined the results in Tables 1 and 2, one can see that our ρ -based models achieve higher performance with minimal increase in computational cost.

| | NAF | ρ -NAF | SAX-NeRF | ρ -SAX-NeRF |
|------------|---------|-------------|----------|------------------|
| GMACs | 0.7559 | 0.9463 | 12.6273 | 12.8177 |
| Params (M) | 14.2667 | 14.2677 | 14.3252 | 14.3262 |

Table 4. Comparison analysis of GMACs and parameters for different models.

Convergence Analysis

Figure 5 shows the loss over training steps for four models: NAF, ρ -NAF, SAX-NeRF, and ρ -SAX-NeRF. The vertical axis represents the loss in log scale, and the horizontal axis indicates the training steps up to 70,000. The inset

zooms highlight the behavior of each model. Models that incorporate attenuation priors (ρ -NAF and ρ -SAX-NeRF) exhibit faster convergence and maintain lower loss values throughout the training compared to the counterparts without these priors (NAF and SAX-NeRF). This suggests that leveraging attenuation priors accelerates the convergence process, allowing the models to reach a stable and lower loss more quickly than those without the priors.




Figure 5. Model convergence with respect to training steps.

5. Conclusion

This work introduces ρ -NeRF, a self-supervised neural framework designed for high-fidelity 3D CT reconstruction from sparse-view X-ray data. By leveraging attenuation priors with a continuous 4D representation, ρ -NeRF advances the performance of CT imaging while reducing reliance on dense data or supervised learning. Key contributions include a novel input schema that integrates initialized attenuation priors, effective feature encoding methods, and robust interpolation techniques. Evaluations on the X3D dataset validate ρ -NeRF's state-of-the-art performance in novel view synthesis and CT reconstruction. This framework not only sets a new benchmark in efficiency and accuracy for sparse-view CT applications but also offers potential for adaptation across a range of Nerf-based methods on CT reconstruction, contributing to advancements in medical and industrial CT imaging.

Supplementary Material

6. Attenuation Initialization Analysis

We present additional experiments to extend the analysis of attenuation initialization strategies discussed in Sec. 4.3 of the main paper. These experiments focus on two key aspects: the role of pre-initialization algorithms and the impact of different interpolation methods. By providing further insights into these components with both quantitative results and qualitative visualizations, we aim to highlight their impacts on the overall CT reconstruction quality and validate the robustness of leveraging attenuation priors.

Pre-initialization Algorithms Analysis

The first part of this analysis, as demonstrated in Table 5 and Figure 6, evaluates the performance of different pre-initialization algorithms—FDK, CGLS, and ASD-POCS—using the nearest interpolation method. The findings reveal that FDK consistently provides superior initialization for both ρ -NAF and ρ -SAX-NeRF models, achieving the best balance between detail preservation and reconstruction quality. While the ASD-POCS excels in some specific scenarios, its smoothing tendency often reduces fine detail accuracy, making FDK the preferred choice for reliable initialization.

| match mode | ρ -NAF | | | ρ -SAX-NeRF | | |
|------------|---------------------|--------------|---------------------|---------------------|--------------|--------------|
| | FDK | CGL | SD-POCS | FDK | CGLS | ASD-POCS |
| = Nearest | | | | | | |
| Jaw | 35.23/0.9512 | 34.87/0.9466 | 34.48/0.9420 | 35.25/0.9522 | 35.20/0.9509 | 34.93/0.9468 |
| Foot | 32.84/0.9506 | 32.20/0.9408 | 32.84/0.9507 | 32.72/0.9481 | 32.23/0.9411 | 32.29/0.9422 |
| Chest | 35.07/0.9733 | 33.13/0.9654 | 33.39/0.9689 | 34.81/0.9747 | 34.76/0.9746 | 34.82/0.9747 |
| Aneurism | 39.43/0.9928 | 39.43/0.9928 | 41.52/0.9973 | 41.41/0.9966 | 41.35/0.9960 | 41.00/0.9958 |
| Bonsai | 34.79/0.9705 | 34.23/0.9627 | 35.31/0.9656 | 36.85/0.9798 | 36.70/0.9780 | 36.79/0.9771 |

Table 5. Quantitative comparison of CT reconstruction results with different pre-initialization algorithms using the nearest interpolation method in terms of PSNR/SSIM, where **bold** fonts indicate the best results.




Figure 6. Visualization of representative results from different pre-initialization algorithms for Jaw, Foot, Chest, Aneurism, and Bonsai (from top to bottom). Zoomed region-of-interests are also provided.

Interpolation Methods Analysis

The second part of this analysis, shown in Table 6 and Figure 7, investigates the impact of interpolation methods—Nearest, Mean, and Trilinear—on reconstruction performance when using FDK as the pre-initialization algorithm.

These experiments evaluate ρ -NAF and ρ -SAX-NeRF models across the same five cases. Nearest interpolation generally demonstrates better performance due to its ability to preserve sharp boundaries, particularly in simpler cases like Foot and Chest. Trilinear interpolation, however, provides competitive or superior results in more complex scenarios such as Aneurism and Bonsai, benefiting from smoother transitions across voxel grids. This analysis underscores the adaptability of interpolation techniques and their critical role in optimizing reconstruction outcomes.

| image_init | ρ -NAF | | | ρ -SAX-NeRF | | |
|------------|---------------------|--------------|--------------|---------------------|--------------|---------------------|
| | Nearest | Mean | Trilinear | Nearest | Mean | Trilinear |
| = FDK | | | | | | |
| Jaw | 35.23/0.9512 | 34.53/0.9421 | 34.85/0.9460 | 35.25/0.9522 | 34.95/0.9477 | 34.89/0.9468 |
| Foot | 32.84/0.9506 | 32.40/0.9488 | 32.56/0.9498 | 32.72/0.9481 | 32.53/0.9480 | 32.73/0.9507 |
| Chest | 35.07/0.9733 | 32.59/0.9605 | 33.49/0.9644 | 34.81/0.9747 | 34.71/0.9741 | 34.80/0.9749 |
| Aneurism | 39.43/0.9928 | 36.98/0.9899 | 38.59/0.9909 | 41.41/0.9966 | 40.67/0.9956 | 42.05/0.9962 |
| Bonsai | 34.79/0.9705 | 35.06/0.9710 | 35.05/0.9720 | 36.85/0.9798 | 36.55/0.9782 | 36.95/0.9805 |

Table 6. Quantitative comparison of representative results with different interpolation methods using the FDK pre-initialization algorithm in terms of PSNR/SSIM, where **bold** fonts indicate the best results.




Figure 7. Visualization of representative results from different interpretation methods for Jaw, Foot, Chest, Aneurism, and Bonsai (from top to bottom). Zoomed region-of-interests are also provided.

References

1. [^]Maire E, Buffiere JY, Salvo L, Blandin JJ, Ludwig W, Letang JM (2001). "On the Application of X-ray Microtomography in the Field of Materials Science". *Advanced Engineering Materials*. 3 (8): 539–546.
2. [^]Singh AK. *Advanced x-ray techniques in research and industry*. IOS press; 2005.
3. [^]Zang G, Aly M, Idoughi R, Wonka P, Heidrich W (2018). "Super-resolution and sparse view CT reconstruction". In: *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 137–153.
4. [^]Morovati B, Lashgari R, Hajihassani M, Shabani H (2023). "Reduced deep convolutional activation features (r-decaf) in histopathology images to improve the classification performance for breast cancer diagnosis". *Journal of Digital Imaging*. 36 (6): 2602–2612.
5. ^{a, b}Wang D, Han S, Xu Y, Wu Z, Zhou L, Morovati B, Yu H (2024). "LoMAE: Simple Streamlined Low-level Masked Autoencoders for Robust, Generalized, and Interpretable Low-dose CT Denoising". *IEEE Journal of Biomedical and Health Informatics*. IEEE.
6. ^{a, b}Herman GT. *Fundamentals of computerized tomography: image reconstruction from projections*. Springer Science & Business Media; 2009.
7. ^{a, b}Feldkamp LA, Davis LC, Kress JW (1984). "Practical cone-beam algorithm". *Josa a*. 1 (6): 612–619.
8. ^{a, b}Gilbert P (1972). "Iterative methods for the three-dimensional reconstruction of an object from projections". *Journal of theoretical biology*. 36 (1): 105–117.
9. ^{a, b}Andersen AH, Kak AC (1984). "Simultaneous algebraic reconstruction technique (SART): a superior implementation of the ART algorithm". *Ultrasonic Imaging*. 6 (1): 81–94.
10. ^{a, b, c}Zeng GL. *Medical image reconstruction*. Springer; 2010.
11. [^]Morovati B, Li M, Han S, Xu Y, Zhou L, Wang G, Yu H. "Impact of network architecture and training strategy for photon counting CT data correction." In: *Developments in X-Ray Tomography XV*. SPIE; 2024. p. 131520D.
12. [^]Han S, Xu Y, Wang D, Morovati B, Zhou L, Maltz JS, Wang G, Yu H (2024). "Physics-informed score-based diffusion model for limited-angle reconstruction of cardiac computed tomography". *arXiv preprint arXiv:2405.14770*.
13. [^]Adler J, Öktem O (2018). "Learned primal-dual reconstruction". *IEEE transactions on medical imaging*. 37 (6): 1322–1332.
14. [^]Zang G, Idoughi R, Li R, Wonka P, Heidrich W (2021). "Intratomo: self-supervised learning-based tomography via sinogram synthesis and prediction". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 1960–1970.
15. ^{a, b, c, d, e}Zha R, Zhang Y, Li H. "NAF: neural attenuation fields for sparse-view CBCT reconstruction." In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer; 2022. p. 442–452.
16. ^{a, b, c, d, e, f}Cai Y, Wang J, Yuille A, Zhou Z, Wang A. "Structure-aware sparse-view x-ray 3d reconstruction." In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024. p. 11174–11183.

17. ^{a, b, c, d, e}Mildenhall B, Srinivasan PP, Tancik M, Barron JT, Ramamoorthi R, Ng R (2021). "Nerf: Representing scenes as neural radiance fields for view synthesis". *Communications of the ACM*. 65 (1): 99–106.
18. [△]Pumarola A, Corona E, Pons-Moll G, Moreno-Noguer F (2021). "D-nerf: Neural radiance fields for dynamic scenes". *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 10318–10327.
19. [△]Wang Z, Wu S, Xie W, Chen M, Prisacariu VA (2021). "NeRF--: Neural radiance fields without known camera parameters". *arXiv preprint arXiv:2102.07064*. [arXiv:2102.07064](https://arxiv.org/abs/2102.07064).
20. [△]Tancik M, Casser V, Yan X, Pradhan S, Mildenhall B, Srinivasan PP, Barron JT, Kretzschmar H (2022). "Block-nerf: Scalable large scene neural view synthesis". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. pp. 8248–8258.
21. ^{a, b}Biguri A, Dosanjh M, Hancock S, Soleimani M (2016). "TIGRE: a MATLAB-GPU toolbox for CBCT image reconstruction". *Biomedical Physics & Engineering Express*. 2 (5): 055010.
22. [△]Defrise M, Noo F, Clackdoyle R, Kudo H (2006). "Truncated Hilbert transform and image reconstruction from limited tomographic data". *Inverse problems*. 22 (3): 1037.
23. [△]Sidky EY, Pan X (2008). "Image reconstruction in circular cone-beam computed tomography by constrained, total-variation minimization". *Physics in Medicine & Biology*. 53 (17): 4777.
24. [△]Dabravolski A, Batenburg KJ, Sijbers J (2014). "Dynamic angle selection in x-ray computed tomography". *Nuclear Instruments and Methods in Physics Research Section B: Beam Interactions with Materials and Atoms*. 324: 17–24.
25. [△]Flores L, Vidal V, Parceros E, Verdú G. Application of a modified LSQR method for CT imaging reconstruction with low doses to patient. In: *2016 9th International Congress on Image and Signal Processing, BioMedical Engineering and Informatics (CISP-BMEI)*. IEEE; 2016. p. 1969–1974.
26. [△]Anirudh R, Kim H, Thiagarajan JJ, Mohan KA, Champley K, Bremer T (2018). "Lose the views: Limited angle CT reconstruction via implicit sinogram completion". *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 6343--6352.
27. [△]Li Z, Cai A, Wang L, Zhang W, Tang C, Li L, Liang N, Yan B (2019). "Promising generative adversarial network based sinogram inpainting method for ultra-limited-angle computed tomography imaging". *Sensors*. 19 (18): 3941.
28. [△]Jiang L, Zhang M, Wei R, Liu B, Bai X, Zhou F (2021). "Reconstruction of 3d ct from a single x-ray projection view using CVAE-GAN". In: *2021 IEEE International Conference on Medical Imaging Physics and Engineering (ICMIPE)*. IEEE. pp. 1–6.
29. [△]Lahiri A, Maliakal G, Klasky ML, Fessler JA, Ravishankar S (2023). "Sparse-view cone beam CT reconstruction using data-consistent supervised and adversarial learning from scarce training data". *IEEE Transactions on Computational Imaging*. 9: 13–28.
30. [△]Zhou L, Wang D, Xu Y, Han S, Morovati B, Fan S, Yu H. "Gradient Guided Co-Retention Feature Pyramid Network for LDCT Image Denoising." In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer; 2024. p. 153–163.

31. [△]Shen C, Gonzalez Y, Chen L, Jiang SB, Jia X (2018). "Intelligent parameter tuning in optimization-based iterative CT reconstruction via deep reinforcement learning". *IEEE Transactions on Medical Imaging*. 37 (6): 1430–1439.
32. [△]Yu Z, Thibault JB, Bouman CA, Sauer KD, Hsieh J (2010). "Fast model-based X-ray CT reconstruction using spatially nonhomogeneous ICD optimization". *IEEE Transactions on Image Processing*. 20 (1): 161–175.
33. [△]Xu C, Yang B, Guo F, Zheng W, Poignet P (2020). "Sparse-view CBCT reconstruction via weighted Schatten p-norm minimization". *Optics Express*. 28 (24): 35469–35482.
34. [△]Genova K, Cole F, Vlastic D, Sarna A, Freeman WT, Funkhouser T (2019). "Learning shape templates with structured implicit functions". *Proceedings of the IEEE/CVF international conference on computer vision*. 7154--7164.
35. [△]Mescheder L, Oechsle M, Niemeyer M, Nowozin S, Geiger A (2019). "Occupancy networks: Learning 3d reconstruction in function space". *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pages 4460–4470.
36. [△]Park JJ, Florence P, Straub J, Newcombe R, Lovegrove S (2019). "Deepsdf: Learning continuous signed distance functions for shape representation". In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 165–174.
37. [△]Sitzmann V, Thies J, Heide F, Nießner M, Wetzstein G, Zollhofer M. Deepvoxels: Learning persistent 3d feature embeddings. In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2019. p. 2437–2446.
38. [△]Lombardi S, Simon T, Saragih J, Schwartz G, Lehrmann A, Sheikh Y (2019). "Neural volumes: Learning dynamic renderable volumes from images". *arXiv preprint arXiv:1906.07751*. Available from: <https://arxiv.org/abs/1906.07751>.
39. [△]Kajiya JT, Von Herzen BP (1984). "Ray tracing volume densities". *ACM SIGGRAPH computer graphics*. 18 (3): 165–174.
40. [△]Barron JT, Mildenhall B, Tancik M, Hedman P, Martin-Brualla R, Srinivasan PP. "Mip-nerf: A multiscale representation for anti-aliasing neural radiance fields." In: *Proceedings of the IEEE/CVF international conference on computer vision*. 2021. p. 5855–5864.
41. [△]Barron JT, Mildenhall B, Verbin D, Srinivasan PP, Hedman P. "Mip-nerf 360: Unbounded anti-aliased neural radiance fields." In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2022. p. 5470–5479.
42. [△]Chen A, Xu Z, Geiger A, Yu J, Su H. Tensorf: Tensorial radiance fields. In: *European conference on computer vision*. Springer; 2022. p. 333–350.
43. [△]Chen Z, Li Z, Song L, Chen L, Yu J, Yuan J, Xu Y (2023). "Neurbf: A neural fields representation with adaptive radial basis functions". In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 4182–4194.
44. [△]Zhang Y, Shao HC, Pan T, Mengke T (2023). "Dynamic cone-beam CT reconstruction using spatial and temporal implicit neural representation learning (STINR)". *Physics in Medicine & Biology*. 68 (4): 045005.
45. [△]Chen Z, Yang L, Lai JH, Xie X (2023). "CuNeRF: Cube-based neural radiance field for zero-shot medical image arbitrary-scale super resolution". *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 2023: 21185–

21195.

46. [△]Hu J, Fan Q, Hu S, Lyu S, Wu X, Wang X. "UMedNeRF: Uncertainty-aware single view volumetric rendering for medical neural radiance fields." In: 2024 IEEE International Symposium on Biomedical Imaging (ISBI). IEEE; 2024. p. 1-4.
47. [△]Shao HC, Mengke T, Pan T, Zhang Y (2024). "Dynamic CBCT imaging using prior model-free spatiotemporal implicit neural representation (PMF-STINR)". *Physics in Medicine & Biology*. 69 (11): 115030.
48. [△]Liu Z, Fang Y, Li C, Wu H, Liu Y, Shen D, Cui Z (2024). "Geometry-Aware Attenuation Learning for Sparse-View CBCT Reconstruction". *IEEE Transactions on Medical Imaging*. IEEE.
49. [△]Hornik K, Stinchcombe M, White H (1989). "Multilayer feedforward networks are universal approximators". *Neural networks*. 2 (5): 359-366.
50. [△]Rahaman N, Baratin A, Arpit D, Draxler F, Lin M, Hamprecht F, Bengio Y, Courville A. "On the spectral bias of neural networks." In: *International conference on machine learning*. PMLR; 2019. p. 5301-5310.
51. [△]Müller T, Evans A, Schied C, Keller A (2022). "Instant neural graphics primitives with a multiresolution hash encoding". *ACM Transactions on Graphics (TOG)*. 41 (4): 1-15.
52. [△]Armato SG III, McLennan G, Bidaut L, McNitt-Gray MF, Meyer CR, Reeves AP, Zhao B, Aberle DR, Henschke CI, Hoffman EA, et al. The lung image database consortium (LIDC) and image database resource initiative (IDRI): a complete reference database of lung nodules on CT scans. *Medical physics*. 38(2):915--931, 2011.
53. [△]Klacansky P. *Scientific visualization datasets* (2022) [Internet]. 2022. Available from: <https://klacansky.com/open-scivis-datasets/>.

Declarations

Funding: No specific funding was received for this work.

Potential competing interests: No potential competing interests to declare.