
Differentiable Voxel-based X-ray Rendering Improves Sparse-View 3D CBCT Reconstruction

Mohammadhossein Momeni^{1,2*} Vivek Gopalakrishnan^{1,2*}

Neel Dey² Polina Golland² Sarah Frisken¹

¹Brigham and Women’s Hospital ²MIT CSAIL

Abstract

We present DiffVox, a self-supervised framework for Cone-Beam Computed Tomography (CBCT) reconstruction by directly optimizing a voxelgrid representation using physics-based differentiable X-ray rendering. Further, we investigate how the different implementations of the X-ray image formation model in the renderer affect the quality of 3D reconstruction and novel view synthesis. When combined with our regularized voxel-based learning framework, we find that using an exact implementation of the discrete Beer-Lambert law for X-ray attenuation in the renderer outperforms both widely used iterative CBCT reconstruction algorithms and modern neural field approaches, particularly when given only a few input views. As a result, we reconstruct high-fidelity 3D CBCT volumes from fewer X-rays, potentially reducing ionizing radiation exposure and improving diagnostic utility.

1 Introduction

CBCT reconstruction estimates internal 3D structure from multiple 2D X-ray projection images, typically acquired in a circular orbit about the subject (Fig. 1A). Current analytical and iterative solvers for CBCT reconstruction [1–3] fail when reconstructing volumes from a limited number of 2D X-rays, typically prescribed to minimize radiation exposure to patients. To this end, several neural field-based methods aim to fill the need for sparse-view CBCT reconstruction [4–6]. However, these methods rely on computationally intensive Transformer- or MLP-hashgrid hybrid architectures [7], which introduce numerous additional hyperparameters and substantially increase reconstruction times. Furthermore, previous neural fields for CBCT reconstruction have predominantly been evaluated on synthetic X-ray datasets [8], making it unclear if these architectures scale to real X-ray images. Despite their high performance on synthetic datasets, neural field-based CBCT reconstruction methods have not seen widespread adoption in real-world tasks.

Instead of parameterizing the unknown volume as a neural network as in previous work [4–6], we propose to perform sparse-view reconstruction by directly optimizing a regularized voxelgrid representation [9] of a CBCT scan. DiffVox, our self-supervised reconstruction framework, has few tunable hyperparameters and is driven by a physics-based differentiable X-ray renderer [10], enabling the fusion of popular voxel-based regularizers (e.g., total variation norm [11]) with gradient-based optimizers (e.g., Adam [12]). We further investigate how different X-ray image formation models affect the quality of CBCT reconstruction: Siddon’s method [13], an exact model of the first-order effects in X-ray physics, and trilinear interpolation, a fast approximation of these phenomena. Whereas most existing reconstruction techniques use trilinear interpolation, we find that optimization with Siddon’s method leads to the highest quality reconstructions and renderings in the sparse-view regime. We evaluate our methods, and many traditional and neural field baselines, on an open dataset of more than 150,000 raw X-ray images from 42 subjects and find that differentiable voxelgrid optimization with DiffVox outperforms existing CBCT reconstruction algorithms.

*Equal contribution. Implementation available at <https://github.com/hosseini-momeni/DiffVox>.

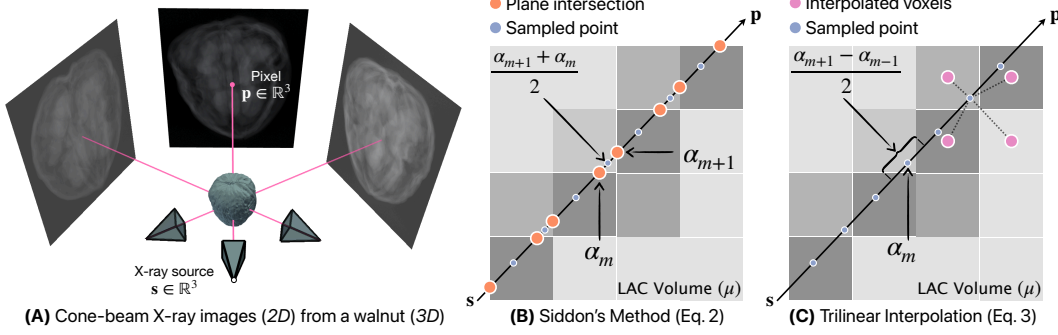


Figure 1: (A) In sparse-view CBCT reconstruction, a small number of X-ray images are acquired in a circular orbit about a subject. We compare two implementations of the X-ray image formation model for reconstruction via differentiable rendering: (B) Siddon’s method and (C) trilinear interpolation.

2 Methods

Image Formation Model (Fig. 1). Let $\vec{r}(\alpha) = \mathbf{s} + \alpha(\mathbf{p} - \mathbf{s})$ with $\alpha \in [0, 1]$ be a beam of photons cast through a heterogeneous medium from an X-ray source $\mathbf{s} \in \mathbb{R}^3$ to a pixel on the detector $\mathbf{p} \in \mathbb{R}^3$. An X-ray image I quantifies the energy-weighted fraction of attenuated incident photons for every pixel in the detector, with the negative log intensity at \mathbf{p} being governed by the Beer-Lambert law:

$$I_{\mu}(\vec{r}) = \int_{\mathbf{x} \in \vec{r}} \mu(\mathbf{x}) d\mathbf{x} = \int_0^1 \mu(\vec{r}(\alpha)) \|\vec{r}'(\alpha)\| d\alpha = \|\mathbf{p} - \mathbf{s}\| \int_0^1 \mu(\mathbf{s} + \alpha(\mathbf{p} - \mathbf{s})) d\alpha, \quad (1)$$

where $\mu : \mathbb{R}^3 \rightarrow [0, \infty)$ represents the medium’s linear attenuation coefficients (LACs), a physical property proportional to the density at every point in space. In all practical settings, μ is discretized onto a finite-resolution voxelgrid. Therefore, CBCT reconstruction via differentiable rendering requires numerical methods to compute Eq. (1) over a discrete volume of optimized LACs, denoted as $\hat{\mu}$, in a manner that is differentiable with respect to $\hat{\mu}$.

Differentiable X-ray Rendering. The first integration technique we consider is Siddon’s method [13], which exactly computes a discretized version of Eq. (1) as the sum of the LAC in every voxel on the path of \vec{r} , weighted by the intersection length of \vec{r} with each voxel:

$$I_{\hat{\mu}}(\vec{r}) = \|\mathbf{p} - \mathbf{s}\| \sum_{m=1}^{M-1} \hat{\mu} \left[\mathbf{s} + \frac{\alpha_{m+1} + \alpha_m}{2} (\mathbf{p} - \mathbf{s}) \right] (\alpha_{m+1} - \alpha_m), \quad (2)$$

where $\{\alpha_1, \dots, \alpha_M\}$ parameterize the intersection of \vec{r} with the parallel planes comprising $\hat{\mu}$, and $\hat{\mu}[\cdot]$ is an indexing operation that returns of the LAC of the intersected voxel (Fig. 1B). It was previously shown that Eq. (2) can be implemented in a completely differentiable manner [10]. Instead of computing every plane intersection, which scales cubically with the resolution of $\hat{\mu}$, we propose to approximate the Beer-Lambert law using interpolatory quadrature via the rectangular rule:

$$I_{\hat{\mu}}(\vec{r}) \approx \|\mathbf{p} - \mathbf{s}\| \sum_{m=1}^{M-1} \hat{\mu} [\mathbf{s} + \alpha_m (\mathbf{p} - \mathbf{s})] \frac{(\alpha_{m+1} - \alpha_{m-1})}{2}, \quad (3)$$

where $\{\alpha_1, \dots, \alpha_M\}$ parameterize M evenly spaced points along \vec{r} and $\hat{\mu}[\cdot]$ represents trilinear interpolation (Fig. 1C). Trilinear interpolation is linear in M and thus faster than Siddon’s method.

Differentiable CBCT Reconstruction. Let \mathcal{R} represent the set of all photon beams cast in a set of 2D X-ray images of a subject, fully specifying the acquisition geometry. For each beam $\vec{r} \in \mathcal{R}$, let $I_{\mu}(\vec{r}) \in [0, \infty)$ correspond to its ground truth pixel intensity. To reconstruct the unknown volume, we directly optimize the voxelgrid $\hat{\mu}$ with the following photometric loss function:

$$\mathcal{L}(\hat{\mu}) = \frac{1}{|\mathcal{R}|} \sum_{\vec{r} \in \mathcal{R}} \|I_{\mu}(\vec{r}) - I_{\hat{\mu}}(\vec{r})\|_1 + \lambda_{\text{TV}} \text{TV}(\hat{\mu}), \quad (4)$$

where $\|\cdot\|_1$ is the L1 norm and $\text{TV}(\cdot)$ is the total variation norm, a commonly used regularizer that encourages reconstructed volumes to be piecewise constant [11]. Note that we implement $I_{\hat{\mu}}(\cdot)$ as a differentiable renderer specifically to minimize \mathcal{L} using gradient-based optimization.

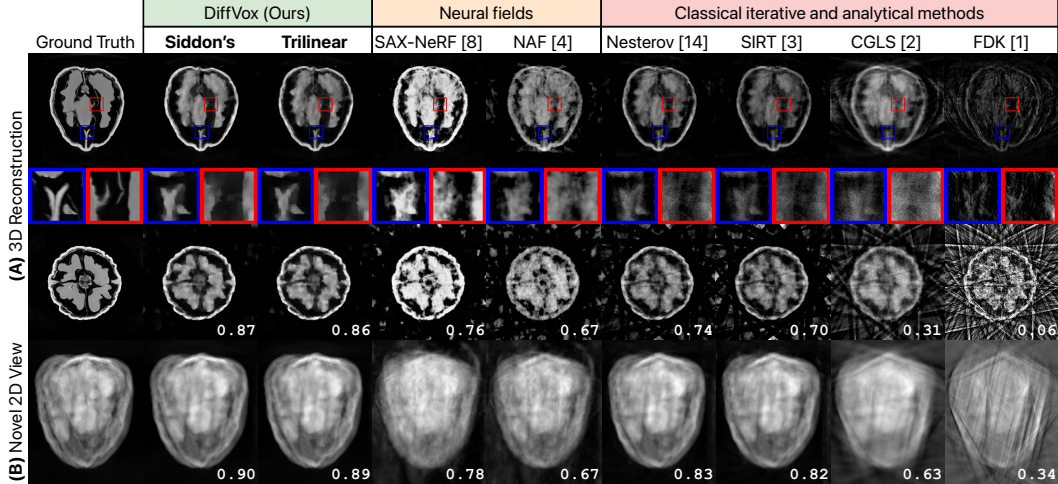


Figure 2: (A) CBCT reconstructions of an exemplar walnut from the test set using 15 input views. The Structural Similarity Index Measure (SSIM) for each 3D reconstruction is annotated. Blue insets highlight where our methods outperform the baselines, with sharper boundaries and fewer artifacts. Red insets indicate areas where all methods struggle, particularly in reconstructing thin structures. (B) Novel views rendered from these estimated volumes are compared to a ground truth X-ray image not seen during reconstruction and similarly annotated with SSIM. Novel views were rendered using the forward model implemented in each method.

3 Experiments

Dataset. Commercial CT scanners typically do not provide access to raw projection data. Therefore, most CBCT reconstruction methods are evaluated on artificial X-ray images simulated from existing 3D scans using computational models [4, 8]. Instead, we evaluated our proposed methods using an open dataset of real high-resolution X-ray images acquired from 42 walnuts [14]. For each walnut, 3,600 X-ray images (972×768 pixels with 0.1496 mm isotropic spacing) were acquired across three circular orbits at different heights. Ground truth CBCT volumes were reconstructed using all images to produce high-resolution $501 \times 501 \times 501$ volumes at 0.1 mm isotropic spacing. At test-time, we reconstructed CBCT volumes of equal resolution using a sparse subset of X-ray images from only the middle orbit to replicate a typical cone-beam acquisition (Fig. 2A). In addition to comparing our 3D reconstructions to the high-resolution ground truth, we also measure how similar novel 2D views rendered from our reconstructed volumes were to real 2D X-ray images from the high and low orbits, which are not seen during reconstruction (Fig. 2B). All hyperparameters were tuned with a grid search over two subjects, and the remaining 40 were used as a held-out test set.

Implementation Details. We initialize $\hat{\mu}$ as a voxelgrid of all zeros. As LACs are physically constrained to be nonnegative, we apply the Softplus activation function ($\beta = 20$) to $\hat{\mu}$ prior to rendering. To optimize $\hat{\mu}$, we use the Adam optimizer with an initial learning rate of 1 that linearly decays to 0 over 50 iterations. In each batch, rays from all projection images are randomly sampled without replacement with uniform probability. The batch size for each renderer was set to the maximum allowable under memory constraints, i.e., $|\mathcal{R}| = 5.5 \times 10^5$ for Siddon’s method and $|\mathcal{R}| = 1.8 \times 10^6$ for trilinear interpolation. Finally, we set $\lambda_{TV} = 25$ for Siddon’s method and $\lambda_{TV} = 15$ for trilinear interpolation. $M = 500$ points per ray were sampled for trilinear interpolation.

Baselines. We compare against GPU-accelerated implementations of four classical CBCT reconstruction algorithms from the ASTRA Toolbox [15]: Feldkamp-Davis-Kress (FDK) [1], Conjugate Gradient Least Squares (CGLS) [2], Simultaneous Iterative Reconstruction Technique (SIRT) [3], and Nesterov-Accelerated Gradient Descent (Nesterov) [14]. Based on grid searches to maximize their performance, we ran CGLS for 20 iterations, SIRT for 500 iterations, and Nesterov for 50 iterations. Additionally, we evaluated two contemporary neural field approaches for CBCT reconstruction: Neural Attenuation Fields (NAF) [4], an MLP-hashgrid hybrid, and Structure-Aware X-ray Neural Radiodensity Fields (SAX-NeRF) [8], a Transformer-hashgrid hybrid, which recently reported state-of-the-art accuracy in sparse-view CBCT reconstruction.

Table 1: Quantitative comparison for 3D reconstruction from 15 views averaged over the 40 test walnuts. Our methods excel across multiple image fidelity metrics, outperforming traditional and neural CBCT reconstruction methods. Additionally, our voxel-based reconstruction method is an order of magnitude faster than existing neural field approaches. Metrics are reported as *mean (se)*.

	SSIM (\uparrow)	PSNR (\uparrow)	MSE $\times 10^{-5}$ (\downarrow)	PCC (\uparrow)	Runtime (\downarrow)
FDK [1]	0.152 (0.014)	21.04 (0.67)	58.9 (0.62)	0.47 (0.01)	0.94 sec
CGLS [2]	0.505 (0.024)	30.01 (0.67)	7.47 (0.08)	0.76 (0.01)	<u>2.36 sec</u>
SIRT [3]	0.809 (0.013)	32.99 (0.70)	3.87 (0.16)	0.88 (0.03)	34.6 sec
Nesterov [14]	0.834 (0.011)	34.35 (0.69)	2.81 (0.09)	0.92 (0.02)	2.3 min
NAF [4]	0.782 (0.015)	31.72 (0.69)	5.09 (0.13)	0.85 (0.02)	123.5 min
SAX-NeRF [8]	0.835 (0.014)	33.43 (0.89)	3.97 (0.32)	0.88 (0.06)	160 min
DiffVox – Trilinear (Ours)	0.912 (0.006)	37.23 (0.67)	1.43 (0.03)	0.96 (0.00)	10.0 min
DiffVox – Siddon’s (Ours)	0.923 (0.006)	37.69 (0.67)	1.28 (0.03)	0.96 (0.00)	11.3 min
DiffVox – Trilinear w/o TV (Ours)	0.827 (0.008)	29.25 (0.56)	10.26 (0.15)	0.73 (0.01)	9.7 min
DiffVox – Siddon’s w/o TV (Ours)	0.812 (0.007)	27.41 (0.56)	15.72 (0.29)	0.67 (0.01)	10.1 min

Results. We tested the performance of DiffVox and relevant baselines in the sparse-view regime, reconstructing 40 walnuts with the number of input views ranging from $\{5, 10, 15, \dots, 60\}$. To evaluate the quality of these 3D reconstructions and subsequent novel 2D views, we computed the following image fidelity scores: Structural Similarity Index Measure (SSIM), Peak Signal-to-Noise Ratio (PSNR), Pearson’s Correlation Coefficient (PCC), and Mean Squared Error (MSE). Additionally, runtimes for five subjects were measured using a single NVIDIA RTX A6000 GPU.

In sparse-view settings, our voxel-based differentiable rendering approach outperforms classical and neural CBCT reconstruction techniques across all image quality metrics (Tab. 1). Visualized slices of reconstructions from 15 input views in Fig. 2A show that our methods successfully recover the highest quality 3D volumes. In comparison, previous methods suffer from stereotypical sparse-view artifacts, such as floaters and aliasing artifacts. Consequently, the novel views rendered from our reconstructed volumes have much higher fidelity (Fig. 2B). When performing reconstructions from 30 and 60 views, SAX-NeRF achieves a higher PSNR than our methods (Fig. 3A, *bottom*). However, across all other number of views and performance metrics, DiffVox outperforms all baseline CBCT reconstruction methods in both reconstruction and novel view synthesis tasks (Fig. 3A,B).

In both 3D reconstruction and 2D novel view synthesis, Siddon’s method provides a modest improvement over trilinear interpolation (Fig. 3A,B) while incurring a slightly higher runtime due to its increased complexity (Fig. 3C). In particular, the advantage of Siddon’s method is evident when rendering 2D projections in the sparse-view setting (5-20 views in Fig. 3B). In this regime, novel views rendered from volumes reconstructed with Siddon’s method have consistently higher fidelity than those produced by either trilinear interpolation or any other baseline method.

For both of our differentiable rendering methods, regularization via TV norm leads to high-quality reconstructions. However, as the ablation study in Tab. 1 shows, removing the TV norm has a far greater impact on Siddon’s method than on trilinear interpolation. This is because, while Siddon’s method exactly computes the line integral in the Beer-Lambert law (which leads to higher fidelity novel views), trilinear interpolation allows rendered pixel intensities to be influenced by voxels not directly intersected by \vec{r} (Fig. 1C). This is advantageous as it enables ground truth 2D pixels to provide self-supervision to many more 3D voxels in the reconstruction, preventing the overfitting observed in Siddon’s method. Furthermore, this behavior is an approximation of second-order effects such as X-ray scatter, something that Siddon’s method does not model.

Compared to previous neural field methods, traditional CBCT reconstruction algorithms have much faster runtimes. The slowest traditional method, Nesterov, takes no more than 5 min for any number of views (Fig. 3C). In contrast, the two slowest methods overall, SAX-NeRF and NAF, take on average 10.3 h and 30.6 h, respectively, to reconstruct a volume from 60 input views. While previous neural fields with hashgrid-augmented neural architectures claim to reconstruct CBCT volumes in tens of minutes [4], these runtimes are achieved when testing on low-resolution synthetic X-rays rendered from subsampled CT scans. Evaluations on real X-rays acquired at a standard resolution for clinical scanners demonstrate that existing neural fields are prohibitively slow for real-world deployment. Comparatively, reconstruction with Siddon’s method and trilinear interpolation in DiffVox is an order of magnitude faster than existing neural fields, requiring only 45.1 min and 39.1 min for 60-view reconstruction, respectively.




Figure 3: Quality of (A) reconstructed 3D volumes and (B) novel 2D views rendered from these volumes, and (C) reconstruction runtimes over a range of input views. DiffVox, using either Siddon’s method or trilinear interpolation as the forward model, achieves the highest quality reconstructions and renderings, with particularly appreciable gains in the sparsest-view settings. Error bars are plotted using the standard error (se).

4 Conclusions

We present DiffVox, a modular framework for performing CBCT reconstruction with differentiable X-ray rendering. In a thorough analysis of a dataset with real X-rays, we found that DiffVox achieved high quality reconstructions given an extremely sparse number of input views (≤ 20), a task for which traditional and neural CBCT reconstruction algorithms proved insufficient. In particular, we observed that Siddon’s method achieved slightly higher accuracy than trilinear interpolation at the expense of moderately increased runtimes. In addition to achieving higher fidelity 3D reconstructions and 2D novel views than existing algorithms, DiffVox was an order of magnitude faster to train than previous neural field methods: DiffVox required tens of minutes to reconstruct clinical-resolution CBCT scans whereas previous methods required hours to days. Furthermore, compared to continuous neural representations, parameterizing the unknown volume as a discrete voxelgrid endowed unique advantages such as fast convergence and the ability to utilize well-validated reconstruction regularizers. Finally, this analysis demonstrated the importance of evaluating reconstruction algorithms for clinical use on real images instead of primarily utilizing synthetic data.

Our preliminary results naturally suggest several potential areas for future work. For example, incorporating a more realistic implementation of X-ray image formation (e.g., Siddon’s method) into neural fields may yield compounding benefits. Additionally, as our physics-based reconstruction framework is fully differentiable, voxelgrid priors beyond TV regularization can be easily incorporated for specialized applications, such as vascular reconstruction [16]. To further exploit the differentiability of our X-ray renderer, it is also possible to jointly optimize the acquisition geometry of the X-ray images along with the unknown volume to address inaccurate or unspecified camera poses [5, 17]. Additionally, instead of initializing our voxelgrid with all zeros, one could instead initialize with the volume estimated by Nesterov or SIRT. Further refining this initial estimate with differentiable rendering could sufficiently reduce the runtime of DiffVox to the point of clinical relevance while retaining its superior reconstruction quality. Note that this initialization scheme is a unique advantage of our discrete voxelgrid representation, and not possible with continuous neural parameterizations.

Acknowledgments. This work was supported in part by NIH NIBIB NAC P41EB015902, NIH NINDS U19NS115388, NIH NIBIB 5T32EB001680-19, and NIH R01EB034223.

References

- [1] Lee A Feldkamp, Lloyd C Davis, and James W Kress. Practical cone-beam algorithm. *Journal of the Optical Society of America A*, 1(6):612–619, 1984.
- [2] Malena Sabaté Landman, Ander Biguri, Sepideh Hatamikia, Richard Boardman, John Aston, and Carola-Bibiane Schönlieb. On Krylov methods for large-scale CBCT reconstruction. *Physics in Medicine & Biology*, 68(15):155008, 2023.
- [3] Jens Gregor and Thomas Benson. Computational analysis and improvement of SIRT. *IEEE Transactions on Medical Imaging*, 27(7):918–924, 2008.
- [4] Ruyi Zha, Yanhao Zhang, and Hongdong Li. NAF: neural attenuation fields for sparse-view CBCT reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 442–452. Springer, 2022.
- [5] Darius Rückert, Yuanhao Wang, Rui Li, Ramzi Idoughi, and Wolfgang Heidrich. NeAT: Neural adaptive tomography. *ACM Transactions on Graphics (TOG)*, 41(4):1–13, 2022.
- [6] Yiqun Lin, Zhongjin Luo, Wei Zhao, and Xiaomeng Li. Learning deep intensity field for extremely sparse-view CBCT reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 13–23. Springer, 2023.
- [7] Thomas Müller, Alex Evans, Christoph Schied, and Alexander Keller. Instant neural graphics primitives with a multiresolution hash encoding. *ACM Transactions on Graphics (TOG)*, 41(4):1–15, 2022.
- [8] Yuanhao Cai, Jiahao Wang, Alan Yuille, Zongwei Zhou, and Angtian Wang. Structure-aware sparse-view X-ray 3D reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11174–11183, 2024.
- [9] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022.
- [10] Vivek Gopalakrishnan and Polina Golland. Fast auto-differentiable digitally reconstructed radiographs for solving inverse problems in intraoperative imaging. In *Workshop on Clinical Image-Based Procedures*, pages 1–11. Springer, 2022.
- [11] Leonid I Rudin, Stanley Osher, and Emad Fatemi. Nonlinear total variation based noise removal algorithms. *Physica D: Nonlinear Phenomena*, 60(1-4):259–268, 1992.
- [12] Diederik P. Kingma and Jimmy Ba. Adam: A method for stochastic optimization. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, 2015.
- [13] Robert L Siddon. Fast calculation of the exact radiological path for a three-dimensional CT array. *Medical Physics*, 12(2):252–255, 1985.
- [14] Henri Der Sarkissian, Felix Lucka, Maureen van Eijnatten, Giulia Colacicco, Sophia Bethany Coban, and Kees Joost Batenburg. A cone-beam X-ray computed tomography data collection designed for machine learning. *Scientific Data*, 6(1):215, 2019.
- [15] Wim Van Aarle, Willem Jan Palenstijn, Jeroen Cant, Eline Janssens, Folkert Bleichrodt, Andrei Dabrovolski, Jan De Beenhouwer, K Joost Batenburg, and Jan Sijbers. Fast and flexible X-ray tomography using the ASTRA toolbox. *Optics Express*, 24(22):25129–25147, 2016.
- [16] Alejandro F Frangi, Wiro J Niessen, Koen L Vincken, and Max A Viergever. 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*, pages 130–137. Springer, 1998.
- [17] Vivek Gopalakrishnan, Neel Dey, and Polina Golland. Intraoperative 2D/3D image registration via differentiable X-ray rendering. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 11662–11672, 2024.