

# Learning Deep Intensity Field for Extremely Sparse-View CBCT Reconstruction

Yiqun  $\operatorname{Lin}^1$ , Zhongjin  $\operatorname{Luo}^2$ , Wei Zhao<sup>3</sup>, and Xiaomeng  $\operatorname{Li}^{1(\boxtimes)}$ 

<sup>1</sup> The Hong Kong University of Science and Technology, Kowloon, Hong Kong eexmli@ust.hk

- $^{2\,}$  The Chinese University of Hong Kong, Shenzhen, China
- <sup>3</sup> Beihang University, Beijing, People's Republic of China

Abstract. Sparse-view cone-beam CT (CBCT) reconstruction is an important direction to reduce radiation dose and benefit clinical applications. Previous voxel-based generation methods represent the CT as discrete voxels, resulting in high memory requirements and limited spatial resolution due to the use of 3D decoders. In this paper, we formulate the CT volume as a continuous intensity field and develop a novel DIF-Net to perform high-quality CBCT reconstruction from extremely sparse  $(\leq 10)$  projection views at an ultrafast speed. The intensity field of a CT can be regarded as a continuous function of 3D spatial points. Therefore, the reconstruction can be reformulated as regressing the intensity value of an arbitrary 3D point from given sparse projections. Specifically, for a point, DIF-Net extracts its view-specific features from different 2D projection views. These features are subsequently aggregated by a fusion module for intensity estimation. Notably, thousands of points can be processed in parallel to improve efficiency during training and testing. In practice, we collect a knee CBCT dataset to train and evaluate DIF-Net. Extensive experiments show that our approach can reconstruct CBCT with high image quality and high spatial resolution from extremely sparse views within 1.6 s, significantly outperforming state-of-the-art methods. Our code will be available at https://github.com/xmed-lab/DIF-Net.

**Keywords:** CBCT Reconstruction  $\cdot$  Implicit Neural Representation  $\cdot$  Sparse View  $\cdot$  Low Dose  $\cdot$  Efficient Reconstruction

# 1 Introduction

Cone-beam computed tomography (CBCT) is a common 3D imaging technique used to examine the internal structure of an object with high spatial resolution and fast scanning speed [20]. During CBCT scanning, the scanner rotates around the object and emits cone-shaped beams, obtaining 2D projections in the detection panel to reconstruct 3D volume. In recent years, beyond dentistry, CBCT

**Supplementary Information** The online version contains supplementary material available at https://doi.org/10.1007/978-3-031-43999-5\_2.

<sup>©</sup> The Author(s), under exclusive license to Springer Nature Switzerland AG 2023 H. Greenspan et al. (Eds.): MICCAI 2023, LNCS 14229, pp. 13–23, 2023. https://doi.org/10.1007/978-3-031-43999-5\_2



**Fig. 1.** (a-b): Comparison of conventional CT and cone-beam CT scanning. (c-d): CBCT reconstruction from a stack of sparse 2D projections.

has been widely used to acquire images of the human knee joint for applications such as total knee arthroplasty and postoperative pain management [3,4,9,15]. To maintain image quality, CBCT typically requires hundreds of projections involving high radiation doses from X-rays, which could be a concern in clinical practice. Sparse-view reconstruction is one of the ways to reduce radiation dose by reducing the number of scanning views ( $10 \times$  fewer). In this paper, we study a more challenging problem, extremely sparse-view CBCT reconstruction, aiming to reconstruct a high-quality CT volume from fewer than 10 projection views.

Compared to conventional CT (e.g., parallel beam, fan beam), CBCT reconstructs a 3D volume from 2D projections instead of a 2D slice from 1D projections, as comparison shown in Fig. 1, resulting in a significant increase in spatial dimensionality and computational complexity. Therefore, although sparse-view conventional CT reconstruction [2, 23, 25, 26] has been developed for many years, these methods cannot be trivially extended to CBCT. CBCT reconstruction can be divided into dense-view ( $\geq 100$ ), sparse-view ( $20 \sim 50$ ), extremely sparseview ( $\leq 10$ ), and single/orthogonal-view reconstructions depending on the number of projection views required. A typical example of dense-view reconstruction is FDK [6], which is a filtered-backprojection (FBP) algorithm that accumulates intensities by backprojecting from 2D views, but requires hundreds of views to avoid streaking artifacts. To reduce required projection views, ART [7] and its extensions (e.g., SART [1], VW-ART [16]) formulate reconstruction as an iterative minimization process, which is useful when projections are limited. Nevertheless, such methods often take a long computational time to converge and cope poorly with extremely sparse projections; see results of SART in Table 1. With the development of deep learning techniques and computing devices, learning-based approaches are proposed for CBCT sparse-view reconstruction. Lahiri et al. [12] propose to reconstruct a coarse CT with FDK and use 2D CNNs to denoise each slice. However, the algorithm has not been validated on medical datasets, and the performance is still limited as FDK introduces extensive streaking artifacts with sparse views. Recently, neural rendering techniques [5, 14, 19, 21, 29] have been introduced to reconstruct CBCT volume by parameterizing the attenuation coefficient field as an implicit neural representation field (NeRF), but they require a long time for per-patient optimization and do not perform well with extremely sparse views due to lack of prior knowledge; see results of NAF in Table 2. For single/orthogonal-view reconstruction, voxel-based approaches [10, 22, 27] are proposed to build 2D-to-3D generation networks that consist of 2D encoders and 3D decoders with large training parameters, leading to high memory requirements and limited spatial resolution. These

methods are special designs with the networks [10, 27] or patient-specific training data [22], which are difficult to extend to general sparse-view reconstruction.

In this work, our goal is to reconstruct a CBCT of high image quality and high spatial resolution from extremely sparse ( $\leq 10$ ) 2D projections, which is an important yet challenging and unstudied problem in sparse-view CBCT reconstruction. Unlike previous voxel-based methods that represent the CT as discrete voxels, we formulate the CT volume as a continuous intensity field, which can be regarded as a continuous function  $q(\cdot)$  of 3D spatial points. The property of a point p in this field represents its intensity value v, i.e., v = q(p). Therefore, the reconstruction problem can be reformulated as regressing the intensity value of an arbitrary 3D point from a stack of 2D projections  $\mathcal{I}$ , i.e.,  $v = q(\mathcal{I}, p)$ . Based on the above formulation, we develop a novel reconstruction framework, namely DIF-Net (Deep Intensity Field Network). Specifically, DIF-Net first extracts feature maps from K given 2D projections. Given a 3D point, we project the point onto the 2D imaging panel of each view, by corresponding imaging parameters (distance, angle, etc.) and query its view-specific features from the feature map of view<sub>i</sub>. Then, K view-specific features from different views are aggregated by a cross-view fusion module for intensity regression. By introducing the continuous intensity field, it becomes possible to train DIF-Net with a set of sparsely sampled points to reduce memory requirement, and reconstruct the CT volume with any desired resolution during testing. Compared with NeRF-based methods [5, 14, 19, 21, 29], the design of DIF-Net shares the similar data representation (i.e., implicit neural representation) but additional training data can be introduced to help DIF-Net learn prior knowledge. Benefiting from this, DIF-Net can not only reconstruct high-resolution CT in a very short time since only inference is required for a new test sample (no retraining), but also performs much better than NeRF-based methods with extremely limited views.

To summarize, the main contributions of this work include 1.) we are the first to introduce the continuous intensity field for supervised CBCT reconstruction; 2.) we propose a novel reconstruction framework DIF-Net that reconstructs CBCT with high image quality (PSNR: 29.3 dB, SSIM: 0.92) and high spatial resolution ( $\geq 256^3$ ) from extremely sparse ( $\leq 10$ ) views within 1.6 s; 3.) we conduct extensive experiments to validate the effectiveness of the proposed sparse-view CBCT reconstruction method on a clinical knee CBCT dataset.

### 2 Method

#### 2.1 Intensity Field

We formulate the CT volume as a continuous intensity field, where the property of a 3D point  $p \in \mathbb{R}^3$  in this field represents its intensity value  $v \in \mathbb{R}$ . The intensity field can be defined as a continuous function  $g : \mathbb{R}^3 \to \mathbb{R}$ , such that v = g(p). Hence, the reconstruction problem can be reformulated as regressing the intensity value of an arbitrary point p in the 3D space from K projections  $\mathcal{I} =$  $\{I_1, I_2, \ldots, I_K\}$ , i.e.,  $v = g(\mathcal{I}, p)$ . Based on the above formulation, we propose a novel reconstruction framework, namely DIF-Net, to perform efficient sparseview CBCT reconstruction, as the overview shown in Fig. 2.



**Fig. 2.** Overview of DIF-Net. (a) Given K projections, a shared 2D encoder is used for feature extraction. (b) For a point p in the 3D space, its view-specific features are queried from feature maps of different views by projection and interpolation. (c) Queried features are aggregated to estimate the intensity value of p. (d) During testing, given input projections, DIF-Net predicts intensity values for points uniformly sampled in 3D space to reconstruct the target CT image.

#### 2.2 DIF-Net: Deep Intensity Field Network

DIF-Net first extracts feature maps  $\{F_1, F_2, \ldots, F_K\} \subset \mathbb{R}^{C \times H \times W}$  from projections  $\mathcal{I}$  using a shared 2D encoder, where C is the number of feature channels and H/W are height/width. In practice, we choose U-Net [18] as the 2D encoder because of its good feature extraction ability and popular applications in medical image analysis [17]. Then, given a 3D point, DIF-Net gathers its view-specific features queried from feature maps of different views for intensity regression.

View-Specific Feature Querying. Considering a point  $p \in \mathbb{R}^3$  in the 3D space, for a projection view<sub>i</sub> with scanning angle  $\alpha_i$  and other imaging parameters  $\beta$  (distance, spacing, etc.), we project p to the 2D imaging panel of view<sub>i</sub> and obtain its 2D projection coordinates  $p'_i = \varphi(p, \alpha_i, \beta) \in \mathbb{R}^2$ , where  $\varphi(\cdot)$  is the projection function. Projection coordinates  $p'_i$  are used for querying view-specific features  $f_i \in \mathbb{R}^C$  from the 2D feature map  $F_i$  of view<sub>i</sub>:

$$f_i = \pi(F_i, p'_i) = \pi(F_i, \varphi(p, \alpha_i, \beta)), \qquad (1)$$

where  $\pi(\cdot)$  is bilinar interpolation. Similar to perspective projection, the CBCT projection function  $\varphi(\cdot)$  can be formulated as

$$\varphi(p,\alpha_i,\beta) = H\left(A(\beta)R(\alpha_i) \begin{bmatrix} p\\1 \end{bmatrix}\right),\tag{2}$$

where  $R(\alpha_i) \in \mathbb{R}^{4 \times 4}$  is a rotation matrix that transforms point p from the world coordinate system to the scanner coordinate system of view<sub>i</sub>,  $A(\beta) \in \mathbb{R}^{3 \times 4}$  is a projection matrix that projects the point onto the 2D imaging panel of view<sub>i</sub>, and  $H : \mathbb{R}^3 \to \mathbb{R}^2$  is the homogeneous division that maps the homogeneous coordinates of  $p'_i$  to its Cartesian coordinates. Due to page limitations, the detailed formulation of  $\varphi(\cdot)$  is given in the supplementary material. **Cross-View Feature Fusion and Intensity Regression.** Given K projection views, K view-specific features of the point p are queried from different views to form a feature list  $F(p) = \{f_1, f_2, \ldots, f_K\} \subset \mathbb{R}^C$ . Then, the cross-view feature fusion  $\delta(\cdot)$  is introduced to gather features from F(p) and generate a 1D vector  $\overline{f} = \delta(F(p)) \in \mathbb{R}^C$  to represent the semantic features of p. In general, F(p) is an unordered feature set, which means that  $\delta(\cdot)$  should be a set function and can be implemented with a pooling layer (e.g., max/avg pooling). In our experiments, the projection angles of the training and test samples are the same, uniformly sampled from 0° to 180° (half rotation). Therefore, F(p) can be regarded as an ordered list ( $K \times C$  tensor), and  $\delta(\cdot)$  can be implemented by a 2-layer MLP ( $K \rightarrow \lfloor \frac{K}{2} \rfloor \rightarrow 1$ ) for feature aggregation. We will compare different implementations of  $\delta(\cdot)$  in the ablation study. Finally, a 4-layer MLP ( $C \rightarrow 2C \rightarrow \lfloor \frac{C}{2} \rfloor \rightarrow \lfloor \frac{C}{8} \rfloor \rightarrow 1$ ) is applied to  $\overline{f}$  for the regression of intensity value  $v \in \mathbb{R}$ .

#### 2.3 Network Training

Assume that the shape and spacing of the original CT volume are  $H \times W \times D$ and  $(s_h, s_w, s_d)$  mm, respectively. During training, different from previous voxelbased methods that regard the entire 3D CT image as the supervision target, we randomly sample a set of N points  $\{p_1, p_2, \ldots, p_N\}$  with coordinates ranging from (0, 0, 0) to  $(s_h H, s_w W, s_d D)$  in the world coordinate system (unit: mm) as the input. Then DIF-Net will estimate their intensity values  $\mathcal{V} = \{v_1, v_2, \ldots, v_N\}$ from given projections  $\mathcal{I}$ . For supervision, ground-truth intensity values  $\hat{\mathcal{V}} =$  $\{\hat{v}_1, \hat{v}_2, \ldots, \hat{v}_N\}$  can be obtained from the ground-truth CT image based on the coordinates of points by trilinear interpolation. We choose mean-square-error (MSE) as the objective function, and the training loss can be formulated as

$$\mathcal{L}(\mathcal{V},\hat{\mathcal{V}}) = \frac{1}{N} \sum_{i=1}^{N} (v_i - \hat{v}_i)^2.$$
(3)

Because background points (62%, e.g., air) occupy more space than foreground points (38%, e.g., bones, organs), uniform sampling will bring imbalanced prediction of intensities. We set an intensity threshold  $10^{-5}$  to identify foreground and background areas by binary classification and sample  $\frac{N}{2}$  points from each area for training.

#### 2.4 Volume Reconstruction

During inference, a regular and dense point set to cover all CT voxels is sampled, i.e., to uniformly sample  $H \times W \times D$  points from (0,0,0) to  $(s_hH, s_wW, s_dD)$ . Then the network will take 2D projections and points as the input and generate intensity values of sampled points to form the target CT volume. Unlike previous voxel-based methods that are limited to generating fixed-resolution CT volumes, our method enables scalable output resolutions by introducing the representation of continuous intensity field. For example, we can uniformly sample  $\lfloor \frac{H}{s} \rfloor \times \lfloor \frac{W}{s} \rfloor \times$ 

 $\lfloor \frac{D}{s} \rfloor$  points to generate a coarse CT image but with a faster reconstruction speed, or sample  $\lfloor sH \rfloor \times \lfloor sW \rfloor \times \lfloor sD \rfloor$  points to generate a CT image with higher resolution, where s > 1 is the scaling ratio.

## 3 Experiments

We conduct extensive experiments on a collected knee CBCT dataset to show the effectiveness of our proposed method on sparse-view CBCT reconstruction. Compared to previous works, our DIF-Net can reconstruct a CT volume with high image quality and high spatial resolution from extremely sparse ( $\leq 10$ ) projections at an ultrafast speed.

### 3.1 Experimental Settings

**Dataset and Preprocessing.** We collect a knee CBCT dataset consisting of 614 CT scans. Of these, 464 are used for training, 50 for validation, and 100 for testing. We resample, interpolate, and crop (or pad) CT scans to have isotropic voxel spacing of (0.8, 0.8, 0.8) mm and shape of  $256 \times 256 \times 256$ . 2D projections are generated by digitally reconstructed radiographs (DRRs) at a resolution of  $256 \times 256$ . Projection angles are uniformly selected in the range of  $180^{\circ}$ .

**Implementation.** We implement DIF-Net using PyTorch with a single NVIDIA RTX 3090 GPU. The network parameters are optimized using stochastic gradient descent (SGD) with a momentum of 0.98 and an initial learning rate of 0.01. The learning rate is decreased by a factor of  $0.001^{1/400} \approx 0.9829$  per epoch, and we train the model for 400 epochs with a batch size of 4. For each CT scan, N = 10,000 points are sampled as the input during one training iteration. For the full model, we employ U-Net [18] with C = 128 output feature channels as the 2D encoder, and cross-view feature fusion is implemented with MLP.

**Baseline Methods.** We compare four publicly available methods as our baselines, including traditional methods FDK [6] and SART [1], NeRF-based method NAF [29], and data-driven denoising method FBPConvNet [11]. Due to the increase in dimensionality (2D to 3D), denoising methods should be equipped with 3D conv/deconvs for a dense prediction when extended to CBCT reconstruction, which leads to extremely high computational costs and low resolution ( $\leq 64^3$ ). For a fair comparison, we use FDK to obtain an initial result and apply the 2D network for slice-wise denoising.

**Evaluation Metrics.** We follow previous works [27–29] to evaluate the reconstructed CT volumes with two quantitative metrics, namely peak signal-to-noise ratio (PSNR) and structural similarity (SSIM) [24]. Higher PSNR/SSIM values represent superior reconstruction quality.

### 3.2 Results

**Performance.** As shown in Table 1, we compare DIF-Net with four previous methods [1, 6, 22, 29] under the setting of reconstruction with different output

**Table 1.** Comparison of DIF-Net with previous methods under measurements of PSNR (dB) and SSIM. We evaluate reconstructions with different output resolutions (Res.) and from different numbers of projection views (K).

Method	Res. $= 128^3$			Res. = $256^3$		
	K = 6	K = 8	K = 10	K = 6	K = 8	K = 10
FDK [6]	14.1/.18	15.7/.22	17.0/.25	14.1/.16	15.7/.20	16.9/.23
SART [1]	25.4/.81	26.6/.85	27.6/.88	24.7/.81	25.8/.84	26.7/.86
NAF [29]	20.8/.54	23.0/.64	25.0/.73	20.1/.58	22.4/.67	24.3/.75
FBPConvNet [11]	26.4/.84	27.0/.87	27.8/.88	25.1/.83	25.9/.83	26.7/.84
DIF-Net (Ours)	28.3/.91	29.6/.92	30.7/.94	27.1/.89	28.3/.90	29.3/.92



Fig. 3. Qualitative comparison of 10-view reconstruction.

resolutions (i.e., 128<sup>3</sup>, 256<sup>3</sup>) and from different numbers of projection views (i.e., 6, 8, and 10). Experiments show that our proposed DIF-Net can reconstruct CBCT with high image quality even using only 6 projection views, which significantly outperforms previous works in terms of PSNR and SSIM values. More importantly, DIF-Net can be directly applied to reconstruct CT images with different output resolutions without the need for model retraining or modification. As visual results are shown in Fig. 3, FDK [6] produces results with many streaking artifacts due to lack of sufficient projection views; SART [1] and NAF [29] produce results with good shape contours but lack detailed internal information; FBPConvNet [11] reconstructs good shapes and moderate details, but there are still some streaking artifacts remaining; our proposed DIF-Net can reconstruct high-quality CT with better shape contour, clearer internal information, and fewer artifacts. More visual comparisons of the number of input views are given in the supplementary material.

**Table 2.** Comparison of different methods in terms of reconstruction quality (PSNR/S-SIM), reconstruction time, parameters, and training memory cost. Default setting: 10-view reconstruction with the output resolution of  $256^3$ ; training with a batch size of 1. †: evaluated with the output resolution of  $128^3$  due to the memory limitation.

Method	PSNR/SSIM	Time (s)	Parameters (M)	Memory Cost (MB)
FDK [6]	16.9/.23	0.3	-	-
SART [1]	26.7/.86	106	-	339
NAF [29]	24.3/.75	738	14.3	3,273
FBPConvNet [11]	26.7/.84	1.7	34.6	3,095
DIF-Net (Ours)	29.3/.92	1.6	31.1	7,617

**Table 3.** Ablation study (10-view) on different cross-view fusion strategies.

Cross-View Fusion	PSNR	SSIM
Avg pooling	27.6	0.88
Max pooling	28.9	0.92
MLP	29.3	0.92

**Table 4.** Ablation study (10-view) on different numbers of training points N.

# Points	PSNR	SSIM
5,000	28.8	0.91
10,000	29.3	0.92
20,000	29.3	0.92

**Reconstruction Efficiency.** As shown in Table 2, FDK [6] requires the least time for reconstruction, but has the worst image quality; SART [1] and NAF [29] require a lot of time for optimization or training; FBPConvNet [11] can reconstruct 3D volumes faster, but the quality is still limited. Our DIF-Net can reconstruct high-quality CT within 1.6 s, much faster than most compared methods. In addition, DIF-Net, which benefits from the intensity field representation, has fewer training parameters and requires less computational memory, enabling high-resolution reconstruction.

Ablation Study. Tables 3 and 4 show the ablative analysis of cross-view fusion strategy and the number of training points N. Experiments demonstrate that 1.) MLP performs best, but max pooling is also effective and would be a general solution when the view angles are not consistent across training/test data, as discussed in Sect. 2.2; 2.) fewer points (e.g., 5,000) may destabilize the loss and gradient during training, leading to performance degradation; 10,000 points are enough to achieve the best performance, and training with 10,000 points is much sparser than voxel-based methods that train with the entire CT volume (i.e.,  $256^3$  or  $128^3$ ). We have tried to use a different encoder like pre-trained ResNet18 [8] with more model parameters than U-Net [18]. However, ResNet18 does not bring any improvement (PSNR/SSIM: 29.2/0.92), which means that U-Net is powerful enough for feature extraction in this task.

# 4 Conclusion

In this work, we formulate the CT volume as a continuous intensity field and present a novel DIF-Net for ultrafast CBCT reconstruction from extremely sparse ( $\leq 10$ ) projection views. DIF-Net aims to estimate the intensity value of an arbitrary point in 3D space from input projections, which means 3D CNNs are not required for feature decoding, thereby reducing memory requirement and computational cost. Experiments show that DIF-Net can perform efficient and high-quality CT reconstruction, significantly outperforming previous state-of-the-art methods. More importantly, DIF-Net is a general sparse-view reconstruction framework, which can be trained on a large-scale dataset containing various body parts with different projection views and imaging parameters to achieve better generalization ability. This will be left as our future work.

Acknowledgement. This work was supported by the Hong Kong Innovation and Technology Fund under Projects PRP/041/22FX and ITS/030/21, as well as by grants from Foshan HKUST Projects under Grants FSUST21-HKUST10E and FSUST21-HKUST11E.

# References

- 1. And ersen, A.H., Kak, A.C.: Simultaneous algebraic reconstruction technique (SART): a superior implementation of the art algorithm. Ultrason. Imaging  ${\bf 6}(1),$  81–94 (1984)
- Anirudh, R., Kim, H., Thiagarajan, J.J., Mohan, K.A., Champley, K., Bremer, T.: Lose the views: limited angle CT reconstruction via implicit sinogram completion. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 6343–6352 (2018)
- 3. Bier, B., et al.: Range imaging for motion compensation in C-arm cone-beam CT of knees under weight-bearing conditions. J. Imaging 4(1), 13 (2018)
- 4. Dartus, J., et al.: The advantages of cone-beam computerised tomography (CT) in pain management following total knee arthroplasty, in comparison with conventional multi-detector ct. Orthop. Traumatol. Surg. Res. **107**(3), 102874 (2021)
- 5. Fang, Y., et al.: SNAF: sparse-view CBCT reconstruction with neural attenuation fields. arXiv preprint arXiv:2211.17048 (2022)
- Feldkamp, L.A., Davis, L.C., Kress, J.W.: Practical cone-beam algorithm. Josa a 1(6), 612–619 (1984)
- Gordon, R., Bender, R., Herman, G.T.: Algebraic reconstruction techniques (art) for three-dimensional electron microscopy and x-ray photography. J. Theor. Biol. 29(3), 471–481 (1970)
- He, K., Zhang, X., Ren, S., Sun, J.: Deep residual learning for image recognition. In: Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition, pp. 770–778 (2016)
- Jaroma, A., Suomalainen, J.S., Niemitukia, L., Soininvaara, T., Salo, J., Kröger, H.: Imaging of symptomatic total knee arthroplasty with cone beam computed tomography. Acta Radiol. 59(12), 1500–1507 (2018)
- Jiang, Y.: MFCT-GAN: multi-information network to reconstruct CT volumes for security screening. J. Intell. Manuf. Spec. Equipment 3, 17–30 (2022)

- Jin, K.H., McCann, M.T., Froustey, E., Unser, M.: Deep convolutional neural network for inverse problems in imaging. IEEE Trans. Image Process. 26(9), 4509– 4522 (2017)
- Lahiri, A., Klasky, M., Fessler, J.A., Ravishankar, S.: Sparse-view cone beam CT reconstruction using data-consistent supervised and adversarial learning from scarce training data. arXiv preprint arXiv:2201.09318 (2022)
- Lechuga, L., Weidlich, G.A.: Cone beam CT vs. fan beam CT: a comparison of image quality and dose delivered between two differing CT imaging modalities. Cureus 8(9) (2016)
- Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: representing scenes as neural radiance fields for view synthesis. Commun. ACM 65(1), 99–106 (2021)
- Nardi, C., et al.: The role of cone beam CT in the study of symptomatic total knee arthroplasty (TKA): a 20 cases report. Br. J. Radiol. 90(1074), 20160925 (2017)
- Pan, J., Zhou, T., Han, Y., Jiang, M.: Variable weighted ordered subset image reconstruction algorithm. Int. J. Biomed. Imaging 2006 (2006)
- Punn, N.S., Agarwal, S.: Modality specific u-net variants for biomedical image segmentation: a survey. Artif. Intell. Rev. 55(7), 5845–5889 (2022)
- Ronneberger, O., Fischer, P., Brox, T.: U-Net: convolutional networks for biomedical image segmentation. In: Navab, N., Hornegger, J., Wells, W.M., Frangi, A.F. (eds.) MICCAI 2015. LNCS, vol. 9351, pp. 234–241. Springer, Cham (2015). https://doi.org/10.1007/978-3-319-24574-4\_28
- Rückert, D., Wang, Y., Li, R., Idoughi, R., Heidrich, W.: Neat: neural adaptive tomography. ACM Trans. Graph. (TOG) 41(4), 1–13 (2022)
- Scarfe, W.C., Farman, A.G., Sukovic, P., et al.: Clinical applications of cone-beam computed tomography in dental practice. J. Can. Dent. Assoc. 72(1), 75 (2006)
- Shen, L., Pauly, J., Xing, L.: NeRP: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. IEEE Trans. Neural Netw. Learn. Syst. (2022)
- Shen, L., Zhao, W., Xing, L.: Patient-specific reconstruction of volumetric computed tomography images from a single projection view via deep learning. Nat. Biomed. Eng. 3(11), 880–888 (2019)
- Tang, C., et al.: Projection super-resolution based on convolutional neural network for computed tomography. In: 15th International Meeting on Fully Three-Dimensional Image Reconstruction in Radiology and Nuclear Medicine, vol. 11072, pp. 537–541. SPIE (2019)
- Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. IEEE Trans. Image Process. 13(4), 600–612 (2004)
- Wu, W., Guo, X., Chen, Y., Wang, S., Chen, J.: Deep embedding-attentionrefinement for sparse-view CT reconstruction. IEEE Trans. Instrum. Meas. 72, 1–11 (2022)
- Wu, W., Hu, D., Niu, C., Yu, H., Vardhanabhuti, V., Wang, G.: Drone: dualdomain residual-based optimization network for sparse-view CT reconstruction. IEEE Trans. Med. Imaging 40(11), 3002–3014 (2021)
- Ying, X., Guo, H., Ma, K., Wu, J., Weng, Z., Zheng, Y.: X2CT-GAN: reconstructing CT from biplanar x-rays with generative adversarial networks. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition, pp. 10619–10628 (2019)

- 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, pp. 1960–1970 (2021)
- Zha, R., Zhang, Y., Li, H.: NAF: neural attenuation fields for sparse-view CBCT reconstruction. In: Wang, L., Dou, Q., Fletcher, P.T., Speidel, S., Li, S. (eds.) MICCAI 2022. LNCS, vol. 13436, pp. 442–452. Springer, Cham (2022). https:// doi.org/10.1007/978-3-031-16446-0\_42