

# MedNeRF: Medical Neural Radiance Fields for Reconstructing 3D-aware CT-Projections from a Single X-ray

Abril Corona-Figueroa<sup>1</sup>, Jonathan Frawley<sup>1</sup>, Sam Bond-Taylor<sup>1</sup>,  
Sarith Bethapudi<sup>2</sup>, Hubert P. H. Shum<sup>1</sup>, Chris G. Willcocks<sup>1</sup>

**Abstract**—Computed tomography (CT) is an effective medical imaging modality, widely used in the field of clinical medicine for the diagnosis of various pathologies. Advances in Multidetector CT imaging technology have enabled additional functionalities, including generation of thin slice multiplanar cross-sectional body imaging and 3D reconstructions. However, this involves patients being exposed to a considerable dose of ionising radiation. Excessive ionising radiation can lead to deterministic and harmful effects on the body. This paper proposes a Deep Learning model that learns to reconstruct CT projections from a few or even a single-view X-ray. This is based on a novel architecture that builds from neural radiance fields, which learns a continuous representation of CT scans by disentangling the shape and volumetric depth of surface and internal anatomical structures from 2D images. Our model is trained on chest and knee datasets, and we demonstrate qualitative and quantitative high-fidelity renderings and compare our approach to other recent radiance field-based methods. Our code and link to our datasets is available at <https://github.com/jonathanfrawley/mednerf>

**Clinical relevance**—Our model is able to infer the anatomical 3D structure from a few or a single-view X-ray, showing future potential for reduced ionising radiation exposure during the imaging process in the evaluation of skeletal dysplasia, osseous trauma or orthopaedic pre-surgical planning.

## I. INTRODUCTION

3D medical imaging often involves joining multiple 2D slices from CT or Magnetic Resonance Imaging (MRI), and part of their workflow consists of specifying values for the position of the patient, the imaging source, and the detector. The quality and accuracy of a CT 3D representation require hundreds of X-ray projections with a thin slice thickness [1]. Moreover, this process exposes patients to more ionising radiation than typical X-rays and requires the patient to remain immobile for up to more than 1 hour, depending on the type of test [2]. Continuous 3D representations would give radiologists optics of every point in the internal anatomy captured. While such representations are useful, there are practical challenges in CT due to the increased radiation exposure, angle-dependent structures, and time consumption [3].

Earlier approaches in medical image reconstruction used analytic and iterative methods [4], [5] on given input data. However, they often encounter mismatches between the

mathematical model and physical properties of the imaging system. Instead, several recent approaches leverage deep learning [6] for sparse view reconstruction [7], [8], [9], 3D CT reconstruction from 2D images [10], and anomaly detection [11]. These deep learning approaches solved the mismatches between the mathematical model and imaging system and reported improved reconstructions by fine-tuning state-of-the-art architectures. However, they require a large amount of training data, which may be difficult to meet in the medical domain where acquiring expert annotations is both cost and time prohibitive.

The Neural Radiance Fields (NeRF) [12] model is a *recent* reformulation for estimating a 3D volumetric representation from images. Such representations encode the radiance field and density of the scene in the parameters of a neural network. The neural network learns to synthesize new views via volume rendering from point samples along cast rays. However, these representations are often captured in controlled settings [13]. First, the scene is taken by a set of fixed cameras within a short time frame. Second, all content in the scene is static and real images often need masking. These constraints prohibit the direct application of NeRF to the medical domain, where the imaging system greatly differs from conventional cameras, and the images are captured over a long time frame hampering the patient’s stillness. Moreover, the overlapping of anatomical structures in medical images hinders the definition of edges which cannot be easily solved with masking. These aspects explain why the NeRF approach especially shows successes for “natural images”.

To address these challenges, we propose MedNeRF, a model that adapts Generative Radiance Fields (GRAF) [14] in the medical domain to render CT projections given a few or even a single-view X-ray. Our approach not only synthesizes realistic images, but also capture the data manifold and provides a continuous representation of how the attenuation and volumetric depth of anatomical structures vary with the viewpoint without 3D supervision. This is achieved via a new discriminator architecture that provides a stronger and more comprehensive signal to GRAF when dealing with CT scans.

Closest to our goal are [8], [9], which both train a coordinate-based network in sinograms of low-dose CT of phantom objects and apply it to the sparse-view tomography reconstruction problem. In contrast to [8], we learn multiple representations in a single model by randomly feeding data of different medical instances instead of separately optimizing

Accepted at IEEE EMBC 2022

<sup>1</sup>Corona-Figueroa, Frawley, Bond-Taylor, Shum and Willcocks are with the Computer Science Department, Durham University, Durham, DH1 3LE, UK [abril.corona-figueroa@durham.ac.uk](mailto:abril.corona-figueroa@durham.ac.uk)

<sup>2</sup>Bethapudi is with the County Durham and Darlington NHS Foundation Trust, Durham, DL3 6HX, UK

for each collection of images. For testing [9] reconstruction ability, they integrate it into reconstruction methods and use at least 60 views. Different from their methods, we do not rely on additional reconstruction algorithms, and we only require multiple views during training.

We render CT projections of our two datasets of digitally reconstructed radiographs (DRR) from chest and knee. We qualitatively and quantitatively demonstrate high-fidelity renderings and compare our approach to other recent radiance field-based methods. Furthermore, we render CT projections of a medical instance given a single-view X-ray and show the effectiveness of our model to cover surface and internal structures.

## II. METHODS

### A. Dataset Preparation

To train our models, we generate DRRs instead of collecting paired X-rays and corresponding CT reconstructions, which would expose patients to more radiation. Furthermore, DRR generation removes patient data and enables control in capture ranges and resolutions. We generated DRRs by using 20 CT chest scans from [15], [16] and five CT knee scans from [17], [18]. These scans cover a diverse group of patients at different contrast types showing both normal and abnormal anatomy. The radiation source and imaging panel are assumed to rotate around the vertical-axis, generating a DRR of  $128 \times 128$  resolution at every five degrees, resulting in 72 DRRs for each object. During training we use the whole set of 72 DRRs (a fifth of all views within a full 360-degree vertical rotation) per patient and let the model render the rest. Our work did not involve experimental procedures on human subjects or animals and thus did not require Institutional Review Board approval.

### B. GRAF Overview

GRAF [14] is a model that builds from NeRF and defines it within an Generative Adversarial Network (GAN). It consists of a generator  $G_\theta$  that predicts an image patch  $\mathbf{P}_{\text{pred}}$  and a discriminator  $D_\phi$  that compares the predicted patch to a patch  $\mathbf{P}_{\text{real}}$  extracted from a real image. GRAF has shown an effective capacity to disentangle 3D shape and viewpoint of objects from 2D images alone, in contrast to the original NeRF [12] and similar approaches such as [19]. Therefore, we aim to translate GRAF's methods to our task, and in subsection II-C we describe our new discriminator architecture, which allows us to disentangle 3D properties from DRRs.

We consider the experimental setting to obtain the radiation attenuation response instead of the color used in natural images. To obtain the attenuation response at a pixel location for an arbitrary projection  $\mathbf{K}$  with pose  $\xi$ , first, we consider a pattern  $\nu = (u, s)$  to sample  $R$  X-ray beams within a  $K \times K$  image-patch  $\mathbf{P}$ . Then, we sample  $N$  3D points  $\mathbf{x}_r^i$  along the X-ray beam  $r$  originating from the pixel location and ordered between the near and far planes of the projection (Fig. 1a).




Fig. 1. An overview of GRAF's generator.

The object representation is encoded in a multi-layer perceptron (MLP) that takes as input a 3D position  $\mathbf{x} = (x, y, z)$  and a viewing direction  $\mathbf{d} = (\theta, \phi)$ , and produces as output a density scalar  $\sigma$  and a pixel value  $c$ . To learn high-frequency features, the input is mapped into a  $2L$ -dimensional representation (Fig. 1b):

$$\gamma(p) = \dots, \cos(2^j \pi p), \sin(2^j \pi p), \dots \quad (1)$$

where  $p$  represents the 3D position or viewing direction, for  $j = 0, \dots, m - 1$ .

For modeling the shape and appearance of anatomical structures, let  $\mathbf{z}_s \sim p_s$  and  $\mathbf{z}_a \sim p_a$  be the latent codes sampled from a standard Gaussian distribution, respectively (Fig. 1c). To obtain the density prediction  $\sigma$ , the shape encoding  $\mathbf{q}$  is transformed to volume density through a density head  $\sigma_\theta$ . Then, the network  $g_\theta(\cdot)$  operates on a shape encoding  $\mathbf{q} = (\gamma(\mathbf{x}), \mathbf{z}_s)$  that is later concatenated with the positional encoding of  $\mathbf{d}$  and appearance code  $\mathbf{z}_a$  (Fig. 1c):

$$(\gamma(\mathbf{x}), \mathbf{z}_s) \mapsto \mathbf{q} \quad (2)$$

$$(\mathbf{q}(\mathbf{x}, \mathbf{z}_s), \gamma(\mathbf{d}), \mathbf{z}_a) \mapsto c \quad (3)$$

$$\mathbf{q}(\mathbf{x}, \mathbf{z}_s) \mapsto \sigma \quad (4)$$

The final pixel response  $c_r$  is computed by the compositing operation (Fig. 1c):

$$c_r = \sum_{i=1}^N c_r^i \alpha_r^i \exp\left(-\sum_{j=1}^{i-1} \sigma_r^j \delta_r^j\right) \quad (5)$$

where  $\alpha_r^i = 1 - \exp(-\sigma_r^i \delta_r^i)$  is the alpha compositing value of sampled point  $i$  and  $\delta_r^i = \|\mathbf{x}_r^{i+1} - \mathbf{x}_r^i\|_2$  is the distance between the adjacent sampled points.

In this way, both the density and pixel values are computed at each sampled point along the beam  $r$  with network  $g_\theta$ . Finally, combining the results of all  $R$  beams, the generator  $G_\theta$  predicts an image patch  $\mathbf{P}_{\text{pred}}$ , as illustrated in Fig. 1d.

### C. MedNeRF

We investigate how we can adapt GRAF to the medical domain and apply it to render a volumetric representation from DRRs. Leveraging a large dataset, GRAF's discriminator  $D_\phi$  is able to continuously provide useful signals to train the generator  $G_\theta$ . However, medical datasets like those

considered in our problem are generally small, which causes two sequential issues:

**The lack of real information to the generator:** In GRAF (and in GAN in general), the only source of features of the training data contributing to the generator is the indirect gradient transferred from the discriminator. We find that the single convolutional feedback from GRAF’s discriminator poorly conveys refined features from DRRs resulting in inaccurate volumetric estimation.

**Brittle adversarial training:** With a limited training dataset, the generator or discriminator may fall into ill-posed settings such as mode collapse, which would lead to generating a limited number of instances and consequently, a suboptimal data distribution estimation. While some works have applied data augmentation techniques to leverage more data in the medical domain, some transformations could mislead the generator to learn the infrequent or even non-existent augmented data distribution [20]. We find that naively applying classic data augmentation works less favorably than our adopted framework.

### 1) Self-supervised Learning for High-Fidelity Synthesis:

To allow *richer* feature-maps covering from the DRRs such that it produces more comprehensive signals to train  $G_\theta$ , we replace GRAF’s discriminator architecture with recent advancements in self-supervised approaches. We allow  $D_\phi$  to learn useful global and local features training it on a pretext task, in particular, the self-supervision method based on auto-encoding [21]. Different from [21], we only use two decoders for the feature-maps on scales:  $f_1$  on  $32^2$  and  $f_2$  on  $8^2$  (Fig. 2a). We find that this choice allows better performance and enables a correct volumetric depth estimation.  $D_\phi$  must therefore not only discriminate  $\mathbf{P}_{\text{pred}}$  predicted from  $G_\theta$  but also extract comprehensive features from real image patches  $\mathbf{P}_{\text{real}}$  that enable the decoders to resemble the data distribution.

To assess global structure in decoded patches from  $D_\phi$ , we use the Learned Perceptual Image Patch Similarity (LPIPS) metric [22]. We compute the weighted pairwise image distance between two VGG16 feature spaces, where the pretrained weights are fit to better match human perceptual judgments. The additional discriminator loss is therefore:

$$\mathcal{L}_r = \mathbb{E}_{\mathbf{f} \sim D(\mathbf{p}), \mathbf{p} \sim \mathcal{P}} \left[ \frac{1}{whd} \|\phi_i(\mathcal{G}(\mathbf{f})) - \phi_i(\mathcal{T}(\mathbf{p}))\|_2 \right] \quad (6)$$

where  $\phi_i(\cdot)$  denotes the  $i$ th layer output of a pretrained VGG16 network, and  $w$ ,  $h$ , and  $d$  stand for the width, height and depth of a feature space, respectively. Let  $\mathcal{G}$  be the processing on the intermediate feature-maps  $\mathbf{f}$  from  $D_\phi$ , and  $\mathcal{T}$  the processing on real image patches. When coupled with this additional reconstruction loss, the network learns representations that transfer across tasks.

### 2) Improving Learning via Data Augmentation:

We improve learning of  $G_\theta$  and  $D_\phi$  by adopting the Data Augmentation Optimized for GAN (DAG) framework [20] in which a data augmentation transformation  $\mathcal{T}_k$  (Fig.




Fig. 2. An overview of our discriminator with self-supervised learning and DAG.

2b) is applied using multiple discriminator heads  $\{D_k\}$ . To further reduce memory usage, we share all layers of  $D_\phi$  except the last layers corresponding to each head (Fig. 2c). Because applying differentiable and invertible data augmentation transformations  $\mathcal{T}_k$  has the Jensen-Shannon (JS) preserving property [20]:

$$\text{JS}(p_d^{\mathcal{T}_k} \parallel p_g^{\mathcal{T}_k}) = \text{JS}(p_d \parallel p_g) \quad (7)$$

where  $p_d^{\mathcal{T}_k}$  is the transformed training data distribution and  $p_g^{\mathcal{T}_k}$  the transformed distribution captured by  $G_\theta$ . By using a total of four transformations combining flipping and rotation, we encourage optimization to the original data distribution, which also brings the most performance boost. These choices allow our model to benefit from not only  $\text{JS}(p_d \parallel p_g)$  but also  $\text{JS}(p_d^{\mathcal{T}_k} \parallel p_g^{\mathcal{T}_k})$ , thereby improving the learning of  $G_\theta$  and generalization of  $D_\phi$ . Furthermore, using multiple discriminators with weight-sharing provides learning regularization of  $D_\phi$ .

Replacing GRAF’s logistic objective with a hinge loss, we then define our overall loss as below:

$$\mathcal{L}(\theta, \{\phi_k\}) = \mathcal{L}(\theta, \phi_0) + \frac{\lambda}{n-1} \sum_{k=1}^n \mathcal{L}(\theta, \phi_k) \quad (8)$$

$$\begin{aligned} \mathcal{L}(\theta, \phi_k) = & \mathbb{E}_{\mathbf{z}_s \sim p_s, \mathbf{z}_a \sim p_a, \xi \sim p_\xi, \nu \sim p_\nu} [f(D_\phi(G_\theta(\mathbf{z}_s, \mathbf{z}_a, \xi, \nu)))] \\ & + \mathbb{E}_{\mathbf{I} \sim p_D, \nu \sim p_\nu} [f(-D_\phi(\mathbf{I}, \nu))] + \mathcal{L}_r \end{aligned} \quad (9)$$

where  $f(u) = \max(0, 1 + u)$ . We optimize this loss with  $n = 4$ , where  $k = 0$  corresponds to the identity transformation and  $\lambda = 0.2$  (as in [20]).

### 3) Volumetric Rendering from a Single View X-ray:

After training a model, we reconstruct the complete X-ray projections within a full vertical rotation of a medical instance given a single view X-ray. We follow the relaxed reconstruction formulation in [23], which fits the generator to a single image. Then, we allow the parameters of the generator  $G_\theta$  to be *slightly* fine-tuned along with the shape and appearance latent vectors  $\mathbf{z}_s$  and  $\mathbf{z}_a$ . The distortion and perception tradeoff is well known in GAN methods [24] and




Fig. 3. Knee renderings from continuous viewpoint rotations showing tissue and bone. Given a single-view X-ray from a CT, we can generate the complete set of CT-projections within a full vertical rotation by *slightly* fine-tuning a pretrained model along with the shape and appearance latent codes.

TABLE I. Quantitative results based on PSNR and SSIM of rendered X-ray projections with single-view X-ray input.

Dataset	$\uparrow$ PSNR (dB) ( $\mu \pm \sigma$ )	$\uparrow$ SSIM ( $\mu \pm \sigma$ )
Knee	<b>30.17</b> $\pm$ 1.93	<b>0.670</b> $\pm$ 0.040
Chest	<b>28.54</b> $\pm$ 0.79	<b>0.462</b> $\pm$ 0.082

therefore we modify our generation objective by adding the distortion Mean Square Error (MSE) loss, which incentivises a balance between blurriness and accuracy:

$$\mathcal{L}_{\text{gen}} = \lambda_1 \mathcal{L}_r(\text{VGG16}) + \lambda_2 \mathcal{L}_{\text{MSE}}(G) + \lambda_3 \mathcal{L}_{\text{NLL}}(z_s, z_a) \quad (10)$$

where NLL corresponds to the negative log-likelihood loss and the tuned hyperparameters  $lr = 0.0005$ ,  $\beta_1 = 0$ ,  $\beta_2 = 0.999$ ,  $\lambda_1 = 0.3$ ,  $\lambda_2 = 0.1$  and  $\lambda_3 = 0.3$ .

Once the model locates an optimal combination of  $z_s$  and  $z_a$ , we replicate them and use them to render the rest of the X-ray projections by continuously controlling the angle viewpoint.

### III. RESULTS

Here we provide an evaluation of MedNeRF on our datasets. We compare our model’s results to the ground truth, two baselines, perform an ablation study, and show qualitative and quantitative evaluations. We train all models for 100,000 iterations with a batch size of 8. Projection parameters  $(u, v)$  are chosen to evenly sample points on the surface of a sphere, specifically a slight horizontal elevation of 70-85 degrees and  $u_{\min} = 0$ ,  $u_{\max} = 1$  for a full 360-degree vertical rotation. However, we only provide a fifth of the views (72-views each at five degrees) during training and let the model render the rest.

#### A. Reconstruction from Single View X-ray

We evaluate our model’s representation for 3D-aware DRR synthesis given a single-view X-ray as input. We find that despite the implicit linear network’s limited capacity, our model can disentangle 3D anatomy identity and attenuation




Fig. 4. Volumetric maps and attenuation renderings on our dataset.

response of different medical instances, which are retrieved through the described reconstruction reformulation in II-C.3. Our model can also facilitate distinguishing bone from tissue via a contrast transformation, as it renders a brighter pixel value for denser structures (e.g. bone) (Fig. 3).

Table I summarises our results based on the peak signal-to-noise ratio (PSNR) and structural similarity (SSIM), which measure the quality of reconstructed signals and human subjective similarity, respectively. We find that our generative loss can achieve a reasonable perception-distortion curve in renderings and show consistency with the location and volumetric depth of anatomical structures at continuous viewpoints compared to the ground truth.

#### B. 2D DRR Rendering

We evaluate our model on the task of 2D rendering and compare it to pixelNeRF [19], and GRAF [14] baseline,

TABLE II. FID and KID analysis comparing other methods.

Method	Chest dataset		Knee dataset	
	↓ FID ( $\mu \pm \sigma$ )	↓ KID ( $\mu \pm \sigma$ )	↓ FID ( $\mu \pm \sigma$ )	↓ KID ( $\mu \pm \sigma$ )
GRAF [14]	68.25 $\pm$ 0.954	0.053 $\pm$ 0.0008	76.70 $\pm$ 0.302	0.058 $\pm$ 0.0001
pixelNeRF [19]	112.96 $\pm$ 2.356	0.084 $\pm$ 0.0012	166.40 $\pm$ 2.153	0.158 $\pm$ 0.0010
Ours	<b>60.26 <math>\pm</math> 0.322</b>	<b>0.041 <math>\pm</math> 0.0005</b>	<b>76.12 <math>\pm</math> 0.193</b>	<b>0.052 <math>\pm</math> 0.0004</b>

TABLE III. FID and KID analysis of ablations of our model.

Ablation	Chest dataset	
	↓ FID ( $\mu \pm \sigma$ )	↓ KID ( $\mu \pm \sigma$ )
MedNeRF - 3 SD, logistic loss, classic DA	84.85 $\pm$ 1.025	0.069 $\pm$ 0.0031
MedNeRF - 2 SD, logistic loss, classic DA	67.73 $\pm$ 0.712	0.051 $\pm$ 0.0006
MedNeRF - 2 SD, hinge loss, classic DA	65.34 $\pm$ 0.353	0.045 $\pm$ 0.0004
MedNeRF - 2 SD, hinge loss, DAG	<b>60.26 <math>\pm</math> 0.322</b>	<b>0.041 <math>\pm</math> 0.0005</b>

wherein the original architecture is used. Our model can more accurately estimate volumetric depth compared to GRAF and pixelNeRF (Fig. 4). For each category, we find an unseen target instance with a similar view direction and shape. Volumetric depth estimation is given by bright colors (far) and dark colors (near). Lacking a perceptual loss, GRAF is not incentivized to produce high-frequency textures. In contrast, we find our model renders a more detailed internal structure with varied attenuation. GRAF produces a consistent attenuation response, but seems to be unable to distinguish the anatomical shape from the background. Our self-supervised discriminator enables the generator to disentangle shape and background by rendering a brighter color for the background and a darker color for the shape, while GRAF renders a bright or dark color for both.

We find pixelNeRF produces blurred attenuation renderings for all datasets, and volumetric maps tend to exhibit strong color shifts (Fig. 4). We believe these artifacts are due to the see-through nature of the dataset, compared to solid-like natural objects on which NeRFs are trained. This data characteristic impairs not only volumetric maps but also fine anatomical structures. In contrast, our model is better able to render both volumetric depth and attenuation response. We also find pixelNeRF is sensitive to slight changes in projection parameters, hampering optimization for the knee category. Our model produces a consistent 3D geometry and does not rely on explicit projection matrices.

Table II compares image quality based on Fréchet Inception Distance (FID) and Kernel Inception Distance (KID) metrics, in which lower values mean better. Optimizing pixelNeRF on our datasets leads to particularly poor results that are unable to compete with the GRAF baseline and our model. In contrast, our model outperforms the baselines on FID and KID metrics for all datasets.

### C. Ablation Study

We evaluate our model with three ablations (Table III): wherein an additional simple decoder (SD) is included; the adversarial logistic loss is replaced by its hinge version; and wherein the non-classical DAG approach is adopted. We find that that the DAG approach brings the most performance boost compared to naively applying classical DA, while the use of a hinge loss performs slightly better than its

logistic version. However, an additional decoder in our self-supervised discriminator can lead to a significant drop in performance.

## IV. CONCLUSION

We have presented a novel Deep Learning architecture based on Neural Radiance Fields for learning a continuous representation of CT scans. We learn a medical category encoding of the attenuation response of a set of 2D DRRs in the weights of a generator. Furthermore, we have found that a stronger and more comprehensive signal from our discriminator allows generative radiance fields to model 3D-aware CT-projections. Experimental evaluation demonstrates significant qualitative and quantitative reconstructions and improvements over other Neural Radiance Field approaches. Whilst the proposed model may not replace CT entirely, the functionality of generating 3D-aware CT-projections from X-rays has great potential for clinical use in osseous trauma, skeletal evaluation in dysplasia and for orthopaedic pre-surgical planning. This could cut down on the radiation dose given to patients, with significant economic implications such as bringing down the cost of investigations.

## ACKNOWLEDGMENT

This work is partially supported by the Mexican Council of Science and Technology (CONACyT).

## REFERENCES

- [1] Paul Suetens, *Visualization for diagnosis and therapy*, p. 190–218, Cambridge University Press, 2 edition, 2009.
- [2] Pechin Lo, Bram van Ginneken, Joseph M. Reinhardt, Tarunashree Yavarna, Pim A. de Jong, Benjamin Irving, Catalin Fetita, Margarete Ortner, Rômulo Pinho, Jan Sijbers, Marco Feuerstein, Anna Fabijanska, Christian Bauer, Reinhard Beichel, Carlos S. Mendoza, Rafael Wiemker, Jaesung Lee, Anthony P. Reeves, Silvia Born, Oliver Weinheimer, Eva M. van Rikxoort, Juerg Tschirren, Ken Mori, Benjamin Odry, David P. Naidich, Ieneke Hartmann, Eric A. Hoffman, Mathias Prokop, Jesper H. Pedersen, and Marleen de Bruijne, “Extraction of airways from ct (exact’09),” *IEEE Transactions on Medical Imaging*, vol. 31, no. 11, pp. 2093–2107, 2012.
- [3] Mary Coffey and Aude Vaandering, “Patient setup for pet/ct acquisition in radiotherapy planning,” *Radiotherapy and Oncology*, vol. 96, no. 3, pp. 298–301, 2010, PET in Radiotherapy Planning.
- [4] Tri Huynh, Yaozong Gao, Jiayin Kang, Li Wang, Pei Zhang, Jun Lian, and Dinggang Shen, “Estimating ct image from mri data using structured random forest and auto-context model,” *IEEE Transactions on Medical Imaging*, vol. 35, no. 1, pp. 174–183, 2016.
- [5] Shoulie Xie, Weimin Huang, Tao Yang, Dajun Wu, and Huiying Liu, “Compressed sensing based image reconstruction with projection recovery for limited angle cone-beam ct imaging,” in *2020 42nd Annual International Conference of the IEEE Engineering in Medicine Biology Society (EMBC)*, 2020, pp. 1307–1310.
- [6] Ge Wang, Jong Chu Ye, Klaus Mueller, and Jeffrey A. Fessler, “Image reconstruction is a new frontier of machine learning,” *IEEE Transactions on Medical Imaging*, vol. 37, no. 6, pp. 1289–1296, 2018.
- [7] Yinsheng Li, Ke Li, Chengzhu Zhang, Juan Montoya, and Guang-Hong Chen, “Learning to reconstruct computed tomography images directly from sinogram data under a variety of data acquisition conditions,” *IEEE Transactions on Medical Imaging*, vol. 38, no. 10, pp. 2469–2481, 2019.
- [8] David B. Lindell, Julien N. P. Martel, and Gordon Wetzstein, “Autoint: Automatic integration for fast neural volume rendering,” in *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2021, pp. 14556–14565.

- [9] Yu Sun, Jiaming Liu, Mingyang Xie, Brendt Wohlberg, and Ulugbek S. Kamilov, "Coil: Coordinate-based internal learning for tomographic imaging," *IEEE Transactions on Computational Imaging*, vol. 7, pp. 1400–1412, 2021.
- [10] Xingde Ying, Heng Guo, Kai Ma, Jian Wu, Zhengxin Weng, and Yefeng Zheng, "X2ct-gan: Reconstructing ct from biplanar x-rays with generative adversarial networks," in *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019, pp. 10611–10620.
- [11] Bao Nguyen, Adam Feldman, Sarath Bethapudi, Andrew Jennings, and Chris G. Willcocks, "Unsupervised region-based anomaly detection in brain mri with adversarial image inpainting," in *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, 2021, pp. 1127–1131.
- [12] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng, "Nerf: Representing scenes as neural radiance fields for view synthesis," in *Computer Vision – ECCV 2020*, Andrea Vedaldi, Horst Bischof, Thomas Brox, and Jan-Michael Frahm, Eds., Cham, 2020, pp. 405–421, Springer International Publishing.
- [13] Ricardo Martin-Brualla, Noha Radwan, Mehdi S. M. Sajjadi, Jonathan T. Barron, Alexey Dosovitskiy, and Daniel Duckworth, "NeRF in the Wild: Neural Radiance Fields for Unconstrained Photo Collections," in *CVPR*, 2021.
- [14] Katja Schwarz, Yiyi Liao, Michael Niemeyer, and Andreas Geiger, "Graf: Generative radiance fields for 3d-aware image synthesis," in *Advances in Neural Information Processing Systems (NeurIPS)*, 2020.
- [15] E B Tsai, S Lungren, M P Hershman, M Roshkovan, L Colak, E Erickson, B J Shih, G Stein, A Kalpathy-Cramer, J Shen, J Hafez, M A F John, S Rajiah, P Pogatchnik, B P Mongan, J T Altinmakas, E Ranschaert, E R Kitamura, and C C. Wu, "Medical imaging data resource center (midrc) - rsna international covid open research database (ricord) release 1b - chest ct covid-[dataset]," *The Cancer Imaging Archive*, 2021.
- [16] K Clark, B Vendt, K Smith, J Freymann, J Kirby, P Koppel, S Moore, S Phillips, D Maffitt, M Pringle, L Tarbox, and F Prior, "The cancer imaging archive (tcia): Maintaining and operating a public information repository," *Journal of Digital Imaging*, vol. 26, no. 6, pp. 1045–1057, December 2013.
- [17] Michael D Harris, Adam J Cyr, Azhar A Ali, Clare K Fitzpatrick, Paul J Rullkoetter, Lorin P Maletsky, and Kevin B Shelburne, "A combined experimental and computational approach to subject-specific analysis of knee joint laxity," *Journal of biomechanical engineering*, vol. 138, no. 8, August 2016.
- [18] Azhar A Ali, Sami S Shalhoub, Adam J Cyr, Clare K Fitzpatrick, Lorin P Maletsky, Paul J Rullkoetter, and Kevin B Shelburne, "Validation of predicted patellofemoral mechanics in a finite element model of the healthy and cruciate-deficient knee," *Journal of biomechanics*, vol. 49, no. 2, pp. 302–309, January 2016.
- [19] Alex Yu, Vickie Ye, Matthew Tancik, and Angjoo Kanazawa, "pixelNeRF: Neural radiance fields from one or few images," in *CVPR*, 2021.
- [20] Ngoc-Trung Tran, Viet-Hung Tran, Ngoc-Bao Nguyen, Trung-Kien Nguyen, and Ngai-Man Cheung, "On data augmentation for gan training," *IEEE Transactions on Image Processing*, vol. 30, pp. 1882–1897, 2021.
- [21] Bingchen Liu, Yizhe Zhu, Kunpeng Song, and Ahmed Elgammal, "Towards faster and stabilized {gan} training for high-fidelity few-shot image synthesis," in *International Conference on Learning Representations*, 2021.
- [22] Richard Zhang, Phillip Isola, Alexei A. Efros, Eli Shechtman, and Oliver Wang, "The unreasonable effectiveness of deep features as a perceptual metric," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.
- [23] Xingang Pan, Xiaohang Zhan, Bo Dai, Dahua Lin, Chen Change Loy, and Ping Luo, "Exploiting deep generative prior for versatile image restoration and manipulation," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, pp. 1–1, 2021.
- [24] Yochai Blau and Tomer Michaeli, "The perception-distortion tradeoff," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, June 2018.