

# PD-INR: Prior-Driven Implicit Neural Representations for TOF-PET Reconstruction

Yuxuan Long<sup>1,†</sup>, Yulin Zhang<sup>2,†</sup>, Hong Wang<sup>3</sup>, Xiaodong Kuang<sup>1</sup>, Hailiang Huang<sup>1</sup>, Fan Rao<sup>1</sup>, Huafeng Liu<sup>2</sup>, Yefeng Zheng<sup>4</sup>, and Wentao Zhu<sup>1,2\*</sup>

<sup>1</sup> Zhejiang Lab, Hangzhou, China

<sup>2</sup> The College of Optical Science and Engineering, Zhejiang University, Hangzhou, China  
zhuwentao.ee@gmail.com

<sup>3</sup> Xi'an Jiaotong University, Xi'an, China

<sup>4</sup> Westlake University, Hangzhou, China

**Abstract.** Positron emission tomography (PET) reconstruction is a challenging inverse problem, where projection data often contain low statistics. While current supervised learning methods offer strong noise suppression abilities, they may suffer from generalization issues and are in many cases not accurate quantitatively. To overcome these challenges, we propose a novel self-supervised Time-of-Flight PET (TOF-PET) reconstruction framework that utilizes Implicit Neural Representations (INR) to model PET images. Specifically, we introduce a differentiable forward projection model based on the imaging mechanism for TOF-PET and reformulate TOF-PET reconstruction problem using INR. To enhance image smoothness, we develop a ray-based total variation (TV) regularization term, distinct from the traditional TV. For the internal structure of our INR, we integrate a multi-resolution hash encoder with our designed prior-image encoder, where the latter provides sufficient image prior and always delivers reliable initial reconstructions for arbitrary network depth. Experiments on brain and chest datasets show that our method outperforms traditional iterative algorithms and self-supervised approaches in noise suppression and contrast recovery. Compared to conventional NeRF-based architectures, our model is more compact and converges faster, providing an efficient solution for TOF-PET reconstruction. The source code repository is hosted on GitHub: <https://github.com/zy1123300/PD-INR.git>

**Keywords:** TOF-PET Reconstruction · Inverse Problem · Implicit Neural Representation · Neural Radiance Fields.

## 1 Introduction

Positron emission tomography (PET) is a widely used functional imaging technique for clinical applications [1, 9, 17], such as tumor detection, by tracking the spatial distribution of radioactive tracers. However, PET systems face challenges like low spatial

<sup>†</sup> Equal contribution

\* Corresponding Author

resolution and high noise sensitivity. Time-of-Flight (TOF) PET addresses these limitations by incorporating the time differences of annihilation photons, enhancing spatial resolution, especially in low-dose scenarios.

Although the TOF kernel’s finite width (typically modeled as a Gaussian) introduces some uncertainty in localization, TOF-PET significantly improves imaging performance. However, effective reconstruction algorithms remain crucial for further enhancing PET image quality.

To address the limitations of PET imaging, iterative reconstruction techniques were introduced. One of the most prominent iterative methods is the Maximum Likelihood Expectation Maximization (MLEM) algorithm [11, 24], which iteratively optimizes the solution by incorporating a non-negativity constraint to prevent negative values in the reconstructed image. Compared to analytical methods, MLEM offers better noise suppression, significantly improving image quality in low-signal regions. As a result, iterative reconstruction has become the standard approach in most modern PET systems.

In recent years, deep learning has achieved significant advancements in PET image reconstruction, particularly in enhancing image quality and accelerating the reconstruction process. Deep learning applications in PET reconstruction can be broadly categorized into three main approaches: (1) image post-processing [7, 14], where traditional iterative algorithms are used to generate a preliminary image that may contain noise or artifacts, and a denoising network is then trained to refine the image; (2) end-to-end deep learning models [10, 29] directly map sinogram data to image space, enabling the networks to learn key features from large datasets and reduce reliance on manual intervention. However, these approaches typically require a large amount of paired data and suffer from significant generalization issues when handling new data. (3) hybrid methods that integrate traditional iterative algorithms with deep learning [5, 6], replacing specific steps such as data fitting and regularization with deep neural networks. This combination retains the stability and theoretical rigor of traditional methods while benefiting from the nonlinear mapping capabilities of deep learning, but still faces issues such as the need for pre-training and generalization problems. In light of these challenges, a self-supervised reconstruction method needs to be proposed, which can optimize the image reconstruction process by leveraging the inherent properties of the data itself, without relying on large amounts of labeled data.

Recently, Implicit Neural Representation (INR) [18] has gained significant attention, particularly through its success in Neural Radiance Fields (NeRF). NeRF uses multilayer perceptrons (MLPs) to model the mapping from spatial coordinates to color and density in continuous volumetric fields, enabling high-quality image reconstruction from sparse or noisy data. This framework has been extended to medical imaging [3, 4, 15, 16, 19, 22, 23, 27], with applications in reconstruction, segmentation, and registration. For example, IREM [19] leverages implicit functions of 3D spatial coordinates to reconstruct high-resolution images from sparse inputs, while D-NAF [16] employs continuous 4D implicit neural representation for dynamic 4D-CBCT reconstruction, achieving superior image quality and efficiency. Other methods, such as CuNeRF [3], enable zero-shot super-resolution through cubic sampling and hierarchical rendering, further advancing medical image reconstruction. However, the application of INR for TOF-PET reconstruction remains unexplored so far.

In this work, our main contributions are summarized as follows:

- We propose a novel self-supervised method for TOF-PET reconstruction. To the best of our knowledge, it is the first exploration of INR’s potential in TOF-PET.
- We develop a novel framework called prior-driven INR (PD-INR), which incorporates image prior into the conventional INR. With only a single-layer MLP, PD-INR could still achieve promising reconstruction quality and convergence.
- We introduce a novel ray-wise TV regularization term in the feature domain, specially customized for our PD-INR.

## 2 Method

### 2.1 TOF-PET Image Reconstruction

In TOF-PET, assuming that the TOF events have been already pre-corrected, the expected true count  $\bar{y}_{ik}$  along line-of-response (LOR)  $i$  at TOF-bin  $k$  can be expressed as  $\bar{y}_{ik} = \sum_j p_{ijk} x_j$ , where  $x_j$  is the activity at voxel  $j$ ,  $p_{ijk}$  is the sensitivity of the detector at  $(i, k)$  for activity in  $j$ . Assuming that TOF-PET data  $\mathbf{y}$  are independently Poisson distributed, the log-likelihood function can be expressed as:

$$\mathcal{L}(\mathbf{y}|\mathbf{x}) = \sum_{i,k} y_{ik} \log \bar{y}_{ik} - \bar{y}_{ik} - \log y_{ik}!. \quad (1)$$

The activity image is usually reconstructed by maximizing the Poisson log-likelihood:

$$\hat{\mathbf{x}} = \arg \max_{\mathbf{x} \geq 0} \mathcal{L}(\mathbf{y}|\mathbf{x}). \quad (2)$$

### 2.2 Prior-Driven INR for TOF-PET Reconstruction

**Mathematical Formulation.** Conventional algorithms reconstruct the image in discretized image space. Here, we adopt INR to model our activity image. Besides the continuity of INR, the internal structure of INR would also bring some implicit prior into the image. For the INR, the activity image is expressed as:

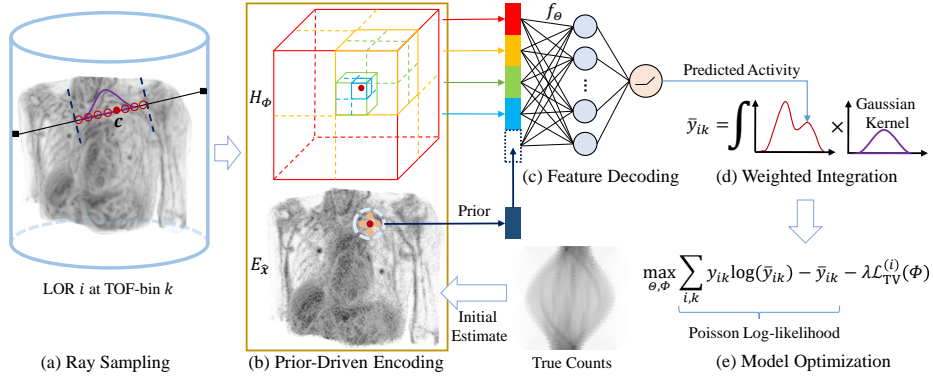
$$x_j = f_{\Theta}(S_{\Phi}(\mathbf{c}_j)) \quad \forall j, \quad (3)$$

where  $\mathbf{c}_j \in \mathbb{R}^3$  is 3D coordinate of the voxel  $j$ ,  $S_{\Phi}(\cdot)$  encodes the spatial coordinates into image features, and  $f_{\Theta}(\cdot)$  decodes the image features into the image intensities. Here,  $\Phi$  and  $\Theta$  are both learnable parameters for our image representation.

To adapt to such image representation, the traditional system matrix may need to be re-defined in order to fully benefit from the expressive power of neural networks. Like NeRF, we would like to reformulate the forward projection in a ray-wise manner, in order to ease batch computation during iterative optimization.

For TOF-PET, the forward projection along the LOR  $i$  at TOF-bin  $k$  is modelled as a line integral, i.e.,

$$\bar{y}_{ik} = \int_{-\infty}^{\infty} \Psi(\mathbf{x}, \mathbf{s}_i + t\mathbf{d}_i) h(t - t_{ik}) dt, \quad (4)$$



**Fig. 1.** Overview of the proposed PD-INR framework. (a) Along LOR  $i$ , numerous points are sampled within TOF-bin  $k$ . (b) For each point, the hash encoder  $H_\Phi$  maps its coordinate  $\mathbf{c}$  to multi-resolution features while  $E_{\hat{\mathbf{x}}}$  extracts the neighboring intensities from the initial estimate  $\hat{\mathbf{x}}$  as the prior feature. (c) All features are concatenated, then decoded by a tiny MLP  $f_\Theta$ . (d) The predicted activities are integrated with the TOF profile, yielding the estimated counts  $\bar{y}_{ik}$ . (e) For all LORs and TOF-bins, the model parameters are optimized and regularized by our designed TV term in feature domain.

where the LOR  $i$  has the starting point  $\mathbf{s}_i \in \mathbb{R}^3$  and the unit direction  $\mathbf{d}_i \in \mathbb{R}^3$ ,  $t_{ik}$  is the TOF variable,  $h(\cdot)$  is the TOF profile often modeled as a truncated Gaussian function, and  $\Psi$  queries the image intensity at the specified spatial position. For conventional forward projection,  $\Psi$  linearly interpolates a discrete image. In the case of INR, we have  $\Psi(\mathbf{x}, \mathbf{s}_i + t\mathbf{d}_i) = f_\Theta(S_\Phi(\mathbf{s}_i + t\mathbf{d}_i))$ . By substituting it into (2) and further introducing a penalty term, the reconstruction problem is finally re-formulated as:

$$\begin{aligned} \max_{\Theta, \Phi} \quad & \mathcal{L}(\mathbf{y}|\Theta, \Phi) - \lambda \sum_i \mathcal{L}_{\text{TV}}^{(i)}(\Phi), \\ \text{s.t.} \quad & \bar{y}_{ik} = \int_{-\infty}^{\infty} f_\Theta(S_\Phi(\mathbf{s}_i + t\mathbf{d}_i))h(t - t_{ik}) dt, \quad \forall i, k, \end{aligned} \quad (5)$$

where  $\mathcal{L}_{\text{TV}}^{(i)}(\Phi)$  is a TV term that enhances the image smoothness, and  $\lambda$  is some small non-negative weight that controls the magnitude of penalty. Here,  $\mathcal{L}_{\text{TV}}^{(i)}(\Phi)$  is defined as:

$$\mathcal{L}_{\text{TV}}^{(i)}(\Phi) = \frac{1}{(P-1)F} \sum_{p=1}^{P-1} \|S_\Phi(\mathbf{v}_{p+1}^{(i)}) - S_\Phi(\mathbf{v}_p^{(i)})\|_1, \quad (6)$$

where  $\{\mathbf{v}_p^{(i)}\}_{p=1}^P$  are the coordinates for the sampled points along LOR  $i$ , with  $\mathbf{v}_p^{(i)} = \mathbf{s}_i + t_p\mathbf{d}_i$ , and the distances  $\{t_p\}_{p=1}^P$  are of ascending order.  $F$  is the feature dimension, i.e.,  $S_\Phi(\mathbf{v}_p^{(i)}) \in \mathbb{R}^F$ . Those sampled points are also used for numerical approximation of the integral in (5), hence avoiding re-computations. Conventional TV penalty is computed in discrete image space, and may be time-consuming for INR as all voxels' activities are to be evaluated. Our TV term (6) is ray-wise (or LOR-wise), implying small computation load. It only penalizes on the difference of neighbouring features

along each LOR. This is nearly equivalent to a ray-wise TV penalty in image domain since all coordinates share the same feature decoder when evaluating the activities.

**Prior-Driven Encoding.** TOF-PET reconstruction problem has been formulated using INR, as shown in (5). Here, we would like to discuss our prior-driven INR (PD-INR) framework in details, especially the prior-driven encoding function  $S_\Phi$ . Specifically,  $S_\Phi$  concatenates the features computed independently from two encoders, i.e.,  $S_\Phi(\mathbf{c}) = [H_\Phi(\mathbf{c})^T, E_{\hat{\mathbf{x}}}(\mathbf{c})^T]^T$ , where  $H_\Phi(\cdot)$  is the learnable hash encoder [20] and  $E_{\hat{\mathbf{x}}}(\cdot)$  is the prior-image encoder. These two encoders are respectively expressed as:

$$H_\Phi(\mathbf{c}) = \oplus_{l=1}^L \Psi \left( H^{(l)}, \mathbf{c} \right), \quad E_{\hat{\mathbf{x}}}(\mathbf{c}) = \oplus_{\mathbf{c}' \in \mathcal{N}(\mathbf{c})} \Psi \left( \hat{\mathbf{x}}, \mathbf{c}' \right), \quad (7)$$

where  $\oplus$  denotes feature concatenation and  $\Psi$  is the linear interpolation function that retrieves the features given the 3D coordinate.

For the hash encoder, there are a total of  $L$  levels of features for concatenation, and  $H^{(l)}$  is the  $l^{\text{th}}$  level cube-like feature map stored in a trainable hash table. The introduction of hash tables not only compresses the number of parameters but also induces the many-to-one property due to hash collisions, which coincides with the sparsity nature of human organs, as claimed in NAF [28]. The hash encoder also concatenates multi-resolution features from each level, which facilitates learning of both low- and high-frequency details, thereby enhancing both image contrast and smoothness.

For the prior-image encoder, as shown in (7),  $\mathcal{N}(\mathbf{c})$  denotes the set of neighbors (including itself) of the coordinate  $\mathbf{c}$ , and  $\hat{\mathbf{x}}$  is the prior image obtained from the projection data.  $E_{\hat{\mathbf{x}}}(\mathbf{c})$  extracts the image intensities from the neighbours and merges them as one feature vector, hence providing the image prior to our PD-INR. It is expected to facilitate convergence and strengthen the expressive power of our spatial encoding.


We design the feature decoder  $f_\Theta(\cdot)$  as a tiny MLP, which is sufficient to handle the reconstruction task. Here, the feature decoder is initialized in the way that our PD-INR start with a prior image. Specifically, the weights of the final layer are initialized as  $[\mathbf{0}_{LF}^T, \mathbf{w}^T]^T$ , such that  $\mathbf{1}^T \mathbf{w} = 1$ ; for the other layers, the weights are simply initialized as identity matrices, while there are no bias in our MLP. And ReLU is used as the activation function in the MLP and also guarantees a non-negative image. For such an MLP of arbitrary number of layers, it could be found that the image intensity  $f_\Theta(S_\Phi(\mathbf{c}))$  is initially set as  $\mathbf{w}^T E_{\hat{\mathbf{x}}}(\mathbf{c})$ , which is just a weighted average of neighbouring intensities in the prior image. In this way, our PD-INR could always start from a prior image since the very beginning, without any need for warm-up fitting.

Since then, our TOF-PET formulation and the internal design of PD-INR have been discussed in details. The overall framework is displayed in Figure 1.

### 3 Experiments

#### 3.1 Data and Settings

**Data.** We conduct experiments on two datasets: 1) Brain data, which are publicly available 3D PET images from BrainWeb [12]; 2) Chest data, which are clinical PET images




**Fig. 2.** TOF-PET reconstruction of different methods on brain data at various doses.

collected from hospitals for patient diagnosis. Specifically, to investigate the algorithm’s performance in tumor recovery, we insert multiple solid spheres with Gaussian textures into the chest images. The projection data are generated using the Joseph3D algorithm from the PyTomography library [21], with the number of TOF bins set to 13, and the TOF time resolution set to 200 ps. For all data, the image size is  $128 \times 128 \times 96$ , with a voxel size of  $1.5 \times 1.5 \times 1.5 \text{ mm}^3$ , and the sinogram size is  $224 \times 449 \times 4096 \times 13$ .

**Baselines.** We compare our **PD-INR** with six baseline methods. **OSEM** [11] is a classical iterative reconstruction techniques for PET. **MAP** [13], is an iterative algorithm with enforced TV penalty. Non-Local Mean (**NLM**) [2] is an image post-processing method that preserves fine details while reducing noise. We also consider self-supervised learning methods, **Noise2Void** [25] and **DIPRecon** [5], which perform image restoration without requiring labeled data. **DenoisingGAN** [8] is used as a post-processing denoising module and serves as a representative of supervised learning methods.

**Implementation details.** We implement our framework in PyTorch and use Adam optimizer to solve (5). All experiments are run on a single NVIDIA V100 GPU with 32GB of RAM. The learning rate is set at  $2 \times 10^{-3}$  and progressively reduced to  $2 \times 10^{-4}$ , with a batch size of 16,384 rays (LORs) per iteration. For our framework, the prior image is obtained by performing an initial OSEM reconstruction on the raw projection data. We take single-layer MLP as the feature decoder and set  $L = 4$  for  $H_\phi$ . Through hyperparameter analysis, as shown on the left of Figure 5, our TV weight  $\lambda$  is set to  $5 \times 10^{-5}$ .

To evaluate the reconstruction quality, we employ PSNR and SSIM metrics [26]. PSNR measures artifact suppression, while SSIM assesses perceptual differences, with higher values indicating more accurate reconstruction. For local lesion recovery, we



**Fig. 3.** TOF-PET reconstruction results of different methods on chest data with enlarged tumor ROI.

**Table 1.** PSNR (dB) and SSIM of different methods on brain data at various doses.

Method	EM [24]	MAP [13]	NLM [2]	N2V [25]	DIP [5]	GAN [8]	PD-INR
100%dose	31.4/.936	30.8/.948	31.5/.942	31.2/.680	30.8/.939	30.4/.954	<b>31.7/.964</b>
10%dose	20.2/.905	23.6/.927	23.8/.916	25.1/.653	25.2/.926	25.3/.916	<b>27.5/.942</b>
1%dose	19.6/.823	22.8/.875	22.5/.871	18.2/.642	24.4/.888	24.0/.844	<b>25.0/.893</b>

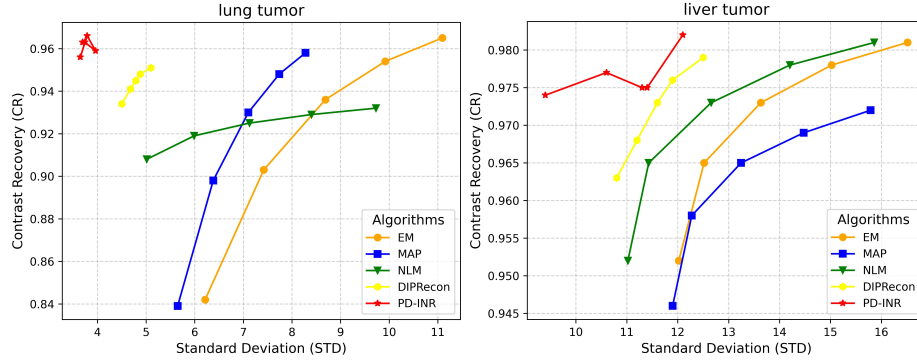
analyze CR and STD curves [5]. CR evaluates the contrast recovery between the tumor region and the background, and STD reflects the variability of background noise, together providing a comprehensive assessment of lesion recovery quality. All metrics are averaged over 20 independent noise realizations to ensure statistical reliability.

### 3.2 Results

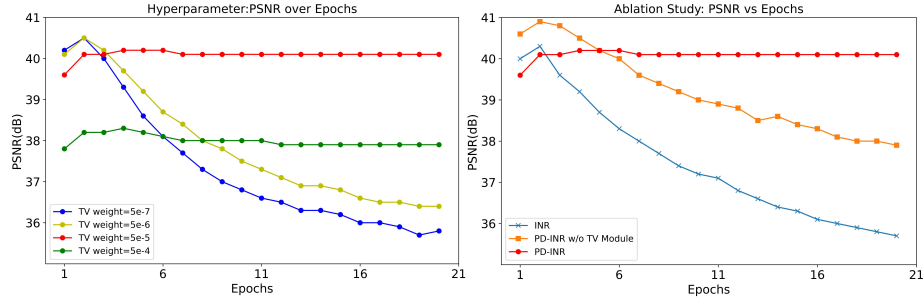
**Quantitative Evaluation.** Table 1 presents the experimental results on the brain data. Our proposed **PD-INR** consistently achieves highest PSNR and SSIM. Notably, as the dose decreases, the performance of classical iterative algorithms (e.g. **EM**) declines significantly, primarily due to the low statistics in the data. In contrast, **PD-INR** demonstrates a clear advantage in the case of low dose, outperforming even the supervised-learning method **DenoisingGAN** (abbreviated as **GAN** in the table). Thanks to our prior-driven encoding and designed TV regularization.

Figure 4 depicts the experimental results for small lesion recovery on the chest data, displaying CR and STD of several iterative algorithms. Results for lung and liver tumors are presented respectively. **PD-INR** shows relatively stable results, where its CR approaches to 1 in few epochs. This is explained by our inserted image prior, which enables fast convergence. Compared with other methods, **PD-INR** better suppresses noise with relatively small STD.

**Qualitative Evaluation.** We also conduct visual comparison of different methods for qualitative evaluation. The reconstruction results are shown in Figure 2 and Figure 3.



**Fig. 4.** Comparison of different methods with CR and STD for tumor recovery.



**Fig. 5.** Results of hyperparameter analysis (left) and ablation studies (right) on chest data using PD-INR.

It could be observed that our **PD-INR** demonstrates superior overall reconstruction quality and more accurate recovery of small lesion. Compared to other methods, **PD-INR** could well preserve the image details while exhibiting robustness to data noises.

**Ablation Studies.** We conducted ablation studies on the chest data. We first found that removal of our prior-image encoder in PD-INR led to failure in convergence, possibly due to lack of expressive power. Then, we compare our PD-INR with following settings: 1) conventional INR framework, i.e. hash encoder + four-layer MLP; 2) removal of the TV regularization from our PD-INR. The results are shown on the right of Figure 5. As expected, removal of regularization (i.e. the orange curve) encouraged PD-INR to overfit to statistical noises, hence PSNR starts to continuously drop after reaching the peak. With TV penalty (i.e. red curve), PD-INR converged to a steady estimate. Besides, it could be observed that our PD-INR (i.e. the orange curve) always performed better than conventional INR (i.e. the blue curve), demonstrating the apparent advantage of our prior-driven encoding.



## 4 Conclusion

This paper proposes a self-supervised learning framework based on INR for TOF-PET reconstruction. We re-formulate TOF-PET reconstruction problem under INR framework, considering the imaging physics behind TOF-PET. To enhance image smoothness while preserving structural details, we introduce ray-based TV regularization that is specially customized for INR. Besides, we propose to employ multi-resolution hash encoder for effective feature learning, combined with our designed image-prior encoder. Experiments on brain and chest data demonstrate that our PD-INR method outperforms existing baselines in noise suppression and contrast recovery, particularly for low-dose PET imaging. In the future, it is highly possible to extend our method to other imaging modalities.

**Acknowledgments.** This work was supported by the Young Scientists Fund of the National Natural Science Foundation of China (Grant No. 62401516), and the National Natural Science Foundation of China (Grant No. 62427807).

**Disclosure of Interests.** The authors have no competing interests to declare that are relevant to the content of this article.

## References

1. Beyer, T., Townsend, D.W., Brun, T., Kinahan, P.E., Charron, M., Roddy, R., Jerin, J., Young, J., Byars, L., Nutt, R.: A combined pet/ct scanner for clinical oncology. *Journal of nuclear medicine* **41**(8), 1369–1379 (2000)
2. Buades, A., Coll, B., Morel, J.M.: A non-local algorithm for image denoising. In: 2005 IEEE computer society conference on computer vision and pattern recognition (CVPR’05). vol. 2, pp. 60–65. Ieee (2005)
3. Chen, Z., Yang, L., Lai, J.H., Xie, X.: Cunerf: Cube-based neural radiance field for zero-shot medical image arbitrary-scale super resolution. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 21185–21195 (2023)
4. Corona-Figueroa, A., Frawley, J., Bond-Taylor, S., Bethapudi, S., Shum, H.P., Willcocks, C.G.: Mednerf: Medical neural radiance fields for reconstructing 3d-aware ct-projections from a single x-ray. In: *44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society*. pp. 3843–3848. IEEE (2022)
5. Gong, K., Catana, C., Qi, J., Li, Q.: Pet image reconstruction using deep image prior. *IEEE transactions on medical imaging* **38**(7), 1655–1665 (2018)
6. Gong, K., Guan, J., Kim, K., Zhang, X., Yang, J., Seo, Y., El Fakhri, G., Qi, J., Li, Q.: Iterative pet image reconstruction using convolutional neural network representation. *IEEE transactions on medical imaging* **38**(3), 675–685 (2018)
7. Gong, K., Guan, J., Liu, C.C., Qi, J.: Pet image denoising using a deep neural network through fine tuning. *IEEE Transactions on Radiation and Plasma Medical Sciences* **3**(2), 153–161 (2018)
8. Gong, Y., Shan, H., Teng, Y., Tu, N., Li, M., Liang, G., Wang, G., Wang, S.: Parameter-transferred wasserstein generative adversarial network (pt-wgan) for low-dose pet image denoising. *IEEE transactions on radiation and plasma medical sciences* **5**(2), 213–223 (2020)
9. Gunn, R.N., Slifstein, M., Searle, G.E., Price, J.C.: Quantitative imaging of protein targets in the human brain with pet. *Physics in Medicine & Biology* **60**(22), R363 (2015)

10. Häggström, I., Schmidtlein, C.R., Campanella, G., Fuchs, T.J.: Deeppet: A deep encoder-decoder network for directly solving the pet image reconstruction inverse problem. *Medical image analysis* **54**, 253–262 (2019)
11. Hudson, H.M., Larkin, R.S.: Accelerated image reconstruction using ordered subsets of projection data. *IEEE transactions on medical imaging* **13**(4), 601–609 (1994)
12. Kwan, R.S., Evans, A.C., Pike, G.B.: Mri simulation-based evaluation of image-processing and classification methods. *IEEE transactions on medical imaging* **18**(11), 1085–1097 (1999)
13. Levitan, E., Herman, G.T.: A maximum a posteriori probability expectation maximization algorithm for image reconstruction in emission tomography. *IEEE transactions on medical imaging* **6**(3), 185–192 (1987)
14. Liu, C.C., Qi, J.: Higher snr pet image prediction using a deep learning model and mri image. *Physics in Medicine & Biology* **64**(11), 115004 (2019)
15. Liu, J., Bai, X.: Volumenerf: Ct volume reconstruction from a single projection view. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 743–753. Springer (2024)
16. Long, Y., Lyu, T., Rao, F., Yuan, P., Chen, Y., Zhu, W.: D-naf: Dynamic neural attenuation fields for 4d cbct reconstruction in pulmonary imaging. In: *2024 IEEE International Conference on Bioinformatics and Biomedicine (BIBM)*. pp. 3549–3554. IEEE (2024)
17. Machac, J.: Cardiac positron emission tomography imaging. In: *Seminars in nuclear medicine*. vol. 35, pp. 17–36. Elsevier (2005)
18. Mildenhall, B., Srinivasan, P.P., Tancik, M., Barron, J.T., Ramamoorthi, R., Ng, R.: Nerf: Representing scenes as neural radiance fields for view synthesis. *Communications of the ACM* **65**(1), 99–106 (2021)
19. Molaei, A., Aminimehr, A., Tavakoli, A., Kazerouni, A., Azad, B., Azad, R., Merhof, D.: Implicit neural representation in medical imaging: A comparative survey. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 2381–2391 (2023)
20. Müller, T., Evans, A., Schied, C., Keller, A.: Instant neural graphics primitives with a multi-resolution hash encoding. *ACM transactions on graphics (TOG)* **41**(4), 1–15 (2022)
21. Polson, L.A., Fedrigo, R., Li, C., Sabouri, M., Dzikunu, O., Ahamed, S., Karakatsanis, N., Kurkowska, S., Sheikhzadeh, P., Esquinas, P., et al.: Pytomography: a python library for medical image reconstruction. *SoftwareX* **29**, 102020 (2025)
22. Ran, Y., Li, Y., Ye, Q., Huo, Y., Bai, Z., Sun, J., Chen, J.: Ct-nerf: Incremental optimizing neural radiance field and poses with complex trajectory. *arXiv preprint arXiv:2404.13896* (2024)
23. Shen, L., Pauly, J., Xing, L.: Nerp: implicit neural representation learning with prior embedding for sparsely sampled image reconstruction. *IEEE Transactions on Neural Networks and Learning Systems* **35**(1), 770–782 (2022)
24. Shepp, L.A., Vardi, Y.: Maximum likelihood reconstruction for emission tomography. *IEEE transactions on medical imaging* **1**(2), 113–122 (2007)
25. Song, T.A., Yang, F., Dutta, J.: Noise2void: unsupervised denoising of pet images. *Physics in Medicine & Biology* **66**(21), 214002 (2021)
26. Wang, Z., Bovik, A.C., Sheikh, H.R., Simoncelli, E.P.: Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing* **13**(4), 600–612 (2004)
27. Xu, J., Moyer, D., Gagoski, B., Iglesias, J.E., Grant, P.E., Golland, P., Adalsteinsson, E.: Nesvor: implicit neural representation for slice-to-volume reconstruction in mri. *IEEE transactions on medical imaging* **42**(6), 1707–1719 (2023)
28. 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*. pp. 442–452. Springer (2022)

29. Zhu, B., Liu, J.Z., Cauley, S.F., Rosen, B.R., Rosen, M.S.: Image reconstruction by domain-transform manifold learning. *Nature* **555**(7697), 487–492 (2018)