TOF-GR: Unsupervised PET Reconstruction via Time-of-Flight Gaussian Representation and Kernel-Based Prior Modeling

Zilan Zhou[1], Yulin Zhang[1], Yiming Wan[1], Min Guo[1,2], Yufei Jin[1], Jiajun Lin[1], Huafeng Liu[1,2*], and Wentao Zhu[1,2*]

[1] State Key Laboratory of Extreme Photonics and Instrumentation, College of Optical Science and Engineering, Zhejiang University, Hangzhou, China
[2] Key Laboratory for Biomedical Engineering of Ministry of Education and Zhejiang Key Laboratory of Intelligent Sensing Technology and Advanced Medical Instrument, College of Biomedical Engineering & Instrument Science, Zhejiang University, Hangzhou, China
liuhf@zju.edu.cn, zhuwentao.ee@gmail.com

Abstract. PET reconstruction is inherently ill-posed due to the extremely low signal-to-noise ratio (SNR) of raw projection data. Although supervised deep learning has achieved superior results in low-count scenarios, its clinical adoption is hindered by the reliance on high-quality paired datasets and limited generalization. While self-supervised Implicit Neural Representations (INRs) show promise for PET reconstruction, their inherent spectral bias causes MLPs to overfit high-frequency noise, yielding artifacts or over-smoothing. We propose a self-supervised framework that transforms the PET reconstruction paradigm toward explicit 3D Gaussian Splatting (3DGS). By representing the radioactivity image as a collection of controllable and anisotropic Gaussian primitives, our method effectively decouples local signal recovery from global noise propagation. We develop a differentiable pipeline incorporating the PET system matrix into Gaussian voxelization. To enhance image fidelity, we introduce a framework integrating Gaussian representation with kernel method, where a kernel matrix derived from MRI priors is incorporated into the voxelization stage. Experiments on brain and thorax datasets demonstrate that our method outperforms traditional iterative algorithms and state-of-the-art self-supervised approaches. By establishing a novel explicit representation for PET, this work provides a robust and interpretable solution for low-SNR clinical scenarios.

Keywords: PET Reconstruction · 3D Gaussian Representation · Kernel Method.

1 Introduction

Positron emission tomography (PET) provides critical functional information for tumor staging by tracking in vivo metabolic activities [26]. However, the

* Corresponding authors: liuhf@zju.edu.cn, zhuwentao.ee@gmail.com

cumulative risks of ionizing radiation have compelled a clinical imaging transition toward low-dose scenarios [21]. Limited by photon starvation, low-dose reconstruction evolves into a highly ill-posed inverse problem, compromising the accurate restoration of sub-centimeter lesions.

To bolster SNR, conventional PET denoising encompasses three distinct strategies. Sinogram filtering prior to FBP induces resolution degradation due to inherent photon sparsity. Iterative optimization, notably MLEM [23] and Maximum A Posteriori (MAP) [12], incorporates structural regularizations into the objective function. Despite sharpening boundaries, excessive iterations or ill-tuned hyperparameters inevitably exacerbate Poisson noise and computational overhead. Image-space algorithms (e.g., BM3D [5], NLM [3]) operate independently of projection data to suppress artifacts, yet their robust denoising proficiency often triggers over-smoothing, systematically obliterating sub-centimeter anatomical configurations.

While deep learning techniques [22, 28] covering from end-to-end direct reconstruction [9] to physics-driven unrolled networks [1, 8] transcend traditional bottlenecks, they rely heavily on large-scale and high-fidelity paired clinical data which is impractical to acquire. Thus, unsupervised paradigms have emerged. Early hand-crafted priors such as total variation (TV) or low-rank sparsity [27] often suffer from over-smoothing. Subsequently, Kernel Methods [7,13,14] guided by high-resolution anatomical priors significantly enhanced structural fidelity. Although Deep Image Prior (DIP) [25] exploits CNNs for denoising, its inherent low-pass filtering prioritizes convergence toward low-frequency backgrounds, blurring lesion edges during early optimization stages [7, 16]. Furthermore, while MLP-based Implicit Neural Representations (INRs) [18] facilitate continuous modeling, they face severe spectral bias [24]. Under strong Poisson noise, their globally coupled mapping easily triggers gradient oscillations, leading to distorted lesion boundaries and prohibitive computational costs [19].

To overcome the representational bottlenecks of implicit networks, 3D Gaussian Splatting (3DGS) [11] has recently emerged as a promising alternative, demonstrating superior detail recovery and computational efficiency compared to NeRFs. By explicitly parameterizing the volume with a set of anisotropic Gaussian kernels, 3DGS facilitates high-frequency structural representation. Consequently, this paradigm has been introduced into the medical imaging domain with promising preliminary attempts, successfully achieving reconstructions in sparse-view CT [15, 29], undersampled MRI [20], and demonstrating superior spatiotemporal detail recovery in dynamic CBCT [4, 10] and vessel reconstruction [17, 30]. However, its application in PET remains largely unexplored. Especially in low-dose scenarios, unconstrained Gaussian representations are prone to overfitting Poisson noise due to the lack of regularization, resulting in morphological collapse and artifacts.

In this study, we propose an unsupervised, multimodal image-guided Gaussian reconstruction framework designed to confine Gaussian primitives within an anatomically plausible manifold. Unlike generative schemes that rely on pre-trained diffusion models [6], our approach is physics-driven and completely cir-

cumvents the need for external pre-trained models. The contributions of our work are as follows:

1. TOF-weighted explicit Gaussian representation: TOF-weighted GS is developed for unsupervised PET reconstruction, capitalizing on explicit local parameterization to precisely capture sub-centimeter lesion details.

2. Anatomical manifold constraint: We propose a kernel regularization strategy guided by multimodal prior images to ensure morphological fidelity.

3. Physics-Consistent Optimization: We formulate a synergistic optimization objective integrating Poisson likelihood with TV penalty, enabling highly efficient and stable convergence under extreme low-count conditions.

2 Method

2.1 PET Gaussian Representation

As shown in Fig.1, we parameterize the continuous radiotracer distribution using a set of learnable 3D Gaussian primitives. Let $\Theta = \{\boldsymbol{\mu}_g, \boldsymbol{\Sigma}_g, \alpha_g\}_{g=1}^{N_g}$ denote the collective attributes of N_g Gaussian kernels, where $\boldsymbol{\mu}_g \in \mathbb{R}^3$, $\boldsymbol{\Sigma}_g \in \mathbb{R}^{3\times3}$, and $\alpha_g \in \mathbb{R}^+$ represent the spatial center, covariance matrix, and intensity coefficient, respectively. The tracer activity at any spatial coordinate $\mathbf{r}_i \in \mathbb{R}^3$ is formulated as the cumulative response of its local Gaussian primitives:

$$x(\mathbf{r}_i|\Theta) = \sum_{g=1}^{N_g} \alpha_g \exp\left(-\frac{1}{2}(\mathbf{r}_i - \boldsymbol{\mu}_g)^\top \boldsymbol{\Sigma}_g^{-1}(\mathbf{r}_i - \boldsymbol{\mu}_g)\right) \tag{1}$$

To enable anisotropic spatial adaptivity for complex anatomical structures, $\boldsymbol{\Sigma}_g$ is derived from a learnable scaling vector $\mathbf{s}_g \in \mathbb{R}^3$ and rotation quaternion $\mathbf{q}_g \in \mathbb{R}^4$. These can be transformed into a diagonal scaling matrix $\mathbf{S}_g \in \mathbb{R}^{3\times3}$ and a rotation matrix $\mathbf{R}_g \in \mathbb{R}^{3\times3}$, such that:

$$\boldsymbol{\Sigma}_g = \mathbf{R}_g \mathbf{S}_g \mathbf{S}_g^\top \mathbf{R}_g^\top \tag{2}$$

By integrating this continuous representation into the physical Time-of-Flight (TOF) PET forward model, the expected projection data expands into a TOF sinogram $\bar{\boldsymbol{y}} \in \mathbb{R}^{M\times T}$, where M is the number of Lines of Response (LORs) and T is the number of TOF bins. The expectation for the m-th LOR and t-th TOF bin, $\bar{y}_{m,t}$, is directly formulated as:

$$\bar{y}_{m,t} = \sum_{i=1}^{N} P_{m,i} h(t - \Delta t_{m,i}) x(\mathbf{r}_i|\Theta) + n_{m,t}$$

Alternatively, this can be expressed in compact matrix notation as:

$$\bar{\boldsymbol{y}} = \boldsymbol{P}_{\mathrm{TOF}} \cdot \boldsymbol{x}(\Theta) + \boldsymbol{n} \tag{3}$$

where $\boldsymbol{P}_{\mathrm{TOF}} \in \mathbb{R}^{MT\times N}$ is the combined spatio-temporal system matrix. The

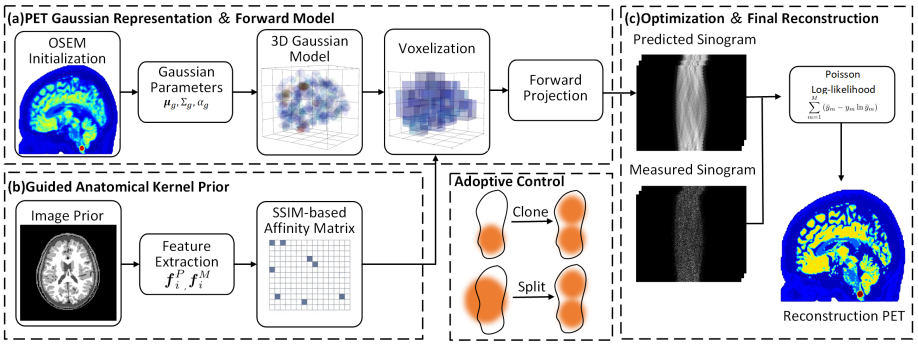

Fig. 1: The overall framework of the proposed anatomically-guided 3D Gaussian PET reconstruction. (a) Parameter initialization derived from the OSEM reconstruction, followed by voxelization and forward projection to acquire the predicted sinogram. (b) Extraction of anatomical priors from MRI/CT images to construct the K matrix. (c) Optimization of the Gaussian parameters using Poisson loss.

spatial component $P_{m,i}$ incorporates detection geometry, attenuation, and normalization. The function $h(\cdot)$ represents the detector TOF profile, typically modeled as a 1D Gaussian corresponding to the system's temporal resolution, and $\Delta t_{m,i}$ denotes the expected photon arrival time difference for an emission at spatial coordinate $\mathbf{r}_i$ along LOR m. Finally, $\boldsymbol{n}$ accounts for the TOF-distributed random and scattered events.

Tractable Optimization Strategy. To facilitate tractable optimization of the Gaussian representation, we implement a tripartite strategy: (i) Rapid voxelization. We employ a rapid voxelization strategy introuduced by [29] to bridge the continuous Gaussian parameterization with the discrete system matrix $\boldsymbol{P}$. (ii) Initialization. To prevent the highly non-convex optimization from collapsing into local minima, the initial spatial coordinates and intensities of the Gaussian primitives are probabilistically sampled from a preliminary OSEM reconstruction , where the latter is further modulated by a scaling factor k to align the initial energy distribution. (iii) Adaptive densification control. Accurate capture of complex sub-centimeter morphologies is achieved through an adaptive cloning and splitting mechanism triggered by positional gradient thresholds.

2.2 Guided Multi-Modal Kernel Prior

To mitigate overfitting to Poisson noise and morphological collapse in ill-posed low-count PET, we regularize 3DGS optimization via a hybrid multi-modal kernel prior. For each voxel i, we extract localized patch features $\boldsymbol{f}_i^P$ and $\boldsymbol{f}_i^M$ from PET and co-registered MRI, respectively. To adaptively fuse functional and structural cues, we derive a voxel-wise reliability weight $\rho_i \propto \max(0, \mathrm{SSIM}_i - C)$

based on their local structural similarity. The affinity $w_{i,j}$ between voxel i and its k-nearest neighbors is formulated via a confidence-weighted distance $D_{i,j}$:

$$w_{i,j} = \exp\left(-\frac{D_{i,j}}{N_f \sigma^2}\right), \quad D_{i,j} = \frac{\bar{\rho}_{ij}^P \|\boldsymbol{f}_i^P - \boldsymbol{f}_j^P\|_2^2 + 3\bar{\rho}_{ij}^M \|\boldsymbol{f}_i^M - \boldsymbol{f}_j^M\|_2^2}{\bar{\rho}_{ij}^P + 3\bar{\rho}_{ij}^M} \tag{4}$$

where $\bar{\rho}_{ij}^P = (1-\rho_i)(1-\rho_j)$ and $\bar{\rho}_{ij}^M = \rho_i \rho_j$ dynamically balance the modality contributions, preventing artifact propagation in unmatched regions. These weights form a sparse, row-normalized affinity matrix $\boldsymbol{K}$. By mapping the Gaussian-parameterized tracer distribution $\boldsymbol{x}(\Theta)$ via $\boldsymbol{x}_{out} = \boldsymbol{K}\boldsymbol{x}(\Theta)$, $\boldsymbol{K}$ acts as a modality-aware projection operator.

2.3 Synergistic Optimization

To jointly optimize the learnable Gaussian attributes Θ, we minimize a composite objective function that couples the negative Poisson log-likelihood:

$$\min_{\Theta} \mathcal{L}(\Theta) = \sum_{m=1}^{M} (\bar{y}_m - y_m \ln \bar{y}_m) + \lambda \|\nabla(\boldsymbol{K}\boldsymbol{x}(\Theta))\|_1 \tag{5}$$

Here, y_m and $\bar{y}_m$ represent the measured and expected projection data, respectively, and λ governs the regularization strength. While the explicit Gaussian primitives retain the spatial agility required for complex morphologies, the MRI/CT-guided matrix $\boldsymbol{K}$ enforces cross-modal anatomical consistency.

3 Experiments and Results

3.1 Data

Brain Data Simulation. The BrainWeb phantom [2] was utilized to simulate 3D ^{18}F-FDG PET images and paired MRI images. The image dimensions were 128×128×96 with 2×2×2 mm^3 voxel size. Eight randomly located lesions (radii: 2–8 mm) generated. 200 ps timing resolution was employed. The system matrix was computed using the Joseph3D algorithm in the PyTomography library, yielding sinograms of size $224 \times 449 \times 4096$.

Real clinical Data. The proposed framework was validated on real clinical TOF-PET/CT studies retrospectively collected from routine oncology examinations at the hospital. All patients underwent whole-body PET/CT scans following the intravenous injection of ^{18}F-FDG. To address GPU memory constraints, original PET 2.1x2.1x2.1 mm^3 and CT 0.875*0.875*1.5 mm^3 volumes were preprocessed and cropped to 128x128x96 voxels.

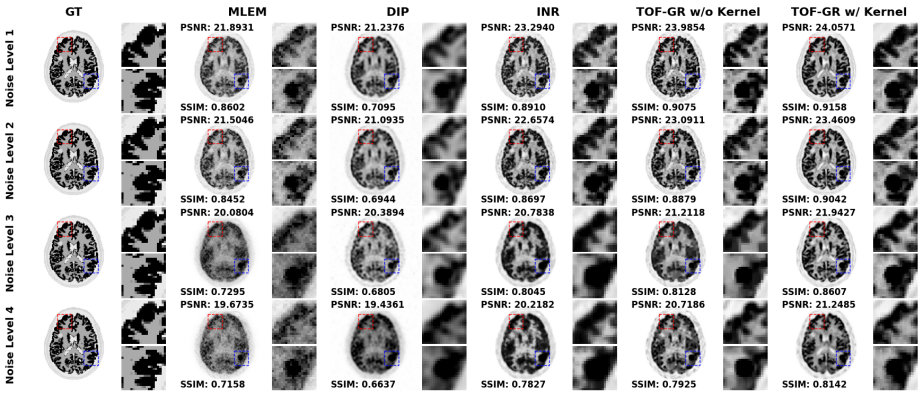

Fig. 2: Visual comparison of the reconstruction results for MLEM, DIP, INR, and proposed TOF-GR on simulated brain data

Table 1: Quantitative comparison of PET reconstruction performance across four noise levels.

Method	**PSNR (dB) ↑**				**SSIM ↑**			
	1e8	5e7	1e7	5e6	1e8	5e7	1e7	5e6
MLEM	22.44	20.24	19.45	18.89	0.8699	0.8123	0.7484	0.7292
DIP	21.72	21.02	21.58	22.19	0.6027	0.6964	0.6814	0.6526
INR	24.12	23.27	23.09	22.76	0.8995	0.8816	0.8561	0.8450
TOF-GR w/o Kernel	26.12	25.29	23.48	22.93	0.9286	0.9164	0.8667	0.8386
TOF-GR w/ Kernel	**26.29**	**25.72**	**24.20**	**23.43**	**0.9423**	**0.9326**	**0.8956**	**0.8633**

Implementation Details. The proposed 3D reconstruction framework was implemented in PyTorch and deployed on a single NVIDIA V100 GPU with 32 GB of memory. The Adam optimizer was utilized with a learning rate of 5×10^{-4}, while the learning rates for position, scale, rotation, and intensity parameters were configured at 0.005, 0.01, 0.001, and 0.005, respectively. In addition, this study calculated two quantitative metrics: the structural similarity index measure (SSIM) and peak signal-to-noise ratio (PSNR). To rigorously quantify the fidelity of subtle lesion recovery, we employ the Contrast Recovery Coefficient (CRC) and background standard deviation (SD). Furthermore, for clinical patient data we use the Contrast-to-Noise Ratio (CNR) and SUV_{max} to evaluate the detectability of lesions.

Baseline. Various baseline methods were compared with our proposed unsupervised framework: Maximum Likelihood Expectation Maximization (MLEM) [23] is a classic traditional reconstruction algorithm. Deep Image Prior (DIP) [25] represents a standard unsupervised deep learning paradigm. Implicit Neural Representation (INR) [18] utilizes a Multilayer Perceptron (MLP) to fit the map-

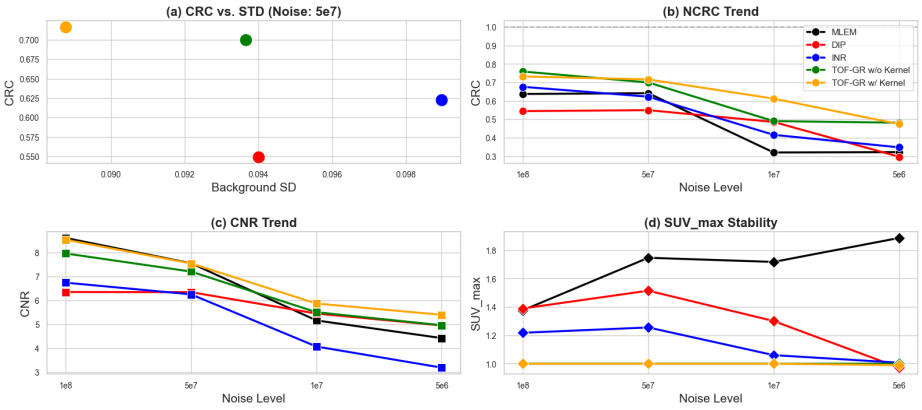

Fig. 3: Quantitative performance comparison of reconstruction methods across multiple noise levels. (a) CRC vs. STD scatter plot at 5×10^7 counts, showing the trade-off between contrast recovery and background noise. (b) Normalized CRC (NCRC) trends across four noise levels, evaluating contrast recovery accuracy. (c) CNR curves, indicating lesion detectability. (d) SUV_{max} stability curves, illustrating the consistency of peak uptake recovery under varying noise conditions.

ping from spatial coordinates to image intensities as an implicit representation method.

3.2 Results

Results on Simulation Data. Fig. 2 demonstrates that our TOF-GR method consistently outperforms MLEM, DIP, and INR across all count levels (5×10^6 to 10^8). TOF-GR effectively recovers intricate anatomical boundaries and sub-centimeter tumors, where baseline methods exhibit noise-induced blurring. Ablation results confirm the Kernel module's role in enhancing local feature representation and anatomical fidelity via Gaussian-parameterized manifolds.

As shown in Fig. 3, TOF-GR achieves a superior bias-variance trade-off. While baseline methods suffer from severe contrast degradation under photon starvation, TOF-GR's explicit spatial parameterization preserves high-frequency lesion boundaries against over-smoothing. Furthermore, NCRC and CNR curves highlight our method's superior stability and noise resilience. Unlike the erratic $\mathrm{SUV}_{\mathrm{max}}$ fluctuations in MLEM and DIP, which suggest poor convergence in ill-posed scenarios, our framework ensures accurate quantification and robustness even under extreme noise, facilitating reliable low-dose imaging.

Results on clinical Data. Fig. 4 compares clinical reconstructions. MLEM exhibits excessive noise, while DIP introduces over-smoothing artifacts that compromise lesion fidelity. Although INR achieves better balance, TOF-GR demonstrates superior edge sharpness and the highest CNR. Specifically, integrating the

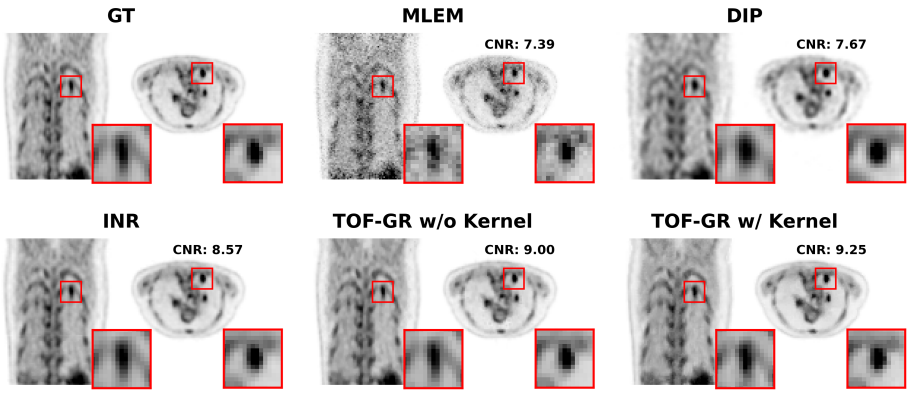

Fig. 4: Qualitative assessment of clinical PET image quality and Contrast-to-Noise Ratio (CNR).

Table 2: Ablation study and parameter sensitivity analysis. Left: Results of TOF-GR w/o Kernel for λ and k. Right: Results of the full TOF-GR w/ Kernel framework for λ and k.

TOF-GR w/o Kernel				**TOF-GR w/ Kernel**			
Method	Parameters	PSNR ↑	SSIM ↑	Method	Parameters	PSNR ↑	SSIM ↑
TOF-GR w/o Kernel	$\lambda = 0.01$	24.81	0.9437	TOF-GR w/ Kernel	$\lambda = 0.01$	**25.57**	**0.9519**
	$\lambda = 0.05$	25.09	0.9469		$\lambda = 0.05$	25.52	0.951
	$\lambda = 0.10$	**25.29**	**0.9488**		$\lambda = 0.10$	25.38	0.9489
	$\lambda = 0.50$	23.93	0.9266		$\lambda = 0.50$	23.91	0.9259
TOF-GR w/o Kernel	$k = 0.20$	24.57	0.9380	TOF-GR w/ Kernel	$k = 0.20$	24.69	0.9402
	$k = 0.50$	**25.29**	**0.9488**		$k = 0.50$	25.52	0.9513
	$k = 0.80$	25.19	0.9477		$k = 0.80$	**25.57**	**0.9519**

Anatomical Kernel Prior further sharpens boundaries and significantly enhances tumor contrast, effectively recovering structural details even in challenging clinical scenarios.

Ablation Studies. Sensitivity analysis in Table 2 underscores the intrinsic stability of our framework. While the baseline TOF-GR performance fluctuates with parameter tuning, TOF-GR w/ Kernel maintains a uniformly higher performance floor across varying regularization coefficients and scaling factors. This demonstrates that the Anatomical Kernel Prior not only enhances peak reconstruction accuracy but also significantly mitigates hyperparameter sensitivity, ensuring robust convergence and practical reliability in clinical settings.

4 Discussion and Conclusion

TOF-GR demonstrates significant potential, yet several frontiers remain for enhancement. Extending the 3D Gaussian primitives to 4D temporal modeling would allow the framework to capture dynamic kinetic distributions. Integrating parametric imaging could further unveil underlying physiological changes. Additionally, changing the projection could further boost resolution. These advancements will broaden TOF-GR's applicability in multi-parametric clinical protocols.

We presented TOF-GR, the first unsupervised PET reconstruction framework leveraging 3D Gaussian Splatting. By replacing grids with explicit primitives, TOF-GR resolves the bias-variance trade-off and preserves sub-centimeter details under extreme noise. Our method consistently outperforms state-of-the-art baselines, offering a robust and dose-efficient solution for clinical diagnostics.

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

References

1. Ai, X., Huang, B., Chen, F., Shi, L., Li, B., Wang, S., Liu, Q.: RED: Residual estimation diffusion for low-dose PET sinogram reconstruction. Med. Image Anal. **102**, 103558 (2025)
2. Aubert-Broche, B., Griffin, M., Pike, G.B., Evans, A.C., Collins, D.L.: Twenty new digital brain phantoms for creation of validation image data bases. IEEE Trans. Med. Imaging **25**(11), 1410–1416 (2006)
3. 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)
4. Fu, Y., Zhang, H., Cai, W., Xie, H., Kuo, L., Cervino, L., Moran, J., Li, X., Li, T.: Spatiotemporal gaussian optimization for 4d cone beam ct reconstruction from sparse projections. arXiv preprint arXiv:2501.04140 (2025)
5. Fumene Feruglio, P., Vinegoni, C., Gros, J., Sbarbati, A., Weissleder, R.: Block matching 3d random noise filtering for absorption optical projection tomography. Physics in Medicine & Biology **55**(18), 5401–5415 (2010)
6. Geng, M., Chen, Z., Hong, R., Li, B., Liu, Q.: Gr-diffusion: 3d gaussian representation meets diffusion in whole-body pet reconstruction. arXiv preprint arXiv:2602.11653 (2026)
7. Gong, K., Catana, C., Qi, J., Li, Q.: Direct reconstruction of linear parametric images from dynamic PET using nonlocal deep image prior. IEEE Trans. Med. Imaging **41**(3), 680–689 (2022)
8. Gong, K., Wu, D., Kim, K., Yang, J., El Fakhri, G., Seo, Y., Li, Q.: EMnet: an unrolled deep neural network for PET image reconstruction. In: Proc. SPIE. vol. 10948, pp. 1203–1208 (2019)
9. 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. Med. Image Anal. **54**, 253–262 (2019)

10. Huang, Y., Singh, I., Joyce, T., Thielemans, K., McClelland, J.R.: Digs: Dynamic cbct reconstruction using deformation-informed 4d gaussian splatting and a low-rank free-form deformation model. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 131–141. Springer (2025)
11. Kerbl, B., Kopanas, G., Leimkühler, T., Drettakis, G.: 3D gaussian splatting for real-time radiance field rendering. arXiv preprint arXiv:2308.04079 (2023)
12. 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 (2007)
13. Li, S., Gong, K., Badawi, R.D., Kim, E.J., Qi, J., Wang, G.: Neural KEM: A kernel method with deep coefficient prior for PET image reconstruction. IEEE Trans. Med. Imaging **42**(3), 785–796 (2023)
14. Li, S., Wang, G.: Deep kernel representation for image reconstruction in PET. IEEE Trans. Med. Imaging **41**(11), 3029–3038 (2022)
15. Li, Y., Fu, X., Li, H., Zhao, S., Jin, R., Zhou, S.K.: 3DGR-CT: Sparse-view CT reconstruction with a 3D gaussian representation. Med. Image Anal. **103**, 103585 (2025)
16. Liu, Q., Tsai, Y.J., Gallezot, J.D., Guo, X., Chen, M.K., Pucar, D., Young, C., Panin, V., Casey, M., Miao, T., Xie, H., Chen, X., Zhou, B., Carson, R., Liu, C.: Population-based deep image prior for dynamic PET denoising. Med. Image Anal. **95**, 103180 (2024)
17. Liu, Z., Zha, R., Zhao, H., Li, H., Cui, Z.: 4drgs: 4d radiative gaussian splatting for efficient 3d vessel reconstruction from sparse-view dynamic dsa images. In: International Conference on Information Processing in Medical Imaging. pp. 361–374. Springer (2025)
18. Long, Y., Zhang, Y., Wang, H., Kuang, X., Huang, H., Rao, F., Liu, H., Zheng, Y., Zhu, W.: PD-INR: Prior-driven implicit neural representations for TOF-PET reconstruction. In: Proc. MICCAI. pp. 459–469. Springer (2025)
19. 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)
20. Peng, T., Zha, R., Li, Z., Liu, X., Zou, Q.: Three-dimensional MRI reconstruction with 3D gaussian representations. IEEE Trans. Med. Imaging pp. 1–1 (2025)
21. Qi, J., Leahy, R.M.: Iterative reconstruction techniques in emission computed tomography. Physics in Medicine & Biology **51**(15), R541–R578 (2006)
22. Reader, A.J., Corda, G., Mehranian, A., da Costa-Luis, C., Ellis, S., Schnabel, J.A.: Deep learning for pet image reconstruction. IEEE Transactions on Radiation and Plasma Medical Sciences **5**(1), 1–25 (2020)
23. Shepp, L.A., Vardi, Y.: Maximum likelihood reconstruction for emission tomography. IEEE Trans. Med. Imaging **1**(2), 113–122 (2007)
24. Tancik, M., Srinivasan, P., Mildenhall, B., Fridovich-Keil, S., Raghavan, N., Singhal, U., Ramamoorthi, R., Barron, J., Ng, R.: Fourier features let networks learn high frequency functions in low dimensional domains. Advances in neural information processing systems **33**, 7537–7547 (2020)
25. Ulyanov, D., Vedaldi, A., Lempitsky, V.: Deep image prior. In: Proc. CVPR. pp. 9446–9454 (2018)
26. Vaquero, J.J., Kinahan, P.: Positron emission tomography: current challenges and opportunities for technological advances in clinical and preclinical imaging systems. Annual review of biomedical engineering **17**(1), 385–414 (2015)
27. Wang, C., Hu, Z., Shi, P., Liu, H.: Low dose PET reconstruction with total variation regularization. In: Proc. EMBC. pp. 1917–1920. IEEE (2014)

28. Wang, T., Lei, Y., Fu, Y., Curran, W.J., Liu, T., Nye, J.A., Yang, X.: Machine learning in quantitative pet: A review of attenuation correction and low-count image reconstruction methods. Physica Medica **76**, 294–306 (2020)
29. Zha, R., Lin, T.J., Cai, Y., Cao, J., Zhang, Y., Li, H.: R^2-gaussian: Rectifying radiative gaussian splatting for tomographic reconstruction. arXiv preprint arXiv:2405.20693 (2024)
30. Zhang, S., Zhao, H., Zhou, Z., Wu, G., Zheng, C., Wang, X., Liu, W.: Togs: Gaussian splatting with temporal opacity offset for real-time 4d dsa rendering. IEEE Journal of Biomedical and Health Informatics (2025)