

MAP-Diff: Multi-Anchor Guided Diffusion for Progressive 3D Whole-Body Low-Dose PET Denoising

Peiyuan Jing^{1,2}, Chun-Wun Cheng³, Liutao Yang², Zhenxuan Zhang², Thiago A. Lima^{1,4,5,6}, Klaus Strobel⁴, Antoine Leimgruber^{4,5,6}, Angelica Aviles-Rivero⁷, Guang Yang², Javier A. Montoya-Zegarra^{1,4,8} ✉

¹ School of Engineering, Zurich University of Applied Sciences, CH

² Bioengineering Department and Imperial-X, Imperial College London, UK

³ DAMTP, University of Cambridge, UK

⁴ Lucerne University Teaching and Research Hospital, CH

⁵ Genolier Swiss Medical Network, CH

⁶ Swiss Theranostics, CH

⁷ Yau Mathematical Sciences Center, Tsinghua University, CN

⁸ Departamento Académico de Ingeniería, Facultad de Ciencias e Ingeniería, Universidad Peruana Cayetano Heredia, PE
javier.montoya@zhaw.ch

Abstract. Low-dose Positron Emission Tomography (PET) reduces radiation exposure but suffers from severe noise and quantitative degradation. Diffusion-based denoising models achieve strong final reconstructions, yet their reverse trajectories are typically unconstrained and not aligned with the progressive nature of PET dose formation. We propose MAP-Diff, a multi-anchor guided diffusion framework for progressive 3D whole-body PET denoising. MAP-Diff introduces clinically observed intermediate-dose scans as trajectory anchors and enforces timestep-dependent supervision to regularize the reverse process toward dose-aligned intermediate states. Anchor timesteps are calibrated via degradation matching between simulated diffusion corruption and real multi-dose PET pairs, and a timestep-weighted anchor loss stabilizes stage-wise learning. At inference, the model requires only ultra-low-dose input while enabling progressive, dose-consistent intermediate restoration. Experiments show consistent improvements over strong deep learning baselines on both internal (Siemens Biograph Vision Quadra) and cross-scanner (United Imaging uEXPLORER) datasets. MAP-Diff improves PSNR from 42.48 dB to 43.71 dB (+1.23 dB) and reduces NMAE from 0.115 to 0.103 compared with 3D DDPM.

Keywords: 3D Whole-body PET Denoising · Low-dose PET · Diffusion Models · Multi-Anchor Supervision

1 Introduction

Positron Emission Tomography (PET) plays a central role in clinical diagnosis, treatment planning, and quantitative functional analysis [5]. However, PET

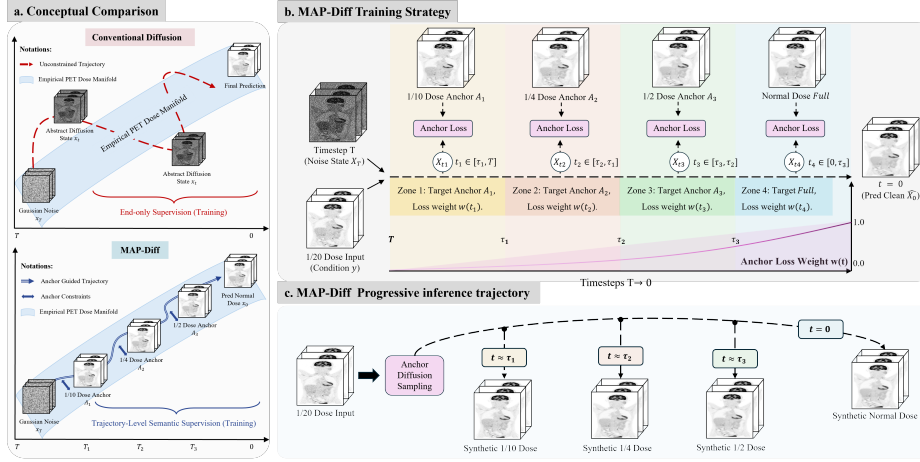


Fig. 1. Overview of MAP-Diff. (a) Unlike conventional diffusion models that supervise only the final target, MAP-Diff aligns intermediate diffusion states with paired PET acquisitions at progressively increasing dose levels. (b) During training, timestep zones are associated with dose-specific anchors through an anchor-weighted loss. (c) Progressive inference trajectory enabling dose-consistent intermediate outputs.

acquisition requires the injection of radioactive tracers, which inevitably introduces radiation exposure to patients [1]. To reduce this risk, clinical protocols often lower the injected dose or shorten the acquisition time, but both strategies substantially degrade image quality due to increased noise and reduced count statistics [6, 16]. This creates a fundamental trade-off between radiation safety and diagnostic reliability. As a result, developing robust low-dose PET denoising methods that can recover high-fidelity images from severely degraded measurements has become an important research problem, with the goal of preserving diagnostic quality while minimizing patient exposure [20, 10, 19].

Deep learning has advanced low-dose PET denoising, progressing through several architectural paradigms. Early CNN-based models [16, 2, 12, 4] achieved effective noise suppression but often produced over-smoothed results with loss of fine anatomical detail. GAN-based approaches [13, 15, 23] improved perceptual realism through adversarial training, yet frequently suffered from training instability and hallucination risk. More recently, transformer-based architectures [21, 10] demonstrated stronger global context modeling, but their high memory and computational cost limit practical deployment for large 3D whole-body PET volumes. Denoising Diffusion Probabilistic Models (DDPM) [8] have therefore emerged as a promising alternative due to their stable optimization and strong generative fidelity. Recent diffusion-based PET denoising frameworks [7, 14, 20, 11] now achieve state-of-the-art performance across multiple benchmarks.

Despite this progress, diffusion-based PET denoising still exhibits two key mismatches between training objectives and the underlying data structure. **First**, standard diffusion models are primarily optimized for final reconstruction accu-

racy at $t = 0$, while the intermediate reverse trajectory remains unconstrained. As illustrated in Fig. 1(a), the denoising path is learned implicitly through noise prediction and is not explicitly aligned with clinically meaningful dose levels, leading to dose-inconsistent intermediate states. As a result, intermediate states may satisfy the diffusion objective but not semantically dose-consistent. **Second**, PET image formation is inherently progressive: image quality improves gradually with increasing count statistics, and intermediate-dose images follow structured, non-arbitrary appearance patterns governed by acquisition physics and tracer kinetics [1, 5]. However, existing diffusion approaches do not explicitly encode this progressive dose structure in the reverse process. Even when multi-dose data are available, supervision is typically applied only at the final target, leaving the reverse trajectory physically unregularized and failing to exploit structured intermediate-dose information. This creates a structural mismatch: the learned generative trajectory is unconstrained with respect to the monotonic and physically grounded dose progression inherent to PET acquisition.

To address these two mismatches, we propose MAP-Diff (Fig. 1(b)), a multi-anchor guided diffusion framework for progressive 3D whole-body low-dose PET denoising. MAP-Diff introduces clinically observed intermediate-dose images as *trajectory anchors* and applies timestep-dependent supervision along the reverse diffusion process. By extending supervision from the final output to the full denoising trajectory, MAP-Diff resolves the objective mismatch. Simultaneously, by enforcing stage-wise alignment with physically grounded dose levels, it regularizes the generative path to reflect the progressive nature of PET acquisition.

As a result, the reverse diffusion trajectory becomes both semantically structured and dose-consistent across clinically meaningful stages, while still requiring only ultra-low-dose input at inference. Our main contributions are threefold:

- We propose MAP-Diff, a multi-anchor guided diffusion framework that constrains the reverse denoising trajectory using clinically observed intermediate-dose anchors, rather than supervising only the final full-dose target.
- We introduce timestep-dependent trajectory supervision that aligns diffusion corruption with real PET dose degradation, enforcing trajectory-level consistency and enabling dose-consistent intermediate reconstructions.
- Comprehensive evaluation on ultra-low-dose 3D whole-body PET across internal and cross-scanner datasets, showing consistent gains in structural fidelity and quantitative accuracy over state-of-the-art baselines.

2 Method

2.1 Problem Setup and Notation

Let $\Omega \subset \mathbb{Z}^3$ denote a 3D voxel grid of size $H \times W \times D$. For each subject, we observe paired multi-dose PET volumes $\{x^{1/20}, x^{1/10}, x^{1/4}, x^{1/2}, x^{\text{full}}\}$. We define the ultra-low-dose image as the conditional input $y = x^{1/20}$ and denote the full-dose reference as $x_0 := x^{\text{full}}$. To model the progressive nature of PET dose formation, we introduce a set of clinically observed **dose anchors** $\mathcal{A} = \{a_1, a_2, a_3, a_4\}$,

where a_1, a_2, a_3 correspond to the 1/10, 1/4, and 1/2 dose volumes, and $a_4 = x_0$. Our goal is to learn a conditional generative model that restores full-dose PET from ultra-low-dose input while enforcing a dose-consistent denoising trajectory. Only y is required at inference time.

2.2 Conditional Diffusion Backbone

We employ a conditional denoising diffusion probabilistic model (DDPM) [8]. Given a variance schedule $\{\alpha_t\}_{t=1}^T$, the forward process admits the closed-form marginal

$$x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}), \quad (1)$$

where $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$. Although PET noise follows Poisson photon-counting statistics [1, 5], we adopt Gaussian diffusion corruption as a tractable surrogate in image space; boundary calibration (Sec. 2.3) aligns statistical degradation levels rather than assuming physical noise equivalence.

A conditional denoiser $\epsilon_\theta(x_t, t | y)$ predicts the injected noise given the noisy volume, timestep, and ultra-low-dose condition. A per-timestep clean estimate is obtained as

$$\hat{x}_0(x_t, t | y) = \frac{1}{\sqrt{\bar{\alpha}_t}} (x_t - \sqrt{1 - \bar{\alpha}_t} \epsilon_\theta(x_t, t | y)). \quad (2)$$

Unlike standard diffusion training that uses $\hat{x}_0$ only implicitly, MAP-Diff uses this estimate for explicit trajectory-level supervision.

2.3 Multi-Anchor Trajectory Supervision

To resolve the objective and structure mismatches discussed in Sec. 1, MAP-Diff (Fig. 1 (b)) introduces trajectory-level supervision using clinically observed dose anchors. Instead of supervising only the final reconstruction target, we constrain the reverse diffusion trajectory itself to follow dose-aligned stages. We partition the timestep axis into four supervision zones with boundaries $0 \leq \tau_3 < \tau_2 < \tau_1 < T$ and define intervals $\mathcal{I}_1 = [\tau_1, T]$, $\mathcal{I}_2 = [\tau_2, \tau_1)$, $\mathcal{I}_3 = [\tau_3, \tau_2)$, $\mathcal{I}_4 = [0, \tau_3)$. Each timestep is assigned an anchor target through a piecewise mapping

$$a(t) = \sum_{j=1}^4 \mathbf{1}[t \in \mathcal{I}_j] a_j.$$

This zoning strategy defines a dose-aligned semantic curriculum along the reverse process: high-noise timesteps are guided toward lower-dose anchors, while later timesteps are constrained toward higher-dose targets. As a result, intermediate states are explicitly regularized to be dose-consistent rather than unconstrained latent denoising states.

Boundary Calibration The timestep boundaries (τ_1, τ_2, τ_3) are calibrated by matching diffusion corruption levels to clinically observed dose degradation. Within a body mask M , we compute degradation signatures relative to full dose using both intensity and structural criteria:

$$d_j^{\text{clin}} = (\text{NMAE}(M \odot a_j, M \odot x_0), 1 - \text{SSIM}(M \odot a_j, M \odot x_0)), \quad (3)$$

$$d^{\text{sim}}(t) = (\text{NMAE}(M \odot x_t, M \odot x_0), 1 - \text{SSIM}(M \odot x_t, M \odot x_0)). \quad (4)$$

Matched timesteps are selected by $t_j^* = \arg \min_t \|d^{\text{sim}}(t) - d_j^{\text{clin}}\|_2$, and final boundaries are obtained from dataset-level averages.

2.4 Training Objective and Inference

Training combines standard diffusion noise prediction with anchor-guided trajectory supervision. Sampling $t \sim \text{Unif}\{1, \dots, T\}$, the standard objective is

$$\mathcal{L}_{\text{noise}} = \mathbb{E}_{t, \epsilon} [\|\epsilon - \epsilon_\theta(x_t, t \mid y)\|_2^2], \quad (5)$$

which ensures correct diffusion modeling but supervises only the final target.

To extend supervision to the trajectory and enforce dose consistency, we add an anchor reconstruction loss on the per-timestep clean estimate:

$$\mathcal{L}_{\text{anch}} = \mathbb{E}_t [w(t) \|\hat{x}_0(x_t, t \mid y) - a(t)\|_2^2]. \quad (6)$$

The timestep weight $w(t) = (1 - t/T)^p$ emphasizes later, lower-noise stages where predictions are more reliable, avoiding over-constraining early noisy states while preserving progressive alignment.

We adopt hard timestep zoning and a polynomial weighting schedule as a simple and stable mechanism for stage-aware supervision. In practice, soft assignments or learned weighting functions could also be employed; however, the piecewise formulation provides clear dose-stage separation and was found empirically sufficient without introducing additional optimization complexity. The final loss is:

$$\mathcal{L} = \mathcal{L}_{\text{noise}} + \lambda \mathcal{L}_{\text{anch}}. \quad (7)$$

The addition of the anchor term modifies the standard DDPM objective. While $\mathcal{L}_{\text{noise}}$ preserves the score-matching formulation, $\mathcal{L}_{\text{anch}}$ acts as a trajectory-level regularizer. Thus, MAP-Diff can be viewed as a regularised conditional score model rather than a strictly likelihood-consistent diffusion model, prioritizing dose-consistent trajectory shaping over exact likelihood optimization.

At inference (Fig. 1 (c)), given only input y , we run the reverse diffusion process to obtain the final full-dose estimate $\hat{x}_0$. Because supervision is applied along the trajectory, intermediate-dose synthesis can be exported by recording $\hat{x}_0(x_t, t \mid y)$ at $t = \tau_1, \tau_2, \tau_3$, yielding synthetic volumes aligned with intermediate dose levels.

Table 1. Quantitative comparison results. Statistical significance is assessed using a two-sided Wilcoxon signed-rank test with Holm correction. * and ** denote statistically significant differences compared with MAP-Diff at $p < 0.05$ and $p < 0.01$, respectively.

Method	Internal (University Hospital of Bern)			External (Shanghai Ruijin Hospital)		
	PSNR $\uparrow$	SSIM $\uparrow$	NMAE $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	NMAE $\downarrow$
Low-dose Input	38.545 $\pm$ 1.798**	0.948 $\pm$ 0.018**	0.199 $\pm$ 0.036**	30.402 $\pm$ 3.475**	0.896 $\pm$ 0.033**	0.257 $\pm$ 0.059**
3D cGAN [15]	40.721 $\pm$ 1.293**	0.977 $\pm$ 0.006**	0.127 $\pm$ 0.015**	33.062 $\pm$ 3.624**	0.927 $\pm$ 0.053**	0.174 $\pm$ 0.031**
3D UNet [4]	41.810 $\pm$ 1.377**	0.972 $\pm$ 0.007**	0.120 $\pm$ 0.015**	33.654 $\pm$ 3.810**	0.904 $\pm$ 0.074**	0.164 $\pm$ 0.031**
RED-CNN [2]	42.185 $\pm$ 1.427**	0.981 $\pm$ 0.005**	0.119 $\pm$ 0.016**	33.989 $\pm$ 3.725**	0.945 $\pm$ 0.025**	0.161 $\pm$ 0.030**
EDCNN [12]	42.355 $\pm$ 1.579**	0.983 $\pm$ 0.006**	0.116 $\pm$ 0.017**	33.841 $\pm$ 3.615**	0.942 $\pm$ 0.023**	0.164 $\pm$ 0.024**
LIT-Former [3]	42.036 $\pm$ 1.537**	0.981 $\pm$ 0.006**	0.121 $\pm$ 0.017**	33.471 $\pm$ 3.723**	0.931 $\pm$ 0.037**	0.176 $\pm$ 0.042**
Restormer [21]	42.198 $\pm$ 1.466**	0.979 $\pm$ 0.007**	0.124 $\pm$ 0.020**	33.657 $\pm$ 3.811**	0.939 $\pm$ 0.027**	0.171 $\pm$ 0.036**
Spach Former [10]	42.290 $\pm$ 1.515**	0.982 $\pm$ 0.005**	0.118 $\pm$ 0.015**	34.109 $\pm$ 3.882**	0.945 $\pm$ 0.030**	0.153 $\pm$ 0.032**
DRMC [19]	42.378 $\pm$ 1.619**	0.982 $\pm$ 0.006**	0.115 $\pm$ 0.018**	34.006 $\pm$ 3.832**	0.942 $\pm$ 0.027**	0.163 $\pm$ 0.034**
3D DDPM [20]	42.483 $\pm$ 1.593**	0.983 $\pm$ 0.005**	0.115 $\pm$ 0.016**	34.117 $\pm$ 3.485**	0.947 $\pm$ 0.021**	0.151 $\pm$ 0.023**
MAP-Diff	43.710$\pm$1.575	0.986$\pm$0.004	0.103$\pm$0.014	34.419$\pm$3.611	0.952$\pm$0.023	0.141$\pm$0.025

3 Experiments and Results

3.1 Experiments

Dataset. We use the Ultra-Low-Dose PET Challenge dataset [18], which provides multi-dose whole-body PET scans from two vendors: University Hospital of Bern (Siemens Biograph Vision Quadra) and Shanghai Ruijin Hospital (United Imaging uEXPLORER). For in-distribution experiments, we select 377 ^{18}F -FDG scans from the Siemens system with paired dose levels (1/20, 1/10, 1/4, 1/2, and normal). Subjects are split into 297/20/60 for training/validation/testing. For cross-scanner evaluation, we use 100 uEXPLORER scans for external testing only. All volumes are converted to SUV, center-cropped to $192 \times 288 \times 520$, and divided into overlapping $96 \times 96 \times 96$ patches for fully 3D training [20].

Evaluation Metrics. We use PSNR [9] and SSIM [17] with normal-dose PET as the golden-rule. PSNR measures voxel-wise fidelity, while SSIM assesses structural and perceptual similarity. We additionally report NMAE [17] to quantify intensity deviations, which is particularly important in PET since all volumes are expressed in SUV units and require accurate quantitative recovery.

Experiment Details. All experiments are implemented in PyTorch and run on NVIDIA A100 GPUs. Diffusion models use $T=1000$ timesteps with a linear noise schedule (β_t from 1×10^{-4} to 2×10^{-2}) and are optimized with Adam [22] at learning rate 1×10^{-4} . All methods are evaluated in fully 3D form for volumetric PET. Baselines originally proposed in 2D [12, 2, 3, 21] are converted to 3D while preserving their original architectures and training settings.

3.2 Results

Overall quantitative and qualitative performance. Table 1 reports quantitative comparisons on the internal (Siemens Quadra) and external (uEXPLORER)




Fig. 2. Qualitative comparison of PET denoising methods on a whole-body PET study reconstructed from 1/20-dose data. Coronal, sagittal, and axial views with enlarged regions of interest are shown. Red arrows highlight structures and uptake patterns that are degraded in competing methods.

cohorts. MAP-Diff achieves the best performance across all metrics on both datasets, with statistically significant improvements over all competing methods ($p < 0.01$). On the internal cohort, MAP-Diff increases PSNR to 43.71 ± 1.58 dB, outperforming the strongest diffusion baseline (3D DDPM [20]) by +1.23 dB. NMAE is reduced from 0.115 to 0.103, while SSIM improves to 0.986, depicting enhanced structural fidelity and quantitative accuracy. On the external cohort, MAP-Diff achieves 34.42 ± 3.61 dB PSNR and reduces NMAE to 0.141, again surpassing 3D DDPM (34.12 dB, 0.151 NMAE) and all CNN- [4, 2, 12, 19], GAN- [15], and Transformer- [3, 21, 10] based baselines. Qualitatively (Fig. 2), MAP-Diff better preserves lesion contrast and anatomical boundaries while effectively suppressing noise without introducing over-smoothing artifacts.

Intermediate-dose reconstruction behavior. To validate the progressive reconstruction property described in Sec. 2.4, we compare intermediate outputs extracted at anchor timesteps with real paired intermediate-dose references. As shown in Fig. 3 (a), MAP-Diff consistently outperforms 3D DDPM at the 1/10, 1/4, and 1/2 dose levels across PSNR, SSIM, and NMAE. The largest margin occurs at the earliest stage (1/10 dose), where noise levels are highest and trajectory guidance is most critical. The error maps in Fig. 3(b) further show that MAP-Diff produces more spatially homogeneous residuals with reduced structured errors compared with 3D DDPM across all stages. This behaviour indicates




Fig. 3. Evaluation of reconstruction fidelity at anchor dose levels. (a) Quantitative comparison against paired real-dose references at the 1/10, 1/4, and 1/2 dose levels. (b) Coronal-view signed error maps with respect to the corresponding dose-specific references.

that anchor-based supervision provides effective stage-local constraints, steering the reverse process toward dose-consistent intermediate states. As a result, MAP-Diff improves not only the final reconstruction but also the quantitative accuracy of intermediate outputs along the trajectory.

Ablation on anchor supervision design. Table 2 presents ablations of the anchor supervision design, varying both the number of anchors and the timestep weighting strategy. Using all three anchors with timestep-weighted loss (S1) yields the best overall performance across all metrics. Reducing the number of anchors leads to gradual degradation, with configurations that include the earliest anchor (1/10 dose) consistently outperforming those that use only later anchors, indicating that early-stage guidance contributes most to trajectory shaping. Replacing timestep-dependent weighting with constant weighting (S8) significantly degrades performance, suggesting that stage-aware loss scaling is important to

Table 2. Ablation study of the anchor supervision design in MAP-Diff. Different anchor configurations and timestep weighting strategies are evaluated.

Set	#A	Anchors	W	PSNR	SSIM	NMAE	Set	#A	Anchors	W	PSNR	SSIM	NMAE
Baseline 3D DDPM [20] (No Anchors): PSNR: 42.483, SSIM: 0.983, NMAE: 0.115													
S1	3	1/10,1/4,1/2	Poly	43.710	0.986	0.103	S2	2	{1/10,1/4}	Poly	43.672	0.984	0.104
S3	2	{1/10,1/2}	Poly	43.521	0.983	0.106	S4	2	{1/4,1/2}	Poly	43.414	0.982	0.107
S5	1	{1/10}	Poly	43.503	0.983	0.106	S6	1	{1/4}	Poly	43.267	0.982	0.109
S7	1	{1/2}	Poly	42.790	0.981	0.111	S8	3	{1/10,1/4,1/2}	Const	41.856	0.905	0.133

stabilize training and avoid over-constraining high-noise timesteps. Together, these results confirm that both multi-anchor coverage and timestep-adaptive weighting are key to effective trajectory supervision.

4 Conclusion

We introduced MAP-Diff, a multi-anchor guided diffusion framework for progressive 3D whole-body low-dose PET denoising. By introducing clinically observed intermediate-dose anchors and timestep-dependent trajectory supervision, MAP-Diff constrains the reverse diffusion process to follow a dose-aligned and stage-consistent restoration path, addressing the mismatch between endpoint-only objectives and progressive PET dose structure. The anchor-based trajectory regularization enables stage-wise semantic consistency while requiring only ultra-low-dose input at inference. Experiments on both internal and cross-scanner datasets demonstrate consistent improvements over strong baselines in visual quality and quantitative accuracy. Future work will explore adaptive anchor selection and extension to other progressive medical image reconstruction settings.

Acknowledgements PJ, CWC, TVL, KS, AL, AIAV, GY, and JAMZ are supported by SNSF (20HW-1 220785). GY was also supported by the UKRI (MC/PC/21013, EP/Z002206/1, MR/U506710/1), and the UKRI Future Leaders Fellowship (MR/V023799/1, UKRI2738).

Disclosure of Interests The authors declare that they have no competing interests.

References

1. Berger, A.: Positron emission tomography. *BMJ: British Medical Journal* **326**(7404), 1449 (2003)
2. Chen, H., Zhang, Y., Kalra, M.K., Lin, F., Chen, Y., Liao, P., Zhou, J., Wang, G.: Low-dose ct with a residual encoder-decoder convolutional neural network. *IEEE transactions on medical imaging* **36**(12), 2524–2535 (2017)
3. Chen, Z., Niu, C., Gao, Q., Wang, G., Shan, H.: Lit-former: Linking in-plane and through-plane transformers for simultaneous ct image denoising and deblurring. *IEEE Transactions on Medical Imaging* **43**(5), 1880–1894 (2024)
4. Çiçek, Ö., Abdulkadir, A., Lienkamp, S.S., Brox, T., Ronneberger, O.: 3d u-net: learning dense volumetric segmentation from sparse annotation. In: *International conference on medical image computing and computer-assisted intervention*. pp. 424–432. Springer (2016)
5. Czernin, J., Phelps, M.E.: Positron emission tomography scanning: current and future applications. *Annual review of medicine* **53**(1), 89–112 (2002)
6. Fletcher, J.W., Djulbegovic, B., Soares, H.P., Siegel, B.A., Lowe, V.J., Lyman, G.H., Coleman, R.E., Wahl, R., Paschold, J.C., Avril, N., et al.: Recommendations on the use of 18f-fdg pet in oncology. *Journal of Nuclear Medicine* **49**(3), 480–508 (2008)

7. Gong, K., Johnson, K., El Fakhri, G., Li, Q., Pan, T.: Pet image denoising based on denoising diffusion probabilistic model. *European Journal of Nuclear Medicine and Molecular Imaging* **51**(2), 358–368 (2024)
8. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
9. Hore, A., Ziou, D.: Image quality metrics: Psnr vs. ssim. In: 2010 20th international conference on pattern recognition. pp. 2366–2369. IEEE (2010)
10. Jang, S.I., Pan, T., Li, Y., Heidari, P., Chen, J., Li, Q., Gong, K.: Spach transformer: Spatial and channel-wise transformer based on local and global self-attentions for pet image denoising. *IEEE transactions on medical imaging* **43**(6), 2036–2049 (2023)
11. Jing, P., Tang, Y., Cheng, C.W., Zhang, Z., Yang, L., Lima, T.V., Strobel, K., Leimgruber, A., Aviles-Rivero, A., Yang, G., et al.: 3d wavelet-based structural priors for controlled diffusion in whole-body low-dose pet denoising. *arXiv preprint arXiv:2601.07093* (2026)
12. Liang, T., Jin, Y., Li, Y., Wang, T.: Edcnn: Edge enhancement-based densely connected network with compound loss for low-dose ct denoising. In: 2020 15th IEEE International conference on signal processing (ICSP). vol. 1, pp. 193–198. IEEE (2020)
13. Luo, Y., Zhou, L., Zhan, B., Fei, Y., Zhou, J., Wang, Y., Shen, D.: Adaptive rectification based adversarial network with spectrum constraint for high-quality pet image synthesis. *Medical image analysis* **77**, 102335 (2022)
14. Lyu, Q., Chen, T., Wang, Y., Guo, E., Zhou, L.: Wid-pet: Pet image reconstruction from low-dose data using a wavelet-informed diffusion model with fast inference. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 684–693. Springer (2025)
15. Wang, Y., Yu, B., Wang, L., Zu, C., Lalush, D.S., Lin, W., Wu, X., Zhou, J., Shen, D., Zhou, L.: 3d conditional generative adversarial networks for high-quality pet image estimation at low dose. *Neuroimage* **174**, 550–562 (2018)
16. Wang, Y.R., Baratto, L., Hawk, K.E., Theruvath, A.J., Pribnow, A., Thakor, A.S., Gatidis, S., Lu, R., Gummidipundi, S.E., Garcia-Diaz, J., et al.: Artificial intelligence enables whole-body positron emission tomography scans with minimal radiation exposure. *European journal of nuclear medicine and molecular imaging* **48**(9), 2771–2781 (2021)
17. 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)
18. Xue, S., Wang, H., Chen, Y., Liu, F., Zhu, H., Viscione, M., Guo, R., Rominger, A., Li, B., Shi, K.: Udpet: Ultra-low dose pet imaging challenge dataset. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 616–623. Springer (2025)
19. Yang, Z., Zhou, Y., Zhang, H., Wei, B., Fan, Y., Xu, Y.: Drmc: A generalist model with dynamic routing for multi-center pet image synthesis. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 36–46. Springer (2023)
20. Yu, B., Ozdemir, S., Dong, Y., Shao, W., Pan, T., Shi, K., Gong, K.: Robust whole-body pet image denoising using 3d diffusion models: evaluation across various scanners, tracers, and dose levels. *European Journal of Nuclear Medicine and Molecular Imaging* **52**(7), 2549–2562 (2025)

21. Zamir, S.W., Arora, A., Khan, S., Hayat, M., Khan, F.S., Yang, M.H.: Restormer: Efficient transformer for high-resolution image restoration. In: Proceedings of the IEEE/CVF conference on computer vision and pattern recognition. pp. 5728–5739 (2022)
22. Zhang, Z.: Improved adam optimizer for deep neural networks. In: 2018 IEEE/ACM 26th international symposium on quality of service (IWQoS). pp. 1–2. Ieee (2018)
23. Zhou, B., Tsai, Y.J., Chen, X., Duncan, J.S., Liu, C.: Mdpet: a unified motion correction and denoising adversarial network for low-dose gated pet. *IEEE transactions on medical imaging* **40**(11), 3154–3164 (2021)