MDAA-Diff: CT-Guided Multi-Dose Adaptive Attention Diffusion Model for PET Denoising

Xiaolong Niu[1,4,5], Zanting Ye[1,4,5], Xu Han[2], Yanchao Huang[3], Hao Sun[1,4,5], Hubing Wu[3], Lijun Lu[1,4,5,6(✉)]

[1]School of Biomedical Engineering, Southern Medical University, Guangzhou, China
ljlubme@gmail.com
[2]School of Biomedical Engineering, Shanghai Jiao Tong University, shanghai, China
[3]Nanfang PET Center, Nanfang Hospital Southern Medical University, Guangzhou, China
[4]Guangdong Provincial Key Laboratory of Medical Image Processing, Southern Medical University, Guangzhou, China
[5]Guangdong Province Engineering Laboratory for Medical Imaging and Diagnostic Technology, Southern Medical University, Guangzhou, China
[6]Pazhou Lab, Guangzhou, China

Abstract. Acquiring high-quality Positron Emission Tomography (PET) images requires administering high-dose radiotracers, which increases radiation exposure risks. Generating standard-dose PET (SPET) from low-dose PET (LPET) has become a potential solution. However, previous studies have primarily focused on single low-dose PET denoising, neglecting two critical factors: discrepancies in dose response caused by inter-patient variability, and complementary anatomical constraints derived from CT images. In this work, we propose a novel CT-Guided Multi-dose Adaptive Attention Denoising Diffusion Model (MDAA-Diff) for multi-dose PET denoising. Our approach integrates anatomical guidance and dose-level adaptation to achieve superior denoising performance under low-dose conditions. Specifically, this approach incorporates a CT-Guided High-frequency Wavelet Attention (HWA) module, which uses wavelet transforms to separate high-frequency anatomical boundary features from CT images. These extracted features are then incorporated into PET imaging through an adaptive weighted fusion mechanism to enhance edge details. Additionally, we propose the Dose-Adaptive Attention (DAA) module, a dose-conditioned enhancement mechanism that dynamically integrates dose levels into channel-spatial attention weight calculation. Extensive experiments on ^{18}F-FDG and ^{68}Ga-FAPI datasets demonstrate that MDAA-Diff outperforms state-of-the-art approaches in preserving diagnostic quality under reduced-dose conditions. Our code is publicly available at https://github.com/Long0121/MDAA-Diff.

Keywords: PET denoising · Multi-dose adaptive · High-frequency CT-guided · Diffusion model.

X. Niu, Z. Ye, X. Han—Co-first authors.

1   Introduction

Positron Emission Tomography (PET) is a molecular imaging technique used to evaluate metabolic activity and functional states of human tissues and organs. PET facilitates precise functional assessments of tissues and organs through tracer-based metabolic visualization imaging, demonstrating indispensable clinical significance in tumor diagnosis, neurodegenerative disease research, and cardiovascular functional analysis[2, 9, 29]. However, PET imaging quality remains highly dependent on radiotracer dose. While low-dose PET (LPET) imaging mitigates the risk of ionizing radiation (particularly for sensitive populations such as children and pregnant women [3, 5, 12]), it inevitably increases image noise, compromising diagnostic accuracy. This inherent trade-off between radiation reduction and diagnostic efficacy remains a major challenge in contemporary PET imaging technology.

Deep learning has emerged as a pivotal approach for radiation dose reduction in PET imaging, enabling diagnostic-quality standard-dose PET (SPET) denoising from LPET. Compared with traditional methods, such as local mean filtering and frequency-domain filtering [1, 4, 24], which often oversmooth critical details, deep neural networks learn complex nonlinear mappings between LPET and SPET through end-to-end training [7, 10, 20, 25, 22, 26, 27], suppressing noise while preserving lesion metabolic activity features. Generative model-based denoising architectures have become a focal point. For instance, Sanaat et al. [17] utilized CycleGAN to achieve cross-domain mapping from LPET to SPET. Zhou et al. [28] combined StyleGAN with segmentation network to generate denoised images with realistic textures. Nevertheless, these deterministic mapping models struggle to quantify uncertainties inherent in the reconstruction process. Denoising Diffusion Probabilistic Model (DDPM) [11] addresses this limitation by iteratively denoising images through a Markov chain, enabling uncertainty modeling while retaining fine-grained details. However, existing approaches predominantly assume single LPET denoising, ignoring the variations in dose response caused by inter-patient variability in clinical practice, such as physiology, metabolism, and tracer uptake patterns.

Preliminary progress has been achieved in research for multi-dose PET denoising. Xie et al. [23] proposed a Unified Noise-aware Network that performs PET denoising across different dose levels. But its weight allocation mechanism compromises recovery performance under high-dose conditions. Tang et al. [19] developed a High-Frequency-Guided Residual Diffusion Model which uses sinusoidal positional encoding to encode the dose, and then input it into the denoising network along with the embedded timesteps. However, the simple concatenation of dose embedding and time embedding may fail to enable effective interaction between dose information and the diffusion process, potentially limiting the model's adaptability to feature variations across different dose levels. Moreover, this method inadequately exploits the advantages of multimodal imaging in clinical applications. In multimodal PET/CT imaging, the high spatial coherence between CT and PET images provides critical anatomical constraints that enhance denoising performance [8]. Chen et al. [6] concatenate multimodal

images as input, ignoring varying modalities contribute differently across image locations. Wang et al. [21] proposed a local adaptive fusion method with position-specific weights, which operates only at the input level and does not exploit deep semantic features, potentially limiting the effectiveness of the multimodal fusion.

In this work, we propose a CT-Guided Multi-Dose Adaptive Attention Diffusion Model (MDAA-Diff) for PET denoising. Specifically, we design a CT-Guided High-frequency Wavelet Attention (HWA) module to extract CT high-frequency anatomical boundary features using multi-scale wavelet decomposition, which are then adaptively fused with PET imaging. Moreover, we construct a Dose-adaptive Attention (DAA) module that utilizes dose factors to guide channel-wise and spatial feature weighting, enhancing the model's perception capability and robustness under different dose conditions.
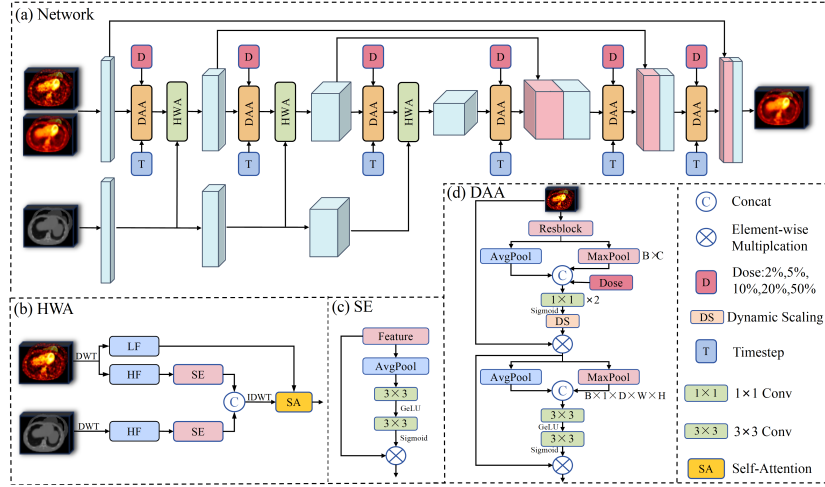

Fig. 1: Overview of the MDAA-Diff. (a) is the overall network architecture of the diffusion model; (b) is a CT-guided modal fusion module; (c) is designed to perform channel-wise attention; and (d) is a multi-dose adaptive attention module.

2 Methodology

2.1 Improved Denoising Diffusion Probabilistic Models

Our model builds upon an Improved Denoising Diffusion Probabilistic Models (IDDPM) [15], which significantly reduces the sampling timesteps compared to DDPM. While DDPM models the reverse denoising process as a conditional probability distribution, defined as:

$$p_\theta(x_{t-1}|x_t) = \mathcal{N}(x_{t-1}; \, \mu_\theta(x_t, t), \, \Sigma_\theta(x_t, t)), \tag{1}$$

where $\mu_\theta(x_t, t)$ represents the learnable value, and the variance $\Sigma_\theta(x_t, t)$ is fixed as $\sigma^2 I$ and cannot be optimized as a learnable parameter. IDDPM addresses this limitation by parameterizing $\Sigma_\theta(x_t, t)$ modeling through the following formulation:

$$\Sigma_\theta(x_t, t) = \exp\left(v \log \beta_t + (1 - v) \log \tilde{\beta}_t\right), \tag{2}$$

where v is a weighting parameter, β_t represents the noise intensity during the forward diffusion process, and $\tilde{\beta}_t$ denotes the adjusted noise intensity in the reverse process. This parameterization enhances both the model's flexibility and the quality of the generated samples. Additionally, it optimizes noise scheduling by selecting specific timesteps $S = \{s_1, s_2, \ldots, s_K\}$, which further improves sampling efficiency.

2.2   MDAA-Diff Model

Our MDAA-Diff model is shown in Fig. 1. First, the noise-perturbed SPET and LPET images are concatenated along the channel dimension and fed into the input model. Multi-scale PET/CT features are extracted by a dedicated encoder to achieve cross-modal information fusion. The dose factor is dynamically integrated into PET feature enhancement through the DAA module, enabling dose-adaptive adjustments.

High-frequency Wavelet Attention Current PET and CT feature fusion methods often rely on channel concatenation or Cross-Attention (CA) mechanisms [14]. The quadratic-complexity operation of CA imposes prohibitive computational demands for high-resolution medical imaging. To address this issue, we propose the HWA module, a fusion module based on 3D wavelet transform (with bior4.4 wavelet) and Squeeze-and-Excitation (SE), to achieve multi-scale cooperative enhancement of cross-modal features. As shown in Fig. 1b, PET and CT features are decomposed into low-frequency and high-frequency components: PET low-frequency features capture global metabolic patterns, while high-frequency features encode lesion edges. CT high-frequency components features capture fine anatomical details, including bone boundaries and organ morphology, which provide accurate spatial localization and structural constraints for PET denoising.

On this basis, the SE module (as shown in Fig. 1c) is utilized to perform dynamic channel weights for both PET and CT high-frequency components, emphasizing key features and reducing redundancy. The weighted PET and CT high-frequency features are concatenated along the channel dimension and processed through convolutional layers for cross-modal interaction and dimensionality reduction. To enable multi-scale information complementarity, HWA modules are embedded across all encoder levels (from shallow local details to deep semantic features) to establish a progressive fusion architecture. Finally, the fused features are reconstructed from the PET low-frequency components and the convolved high-frequency components through inverse wavelet transform.

Dose-Adaptive Attention Module The DAA module adaptively enhances multi-dose features through the synergistic mechanism of the channel enhancement branch and the spatial enhancement branch. The details of the HWA module are shown in Fig. 1d. The module dynamically adjusts the weights of the input features to accommodate different dose-level distributions and improves the robustness of the model.

The channel feature enhancement branch. The branch primarily focuses on dynamically adjusting the channel weights of the input features by capturing global and local information, while integrating the dose factor D to achieve adaptive optimization. Specifically, D is projected through a multi-layer perceptron (MLP) to generate the dose embedding. Global average pooling and max pooling operations are applied to the input features, and the results are concatenated with the dose embedding. The channel weights are then generated through a convolutional network, defined as:

$$W_{\text{channel}} = \text{sigmoid}\Big(W_2 \cdot \text{ReLU}\big(W_1 \cdot \text{Concat}(\text{Pool}(X), D_{\text{emb}})\big)\Big) \qquad (3)$$

where $\text{Pool}(X) \in \{\text{AvgPool}(X), \text{MaxPool}(X)\}$, W_1 and W_2 are 1×1 convolutional weight matrices. The channel weights are further scaled by the factor α and shifted by the bias factor β, which are used to normalize and adjust the input features, achieving dynamic channel enhancement, defined as:

$$X' = \text{Norm}(X) \odot (1 + \alpha) + \beta \qquad (4)$$

where $\text{Norm}(X)$ represents a normalization operation, and $\odot$ denotes element-wise multiplication.

The spatial feature enhancement branch. The branch aims to extract the structural detail and local texture information of the input features along the spatial dimension. Specifically, the input features are pooled along the channel dimension using global average pooling and max pooling. The results are concatenated and passed through a convolutional layer to generate the spatial attention map, defined as:

$$X' = X \odot \text{sigmoid}\left(\text{Conv}\left(\text{Concat}(\text{Mean}(X), \text{Max}(X))\right)\right), \qquad (5)$$

where Conv is a convolutional block composed of a 3×3 kernel and a dilated convolution (dilation rate set to 2). Finally, the spatial attention map is multiplied element-wise with the input features to achieve context enhancement along the spatial dimension.

3   Experiments

3.1   Datasets

We evaluated the performance of MDAA-Diff using two independent PET/CT datasets comprising 51 patients with ^{18}F-FDG tracer and 60 patients with ^{68}Ga-FAPI tracer. The datasets were provided by the Nanfang PET Center, Nanfang

Hospital, Southern Medical University. SPET images were reconstructed from 300 seconds (100%) acquisitions obtained 60 minutes post-injection. LPET images were reconstructed based on time windows resampled during acquisition, simulating reduced acquisition durations of 6 seconds (2%), 15 seconds (5%), 30 seconds (10%), 60 seconds (20%), and 150 seconds (50%). All PET images were reconstructed with resolution of 192×192×673. To address PET-CT mismatches, CT images were resampled and registered to the PET coordinate system using mutual information optimization.

3.2   Implementation Details

Our MDAA-Diff model was implemented in PyTorch and trained on an NVIDIA GeForce RTX 4090 GPU with 24GB of memory. A five-fold cross-validation strategy was employed in the experimental design. The number of diffusion steps during the training phase was set to 1000, and the number of sampling steps during the inference phase was set to 50. We used an AdamW optimizer with an initial learning rate of 1e-4, utilizing the cosine annealing scheduling strategy. The generated image quality was evaluated using the Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index (SSIM) metrics. Visual comparisons were conducted to demonstrate the model's ability to improve anatomical fidelity and effectively suppress noise in critical regions.

Table 1: Quantitative comparison of MDAA-Diff with five state-of-the-art methods. PSNR and SSIM are reported for both tracers at different dose levels.

^{18}F-FDG				
PSNR↑/SSIM↑ 6s	15s	30s	60s	150s
CNN [18] 39.858/0.969	40.082/0.972	40.269/0.974	40.459/0.977	40.614/0.978
SwinUnetr [16] 44.408/0.976	44.979/0.979	45.417/0.981	46.059/0.983	47.115/0.986
MambaMIR [13] 44.349/0.978	45.801/0.982	45.830/0.983	46.043/0.985	46.774/0.987
UNN [23] 44.504/0.972	45.185/0.977	45.663/0.980	46.071/0.982	46.646/0.984
HF-ResDiff [19] 44.728/0.979	45.243/0.982	45.677/0.984	46.174/0.986	47.004/0.984
Ours **46.035/0.984**	**46.523/0.987**	**46.889/0.988**	**47.315/0.989**	**48.042/0.990**
^{68}Ga-FAPI				
PSNR↑/SSIM↑ 6s	15s	30s	60s	150s
CNN [18] 39.570/0.967	39.832/0.961	40.049/0.963	40.276/0.965	40.680/0.968
SwinUnetr [16] 44.119/0.966	44.560/0.969	44.866/0.971	45.171/0.973	45.783/0.976
MambaMIR [13] 44.028/0.968	44.429/0.970	44.810/0.972	45.207/0.974	45.844/0.977
UNN [23] 44.759/0.969	45.019/0.973	45.432/0.975	45.765/0.977	46.496/0.980
HF-ResDiff [19] 44.219/0.964	44.632/0.968	45.081/0.971	45.493/0.974	46.287/0.977
Ours **46.689/0.980**	**47.009/0.981**	**47.160/0.981**	**47.362/0.982**	**47.921/0.983**

3.3   Comparison with State-of-the-Art Methods

We compared MDAA-Diff with five state-of-the-art methods, including single-dose approaches (CNN [18], SwinUnetr [16], MambaMIR [13]) and multi-dose
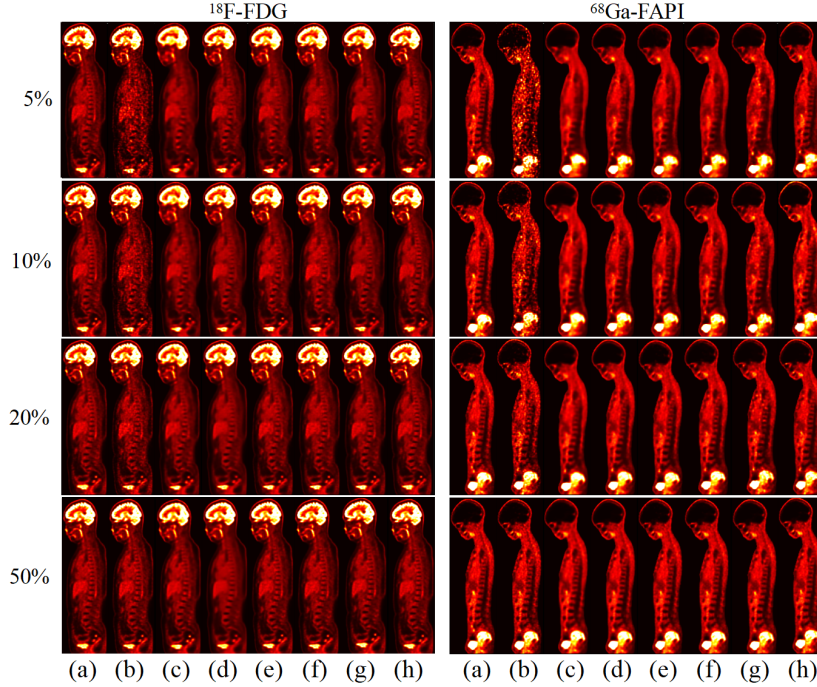
Fig. 2: The comparison results of MDAA-Diff with several state-of-the-art methods at different dose levels. (a): SPET; (b): LPET; (c): CNN; (d): UNN; (e): SwinUnetr; (f): MambaMIR; (g): HF-ResDiff; (h): Ours.

approaches (HF-ResDiff [19], UNN [23]). All comparison methods take concatenated PET and CT images as input. Quantitative analysis (as shown in Table 1) demonstrated that MDAA-Diff outperformed all comparison models across both PSNR and SSIM metrics. Moreover, its performance improved significantly at lower dose levels (e.g., 2% and 5% doses). For instance, in 2% dose imaging for the ^{18}F-FDG tracer, MDAA-Diff achieved an SSIM of 98.4%, which is 0.5 percentage points higher than the second-best model (HF-ResDiff: 97.9%). This finding highlights its robustness to complex noise distributions.

Qualitative analysis further validated the clinical value of the proposed method. As illustrated in Fig. 2, MDAA-Diff exhibited significant improvements in boundary sharpness for the spine region. Particularly in the 2% dose experiments, its structural fidelity was noticeably superior to that of other comparison methods. For denoising tasks of small lesions, MDAA-Diff effectively suppressed noise artifacts (as indicated by the blue arrows in Fig. 3) while preserving the fine texture details of the lesions.
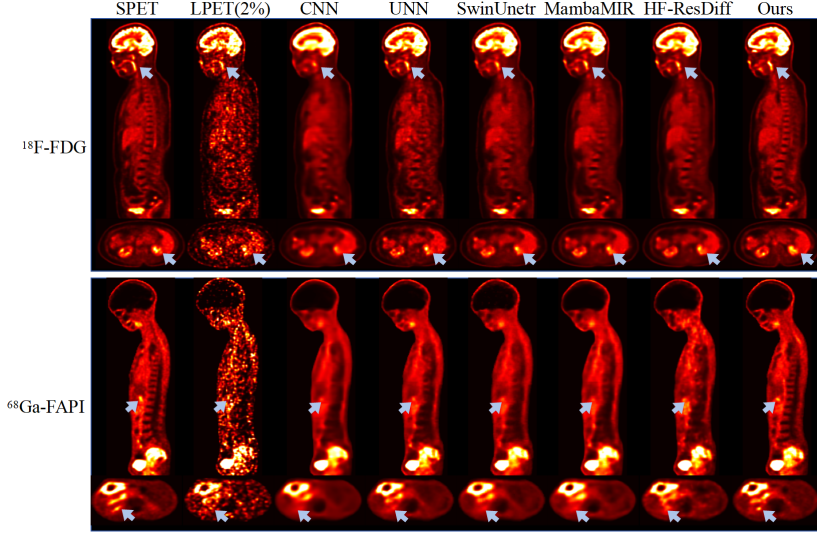
Fig. 3: Quantitative comparison with state-of-the-art methods at 2% dose level.

Table 2: Quantitative ablation results on the ^{18}F-FDG and ^{68}Ga-FAPI tracer datasets.

^{18}F-FDG					
PSNR↑/SSIM↑	6s	15s	30s	60s	150s
LPET	43.151/0.931	43.997/0.944	44.807/0.954	45.674/0.962	47.096/0.969
IDDPM	44.944/0.978	45.108/0.981	45.623/0.983	46.232/0.985	47.341/0.987
IDDPM+HWA	45.724/0.981	46.068/0.983	46.273/0.985	46.739/0.987	47.805/0.988
Ours	**46.035/0.984**	**46.523/0.987**	**46.889/0.988**	**47.315/0.989**	**48.042/0.990**
^{68}Ga-FAPI					
PSNR↑/SSIM↑	6s	15s	30s	60s	150s
LPET	40.443/0.873	41.382/0.892	42.302/0.908	43.224/0.922	44.689/0.935
IDDPM	44.790/0.973	45.363/0.974	45.825/0.976	46.154/0.977	46.796/0.979
IDDPM+HWA	45.992/0.976	46.187/0.978	46.369/0.980	46.648/0.980	47.287/0.982
Ours	**46.689/0.980**	**47.009/0.981**	**47.160/0.981**	**47.362/0.982**	**47.921/0.983**

3.4   Ablation Study

To validate the effectiveness of each proposed module, we conducted ablation studies with three sub-models: (1) IDDPM: the standard diffusion model; (2) IDDPM-HWA: IDDPM integrated with the CT-guided HWA; (3) MDAA-Diff (Ours): IDDPM incorporating both HWA and DAA module. All approaches were kept consistent in the experimental setup, ensuring a fair comparison (as shown in Table 2). The comparison results showed that IDDPM-HWA significantly improved denoising performance over IDDPM by leveraging CT fusion. MDAA-Diff further enhanced denoising capability through multi-dose adaptation and CT guidance, achieving clearer spinal and pulmonary boundaries with enhanced

structural fidelity. These findings confirmed that the HWA and DAA module effectively improve PET denoising performance, validating their critical role in modeling multi-dose characteristics and anatomical constraints.

4   Conclusion

In this study, we propose a CT-Guided Multi-Dose Adaptive Attention Diffusion Model (MDAA-Diff) for LPET denoising. It utilizes high-frequency wavelet transforms from CT images to integrate multimodal information. This approach address detail blurring limitations in traditional single-modal denoising models. Additionally, MDAA-Diff incorporates DAA module to explicitly model dose-dependent relationships, providing a robust and reliable solution for LPET imaging in clinical applications.

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

References

1. Arabi, H., Zaidi, H.: Improvement of image quality in pet using post-reconstruction hybrid spatial-frequency domain filtering. Physics in Medicine & Biology **63**(21), 215010 (2018)
2. Boellaard, R.: Standards for pet image acquisition and quantitative data analysis. Journal of nuclear medicine **50**(Suppl 1), 11S–20S (2009)
3. Catana, C.: The dawn of a new era in low-dose pet imaging (2019)
4. Chan, C., Fulton, R., Barnett, R., Feng, D.D., Meikle, S.: Postreconstruction nonlocal means filtering of whole-body pet with an anatomical prior. IEEE Transactions on medical imaging **33**(3), 636–650 (2013)
5. Chawla, S.C., Federman, N., Zhang, D., Nagata, K., Nuthakki, S., McNitt-Gray, M., Boechat, M.I.: Estimated cumulative radiation dose from pet/ct in children with malignancies: a 5-year retrospective review. Pediatric radiology **40**, 681–686 (2010)
6. Chen, K.T., Gong, E., de Carvalho Macruz, F.B., Xu, J., Boumis, A., Khalighi, M., Poston, K.L., Sha, S.J., Greicius, M.D., Mormino, E., et al.: Ultra–low-dose 18f-florbetaben amyloid pet imaging using deep learning with multi-contrast mri inputs. Radiology **290**(3), 649–656 (2019)
7. Cui, J., Wang, Y., Zhou, L., Fei, Y., Zhou, J., Shen, D.: 3d point-based multimodal context clusters gan for low-dose pet image denoising. IEEE Transactions on Circuits and Systems for Video Technology (2024)
8. Fallahpoor, M., Chakraborty, S., Pradhan, B., Faust, O., Barua, P.D., Chegeni, H., Acharya, R.: Deep learning techniques in pet/ct imaging: A comprehensive review from sinogram to image space. Computer methods and programs in biomedicine **243**, 107880 (2024)

9. 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)

10. Fu, Y., Dong, S., Huang, Y., Niu, M., Ni, C., Yu, L., Shi, K., Yao, Z., Zhuo, C.: Mpgan: multi pareto generative adversarial network for the denoising and quantitative analysis of low-dose pet images of human brain. Medical Image Analysis **98**, 103306 (2024)

11. Fu, Y., Dong, S., Huang, Y., Niu, M., Ni, C., Yu, L., Shi, K., Yao, Z., Zhuo, C.: Mpgan: multi pareto generative adversarial network for the denoising and quantitative analysis of low-dose pet images of human brain. Medical Image Analysis **98**, 103306 (2024)

12. Huang, B., Law, M.W.M., Khong, P.L.: Whole-body pet/ct scanning: estimation of radiation dose and cancer risk. Radiology **251**(1), 166–174 (2009)

13. Huang, J., Yang, L., Wang, F., Nan, Y., Aviles-Rivero, A.I., Schönlieb, C.B., Zhang, D., Yang, G.: Mambamir: An arbitrary-masked mamba for joint medical image reconstruction and uncertainty estimation. arXiv preprint arXiv:2402.18451 (2024)

14. Jiang, C., Pan, Y., Liu, M., Ma, L., Zhang, X., Liu, J., Xiong, X., Shen, D.: Pet-diffusion: Unsupervised pet enhancement based on the latent diffusion model. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 3–12. Springer (2023)

15. Nichol, A.Q., Dhariwal, P.: Improved denoising diffusion probabilistic models. In: International conference on machine learning. pp. 8162–8171. PMLR (2021)

16. Sanaat, A., Boccalini, C., Mathoux, G., Perani, D., Frisoni, G.B., Haller, S., Montandon, M.L., Rodriguez, C., Giannakopoulos, P., Garibotto, V., et al.: A deep learning model for generating [18f] fdg pet images from early-phase [18f] florbetapir and [18f] flutemetamol pet images. European Journal of Nuclear Medicine and Molecular Imaging **51**(12), 3518–3531 (2024)

17. Sanaat, A., Shiri, I., Arabi, H., Mainta, I., Nkoulou, R., Zaidi, H.: Deep learning-assisted ultra-fast/low-dose whole-body pet/ct imaging. European journal of nuclear medicine and molecular imaging **48**, 2405–2415 (2021)

18. Schaefferkoetter, J., Yan, J., Ortega, C., Sertic, A., Lechtman, E., Eshet, Y., Metser, U., Veit-Haibach, P.: Convolutional neural networks for improving image quality with noisy pet data. EJNMMI research **10**, 1–11 (2020)

19. Tang, Z., Jiang, C., Cui, Z., Shen, D.: Hf-resdiff: High-frequency-guided residual diffusion for multi-dose pet reconstruction. In: International Conference on Medical Image Computing and Computer-Assisted Intervention. pp. 372–381. Springer (2024)

20. 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)

21. Wang, Y., Zhou, L., Yu, B., Wang, L., Zu, C., Lalush, D.S., Lin, W., Wu, X., Zhou, J., Shen, D.: 3d auto-context-based locality adaptive multi-modality gans for pet synthesis. IEEE transactions on medical imaging **38**(6), 1328–1339 (2018)

22. Xiang, L., Qiao, Y., Nie, D., An, L., Lin, W., Wang, Q., Shen, D.: Deep auto-context convolutional neural networks for standard-dose pet image estimation from low-dose pet/mri. Neurocomputing **267**, 406–416 (2017)

23. Xie, H., Liu, Q., Zhou, B., Chen, X., Guo, X., Wang, H., Li, B., Rominger, A., Shi, K., Liu, C.: Unified noise-aware network for low-count pet denoising with varying count levels. IEEE Transactions on Radiation and Plasma Medical Sciences **8**(4), 366–378 (2023)

24. Xu, J., Gong, E., Pauly, J., Zaharchuk, G.: 200x low-dose pet reconstruction using deep learning. arXiv preprint arXiv:1712.04119 (2017)
25. Xue, S., Guo, R., Bohn, K.P., Matzke, J., Viscione, M., Alberts, I., Meng, H., Sun, C., Zhang, M., Zhang, M., et al.: A cross-scanner and cross-tracer deep learning method for the recovery of standard-dose imaging quality from low-dose pet. European journal of nuclear medicine and molecular imaging pp. 1–14 (2022)
26. 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)
27. Zhou, L., Schaefferkoetter, J.D., Tham, I.W., Huang, G., Yan, J.: Supervised learning with cyclegan for low-dose fdg pet image denoising. Medical image analysis **65**, 101770 (2020)
28. Zhou, Y., Yang, Z., Zhang, H., Chang, E.I.C., Fan, Y., Xu, Y.: 3d segmentation guided style-based generative adversarial networks for pet synthesis. IEEE Transactions on Medical Imaging **41**(8), 2092–2104 (2022)
29. Zhuang, H., Codreanu, I.: Growing applications of fdg pet-ct imaging in non-oncologic conditions. Journal of biomedical research **29**(3), 189 (2015)