



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

PET-Adapter: Test-Time Domain Adaptation for Full and Limited-Angle PET Image Reconstruction

Rüveyda Yilmaz¹[0009-0007-7351-698X], Yuli Wu²[0000-0002-6216-4911],
Johannes Stegmaier²[0000-0003-4072-3759], and
Volkmar Schulz¹[0000-0003-1341-9356]

¹ Institute of Imaging and Computer Vision, RWTH Aachen University, Germany

² Machine Learning for Medical Data, Heinrich Heine University Düsseldorf, Germany
Email:rueveyda.yilmaz@lfb.rwth-aachen.de

Abstract. Positron Emission Tomography (PET) image reconstruction is inherently challenged by Poisson noise and physical degradation factors, which are further exacerbated in limited-angle acquisitions. While deep learning methods demonstrate promising performance, their generalization to unseen clinical data distributions remains limited without extensive retraining. We propose PET-Adapter, a test-time domain adaptation framework for generative PET reconstruction models pretrained solely on phantom data. Our method enables adaptation to clinical datasets with varying anatomies, tracers, and scanner configurations without requiring paired ground truth. PET-Adapter introduces layer-wise low-rank anatomical conditioning during adaptation and Ordered Subset Expectation Maximization-based warm-starting that initializes the generation from physics-informed reconstructions, reducing diffusion steps from 50 to 2 without compromising quality. Experiments across multiple clinical datasets demonstrate superior 3D reconstruction performance in both full-angle and limited-angle settings, highlighting the clinical feasibility and computational efficiency of the proposed approach. The code base for PET-Adapter will be made publicly available under <https://github.com/ruveydayilmaz0/PetAdapter>.

Keywords: Limited-Angle PET · PET Imaging · Diffusion Models

1 Introduction

Positron Emission Tomography (PET) is a powerful medical imaging modality widely used for the diagnosis, treatment planning, and therapeutic monitoring of diseases such as cancer, epilepsy, and Alzheimer’s disease [4, 6, 33]. Despite its clinical value, PET image reconstruction remains challenging due to Poisson noise, Compton scattering, random coincidences, and detector gaps [4]. Conventional iterative reconstruction methods [3, 8] partially mitigate these effects; however, they are limited in their ability to produce high-quality images and are increasingly being challenged by Deep Learning (DL) approaches [28].

DL-based post-processing methods can be applied in the image domain to denoise reconstructions from conventional algorithms [9, 18]. Although effective,

these approaches are inherently constrained by the quality of the initial reconstruction and do not leverage raw measurement data. In contrast, end-to-end models are trained on measurement-ground truth (GT) pairs and act as standalone reconstructors [10, 29]. While these methods often produce sharper images, they may hallucinate structures that are inconsistent with the measured data [6]. To address this, DL-based iterative methods integrate neural networks as regularizers or learned priors within classical optimization frameworks (e.g., ADMM or unrolled FBSEM), enforcing stronger data consistency [12, 20, 28]. To reduce the reliance on large, paired training datasets, recent work has shifted toward unsupervised methods. Deep Image Prior (DIP)-based methods have been adapted to PET sinogram and list-mode data, demonstrating improved preservation of high-resolution information [7, 14]. More recently, generative modeling has advanced toward diffusion models with PET-specific adaptations and joint PET-magnetic resonance (MR) image reconstruction [20, 30]. Beyond standard reconstruction challenges, limited-angle (LA) PET introduces additional difficulties due to incomplete measurements from dual-panel geometries or partial-ring scanners [11, 16]. Prior approaches to LA artifact mitigation include supervised U-Nets to map incomplete data to full-ring equivalents [16], conditional diffusion models for sinogram inpainting [31], or implicit neural representations [11]. However, these approaches for both full-angle (FA) and LA PET image reconstruction often exhibit limited generalization to previously unseen data distributions, thereby constraining their practical applicability. As a result, there remains a need for methods capable of adapting to variations in scanners, tracers, and anatomical characteristics without requiring extensive retraining on target data with paired reference GT images. Hashimoto et al. [5] and Webber et al. [27] address this problem by adapting a Deep Diffusion Image Prior and a Score-Based Generative Model, respectively, on a per-scan basis. We introduce PET-Adapter, a dataset-wide adaptation strategy for Test-Time Domain Adaptation (TTDA) that generalizes to unseen clinical datasets using a model pretrained on phantom images. Our main contributions are as follows: (1) We present an unsupervised adaptation framework with a tailored loss function to transfer generative PET reconstruction models trained on phantom data to clinical datasets with varying anatomies, pathologies, and tracers, thereby improving 3D reconstruction performance in both FA and LA settings. (2) We propose a strategy to incorporate MR images into unconditional pretrained models during adaptation via layer-wise conditioning [15, 23], providing anatomical guidance. (3) We introduce a warm-starting strategy via diffusion inversion [21] to initialize the inference with reconstructions from the Ordered Subset Expectation Maximization (OSEM) method [8], reducing the diffusion sampling steps from 50 to 2 while maintaining quality.

2 Method

PET image reconstruction is an ill-posed inverse problem, primarily due to the ill-conditioned nature of the system matrix and the stochastic Poisson




Fig. 1: The architecture of PET-Adapter. (a) Pretraining on phantom data is followed by (b) test-time domain adaptation on clinical target domains via (c) layer-wise adaptation blocks. The final reconstruction is initialized with (d) warm-starting (WS) and (e) optimized through T alternating steps of prior refinement and data-consistency updates.

noise inherent to radioactive decay. This condition is further exacerbated by physical degrading factors such as photon attenuation, scattering, and random coincidences [20]. The statistical model of the acquisition process is defined as:

$$\mathbf{y} \sim \text{Poisson}(\mathbf{Ax} + \mathbf{b}), \quad (1)$$

where $\mathbf{y} \in \mathbb{N}^m$ represents the measured projection data (prompts). $\mathbf{A} \in \mathbb{R}^{m \times n}$ is the system matrix (or forward projector), modeling the probability that an emission from voxel j , where $j \in \{0, 1, \dots, n-1\}$, is detected by the detector pair i , where $i \in \{0, 1, \dots, m-1\}$. $\mathbf{x} \in \mathbb{R}_{\geq 0}^n$ denotes the unknown radiotracer activity distribution, and $\mathbf{b} \in \mathbb{R}^m$ accounts for the expected additive background events; namely, scattered and random coincidences. PET image reconstruction aims to estimate the activity distribution $\hat{\mathbf{x}}$ from the noisy measurements $\mathbf{y}$.

Score-based generative models can generate samples by learning to reverse a stochastic noising process described by Stochastic Differential Equations (SDEs) [22]. While the standard reverse process requires many iterative steps, the Denoising Diffusion Implicit Model (DDIM) formulation allows for accelerated sampling [21]. A generalized update step at timestep $t_k \in (0, 1]$, which can interpolate between deterministic DDIM sampling and stochastic variance-preserving SDE, is formulated as:

$$\mathbf{x}_{t_{k-1}} = \gamma_{t_{k-1}} \hat{\mathbf{x}}_0(\mathbf{x}_{t_k}) - v_{t_k} \sqrt{v_{t_{k-1}}^2 - \eta_{t_k}^2} s_\theta(\mathbf{x}_{t_k}, t_k) + \eta_{t_k} \mathbf{z} \text{ with } \mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}), \quad (2)$$

where $\hat{\mathbf{x}}_0(\mathbf{x}_{t_k}) := (\mathbf{x}_{t_k} + v_{t_k}^2 s_\theta(\mathbf{x}_{t_k}, t_k)) / \gamma_{t_k}$ denotes Tweedie's estimate of the clean image derived from the noisy state $\mathbf{x}_{t_k}$, and s_θ is the learned score function.

Table 1: Quantitative results for each reconstruction method across clinical test-time datasets for (a) full-angle (FA) and (b) limited-angle (LA) acquisitions.

	Method	CERMEP		AGIEF		NX		Huntington	
		PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
a) FA	OSEM	29.2	0.85	28.9	0.82	30.7	0.89	26.7	0.88
	Bowsher	30.4	0.89	29.5	0.83	32.1	0.94	<u>28.0</u>	<u>0.90</u>
	PET-PIR ($T = 50$)	24.5	0.84	27.2	0.77	<u>36.4</u>	0.96	25.5	0.84
	PET-DDS ($T = 50$)	27.4	0.83	<u>32.1</u>	<u>0.91</u>	35.7	0.94	24.9	0.81
	LiSch-SCD ($T = 50$)	<u>32.7</u>	<u>0.92</u>	29.4	0.85	35.7	<u>0.95</u>	28.3	0.87
	Ours ($T = 2$)	35.3	0.96	35.4	0.96	37.6	<u>0.95</u>	30.7	0.91
b) LA	OSEM	25.5	0.74	24.9	0.69	23.8	0.67	25.0	<u>0.75</u>
	Bowsher	26.2	0.77	25.6	0.71	25.7	0.75	25.9	0.78
	PET-PIR ($T = 50$)	22.9	0.67	21.4	0.64	30.5	<u>0.90</u>	22.7	0.69
	PET-DDS ($T = 50$)	22.2	0.70	<u>28.4</u>	<u>0.79</u>	30.8	0.89	22.3	0.67
	LiSch-SCD ($T = 50$)	<u>26.4</u>	<u>0.81</u>	27.3	0.78	<u>30.9</u>	<u>0.90</u>	23.4	0.70
	Ours ($T = 2$)	32.2	0.91	31.7	0.87	32.0	0.92	<u>25.2</u>	0.78

The coefficients γ_t , v_t , and η_t are time-dependent scalars that define the signal schedule (mean and variance) and the degree of stochasticity, respectively [20]. s_θ is trained with the standard denoising score matching (DSM) loss L_{DSM} [25]. To adapt this sampling approach for PET image reconstruction, Singh et al. [20] propose a framework with a Maximum A Posteriori (MAP) objective that enforces data consistency, producing $\hat{\mathbf{x}}^{\text{MAP}}$. This is achieved by splitting the update into a prior term (U-Net-based diffusion) and a likelihood term. Specifically, the intermediate estimate $\hat{\mathbf{x}}_0$ is refined by maximizing the penalized Poisson log-likelihood (MAP objective) using preconditioned gradient ascent with the standard expectation maximization preconditioner. For our experiments in Section 3, we employ this approach as the phantom-pretrained model (see Fig. 1a) to showcase the adaptation capabilities of PET-Adapter on novel datasets.

Test-Time Domain Adaptation (TTDA): For PET image reconstruction, anatomical guidance from computed tomography (CT) or MR images is commonly used to regularize the reconstruction and improve image quality [6, 7]. Given the superior soft-tissue contrast of magnetic resonance imaging (MRI) data, we focus on MRI-guided PET image reconstruction. We condition the reconstruction model on MR images at test time by introducing a conditional branch that modulates the original network blocks via affine transformations [15, 23]. To integrate this conditioning into the convolutional blocks of the U-Net-based pretrained score function s_θ , we perform conditional low-rank adaptation [23]. Specifically, for a layer l with input $\mathbf{x}^l \in \mathbb{R}^k$, the adapted output $\mathbf{h}^l$ is defined as follows:

$$\mathbf{h}^l = \mathbf{W}_0 \mathbf{x}^l + \Delta \mathbf{W} \mathbf{x}^l = \mathbf{W}_0 \mathbf{x}^l + \mathbf{U} \phi(\mathbf{D} \mathbf{x}^l \mid g(\mathbf{x}_{\text{MR}})), \quad (3)$$

with $\phi(\mathbf{z} \mid g(\mathbf{x}_{\text{MR}})) = \gamma_\phi(g(\mathbf{x}_{\text{MR}})) \odot \mathbf{z} + \beta_\phi(g(\mathbf{x}_{\text{MR}}))$

Here, $\mathbf{W}_0, \Delta\mathbf{W} \in \mathbb{R}^{d \times k}$ denote the frozen pretrained weights and their adaptation, respectively. The matrices $\mathbf{D} \in \mathbb{R}^{r \times k}$ and $\mathbf{U} \in \mathbb{R}^{d \times r}$ are learnable low-rank adaptation matrices with rank $r \ll \min(d, k)$. The learnable mapper g projects the conditioning input $\mathbf{x}_{\text{MR}}$ to the latent adaptation space, while γ_ϕ and β_ϕ produce layer-wise scale and shift parameters from $g(\mathbf{x}_{\text{MR}})$ (see Fig. 1b, c). At test time, we adapt the pretrained model s_θ to target datasets by learning only these additional parameters. We define an adaptation objective L_{adapt} combining a measurement-based loss (L_p) and an MRI-driven anatomical loss (L_a). For L_p , we minimize the Poisson NLL between the acquired measurements $\mathbf{y}$ and the forward projection $\hat{\mathbf{y}} = \mathbf{A}\hat{\mathbf{x}}^{\text{MAP}} + \mathbf{b}$ [2]. The MRI-driven anatomical loss L_a is weighted by the total variation (TV) of $\mathbf{x}_{\text{MR}}$, penalizing gradients more strongly in anatomically homogeneous regions and thereby encouraging the reconstruction to conform to the underlying anatomy. From each MR image $\mathbf{x}_{\text{MR}} \in \mathbb{R}_{\geq 0}^n$, we derive weights w_p , where $p \in \{0, 1, \dots, n-1\}$, to encourage smoothing within homogeneous tissue regions while preserving edges aligned with anatomical boundaries. We define

$$w_j = \exp\left(-\frac{\|(\nabla\mathbf{x}_{\text{MR}})_j\|}{\sigma}\right), \quad (4)$$

where $\|(\nabla\mathbf{x}_{\text{MR}})_j\|$ is the gradient magnitude at pixel j , and σ controls sensitivity to anatomical edges. The final adaptation objective is $L_{\text{adapt}} = L_p + \lambda L_a$ where λ controls the trade-off between data fidelity and anatomical regularization. Unlike prior TTDA approaches that perform per-image adaptation [5, 27], we perform dataset-level global adaptation for all reconstructions, reducing overhead while improving generalization and limiting per-image overfitting.

Warm-starting (WS): Classical analysis of maximum-likelihood expectation-maximization (MLEM) or OSEM shows that low-frequency spatial components associated with large singular values of the system matrix converge rapidly, placing intermediate reconstructions close to the anatomical image manifold [32]. Due to the ill-posed nature of the inverse problem, continued iterations tend to fit high-frequency Poisson noise rather than recover meaningful anatomical details, progressively driving the reconstruction away from the clean PET image manifold. Guided by this geometric insight, we use the OSEM reconstructions to bypass the redundant low-frequency generation steps and later rely on the model s_θ to predict the high-frequency structures. Specifically, we predict the noisy image corresponding to the OSEM images via diffusion inversion [21] and initiate denoising from this noise representation instead of random noise (see Fig. 1d). This reduces the number of required diffusion steps substantially, without introducing additional computational overhead for computing the OSEM images, as these images are already required for measurement-based normalization regardless of the initialization strategy [20, 27] (see Section 3). During inference, after WS, we perform T alternating steps of prior refinements using s_θ and data-consistency updates to compute $\hat{\mathbf{x}}^{\text{MAP}}$ (see Fig. 1e).




Fig. 2: Qualitative results for each reconstruction method and the reference GT volumes from axial (left), coronal (top right), and sagittal (bottom right) views.

3 Experiments and Results

Datasets: We work exclusively with joint PET–MRI datasets for anatomically guided PET image reconstruction, using brain phantom images for training and clinical datasets for domain adaptation experiments. The BrainWeb dataset consists of 20 simulated FDG-PET volumes with coregistered MR images and represents a controlled, healthy human anatomy [1]. CERMEP [13] provides clinical brain PET–MRI data from healthy adults, which we use to evaluate the adaptation performance of the model trained on simulated phantoms on real data. The AGIEF dataset [33] includes scans from both healthy subjects and epilepsy patients, with unseen pathological metabolic patterns—such as focal

Table 2: Ablation results comparing the full adaptation pipeline (PET-Adapter) to no warm-starting (*w/o* WS) or no adaptation (*w/o* TTDA) for (a) full-angle (FA) and (b) limited-angle (LA) acquisitions.

	Method		CERMEP		AGIEF		NX		Huntington	
			PSNR	SSIM	PSNR	SSIM	PSNR	SSIM	PSNR	SSIM
a) FA	PET-Adapter	($T = 2$)	35.3	0.96	35.4	0.96	37.6	0.95	30.7	0.91
	<i>w/o</i> WS	($T = 2$)	29.6	0.85	29.1	0.84	30.9	0.90	26.7	0.79
	<i>w/o</i> WS	($T = 50$)	34.9	0.95	36.4	0.97	35.2	0.95	31.2	0.95
	<i>w/o</i> TTDA	($T = 2$)	26.2	0.74	26.3	0.73	34.7	0.91	24.7	0.78
b) LA	PET-Adapter	($T = 2$)	32.2	0.91	31.7	0.87	32.0	0.92	25.2	0.78
	<i>w/o</i> WS	($T = 2$)	25.6	0.75	28.4	0.84	28.8	0.84	23.2	0.75
	<i>w/o</i> WS	($T = 50$)	32.6	0.93	30.4	0.86	30.8	0.90	26.2	0.84
	<i>w/o</i> TTDA	($T = 2$)	21.8	0.65	23.3	0.62	28.3	0.79	20.6	0.62

hypometabolism and altered oxygen–glucose indices—not present in BrainWeb’s healthy simulations. The NeuroExplorer (NX) project provides multi-tracer data acquired using a long axial field-of-view PET/CT scanner, together with corresponding MR images, introducing previously unseen conditions arising from both novel hardware configurations and tracer distributions [26]. To probe extreme anatomical domain shifts, we also use animal models of Huntington’s disease [17]. Specifically, mouse brain MRI data are used to assign theoretical [^{18}F]FDG uptake values to segmented gray and white matter, thereby simulating reference PET images [12]. Unlike the NX dataset, the PET images provided in CERMEP and AGIEF are of low quality; therefore, for these datasets, we first generate high-SNR reference GT volumes by averaging the PET values within each brain region defined by the MRI segmentations. Next, for each dataset, we simulate PET measurements from the GT images using the Parallelproj framework [19], modeling PET physics and noise, with true counts per emission volume set to 10 and 40% of the detectors disabled during the LA simulations.

Training: We train the score model s_θ on 2D slices of 3D images due to the limited number of training volumes. s_θ is still used within a 3D reconstruction framework, where inter-slice consistency is enforced through warm-start initialization from 3D OSEM reconstructions. The same model weights are used for both FA and LA experiments, as the prediction of the data prior is independent of the measurement geometry. The network is trained for 2000 epochs with a batch size of 32 and a learning rate of 1×10^{-4} on an NVIDIA RTX 4000 GPU. To account for large intensity variations across PET scans, image-wise normalization is applied to set the mean emission value of each training image to 1.0 [20].

TTDA: To adapt s_θ to out-of-distribution (OOD) data with additional anatomical conditioning, we set the hyperparameters based on two held-out samples from each target dataset [2]. For LA experiments, the contributions from missing sinogram regions are masked out when computing the loss L_p .

Inference: During inference, we set the number of generation steps T to 2,

which is substantially fewer than in other generative PET reconstruction methods [20, 27], enabled by warm-start initialization. The prior predictions ($\hat{\mathbf{x}}_0$) from the score function s_θ are denormalized to the original image intensity range before the data consistency step to ensure a proper comparison between $\mathbf{y}$ and $\hat{\mathbf{y}}$ [20]. Since GT images are unavailable during adaptation and inference, the intensity range is estimated from OSEM reconstructions with 30 subsets and 2 iterations following [20]. For warm-starting, we reverse the SDE to step $t_r = 0.4$.

Comparative Methods: We evaluate our approach against established PET reconstruction methods using Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM). Baseline comparisons include classical algorithms, namely, OSEM [8] and Maximum a Posteriori Expectation Maximization with a Bowsher prior (MAPEM-Bowsher) [3], where MR images are used to construct the anatomical prior. We additionally compare against recent generative-model-based PET reconstruction methods. These include PET-DDS [20], which applies classifier-free guidance (CFG) using MR images, and DiffPIR with feature injection conditioning [24], which we adapt to incorporate structural MRI information and PET-specific measurement updates, and rename as PET-PIR. Finally, we compare to PET-LiSch-SCD (LiSch-SCD in short) [27], which applies the Steerable Conditional Diffusion framework [2] to PET image reconstruction. Following the published description, we implemented this method based on CFG-guided PET-DDS, as the original code is not publicly available.

Results: We report the quantitative results in terms of PSNR and SSIM in Table 1. Notably, PETAdapter achieves successful TTDA using only $T = 2$ generation steps, compared to the other diffusion-based methods listed in the table, which require $T = 50$ steps. As expected, the LA setting (1b) yields lower quality than the FA setting (1a) due to incomplete measurements. Figure 2 presents qualitative comparisons in axial, coronal, and sagittal views for the FA setting (the LA reconstructions are not shown due to space constraints). The visual results demonstrate that PETAdapter produces reconstructions with superior structural fidelity and intensity consistency compared to the other methods.

Ablations: To evaluate the contribution of different components in PETAdapter, we conduct ablation studies by omitting WS and TTDA. In the absence of WS, test-time adaptation is still performed; however, the reconstruction is initialized from pure noise samples instead of OSEM reconstructions. For this setting, we consider two configurations with $T = 2$ and $T = 50$ iterations. With $T = 2$, the reconstruction quality degrades substantially across all datasets compared to the full model with WS, as the limited number of iterations is insufficient for adequate prior refinement and measurement updates (see Table 2). Increasing the number of iterations to $T = 50$ restores the performance to a level comparable to $T = 2$ with WS, highlighting the effectiveness of WS in reducing the required number of refinement steps. To assess the impact of TTDA, we retain WS, set $T = 2$, and omit adaptation to the target dataset. This results in a significant performance drop due to domain shift, as shown in Table 2.

4 Conclusion

In this work, we introduced PET-Adapter for TTDA of generative PET image reconstruction models. By enabling robust adaptation to novel clinical datasets, it improves the generalization across varying scanners, tracers, anatomies, and pathologies without extensive retraining. Experiments in both FA and LA settings demonstrate that PET-Adapter achieves superior performance while reducing generative sampling steps from 50 to 2 via OSEM warm-starting. Coupled with dataset-level adaptation, this reduction highlights its potential for immediate, computationally efficient clinical deployment.

Acknowledgements: We acknowledge the use of the CERMEP-iDB-MRXFDG dataset. © Copyright CERMEP – Imagerie du vivant, www.cermep.fr and Hospices Civils de Lyon. All rights reserved.

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

References

1. 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 Transactions on Medical Imaging* **25**(11), 1410–1416 (2006)
2. Barbano, R., Denker, A., Chung, H., Roh, T.H., Arridge, S., Maass, P., Jin, B., Ye, J.C.: Steerable conditional diffusion for out-of-distribution adaptation in medical image reconstruction. *IEEE Transactions on Medical Imaging* **44**(5), 2093–2104 (2025)
3. Bowsher, J.E., Yuan, H., Hedlund, L.W., Turkington, T.G., Akabani, G., Badea, A., Kurylo, W.C., Wheeler, C.T., Cofer, G.P., Dewhurst, M.W., et al.: Utilizing MRI information to estimate F18-FDG distributions in rat flank tumors. In: *IEEE Symposium Conference Record Nuclear Science 2004*. vol. 4, pp. 2488–2492 (2004)
4. Cherry, S.R., Dahlbom, M.: *PET: Physics, Instrumentation, and Scanners*. Springer (2006)
5. Hashimoto, F., Gong, K.: Pet image reconstruction using deep diffusion image prior. *IEEE Transactions on Medical Imaging* (2026)
6. Hashimoto, F., Onishi, Y., Ote, K., Tashima, H., Reader, A.J., Yamaya, T.: Deep learning-based PET image denoising and reconstruction: a review. *Radiological Physics and Technology* **17**(1), 24–46 (2024)
7. Hashimoto, F., Onishi, Y., Ote, K., Tashima, H., Yamaya, T.: Fully 3D implementation of the end-to-end deep image prior-based PET image reconstruction using block iterative algorithm. *Physics in Medicine & Biology* **68**(15), 155009 (2023)
8. 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)
9. 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)
10. Liu, Z., Ye, H., Liu, H.: Deep-learning-based framework for PET image reconstruction from sinogram domain. *Applied Sciences* **12**(16), 8118 (2022)

11. Makkar, S., Ye, S., Béguin, M., Dissertori, G., Hrbacek, J., Lomax, A., McNamara, K., Ritzer, C., Weber, D.C., Xing, L., et al.: Partial-ring PET image correction using implicit neural representation learning. In: 20th International Conference on the use of Computers in Radiation Therapy, Lyon, France. pp. 8–11 (2024)
12. Mehranian, A., Reader, A.J.: Model-based deep learning PET image reconstruction using forward–backward splitting expectation–maximization. *IEEE Transactions on Radiation and Plasma Medical Sciences* **5**(1), 54–64 (2021)
13. Mérida, I., Jung, J., Bouvard, S., Le Bars, D., Lancelot, S., Lavenne, F., Bouillot, C., Redouté, J., Hammers, A., Costes, N.: CERMEP-IDB-MRXFDG: a database of 37 normal adult human brain [18F] FDG PET, T1 and FLAIR MRI, and CT images available for research. *EJNMMI Research* **11**(1), 91 (2021)
14. Ote, K., Hashimoto, F., Onishi, Y., Isobe, T., Ouchi, Y.: List-mode PET image reconstruction using deep image prior. *IEEE Transactions on Medical Imaging* **42**(6), 1822–1834 (2023)
15. Perez, E., Strub, F., De Vries, H., Dumoulin, V., Courville, A.: Film: Visual reasoning with a general conditioning layer. In: Proceedings of the AAAI Conference on Artificial Intelligence. vol. 32 (2018)
16. Raj, J., Millardet, M., Krishnamoorthy, S., Karp, J.S., Surti, S., Matej, S.: Recovery of the spatially-variant deformations in dual-panel PET reconstructions using deep-learning. *Physics in Medicine & Biology* **69**(5), 055028 (2024)
17. Sawiak, S.J., Morton, A.J.: The cambridge MRI database for animal models of huntington disease. *NeuroImage* **124**, 1260–1262 (2016)
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), 105 (2020)
19. Schramm, G., Thielemans, K.: Parallelproj—an open-source framework for fast calculation of projections in tomography. *Frontiers in Nuclear Medicine* **3**, 1324562 (2024)
20. Singh, I.R., Denker, A., Barbano, R., Kereta, Ž., Jin, B., Thielemans, K., Maass, P., Arridge, S., et al.: Score-based generative models for PET image reconstruction. *Machine Learning for Biomedical Imaging* **2**, 547–585 (2024)
21. Song, J., Meng, C., Ermon, S.: Denoising diffusion implicit models. In: International Conference on Learning Representations (2021)
22. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: International Conference on Learning Representations (2021)
23. Stracke, N., Baumann, S.A., Susskind, J., Bautista, M.A., Ommer, B.: CTRLorAL-Ter: conditional loradapter for efficient 0-shot control and altering of t2i models. In: European Conference on Computer Vision. pp. 87–103. Springer (2024)
24. Tumanyan, N., Geyer, M., Bagon, S., Dekel, T.: Plug-and-play diffusion features for text-driven image-to-image translation. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. pp. 1921–1930 (2023)
25. Vincent, P.: A connection between score matching and denoising autoencoders. *Neural Computation* **23**(7), 1661–1674 (2011)
26. Volpi, T., Toyonaga, T., Khattar, N., Gallezot, J.D., Naganawa, M., Vanderlinden, G., Honhar, P., Zeng, T., Fontaine, K., Mulnix, T., et al.: Exceptional brain PET images from the NeuroEXPLORER: scans with targeted radiopharmaceuticals and comparison to HRRT. *European Journal of Nuclear Medicine and Molecular Imaging* pp. 1–12 (2025)

27. Webber, G., Hammers, A., King, A., Reader, A.: Steerable conditional diffusion for domain adaptation in PET image reconstruction. In: 2025 IEEE Nuclear Science Symposium (NSS), Medical Imaging Conference (MIC) and Room Temperature Semiconductor Detector Conference (RTSD). pp. 1–1. IEEE (2025)
28. Webber, G., Reader, A.J.: Diffusion models for medical image reconstruction. *BJR| Artificial Intelligence* **1**(1), ubae013 (2024)
29. Whiteley, W., Luk, W.K., Gregor, J.: Directpet: full-size neural network PET reconstruction from sinogram data. *Journal of Medical Imaging* **7**(3), 032503–032503 (2020)
30. Xie, T., Cui, Z.X., Luo, C., Wang, H., Liu, C., Zhang, Y., Wang, X., Zhu, Y., Chen, G., Liang, D., et al.: Joint diffusion: mutual consistency-driven diffusion model for PET-MRI co-reconstruction. *Physics in Medicine & Biology* **69**(15), 155019 (2024)
31. Yilmaz, R., Thull, J., Stegmaier, J., Schulz, V.: Diffusion-based sinogram interpolation for limited angle PET. In: 2025 IEEE Nuclear Science Symposium (NSS), Medical Imaging Conference (MIC) and Room Temperature Semiconductor Detector Conference (RTSD). pp. 1–1 (2025)
32. Zeng, G.L.: Maximum-likelihood expectation-maximization algorithm versus windowed filtered backprojection algorithm: a case study. *Journal of Nuclear Medicine Technology* **46**(2), 129–132 (2018)
33. Zhang, M., Qin, Q., Zhang, S., Liu, W., Meng, H., Xu, M., Huang, X., Lin, X., Lin, M., Herman, P., et al.: Aerobic glycolysis imaging of epileptic foci during the inter-ictal period. *EBioMedicine* **79** (2022)