Distribution-Guided Multi-Tracer Brain PET Synthesis from Structural MRI with Class-Conditioned Weighted Diffusion

Minhui Yu[1,2], David S. Lalush[2], Derek C. Monroe[3], Kelly S. Giovanello[4], Weili Lin[1], Pew-Thian Yap[1], Jason P Mihalik[3], and Mingxia Liu[1*]

[1] Department of Radiology and Biomedical Research Imaging Center, University of North Carolina at Chapel Hill (UNC-CH), Chapel Hill, NC 27599, USA
[2] Lampe Joint Department of Biomedical Engineering, UNC-CH and North Carolina State University, Chapel Hill, NC 27599, USA
[3] Department of Exercise and Sport Science, UNC-CH, Chapel Hill, NC, USA
[4] Department of Psychology and Neuroscience, UNC-CH, Chapel Hill, NC, USA
*Corresponding author (mingxia_liu@med.unc.edu)

Abstract. Multi-tracer positron emission tomography (PET), which assesses key neurological biomarkers such as tau pathology, neuroinflammatory, β-amyloid deposition, and glucose metabolism, plays a vital role in diagnosing neurological disorders by providing complementary insights into the brain's molecular and functional state. Acquiring multi-tracer PET scans remains challenging due to high costs, radiation exposure, and limited tracer availability. Recent studies have attempted to synthesize multi-tracer PET images from structural MRI. However, these approaches typically either rely on direct mappings to individual tracers or lack distributional constraints, leading to inconsistencies in image quality across tracers. To this end, we propose a normalized diffusion framework (NDF) to generate high-quality multi-tracer PET images from a single MRI through a distribution-guided class-conditioned weighted diffusion model. Specifically, a diffusion model conditioned on MRI and tracer-specific class labels is trained to synthesize PET images of multiple tracers, and a pre-trained normalizing flow model refines these outputs by mapping them into a shared distribution space. This mapping ensures that the subject-specific high-level features across different PET tracers are preserved, resulting in more consistent and accurate synthesis. Experiments on a total of 425 subjects with multi-tracer PET scans demonstrate that our NDF outperforms current state-of-the-art methods, indicating its potential for advancing multi-tracer PET synthesis.

Keywords: Multi-Tracer PET · Synthesis · MRI · Diffusion

1 Introduction

Neurological disorders, such as Alzheimer's disease (AD) and other neurodegenerative conditions, involve complex pathological processes that require advanced

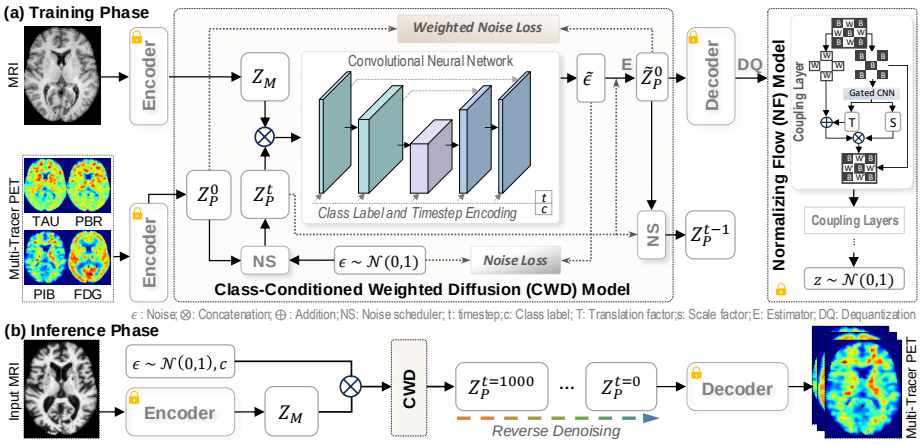

Fig. 1. Illustration of the proposed normalized diffusion framework (NDF) that synthesizes multi-tracer PET images from a single T1-weighted MRI input.

imaging techniques for accurate diagnosis and progression monitoring [1–5]. Multi-tracer positron emission tomography (PET) is a powerful tool that enables the assessment of multiple molecular and functional biomarkers, providing complementary insights into brain pathology. For example, PET imaging is widely used to study *Tau* pathology with ^{18}F-T807 [6], neuroinflammatory with ^{18}F-PBR111 [7], β-amyloid deposition with ^{11}C-PIB [8] and glucose metabolism with ^{18}F-FDG [9]. While each tracer highlights distinct pathological processes, they collectively contribute to comprehensive understanding of neurodegeneration.

Despite the advantages, multi-tracer PET imaging faces significant challenges in clinical and research settings due to high costs, radiation exposure, and limited tracer availability. Recent studies have explored deep learning-based methods to synthesize PET images from structural MRI as a cost-effective and non-invasive alternative. However, existing approaches often rely on direct MRI-to-PET mappings for single-tracer PET synthesis [10–14]. Even though several studies [15,16] have proposed to generate multi-tracer PET images, they typically lack explicit distributional constraints, which could lead to inconsistent image quality and reduced reliability across synthesized multi-tracer PET images.

To this end, we propose a normalized diffusion framework (NDF) to generate multi-tracer PET images from a single T1-weighted (T1) MRI through a distribution-guided diffusion model. As shown in Fig. 1 (a), the input 3D MRI and PET images are first encoded to a latent space. Then, a class-conditioned weighted diffusion (CWD) model simulates Markov chain transitions to generate latent PET features by iteratively refining a Gaussian noise, conditioned on the latent MRI representation and tracer-specific class labels. The synthesized latent PET features are then decoded into PET images. A pretrained normalizing flow (NF) model simultaneously constrains the training process, guiding the decoded images across tracers to align with a normal distribution. This ensures that synthetic PET images match the target tracer distribution. Experimental results

on a total of 425 subjects with multi-tracer PET scans validate the superiority of our method over the current state-of-the-art (SOTA) methods.

2 Materials and Methodology

Subjects and Image Preprocessing. Two datasets are involved in this work. (1) NFL-LONG [17], which includes 132 all-male subjects with an average age of 60.49 ± 6.24. Random tracers are missing for some subjects due to various reasons, resulting in different numbers of subjects for each modality pairing: 114 subjects had both T1 MRI and PBR-PET images, 90 subjects had both T1 MRI and PIB-PET images, and 97 subjects had both T1 MRI and TAU-PET images. PET preprocessing involves initial rigid alignment of PET scans to each subject's T1 MR image (subject space), followed by nonlinear transformation of images to a standardized MNI space. All PET images from this dataset are processed to have the resolution of $2 \times 2 \times 2\ mm^3$. A common MNI brain mask is then applied to each image for skull stripping. (2) The public ADNI dataset [18] with 293 cognitively normal subjects (141 female, 152 male; average age 74.21 ± 5.97) with paired MRI and FDG-PET data. MR image preprocessing involves skull stripping, intensity normalization, and nonlinear registration to the MNI space. PET images undergo initial skull stripping, followed by linear alignment to their corresponding MR images, and are subsequently registered to MNI using the MRI-derived deformation fields. To exclude non-informative background regions while preserving the complete brain, images from both datasets are uniformly cropped to $80 \times 96 \times 80$. Image intensities are normalized to the range $[-1, 1]$ during training, and $[0, 1]$ for quantitative and qualitative evaluation.

Proposed Method. Our normalized diffusion framework (NDF) synthesizes multi-tracer PET images from a single MRI input. As illustrated in Fig. 1, it consists of three key components: (1) a latent space encoding and decoding, (2) a class-conditioned weighted diffusion (CWD) model for latent PET generation, and (3) a pretrained normalizing flow (NF) model for distribution alignment.

(1) ***Latent Space Encoding and Decoding***. To ensure computation efficiency, an autoencoder with an encoder and a decoder is pretrained to transform PET images into a compact latent representation, reducing the spatial resolution from $80\times96\times480$ to $10\times12\times10$. The encoder maps the input into latent space, while the decoder reconstructs it back into a 3D PET image. The autoencoder training is initially guided by an l_1 loss, a perceptual loss, and a Kullback-Leibler divergence loss [19]. After a 15-epoch warm-up, a patch-based adversarial loss is introduced to further refine the model through adversarial training.

(2) ***Latent PET Generation with CWD***. A class-conditioned weighted diffusion (CWD) model is used to synthesize latent PET from latent MRI [13]. The proposed CWD model learns to progressively denoise Gaussian noise into meaningful PET features. This process follows a Markov chain transition model, conditioned on the encoded MRI latent features and tracer-specific class labels (with each tracer corresponding to a specific class). In the training phase, the

pretrained encoder maps both MRI and PET images into the latent space, producing corresponding latent features Z_M and Z_P^0 as inputs to the CWD model. At each training step, a timestep t is randomly chosen from 1 to $1,000$, and a Gaussian noise $\epsilon \sim N(0,1)$ is introduced to Z_P^0 to create the noisy Z_P^t at timestep t according to a noise scheduler (NS), formulated as:

$$Z_P^t = \sqrt{\bar{\alpha}^t} Z_P^0 + \epsilon \sqrt{1-\bar{\alpha}^t}. \tag{1}$$

The corrupted PET latent representation is then concatenated with MRI (Z_M) and input into a convolutional neural network to predict the added noise ϵ. With the predicted noise $\tilde{\epsilon}$, we can reformulate Eq. (1) to derive the l_1 difference between predicted latent PET $\tilde{Z}_P^0$ and actual latent PET Z_P^0:

$$\mathcal{L}_W = Z_P^0 - \tilde{Z}_P^0 = \frac{\sqrt{1-\bar{\alpha}^t}}{\sqrt{\bar{\alpha}^t}} |\tilde{\epsilon} - \epsilon|, \tag{2}$$

where $\bar{\alpha}^t := \prod_{s=1}^{t} \alpha^s$ is a time-dependent hyperparameter and $|\cdot|$ is the l_1 loss. Since the value of $\frac{\sqrt{1-\bar{\alpha}^t}}{\sqrt{\bar{\alpha}^t}}$ increases monotonically with t, this loss $\mathcal{L}_W$ provides an adaptive weight for the noise loss $\tilde{\epsilon} - \epsilon$, leading to stronger penalty at larger timesteps. This dynamic weighting improves training stability and image fidelity since noise removal becomes increasingly challenging at later timesteps. The loss for CWD contains the standard noise loss $\mathcal{L}_N$ and weighted noise loss $\mathcal{L}_W$:

$$\mathcal{L}_1 = \mathcal{L}_N + \mathcal{L}_W = |\tilde{\epsilon} - \epsilon| + \frac{\sqrt{1-\bar{\alpha}^t}}{\sqrt{\bar{\alpha}^t}} |\tilde{\epsilon} - \epsilon|. \tag{3}$$

During inference, a random noise serves as CWD's input at the largest timestep $Z_P^{t=1000}$ and is fed into the CWD model together with the encoded MRI to predict the noise at the previous step. The predicted noise $\tilde{\epsilon}$ is then removed by inverse NS to estimate the denoised PET latent representation Z_P^{t-1}. This process is repeated $1,000$ times, gradually removing noise from PET representation. We obtain $\tilde{Z}_P^0$ at the final timestep, which is the estimated PET latent features and is then decoded by the pre-trained decoder to generate a synthesized PET.

(3) ***Distribution Alignment with NF***. We aim for synthesized PET images to be both realistic and consistent with true PET distributions, which is an objective that prior related works have not explicitly addressed. To achieve this, we introduce a separately trained normalizing flow (NF) module [20] that learns a bijective mapping from real PET images to a Gaussian latent space. During training, synthesized PET images are passed through the NF, and their log-likelihoods are computed based on how closely their transformed representations align with the Gaussian prior. Unrealistic or distributionally inconsistent outputs are assigned lower likelihoods and penalized accordingly. In this way, the NF module serves as a distributional constraint, guiding the model to generate multi-tracer PET images with improved fidelity and consistency. Technically, we first dequantize the synthesized PET and employ a sequence of eight coupling layers to construct a bijective transformation function. These coupling layers split the input into two partitions, where one remains unchanged while the other undergoes an affine transformation parameterized by neural networks

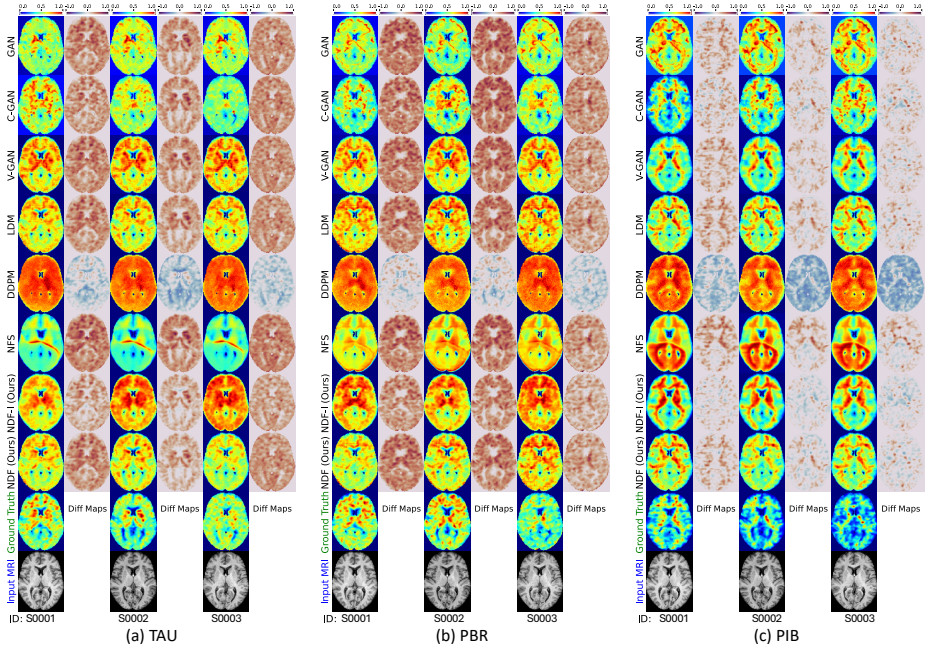

Fig. 2. Visualization of test PET images of 3 tracers – TAU (a), PBR (b), PIB (c) – synthesized by 8 methods, and difference (Diff) maps. The ground-truth PET images and input T1 MRI are displayed at the bottom with the corresponding subject ID.

that generate shift and scale factors. During pretraining, the NF model is optimized to maximize the likelihood that the mapped output follows a normal distribution. Once pretrained, NF is frozen and serves as a guiding constraint in NDF, where the CWD model is trained to maximize this likelihood, ensuring that the synthesized PET images align with the target tracer distribution.

Model Extension. We further extend NDF as NDF-I, which directly takes 3D MRI and PET scans as inputs instead of their latent representations. Different from the weighed loss in Eq. (2), we create an image-level constraint, formulated as $\mathcal{L}_I = |X_P^0 - \tilde{X}_P^0|$, where X_P^0 and $\tilde{X}_P^0$ denote the original PET image and the generated one at each step. This modification enables the model to operate in the original image space rather than a compressed latent space, potentially preserving more fine-grained details in synthesized PET images. NDF-I is optimized by minimizing this objective function: $\mathcal{L}_2 = \mathcal{L}_N + \mathcal{L}_I$.

Implementation Details. We design a four-stage training strategy for NDF: (1) autoencoder pertaining (200 epochs) with PET images from all tracers to learn a compact latent representation, (2) CWD model training (100 epochs) with the frozen autoencoder to establish initial weights, (3) NF model pretraining (200 epochs) using all training PET data to learn a distribution mapping for

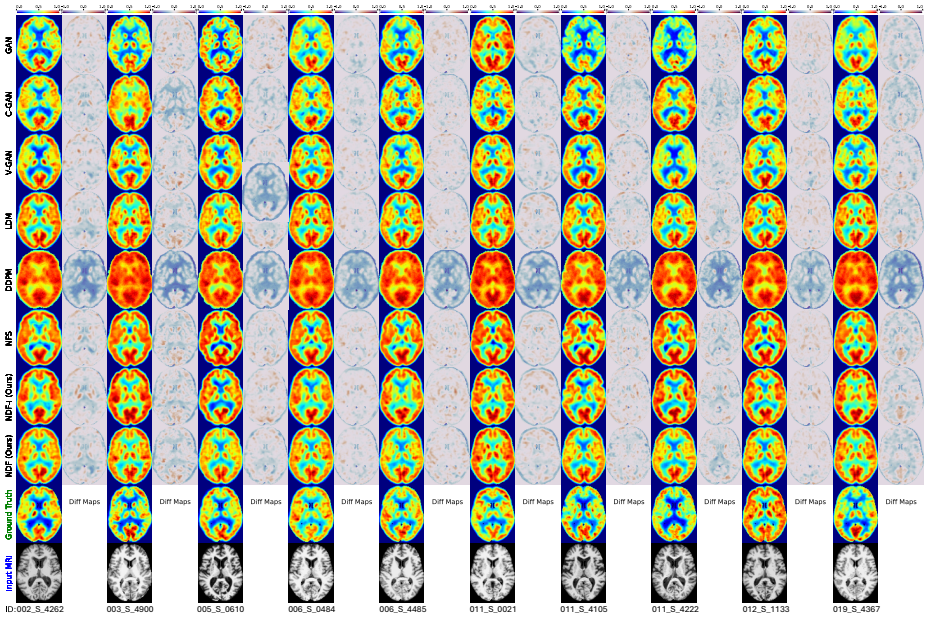

Fig. 3. Test images of FDG-PET synthesized by 8 methods and difference (Diff) maps.

PET images, and (4) CWD model retraining (100 epochs) with the pretrained NF model as a distribution constraint and the autoencoder weights for feature extraction. NDF-I follows only the last two stages. We use 90% of the data for training and 10% for testing. Subjects with complete data (T1 MRI, PBR-, PIB-, and TAU-PET in the private dataset; T1 MRI and FDG-PET in ADNI) are included in the test set. FDG-PET data, three times larger than other tracers, has its loss down-weighted by ten to avoid dominating training and ensure balanced optimization across tracers. NDF and NDF-I are implemented in PyTorch and trained on a cluster with 4 NVIDIA H100 GPUs (each with 80 GB of memory). The diffusion process consists of 1,000 iterative steps during inference, yet NDF achieves an efficient inference speed of approximately 0.5 seconds per image.

3 Experiment

Experimental Settings. Four evaluation metrics are employed: peak signal-to-noise ratio (PSNR), structural similarity index (SSIM), mean absolute error (MAE), and normalized mutual information (NMI). Before evaluation, we restore the synthesized images to their original dimensions through padding. Six methods are used for comparison. (1) GAN [14]: A generative adversarial network trained with adversarial and l_1 losses, producing a four-channel output for multi-tracer PET synthesis. (2) C-GAN [16]: A conditional GAN that concatenates tracer labels with the input and employs cycle-consisitency loss alongside adversarial and l_1 losses. (3) V-GAN [21]: A variational autoencoder-GAN

Table 1. Quantitative results achieved by NDF, NDF-I and six competing methods for TAU-, PBR-, PIB-, and FDG-PET generation, with best results shown in bold.

Method	Synthesized TAU-PET Image				Synthesized PBR-PET Image			
	PSNR↑	SSIM↑	MAE↓	NMI↑	PSNR↑	SSIM↑	MAE↓	NMI↑
GAN	27.932±2.454	0.851±0.021	0.017±0.006	0.579±0.008	28.339±2.519	0.831±0.018	0.016±0.006	0.585±0.008
C-GAN	27.038±2.763	0.831±0.037	0.019±0.007	0.594±0.011	28.065±2.687	0.833±0.022	0.017±0.007	0.594±0.007
V-GAN	28.451±2.591	0.850±0.032	0.016±0.006	0.603±0.016	29.701±1.953	0.841±0.019	0.013±0.004	0.594±0.013
LDM	28.887±2.756	0.861±0.027	0.016±0.006	0.602±0.014	29.848±1.608	0.855±0.014	**0.013**±0.003	0.594±0.007
DDPM	14.082±1.854	0.735±0.024	0.100±0.020	0.595±0.011	15.544±0.841	0.749±0.012	0.083±0.008	0.599±0.013
NFS	28.425±4.210	0.849±0.045	0.017±0.011	0.592±0.010	29.010±1.880	0.842±0.018	0.014±0.004	0.581±0.004
NDF-I	25.259±2.729	0.853±0.025	0.026±0.009	**0.614**±0.200	27.747±1.607	**0.858**±0.012	0.018±0.004	**0.610**±0.013
NDF	**29.374**±2.349	**0.866**±0.022	**0.014**±0.005	0.603±0.016	**29.870**±1.549	0.855±0.016	**0.013**±0.003	0.592±0.009

Method	Synthesized PIB-PET Image				Synthesized FDG-PET Image			
	PSNR↑	SSIM↑	MAE↓	NMI↑	PSNR↑	SSIM↑	MAE↓	NMI↑
GAN	25.357±0.777	0.816±0.013	0.021±0.002	0.630±0.017	26.075±1.168	0.887±0.021	0.022±0.003	0.804±0.040
C-GAN	24.951±1.373	0.815±0.018	0.022±0.004	0.633±0.020	26.174±1.428	0.885±0.021	0.022±0.004	0.808±0.033
V-GAN	**26.713**±1.198	0.833±0.017	**0.018**±0.003	0.664±0.025	25.703±0.848	0.870±0.019	0.022±0.002	0.777±0.025
LDM	25.637±1.625	0.830±0.016	0.021±0.004	0.651±0.021	26.533±1.016	0.875±0.019	**0.020**±0.003	0.806±0.029
DDPM	16.699±1.279	0.743±0.018	0.071±0.012	0.640±0.022	18.199±1.029	0.795±0.028	0.061±0.008	0.799±0.033
NFS	24.522±1.304	0.805±0.017	0.024±0.004	0.614±0.015	**26.706**±1.065	0.877±0.023	**0.020**±0.003	0.809±0.032
NDF-I	25.452±1.749	**0.847**±0.020	0.022±0.006	**0.673**±0.030	26.572±1.106	**0.894**±0.022	**0.020**±0.003	**0.813**±0.035
NDF	26.038±0.832	0.828±0.013	0.020±0.003	0.645±0.018	26.516±1.180	0.875±0.020	0.021±0.003	0.808±0.033

hybrid, where a VAE serves as the generator, and a discriminator is used for adversarial training. (4) DDPM [22]: A denoising diffusion probabilistic model (DDPM) that learns a reverse denoising process within a Markov chain, conditioned by tracer labels and MRI inputs. (5) LDM [19, 23]: A latent diffusion model that encodes images into a compact latent space ($10 \times 12 \times 10$) before applying diffusion-based modality translation. The pretrained encoder and decoder are the same as those used in our NDF. (6) NFS [20]: A normalizing flow-based model trained end-to-end, leveraging normalizing flows to guide PET synthesis.

Qualitative Results. Figure 2 and Figure 3 present axial slices of PET images generated by NDF, NDF-I, and six comparison methods on the two datasets, alongside the ground-truth PET and input T1 MRI displayed at the bottom. The difference (Diff) maps between each synthesized PET and its ground truth are displayed beside each image. As can be observed from Figs. 2-3, the proposed NDF and NDF-I generally produce higher-quality PET images compared to the competing methods. While the GAN model generates visually realistic images, the results generally lack variability, resulting in highly similar synthesized PET images across different subjects for the same tracer. This issue is also observed in the results of V-GAN and NFD, suggesting that those methods struggle to fully capture subject-specific characteristics. C-GAN and LDM produce generally good results. However, their difference maps indicate less precise synthesis compared to our NDF and NDF-I. Meanwhile, DDPM performs the worst, exhibiting the most noticeable discrepancies and the least reliable synthesis.

Quantitative Results. Table 1 reports the quantitative results of eight methods for multi-tracer PET synthesis. This table shows that our method consistently outperforms the competing methods, achieving the highest PSNR and lowest MAE across most tracers. In TAU-PET synthesis, NDF improves PSNR by 1.7% and reduces MAE by 12.5% compared to the second-best method LDM,

demonstrating superior synthesizing accuracy. For PBR- and PIB-PET synthesis, NDF-I attains the highest SSIM and NMI, indicating enhanced structural fidelity. In FDG-PET synthesis, NDF-I achieves the highest SSIM and NMI, while NDF maintains strong PSNR performance, showing the robustness of both approaches. Additionally, GAN-based approaches (GAN, C-GAN, V-GAN) perform worse than diffusion-based models (LDM, DDPM, NDF-I, and NDF), while DDPM exhibits the weakest performance. NFS achieves competitive results but is still outperformed by NDF and NDF-I, demonstrating the advantage of normalizing flow-guided diffusion models in PET synthesis.

Ablation Study. To assess the impact of key components in our framework, we conduct an ablation study on both NDF and NDF-I, evaluating their performance under three different settings: (1) **w/oTC** having the tracer label input removed, making the model rely solely on MRI input; (2) **w/oWL** having the dynamic weighting noise loss omitted; and (3) **w/oNF** having the normalizing flow guiding removed during training. For a fair comparison, these variants share the same architecture as NDF. The quantitative results are reported in Table 2.

Results in Table 2 suggest that removing tracer labels leads to a notable drop in PSNR and a rise in MAE, especially in TAU- and PBR-PET synthesis, highlighting its importance for tracer-specific PET synthesis. Excluding WL significantly degrades performance, particularly in NDF-I, confirming its role in stabilizing training. Removing NF generally lowers synthesis quality, especially in TAU and PBR for NDF-I, showing its importance in output regularization. The greater improvement observed in NDF-I compared to NDF when incorporating NF is likely due to the distribution regulation being applied at the image level. Since NDF-I directly generates images, whereas NDF needs to decode synthesized latent PET features into images, the normalization flow may have a more direct influence on preserving distributional consistency and enhancing image fidelity in NDF-I. Overall, the NDF and NDF-I models achieve the best results, with NDF performing better in TAU-, PBR-, and PIB-PET synthesis and NDF-I performing better in FDG-PET generation.

Influence of NF Quality. To evaluate the dependence of NDF on the quality of the pretrained NF model, we conducted experiments using NFs trained for 10, 50, 100, 150, and 200 epochs. Across all tracers, the synthesis performance remains stable, with only slight reductions in variance as the number of NF training epochs increased. These results indicate that NDF is robust to variations in NF quality. Briefly trained NFs are sufficient to guide effective PET synthesis. Nonetheless, longe r-trained NFs yield marginally more consistent outputs.

4 Conclusion and Future Work

In this work, we propose NDF, a novel distribution-guided diffusion framework for multi-tracer PET synthesis from a single MRI input. By leveraging a class-

Table 2. Quantitative results achieved by NDF and NDF-I, as well as their degraded variants for multi-tracer PET generation, with best results shown in bold.

Tracer & Method		NDF		NDF-I	
		PSNR↑	MAE↓	PSNR↑	MAE↓
TAU	w/oTC	25.6475±2.1807	0.0227±0.0071	22.1891±3.0992	0.0384±0.0154
	w/oWL	27.7529±2.3563	0.0181±0.0062	18.8648±2.0542	0.0564±0.0142
	w/oNF	28.8266±2.5574	0.0156±0.0055	23.5384±3.0300	0.0325±0.0119
	Ours	**29.3735**±2.3491	**0.0145**±0.0048	**25.2588**±2.7285	**0.0258**±0.0093
PBR	w/oTC	25.9791±1.5066	0.0214±0.0045	22.6238±2.6865	0.0354±0.0107
	w/oWL	29.7926±1.5012	0.0129±0.0024	19.6417±1.4108	0.0506±0.0089
	w/oNF	**29.9082**±1.8321	**0.0128**±0.0033	24.8161±1.7490	0.0266±0.0063
	Ours	29.8696±1.5485	**0.0128**±0.0025	**27.7473**±1.6069	**0.0178**±0.0042
PIB	w/oTC	25.5372±1.6282	0.0211±0.0043	23.2225±1.7521	0.0298±0.0076
	w/oWL	25.8981±1.1692	0.0199±0.0025	18.6655±1.3171	0.0534±0.0101
	w/oNF	25.7697±1.3491	0.0202±0.0028	25.3973±2.2005	0.0233±0.0080
	Ours	**26.0384**±0.8315	**0.0198**±0.0025	**25.4519**±1.7489	**0.0224**±0.0055
FDG	w/oTC	26.5687±1.0881	**0.0203**±0.0029	20.7926±1.2528	0.0413±0.0070
	w/oWL	**26.5726**±1.0236	**0.0203**±0.0027	20.0019±1.1389	0.0456±0.0063
	w/oNF	26.2243±1.0349	0.0213±0.0029	25.8093±1.8916	0.0237±0.0055
	Ours	26.5156±1.1800	0.0206±0.0032	**26.5721**±1.1058	**0.0204**±0.0029

conditioned latent diffusion model and a pretrained normalizing flow model, NDF effectively captures tracer-specific PET features while ensuring anatomical consistency. Extensive experiments demonstrate that NDF consistently outperforms state-of-the-art methods, achieving superior performance across multiple PET tracers. For future work, we could explore synthesizing between PET tracers to enhance the model by incorporating molecular information.

Acknowledgements. This work was supported in part by NIH grants (Nos. AG073297, AG082938, EB035160, and NS134841).

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

References

1. DeKosky, S.T., Marek, K.: Looking backward to move forward: early detection of neurodegenerative disorders. Science **302**(5646) (2003) 830–834
2. Veitch, D.P., Weiner, M.W., Aisen, P.S., Beckett, L.A., Cairns, N.J., Green, R.C., Harvey, D., Jack Jr, C.R., Jagust, W., Morris, J.C., et al.: Understanding disease progression and improving Alzheimer's disease clinical trials: Recent highlights from the Alzheimer's Disease Neuroimaging Initiative. Alzheimer's & Dementia **15**(1) (2019) 106–152
3. Lian, C., Liu, M., Wang, L., Shen, D.: Multi-task weakly-supervised attention network for dementia status estimation with structural MRI. IEEE Transactions on Neural Networks and Learning Systems **33**(8) (2021) 4056–4068
4. Guan, H., Liu, M.: Domain adaptation for medical image analysis: A survey. IEEE Transactions on Biomedical Engineering **69**(3) (2021) 1173–1185
5. Zhang, L., Wu, J., Wang, L., Wang, L., Steffens, D.C., Qiu, S., Potter, G.G., Liu, M.: Brain anatomy prior modeling to forecast clinical progression of cognitive impairment with structural mri. Pattern Recognition **165** (2025) 111603

6. Xia, C.F., Arteaga, J., Chen, G., Gangadharmath, U., Gomez, L.F., Kasi, D., Lam, C., Liang, Q., Liu, C., Mocharla, V.P., et al.: [18F] T807, a novel tau positron emission tomography imaging agent for Alzheimer's disease. Alzheimer's & Dementia **9**(6) (2013) 666–676
7. Colasanti, A., Guo, Q., Muhlert, N., Giannetti, P., Onega, M., Newbould, R.D., Ciccarelli, O., Rison, S., Thomas, C., Nicholas, R., et al.: In vivo assessment of brain white matter inflammation in multiple sclerosis with 18F-PBR111 PET. Journal of Nuclear Medicine **55**(7) (2014) 1112–1118
8. Antoni, G., Lubberink, M., Estrada, S., Axelsson, J., Carlson, K., Lindsjö, L., Kero, T., Lrangström, B., Granstam, S.O., Rosengren, S., et al.: In vivo visualization of amyloid deposits in the heart with 11C-PIB and PET. Journal of Nuclear Medicine **54**(2) (2013) 213–220
9. Chételat, G., Arbizu, J., Barthel, H., Garibotto, V., Law, I., Morbelli, S., van de Giessen, E., Agosta, F., Barkhof, F., Brooks, D.J., et al.: Amyloid-PET and 18F-FDG-PET in the diagnostic investigation of Alzheimer's disease and other dementias. The Lancet Neurology **19**(11) (2020) 951–962
10. Hu, S., Lei, B., Wang, S., Wang, Y., Feng, Z., Shen, Y.: Bidirectional mapping generative adversarial networks for brain MR to PET synthesis. IEEE Transactions on Medical Imaging **41**(1) (2021) 145–157
11. Hussein, R., Zhao, M.Y., Shin, D., Guo, J., Chen, K.T., Armindo, R.D., Davidzon, G., Moseley, M., Zaharchuk, G.: Multi-task deep learning for cerebrovascular disease classification and MRI-to-PET translation. In: 2022 26th International Conference on Pattern Recognition (ICPR), IEEE (2022) 4306–4312
12. Zhang, L., Xiao, Z., Zhou, C., Yuan, J., He, Q., Yang, Y., Liu, X., Liang, D., Zheng, H., Fan, W., et al.: Spatial adaptive and transformer fusion network (STFNet) for low-count PET blind denoising with MRI. Medical Physics **49**(1) (2022) 343–356
13. Yu, M., Wu, M., Yue, L., Bozoki, A., Liu, M.: Functional imaging constrained diffusion for brain pet synthesis from structural mri. arXiv preprint arXiv:2405.02504 (2024)
14. Pan, Y., Liu, M., Xia, Y., Shen, D.: Disease-image-specific learning for diagnosis-oriented neuroimage synthesis with incomplete multi-modality data. IEEE Transactions on Pattern Analysis and Machine Intelligence **44**(10) (2021) 6839–6853
15. Sharma, A., Hamarneh, G.: Missing MRI pulse sequence synthesis using multi-modal generative adversarial network. IEEE Transactions on Medical Imaging **39**(4) (2019) 1170–1183
16. Zhou, B., Wang, R., Chen, M.K., Mecca, A.P., O'Dell, R.S., Van Dyck, C.H., Carson, R.E., Duncan, J.S., Liu, C.: Synthesizing multi-tracer PET images for alzheimer's disease patients using a 3D unified anatomy-aware cyclic adversarial network. In: Medical Image Computing and Computer Assisted Intervention–MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI 24, Springer (2021) 34–43
17. Walton, S.R., Kerr, Z.Y., Mannix, R., Brett, B.L., Chandran, A., DeFreese, J.D., McCrea, M.A., Guskiewicz, K.M., Meehan, W.P., Echemendia, R.J.: Subjective concerns regarding the effects of sport-related concussion on long-term brain health among former NFL players: An NFL-LONG study. Sports Medicine (2022) 1–15
18. Jack Jr., C.R., Bernstein, M.A., Fox, N.C., Thompson, P., Alexander, G., Harvey, D., Borowski, B., Britson, P.J., L. Whitwell, J., Ward, C., Dale, A.M., et al.: The Alzheimer's disease neuroimaging initiative (ADNI): MRI methods. Journal of Magnetic Resonance Imaging: An Official Journal of the International Society for Magnetic Resonance in Medicine **27**(4) (2008) 685–691

19. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition. (2022) 10684–10695
20. Beizaee, F., Desrosiers, C., Lodygensky, G.A., Dolz, J.: Harmonizing Flows: Unsupervised MR harmonization based on normalizing flows. In: International Conference on Information Processing in Medical Imaging, Springer (2023) 347–359
21. Larsen, A.B.L., S**o**nderby, S.K., Larochelle, H., Winther, O.: Autoencoding beyond pixels using a learned similarity metric. In: International Conference on Machine Learning, PMLR (2016) 1558–1566
22. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. Advances in Neural Information Processing Systems **33** (2020) 6840–6851
23. Wu, M., Zhang, L., Yap, P.T., Lin, W., Zhu, H., Liu, M.: Structural MRI harmonization via disentangled latent energy-based style translation. In: International Workshop on Machine Learning in Medical Imaging, Springer (2023) 1–11