

# 000 001 002 003 004 005 006 007 RELA-DIFFUSION: RELATIVISTIC ADVERSARIAL DIF- 008 FUSION FOR MULTI-TRACER PET SYNTHESIS FROM 009 MULTI-SEQUENCE MRI 010 011

012 **Anonymous authors**  
013 Paper under double-blind review  
014  
015  
016  
017  
018  
019  
020  
021  
022  
023  
024  
025  
026  
027  
028  
029  
030  
031  
032  
033

## ABSTRACT

034 Multi-tracer positron emission tomography (PET) provides critical insights into  
035 diverse neuropathological processes such as tau accumulation, neuroinflammation,  
036 and  $\beta$ -amyloid deposition in the brain, making it indispensable for comprehensive  
037 neurological assessment. However, routine acquisition of multi-tracer PET is limited by high costs, radiation exposure, and restricted tracer availability.  
038 Recent efforts have explored deep learning approaches for synthesizing PET images  
039 from structural MRI. While some methods rely solely on T1-weighted MRI, others  
040 incorporate additional sequences such as T2-FLAIR to improve pathological  
041 sensitivity. However, existing methods often struggle to capture fine-grained  
042 anatomical and pathological details, resulting in artifacts and unrealistic outputs.  
043 To address these limitations, we propose RelA-Diffusion, a Relativistic Adversarial  
044 Diffusion framework for multi-tracer PET synthesis from multi-sequence MRI.  
045 By leveraging both T1-weighted and T2-FLAIR scans as complementary inputs,  
046 RelA-Diffusion captures richer structural information to guide image generation.  
047 To improve synthesis fidelity, we introduce a gradient-penalized relativistic ad-  
048 versarial loss to the intermediate clean predictions of the diffusion model. This  
049 loss compares real and generated images in a relative manner, encouraging the  
050 synthesis of more realistic local structures. Both the relativistic formulation and  
051 the gradient penalty contribute to stabilizing the training, while adversarial feed-  
052 back at each diffusion timestep enables consistent refinement throughout the gen-  
053 eration process. Extensive experiments on two datasets demonstrate that RelA-  
054 Diffusion outperforms existing methods in both visual fidelity and quantitative  
055 metrics, highlighting its potential for accurate synthesis of multi-tracer PET.

## 1 INTRODUCTION

056 Accurate assessment and monitoring of neurodegenerative disorders rely heavily on the ability to  
057 capture diverse and subtle pathological processes within the brain. Brain positron emission tomog-  
058 raphy (PET) is a powerful imaging modality that enables the *in vivo* visualization of various molec-  
059 ular targets, including  $\beta$ -amyloid deposition with  $^{11}\text{C}$ -PIB (PIB), tau pathology with  $^{18}\text{F}$ -AV1451  
060 (TAU) Johnson et al. (2016); Schwarz et al. (2016); Antoni et al. (2013); Xia et al. (2013), and  
061 neuroinflammatory activity with  $^{18}\text{F}$ -PBR111 (PBR) Colasanti et al. (2014). This multi-tracer PET  
062 imaging can provide complementary information, offering a more comprehensive understanding of  
063 disease mechanisms. However, the use of multi-tracer PET imaging in clinical and research settings  
064 is hindered by several practical constraints, including high scanning costs, radiation exposure, and  
065 limited access to specific tracers. These limitations have motivated the development of alternative  
066 approaches that can synthesize PET data without requiring tracer administration. Recent research  
067 has focused on developing deep learning techniques to synthesize PET images from commonly ac-  
068 cessible structural MRI scans, offering a non-invasive and cost-effective option.

069 Earlier methods for MRI-based PET image synthesis primarily employed GAN-based models. Wei  
070 et al. (2020) propose a conditional flexible self-attention GAN model to predict PET parametric  
071 maps from multi-sequence MRI, training the GAN using a Sketcher–Refiner approach. Hu et al.  
072 (2021) introduce a bidirectional GAN that learns both PET generation and inversion back to la-  
073 tent representations, aiming to effectively capture semantic content within the latent space. Zhang

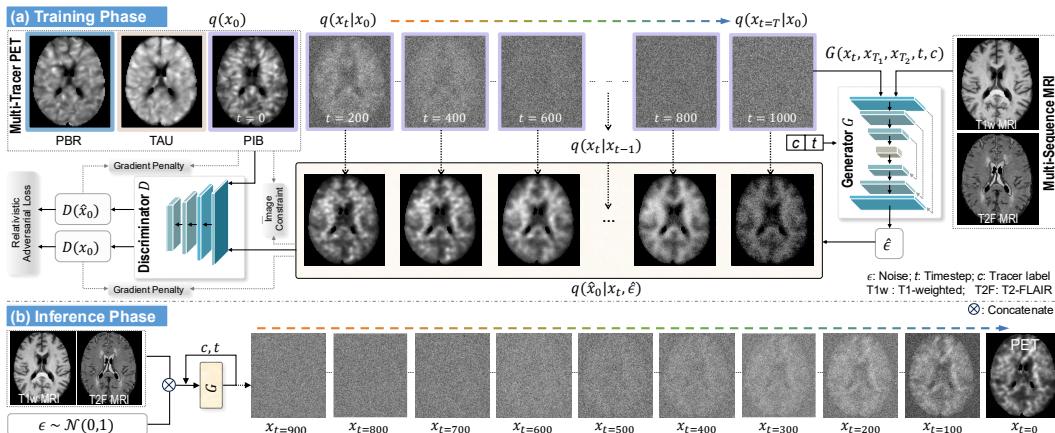


Figure 1: Illustration of the proposed relativistic adversarial diffusion (RelA-Diffusion) framework that synthesizes multi-tracer PET images from multi-sequence MRI input such as T1-weighted (T1w) and T2-FLAIR (T2F) MRIs.

et al. (2022) develop BPGAN, which utilizes a 3D multi-convolution U-Net generator with gradient profile and structural similarity index measure constraints to better preserve structural details in synthesized PET images. Several studies have applied CycleGAN-based frameworks for paired MRI-to-PET synthesis. For instance, Pan et al. (2021) introduce a method that translates between T1-weighted MRI and fluorodeoxyglucose-PET while integrating disease-specific feature disentanglement and diagnosis-guided supervision to enhance clinical relevance. Zhou et al. (2021) explore translation between multiple PET tracers, conditioning the synthesis process on anatomical context derived from structural MRI. Despite their effectiveness, GAN-based approaches frequently encounter optimization challenges, such as mode collapse, which can significantly limit model performance Goodfellow (2016).

More recent strategies have shifted toward diffusion models, which offer significant advantages in generating high-quality medical images. For example, a Joint Diffusion Attention Model (JDAM) (Xie et al., 2024) is designed to generate synthetic PET images from high-field and ultra-high-field MRI by learning the joint probability distribution between MRI input and the noisy PET output. Another study (Yu et al., 2024) presents a functional imaging constrained diffusion (FICD) method, incorporating a voxel-wise alignment loss to improve the quantitative fidelity of PET predictions by enforcing agreement with the ground truth. Additionally, a diffusion model MTGD (Zhong et al., 2025) is designed to fuse multi-sequence MRI to facilitate multi-tracer PET synthesis, showing improved performance in combining complementary MRI contrasts.

To leverage the strengths of both paradigms, several approaches have integrated GAN objectives into diffusion frameworks. The study by (Xiao et al., 2021) models each denoising step of a diffusion process using a multi-modal conditional GAN, allowing for significantly faster sampling while maintaining high sample quality and diversity, and improving mode coverage and training stability compared to traditional GANs. Another study (Wang et al., 2023) introduces Diffusion-GAN, a GAN framework that injects instance noise generated by a forward diffusion chain into the discriminator’s input, thereby addressing GAN training instability and mode collapse by providing more stable and diverse learning signals. While these hybrid frameworks mitigate several limitations of standalone GANs and diffusion models, they often introduce complex training dynamics due to the competing nature of adversarial and denoising objectives. Consequently, such models may struggle to balance global structural coherence with fine-grained detail, leading to suboptimal synthesis fidelity in some cases.

To address these critical limitations, we propose RelA-Diffusion, a novel relativistic adversarial diffusion framework specifically developed to synthesize realistic multi-tracer PET images from multi-sequence MRI scans, including T1-weighted (T1w) and T2-FLAIR (T2F) MRIs. As illustrated in Figure 1, RelA-Diffusion employs a conditional denoising diffusion probabilistic model that iteratively transforms Gaussian noise into realistic PET images, conditioned on multi-sequence MRI inputs. To improve training dynamics and enhance image fidelity, we introduce the gradient-

108 penalized relativistic adversarial loss during training. This loss directly compares real and synthesized  
 109 PET images to encourage sharper anatomical boundaries and more realistic local structures.  
 110 Both the gradient penalty and the relativistic formulation help stabilize adversarial learning and  
 111 regularize model behavior, reducing artifacts and promoting detailed structural representation. Ex-  
 112 tensive evaluations on two multi-tracer PET datasets demonstrate that RelA-Diffusion consistently  
 113 outperforms state-of-the-art (SOTA) methods in terms of visual realism, anatomical accuracy, and  
 114 quantitative performance metrics, establishing its effectiveness and reliability for multi-tracer PET  
 115 image synthesis. The main contributions of this work are summarized as follows:

- 116 • We introduce RelA-Diffusion, a diffusion-based framework tailored for multi-tracer PET  
 117 synthesis from multi-sequence MRI, leveraging both T1-weighted and T2-FLAIR MRI  
 118 inputs for enhanced pathological sensitivity.
- 119 • We propose to apply the gradient-penalized relativistic adversarial feedback to the interme-  
 120 diate clear image estimated during the training of the diffusion model, enabling consistent  
 121 supervision and improved detail preservation, while stabilizing the training.
- 122 • Extensive experiments on two real-world datasets demonstrate that RelA-Diffusion consis-  
 123 tently outperforms existing state-of-the-art methods in both visual fidelity and quantitative  
 124 metrics.

## 126 2 PROPOSED METHOD

127 We propose **RelA-Diffusion**, a conditional diffusion framework for synthesizing high-fidelity multi-  
 128 tracer PET images from multi-sequence MRI. The framework introduces a stabilized adversarial  
 129 supervision mechanism to diffusion training by leveraging a gradient-penalized relativistic discrim-  
 130 inator to provide feedback on intermediate clean predictions  $\hat{x}_0$  at each training step. This feedback  
 131 guides the learning of fine-grained tracer uptake patterns while ensuring stable optimization dynam-  
 132 ics throughout the denoising process.

### 133 2.0.1 OVERALL FRAMEWORK.

134 RelA-Diffusion consists of two main components: a *conditional denoising diffusion model* as the  
 135 generator  $G$ , and a *relativistic discriminator*  $D$  used only during training. As illustrated in Fig-  
 136 ure 1, the generator predicts the noise component  $\hat{\epsilon} = G(x_t, x_{T_1}, x_{T_2}, t, c)$  at each denoising step  
 137  $t$ , where  $x_t$  is the noisy image,  $x_{T_1}$  and  $x_{T_2}$  are the T1w and T2F MRI inputs, and  $c$  denotes the  
 138 tracer label. The clean image estimate  $\hat{x}_0$  is then computed from the predicted noise and  $x_t$  using  
 139 the DDPM formulation (Ho et al., 2020). The discriminator  $D(\hat{x}_0, x_0)$  compares the predicted  $\hat{x}_0$   
 140 with the ground-truth PET image  $x_0$  in a relativistic manner to provide adversarial feedback on the  
 141 relative realism of the intermediate prediction. At inference time, only the trained generator  $G$  is  
 142 used to synthesize PET images through the learned reverse diffusion process, without involving the  
 143 discriminator.

### 144 2.0.2 DIFFUSION PROCESS.

145 Our generative model follows a conditional denoising diffusion framework (Ho et al., 2020), which  
 146 defines a forward noising process as a fixed Markov chain that progressively corrupts a clean PET  
 147 image  $x_0$  into a sequence of noisy variables  $\{x_t\}_{t=1}^T$  over  $T$  timesteps. This is achieved by adding  
 148 Gaussian noise at each step according to a predefined noise schedule Ho et al. (2020). The re-  
 149 verse denoising process is then learned to invert this degradation, reconstructing the original data by  
 150 estimating the noise added at each step.

### 151 2.0.3 FORWARD DIFFUSION PROCESS.

152 Given a real PET image  $x_0 \sim q(x_0)$ , the forward process adds Gaussian noise over  $T$  timesteps,  
 153 creating a series of noisy samples  $\{x_1, \dots, x_T\}$ . This process is defined by a fixed Markov chain:

$$154 q(x_{1:T}|x_0) = \prod_{t=1}^T q(x_t|x_{t-1}), \quad (1)$$

$$155 q(x_t|x_{t-1}) = \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t \mathbf{I}), \quad (2)$$

162 where  $\mathcal{N}$  is Gaussian distribution with mean  $\mu_t$  and variance  $\sigma_t$ ,  $\beta_t$  is a time-dependent hyperparameter that controls the noise level at each timestep, and  $\mathbf{I}$  is the identity matrix indicating isotropic variance. The marginal distribution at timestep  $t$  can be expressed as:

$$166 \quad q(x_t|x_0) = \mathcal{N}(x_t; \sqrt{\bar{\alpha}_t}x_0, (1 - \bar{\alpha}_t)\mathbf{I}), \quad (3)$$

167 where  $\alpha_t = 1 - \beta_t$  and  $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$ .

#### 169 2.0.4 REVERSE DENOISING PROCESS.

171 Starting from a noisy image  $x_t$ , the reverse denoising process aims to progressively estimate and  
172 remove noise to recover the cleaner image  $x_{t-1}$  in the previous timestep, and ultimately aims to  
173 reconstruct the original data  $x_0$ . The trainable component for this process is a generator  $G$  that  
174 predicts the noise  $\epsilon$  added to  $x_0$  during the forward diffusion process to obtain  $x_t$ , denoted as  $\hat{\epsilon}$ . As  
175 illustrated in the top right of Figure 1, the generator is conditioned on multi-sequence MRI inputs  
176 (i.e.,  $x_{T_1}$  for T1w MRI and  $x_{T_2}$  for T2F MRI) via input channel concatenation, and conditioned  
177 on the tracer label  $c$  and timestep  $t$  via cross-attention mechanisms, allowing the model to leverage  
178 anatomical information from MRI and the target tracer type to guide PET image synthesis.

179 In the *training phase*, a randomly selected timestep  $t \in \{1, \dots, T\}$  is sampled uniformly, and a  
180 noisy image  $x_t$  is generated by adding Gaussian noise to the ground-truth PET image  $x_0$  according  
181 to the forward diffusion process:

$$182 \quad x_t = \sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, \quad \epsilon \sim \mathcal{N}(0, \mathbf{I}). \quad (4)$$

183 The generator  $G$  is then trained to predict the added noise  $\epsilon$  from the noisy image  $x_t$ . The training  
184 objective is to minimize the mean squared error (MSE) between the predicted noise  $\hat{\epsilon}$  and the ground-  
185 truth noise  $\epsilon$ , defined as:

$$187 \quad \mathcal{L}_N = \frac{1}{n} \sum_{i=1}^n (\epsilon - \hat{\epsilon}(x_t, t, x_{T_1}, x_{T_2}))^2, \quad (5)$$

189 where  $n$  denotes the number of samples. Given this predicted noise, an estimate of the clean image  
190  $\hat{x}_0$  at timestep  $t$  can be computed as:

$$192 \quad \hat{x}_0(x_t, t) = \frac{x_t - \sqrt{1 - \bar{\alpha}_t}\hat{\epsilon}}{\sqrt{\bar{\alpha}_t}}. \quad (6)$$

194 In the *inference phase*, the generator iteratively applies its noise predictions to progressively denoise  
195 an initial pure Gaussian noise input, starting from  $x_T$  and moving step-by-step towards  $x_0$ . At each  
196 step, the generator predicts the noise  $\hat{\epsilon}$ , which is then utilized to estimate a cleaner version:

$$197 \quad x_{t-1} = \frac{1}{\sqrt{\bar{\alpha}_t}} \left( x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}}\hat{\epsilon} \right) + \sigma_t z, \quad (7)$$

200 where  $\sigma_t$  is the standard deviation of added noise and  $z \sim \mathcal{N}(0, 1)$ . This iterative sampling procedure,  
201 guided by the estimated  $\hat{x}_0$  and the conditional multi-sequence MRI and tracer label inputs,  
202 generating the multi-tracer PET image.

#### 203 2.0.5 SUPERVISION ON INTERMEDIATE CLEAN PREDICTIONS.

205 To enhance output fidelity, we introduce the following two constraints on the intermediate clean  
206 image  $\hat{x}_0$  by comparing them against the true  $x_0$  in RelA-Diffusion.

207 (1) **Image Constraint:** The image-level constraint aims to minimize the  $l_1$  loss between  $\hat{x}_0$  and  $x_0$ ,  
208 formulated as:

$$210 \quad \mathcal{L}_I = \frac{1}{n} \sum_{i=1}^n |x_0 - \hat{x}_0|, \quad (8)$$

212 where  $n$  denotes the number of samples. This design explicitly guides  $\hat{x}_0$  during training to promote  
213 more reliable outputs across the entire reverse process. It is essential since denoising becomes in-  
214 creasingly difficult at larger timesteps, where the input  $x_t$  is heavily corrupted. Supervising  $\hat{x}_0$  at all  
215 timesteps implicitly places stronger regularization on the model’s behavior in high-noise diffusion  
steps, enabling the network to recover meaningful details even from severely degraded inputs. This

effect has been theoretically discussed in prior works Hang et al. (2023); Yu et al. (2024) and provides a foundation for the subsequent adversarial feedback, ensuring that the discriminator receives informative and reliable intermediate predictions at all diffusion steps.

(2) **Gradient-Penalized Relativistic Adversarial Constraint:** To further enhance realism and detail of synthesized PET images, we integrate a *relativistic adversarial loss* Jolicoeur-Martineau (2019); Huang et al. (2024). Unlike traditional GANs where the discriminator judges whether an image is real or fake in an absolute manner, a relativistic discriminator assesses whether a generated image is more realistic than a real image, or vice versa. The discriminator  $D$  is a classifier that receives both intermediate clean predictions  $\hat{x}_0$  produced from noisy PET and the corresponding real PET image  $x_0$ . We perform pairwise comparison to estimate the relative authenticity between real and generated samples. Denote  $\mathbb{P}$  and  $\mathbb{Q}$  as the probability distributions of real and synthesized PET images, respectively. The relativistic adversarial (RA) loss is formulated as:

$$\mathcal{L}_R = \mathbb{E}_{(x_0, \hat{x}_0) \sim (\mathbb{P}, \mathbb{Q})} [f(D(\hat{x}_0) - D(x_0))] + \mathbb{E}_{(x_0, \hat{x}_0) \sim (\mathbb{P}, \mathbb{Q})} [f(D(x_0) - D(\hat{x}_0))], \quad (9)$$

where  $D(x)$  denotes the discriminator’s scalar output for image  $x$ ,  $D(\hat{x})$  denotes the output for the synthesized image  $\hat{x}$ ,  $f(y) = -\log(1 + e^{-y})$  is the activation function Nowozin et al. (2016), and the expectations are computed over the batch of real and synthesized images. The first term in Eq. equation 9 encourages the generator to increase the relative realism of its outputs compared to real PET images. The second term is designed to estimate how much more realistic a real image is compared to a generated one, rather than assigning a realism score to each image in isolation.

Compared to conventional GANs, which can suffer from gradient saturation and mode collapse due to the discriminator’s binary real/fake decisions, the relativistic mechanism has been shown to provide a smoother learning signal that stabilizes adversarial training. It can also improve image quality by preventing the discriminator from becoming overly confident in real-vs-fake discrimination, to maintain meaningful gradients throughout training. Applied jointly with the diffusion loss, it promotes more consistent improvements in the perceptual quality of the diffusion model and guides the model toward producing more realistic outputs.

To facilitate model training, we also incorporate a zero-centered *gradient penalty* on both real and generated data Huang et al. (2024). Specifically, we penalize the squared norm of the discriminator’s gradients with respect to its inputs. This penalty constrains the gradient norm of the discriminator  $D$  on the real PET image  $x_0$  and the generated image  $\hat{x}_0$ , which is formulated as:

$$\mathcal{L}_G = \mathbb{E}_{x_0 \sim \mathbb{P}} [\|\nabla_{x_0} D(x_0)\|^2] + \mathbb{E}_{\hat{x}_0 \sim \mathbb{Q}} [\|\nabla_{\hat{x}_0} D(\hat{x}_0)\|^2], \quad (10)$$

where  $\nabla$  is the gradient operation. This gradient penalty (GP) is designed to promote smoother discriminator behavior on both real and synthesized PET images, improving convergence and ensuring stability throughout training Huang et al. (2024); Nagarajan & Kolter (2017).

#### 2.0.6 OVERALL TRAINING OBJECTIVE.

The RelA-Diffusion framework is trained by minimizing a hybrid objective function that balances the diffusion model’s noise prediction accuracy with the image constraint and the gradient-penalized adversarial guidance for realism. The hybrid loss  $\mathcal{L}$  for RelA-Diffusion is finally formulated as:

$$\mathcal{L} = \mathcal{L}_N + \mathcal{L}_I + \mathcal{L}_R + \mathcal{L}_G. \quad (11)$$

#### 2.0.7 IMPLEMENTATION.

The proposed RelA-Diffusion is implemented using the MONAI framework Cardoso et al. (2022) on PyTorch and trained on a computing cluster with multiple NVIDIA H100 GPUs (each with 80GB memory). We employ the Adam optimizer for both the generator and the discriminator. Empirically, the learning rate for the generator is set to  $5 \times 10^{-5}$  and for the discriminator to  $5 \times 10^{-6}$ . Training is performed with a batch size of 3. The diffusion process is set to  $T = 1000$  timesteps with a linear noise schedule for  $\beta_t$  values, ranging from  $\beta_1 = 0.0005$  to  $\beta_T = 0.0195$ . The model is trained for 100 epochs. The adversarial loss  $\mathcal{L}_R + \mathcal{L}_G$  is weighted by 0.1 throughout all experiments.

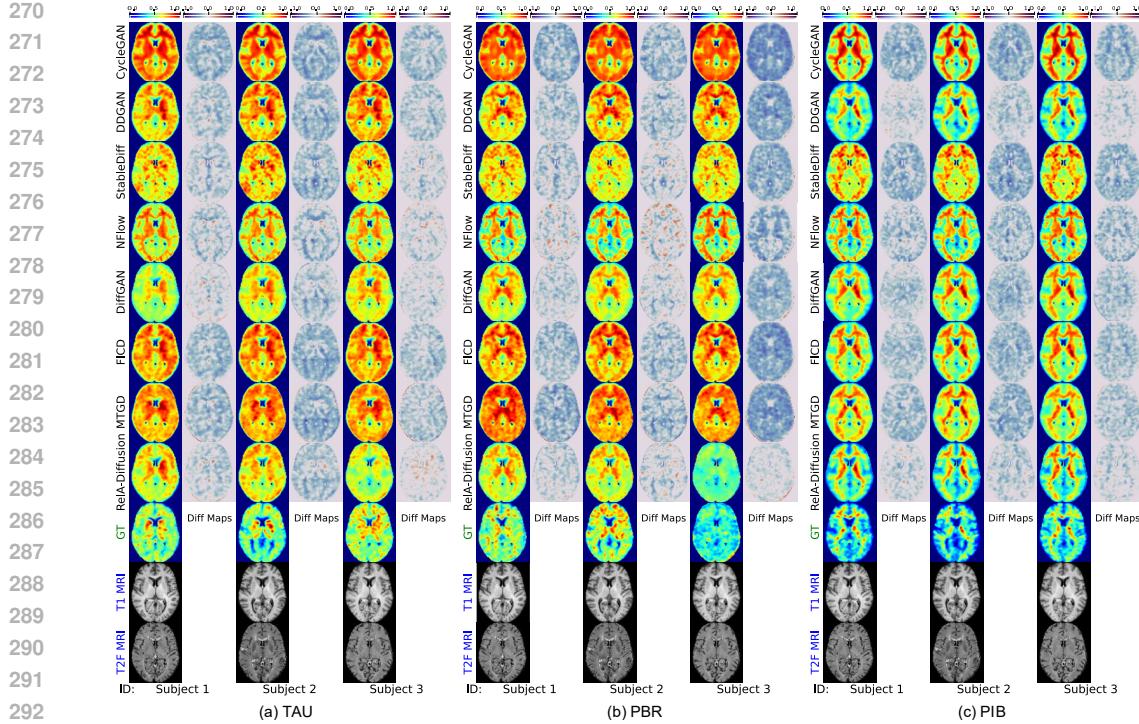


Figure 2: Visualization of test PET images for (a) TAU, (b) PBR, and (c) PIB, synthesized by 8 methods, along with corresponding difference (Diff) maps. Ground-truth (GT) and input T1w/T2F MRIs with subject IDs are shown at the bottom.

### 3 EXPERIMENTS

#### 3.0.1 MATERIALS AND IMAGE PRE-PROCESSING.

Two datasets are utilized in this study: (1) NFL-LONG Walton et al. (2022) comprises 152 participants with a mean age of  $60.15 \pm 6.25$  years. While all subjects have T1-weighted (T1w) and T2-FLAIR (T2F) MRI scans, the availability of PET tracers varies across individuals: 142 subjects have PBR-PET, 105 subjects have PIB-PET, and 124 subjects have TAU-PET using the AV1451 tracer. A subset of 80 subjects possesses all imaging modalities. From this complete cohort, 10 subjects are used for testing across all three tracer types, while the remaining data are used for training. Multi-tracer PET from the same session is rare in public datasets. (2) ADNI Jack Jr. et al. (2008), a public dataset, contains 502 cognitively normal subjects with available T1w MRI, T2F MRI, and TAU-PET scans (with the AV1451 tracer) in the baseline visit. We use all 502 subjects as an external evaluation cohort to assess the generalizability of our model. The mean age of these ADNI subjects is  $69.83 \pm 8.80$ . The subject IDs used are provided in the *Supplementary Materials* for reproducibility. T1w MRI scans are registered to the MNI space. The PET images are rigidly aligned to the corresponding T1w MRI in subject space and then nonlinearly warped to the MNI space using the MRI-derived deformation field. The T2-FLAIR images follow the same pipeline as PET, with an additional bias correction applied prior to rigid alignment. All volumes are resampled to an isotropic resolution of  $1 \times 1 \times 1 \text{ mm}^3$ , and a standard MNI brain mask is applied for skull removal. To exclude irrelevant background, all images are center-cropped to  $160 \times 180 \times 160$ . During training, voxel intensities are scaled to the range  $[-1, 1]$ , while evaluation uses the normalized range  $[0, 1]$  for both visual and quantitative analyses.

#### 3.0.2 EXPERIMENTAL SETTINGS.

We quantitatively evaluate image quality using three metrics: peak signal-to-noise ratio (PSNR), structural similarity index (SSIM), and mean absolute error (MAE). Prior to evaluation, all synthesized PET images are intensity-normalized to  $[0, 1]$  and padded to match their original spatial dimensions.

324  
 325 Table 1: Quantitative results of eight methods for TAU-, PBR- and PIB-PET generation on NFL-  
 326 LONG (best results in bold). The column *sig* indicates statistically significant ( $p < 0.05$ ).

Method	Synthesized TAU-PET Image				Synthesized PBR-PET Image				Synthesized PIB-PET Image			
	PSNR $\uparrow$	SSIM $\uparrow$	MAE $\downarrow$	Sig.	PSNR $\uparrow$	SSIM $\uparrow$	MAE $\downarrow$	Sig.	PSNR $\uparrow$	SSIM $\uparrow$	MAE $\downarrow$	Sig.
CycleGAN (Zhou et al., 2021)	27.506 $\pm$ .598	0.869 $\pm$ .003	<b>0.017</b> $\pm$ .004	*	28.065 $\pm$ .2687	0.833 $\pm$ .0002	0.017 $\pm$ .0007	*	24.954 $\pm$ .2274	0.846 $\pm$ .0017	0.024 $\pm$ .0010	*
DDGAN (Xiao et al., 2021)	25.448 $\pm$ .2373	0.892 $\pm$ .005	0.025 $\pm$ .0009	*	25.201 $\pm$ .3044	0.884 $\pm$ .0017	0.027 $\pm$ .0011	*	<b>26.382</b> $\pm$ .2261	<b>0.870</b> $\pm$ .0016	<b>0.021</b> $\pm$ .0010	*
StableDiff (Rombach et al., 2022)	26.483 $\pm$ .4426	0.881 $\pm$ .001	0.023 $\pm$ .0015	*	<b>28.845</b> $\pm$ .2681	0.892 $\pm$ .0008	<b>0.015</b> $\pm$ .0006	*	25.245 $\pm$ .1735	0.852 $\pm$ .0013	0.023 $\pm$ .0007	*
NFlow (Beizaee et al., 2023)	27.615 $\pm$ .490	0.881 $\pm$ .0020	0.019 $\pm$ .0010	*	28.844 $\pm$ .2030	0.887 $\pm$ .0013	<b>0.015</b> $\pm$ .0004	*	24.522 $\pm$ .1304	0.805 $\pm$ .0017	0.024 $\pm$ .0004	*
DiffGAN (Wang et al., 2023)	24.980 $\pm$ .230	0.884 $\pm$ .0013	0.026 $\pm$ .0009	*	23.112 $\pm$ .2938	0.870 $\pm$ .0012	0.034 $\pm$ .0012	*	23.491 $\pm$ .2440	0.855 $\pm$ .0021	0.031 $\pm$ .0013	*
FICD (Yu et al., 2024)	24.236 $\pm$ .2877	0.883 $\pm$ .0015	0.030 $\pm$ .0012	*	25.594 $\pm$ .2945	0.885 $\pm$ .0017	0.025 $\pm$ .0010	*	24.229 $\pm$ .2360	0.860 $\pm$ .0021	0.029 $\pm$ .0012	*
MTGD (Zhong et al., 2025)	24.998 $\pm$ .1937	0.872 $\pm$ .0012	0.026 $\pm$ .0008	*	23.847 $\pm$ .2379	0.859 $\pm$ .0012	0.031 $\pm$ .0010	*	25.677 $\pm$ .2326	0.846 $\pm$ .0013	0.023 $\pm$ .0010	*
RelA-Diffusion (Ours)	<b>28.314</b> $\pm$ .3392	<b>0.898</b> $\pm$ .0017	<b>0.017</b> $\pm$ .0009	—	<b>29.324</b> $\pm$ .2437	<b>0.898</b> $\pm$ .0017	<b>0.015</b> $\pm$ .0006	—	<b>26.270</b> $\pm$ .1687	0.861 $\pm$ .0012	<b>0.020</b> $\pm$ .0006	—

### 333 3.0.3 COMPETING METHODS.

334  
 335 We compare our method against seven SOTA methods for medical image synthesis. (1) **Cycle-  
 336 GAN** Zhou et al. (2021): A conditional GAN that concatenates tracer labels with the input and  
 337 employs cycle-consistency loss alongside adversarial and  $l_1$  losses. (2) Denoising Diffusion GAN  
 338 (**DDGAN**) Xiao et al. (2021): This method models each denoising step of a diffusion process  
 339 using a multi-modal conditional GAN, allowing for faster sampling while maintaining high sam-  
 340 ple quality and diversity by improving mode coverage and training stability. (3) Stable Diffusion  
 341 (**StableDiff**) Rombach et al. (2022): A latent diffusion model that encodes 3D images into a  
 342 compact latent space before applying diffusion-based modality translation, and then uses a decoder to  
 343 recover the image from the latent space. Its encoder and decoder components are pretrained on  
 344 all training PET images in this study and kept frozen during the diffusion model’s training for the  
 345 modality translation task. (4) Normalizing Flow (**NFlow**) Beizaee et al. (2023): A normalizing flow-  
 346 based model trained end-to-end, leveraging normalizing flows to guide PET synthesis. (5) Diffusion  
 347 GAN (**DiffGAN**) Wang et al. (2023): This GAN framework injects instance noise, generated by a  
 348 forward diffusion chain, into the discriminator’s input. This approach aims to address GAN training  
 349 instability and mode collapse by providing more stable and diverse learning signals. (6) **FICD** Yu  
 350 et al. (2024): A denoising diffusion probabilistic model that learns a reverse denoising process  
 351 within a Markov chain and constrained by the clear intermediate prediction. (7) **MTGD** Zhong  
 352 et al. (2025): A diffusion-based model for multi-tracer PET synthesis from multi-sequence MRI  
 353 (T1w, T2w, T2F), using a Multi-Sequence Attention Encoder and cross-attention fusion. The fused  
 354 representation, combined with timestep and tracer label, conditions the U-Net via its bottleneck.  
 355 To ensure a fair comparison, all methods are implemented on 3D images, with architectures and  
 356 hyperparameters kept as consistent as possible. All diffusion models operated at the image level  
 357 (i.e., FICD, DDDGAN, DiffGAN, and MTGD) implement the image constraint  $\mathcal{L}_I$  supervision on  
 358 intermediate clean predictions as standard DDPM often fails to capture fine structural details in 3D  
 359 settings. All diffusion models (i.e., DDDGAN, StableDiff, DiffGAN, FICD, and MTGD) receive  
 360 input by concatenating T1w and T2F MRI and noisy PET along the channel dimension, and syn-  
 361 thesize tracer-specific PET based on the input label. The tracer label is embedded together with the  
 362 diffusion timestep and injected into the network. CycleGAN and Nflow input concatenated T1w  
 363 and T2F MRI and generate three-channel outputs, each representing a tracer-specific PET image.

### 364 3.0.4 QUALITATIVE RESULTS.

365  
 366 Figure 2 displays axial PET slices synthesized by RelA-Diffusion and seven competing methods  
 367 on the NFL-LONG dataset, alongside the corresponding ground-truth PET images and input T1w  
 368 and T2F MRIs shown at the bottom. For each method, we also visualize the difference (Diff) maps  
 369 between its synthesized PET and the ground truth to highlight voxel-level deviations. As shown  
 370 in Figure 2, our RelA-Diffusion consistently produces PET images with higher visual quality and  
 371 closer resemblance to the ground truth compared to the competing approaches. This improvement  
 372 is particularly evident in the PBR and PIB tracers, where our method effectively captures subject-  
 373 specific patterns of neuroinflammation and amyloid deposition in the brain. In contrast, the compet-  
 374 ing methods tend to generate similar outputs across different subjects, failing to reflect individual  
 375 biological variation. Additionally, hybrid methods that combine diffusion with adversarial learning  
 376 (e.g., DiffGAN and Rela-Diffusion) outperform those relying solely on GANs (e.g., CycleGAN) or  
 377 diffusion models (e.g., FICD, MTGD, StableDiff). This highlights the benefit of integrating adver-  
 378 sarial supervision into the diffusion framework to improve PET image synthesis. Among the hybrid  
 379 methods, our RelA-Diffusion yields the best visual outputs.

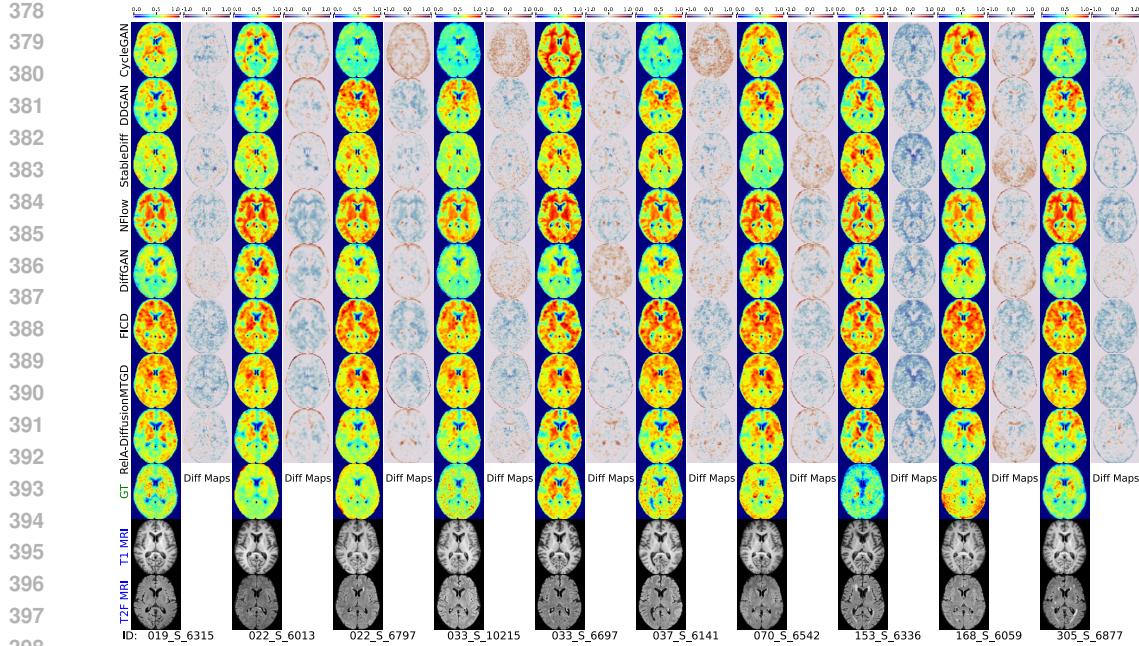


Figure 3: Visualization of test ADNI TAU-PET images synthesized by 8 methods, and difference (Diff) maps. The ground-truth (GT) PET images along with the input T1w and T2F MRIs are displayed at the bottom with the corresponding subject IDs.

Table 2: Quantitative results of eight methods for synthesized ADNI TAU-PET images (best results in bold).

Method	PSNR↑	SSIM↑	MAE↓
CycleGAN (Zhou et al., 2021)	23.766±3.170	0.843±0.034	0.030±0.015
DDGAN (Xiao et al., 2021)	20.563±3.126	0.839±0.028	0.045±0.018
StableDiff (Rombach et al., 2022)	22.179±3.279	0.838±0.035	0.038±0.018
NFlow (Bezaizea et al., 2023)	<b>25.449±6.635</b>	0.851±0.033	<b>0.024±0.010</b>
DiffGAN (Wang et al., 2023)	24.797±2.887	0.862±0.029	0.026±0.010
FICD (Yu et al., 2024)	18.624±2.934	0.831±0.028	0.058±0.021
MTGD (Zhong et al., 2025)	17.632±2.374	0.818±0.023	0.064±0.019
RelA-Diffusion (Ours)	23.781±3.993	<b>0.864±0.034</b>	0.032±0.015

### 3.0.5 QUANTITATIVE RESULTS.

Table 1 reports the quantitative performance of all methods on the NFL-LONG cohort. Across all three PET tracers, RelA-Diffusion consistently achieves the best or highly competitive results. For TAU-PET synthesis, it achieves the highest PSNR (28.314) and SSIM (0.898), and ties with CycleGAN for the lowest MAE (0.017), indicating accurate voxel-wise reconstruction and strong anatomical preservation. For PBR-PET, RelA-Diffusion again achieves the best performance with a PSNR of 29.324, SSIM of 0.898, and MAE of 0.015, outperforming StableDiff and NFlow. While the latter show competitive MAE values, their slightly lower PSNR and SSIM indicate that RelA-Diffusion captures tracer patterns more faithfully. For PIB-PET synthesis, while DDGAN obtains the highest SSIM (0.870), RelA-Diffusion still delivers the best MAE (0.020) and a competitive PSNR (26.270), outperforming other diffusion-based methods (e.g., FICD and MTGD).

## 4 DISCUSSION

#### 4.0.1 GENERALIZATION TO EXTERNAL DATA

430 To assess the generalization capability of our method, we perform external evaluation on  
 431 ADNI Jack Jr. et al. (2008), which includes paired T1w and T2F MRIs along with TAU-PET scans  
 432 acquired using the AV1451 tracer. Without any fine-tuning, the model trained on NFL-LONG is

432

433 Table 3: Quantitative results achieved by RelA-Diffusion, as well as their degraded variants for  
434 multi-tracer PET generation on NFL-LONG, with best results shown in bold.

Method	Synthesized TAU-PET Image			Synthesized PBR-PET Image			Synthesized PIB-PET Image		
	PSNR↑	SSIM↑	MAE↓	PSNR↑	SSIM↑	MAE↓	PSNR↑	SSIM↑	MAE↓
w/oRA	28.156 $\pm$ 4.304	<b>0.908</b> $\pm$ 0.023	0.019 $\pm$ 0.012	29.311 $\pm$ 2.547	<b>0.912</b> $\pm$ 0.018	0.015 $\pm$ 0.006	25.408 $\pm$ 2.182	<b>0.877</b> $\pm$ 0.010	0.022 $\pm$ 0.006
w/oGP	18.401 $\pm$ 2.780	0.780 $\pm$ 0.027	0.061 $\pm$ 0.021	20.947 $\pm$ 2.732	0.786 $\pm$ 0.029	0.045 $\pm$ 0.013	16.855 $\pm$ 0.953	0.764 $\pm$ 0.008	0.066 $\pm$ 0.007
w/oRAGP	24.956 $\pm$ 4.644	0.877 $\pm$ 0.038	0.029 $\pm$ 0.016	29.261 $\pm$ 4.173	0.900 $\pm$ 0.029	0.016 $\pm$ 0.010	23.721 $\pm$ 3.166	0.850 $\pm$ 0.020	0.029 $\pm$ 0.010
w/oT2F	28.109 $\pm$ 4.457	0.898 $\pm$ 0.029	0.019 $\pm$ 0.013	29.490 $\pm$ 2.380	0.905 $\pm$ 0.017	<b>0.014</b> $\pm$ 0.005	25.366 $\pm$ 1.814	0.863 $\pm$ 0.012	0.022 $\pm$ 0.005
w/oT1w	26.022 $\pm$ 4.721	0.887 $\pm$ 0.035	0.026 $\pm$ 0.016	<b>29.894</b> $\pm$ 2.839	0.907 $\pm$ 0.020	<b>0.014</b> $\pm$ 0.006	23.695 $\pm$ 2.748	0.851 $\pm$ 0.014	0.029 $\pm$ 0.008
RelA-Diffusion	<b>28.314</b> $\pm$ 3.392	0.898 $\pm$ 0.017	<b>0.017</b> $\pm$ 0.009	29.324 $\pm$ 2.437	0.898 $\pm$ 0.017	0.015 $\pm$ 0.006	<b>26.270</b> $\pm$ 1.687	0.861 $\pm$ 0.012	<b>0.020</b> $\pm$ 0.006

441

442

443 Table 4: Quantitative results and inference time of RelA-Diffusion and RelA-Diffusion-Latent for  
444 TAU-, PBR-, and PIB-PET generation on NFL-LONG.

Method	Synthesized TAU-PET Image			Synthesized PBR-PET Image			Synthesized PIB-PET Image			Time (s/vol)↓
	PSNR↑	SSIM↑	MAE↓	PSNR↑	SSIM↑	MAE↓	PSNR↑	SSIM↑	MAE↓	
RelA-Diffusion (Ours)	28.314 $\pm$ 3.392	0.898 $\pm$ 0.017	0.017 $\pm$ 0.009	29.324 $\pm$ 2.437	0.898 $\pm$ 0.017	0.015 $\pm$ 0.006	26.270 $\pm$ 1.687	0.861 $\pm$ 0.012	0.020 $\pm$ 0.006	126.778
RelA-Diffusion-Latent (Ours)	27.921 $\pm$ 3.902	0.900 $\pm$ 0.022	0.019 $\pm$ 0.011	28.823 $\pm$ 1.395	0.901 $\pm$ 0.012	0.015 $\pm$ 0.004	24.523 $\pm$ 2.076	0.865 $\pm$ 0.010	0.024 $\pm$ 0.005	0.778

445

446

447 directly applied to 30 ADNI subjects for TAU-PET synthesis. Table 2 and Figure 3 present the  
448 qualitative examples and quantitative metrics for multi-sequence MRI to TAU-PET synthesis on  
449 ADNI.

450

451

452 As shown in Table 2, our method consistently achieves the best performance across all metrics  
453 when compared with the seven SOTA methods, demonstrating superior generalization capability.  
454 Diffusion-only methods (e.g., FICD and MTGD) cannot yield good results compared to hybrid  
455 approaches such as DiffGAN, suggesting limited generalizability to unseen data. The qualitative  
456 results in Figure 3 further support this observation, where RelA-Diffusion produces images that  
457 closely resemble the ground truth, and the corresponding difference (Diff) maps are visibly lighter.  
458 These findings collectively highlight the strong generalizability of our method when applied to external  
459 datasets acquired from different scanners and scanning protocols, suggesting its greater potential  
460 for clinical deployment across diverse imaging environments and patient populations.

461

462

## 463 4.0.2 LATENT RELA-DIFFUSION.

464

465 We further develop a latent-space variant of our framework, termed RelA-Diffusion-Latent, to im-  
466 prove computational efficiency. Instead of directly operating in the voxel space, both the conditional  
467 MRI inputs and the target PET images are first encoded into a latent space ( $10 \times 12 \times 10$ ) using a pre-  
468 trained frozen autoencoder. The diffusion model operates entirely in this latent space and predicts the  
469 noise added to the PET latent. The diffusion processes and the relativistic adversarial supervision re-  
470 main the same as RelA-Diffusion. We apply the image constraint  $L_1$  directly on the predicted clean  
471 latent  $\hat{z}_0$  to enforce accurate reconstruction in the latent domain. In contrast, the gradient-penalized  
472 relativistic adversarial constraint is applied on the *decoded* PET image  $\hat{x}_0$ , obtained by passing  $\hat{z}_0$   
473 through the decoder, allowing the discriminator to assess perceptual realism in the image space. This  
474 hybrid latent–image supervision strategy preserves the computational advantages of latent diffusion  
475 while retaining strong adversarial guidance for high-fidelity PET synthesis. Table 4 compares RelA-  
476 Diffusion and RelA-Diffusion-Latent, showing comparable synthesis performance while the latent  
477 version significantly reduces inference time. Notably, RelA-Diffusion-Latent achieves an inference  
478 speed of 0.778 s/vol, which is even faster than the GAN baseline (CycleGAN: 1.47 s/vol), while  
479 diffusion-only methods such as FICD remain much slower (134.8 s/vol). This demonstrates that the  
latent variant offers an excellent efficiency–fidelity trade-off for practical PET synthesis.

480

481

## 482 4.0.3 ABLATION STUDY.

483

484

485 To evaluate the contributions of key components in our RelA-Diffusion framework, we conduct an  
486 ablation study on the NFL-LONG dataset by comparing it to four ablated variants: (1) **w/oRA**, which  
487 replaces the relativistic adversarial loss with a standard non-relativistic GAN loss; (2) **w/oGP**, which  
488 removes gradient penalties; (3) **w/oRAGP**, which replaces the relativistic adversarial loss with a  
489 standard non-relativistic GAN loss and removes gradient penalties; (4) **w/oT2F**, which excludes T2-

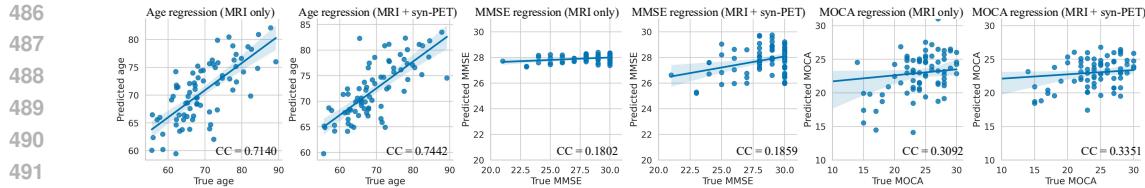


Figure 4: Downstream regression performance for age, MMSE, and MoCA prediction using synthesized PET (syn-PET) generated from ADNI TAU images together with MRI, in comparison with MRI-only inputs.

FLAIR input and uses only T1w MRI as condition; and (5) **w/oT1w**, which uses only T2-FLAIR MRI and discards T1w input. All variants retain the same network architecture as the full model to ensure a fair comparison. The quantitative results are reported in Table 3.

Ablation results in Table 3 demonstrate that removing gradient penalty (GP) leads to the largest performance drop across all tracers, confirming that GP is essential for stabilizing adversarial training and preventing degradation in diffusion-GAN synthesis. Replacing the relativistic adversarial loss (RA) with a standard GAN loss also reduces fidelity, indicating that RA further enhances structural consistency. When both replacing RA with a standard GAN loss and removing GP (w/oRAGP), there is a consistent drop in PSNR and increased MAE across all tracer types, demonstrating that GP and RA are complementary, and their combination is critical for achieving high-quality PET synthesis. Excluding T1w input (w/oT1w) degrades performance, particularly for TAU- and PBR-PET, indicating that the complementary anatomical context captured by T2-FLAIR and T1w MRIs is critical for accurate tracer distribution modeling. Similarly, removing T2-FLAIR (w/oT2F) also causes a performance drop, but the effect is slightly less pronounced than w/oT1w, suggesting that while both modalities contribute valuable information, T1w may offer more specific pathological contrast for PET synthesis. Overall, the full model leveraging both T1w and T2F MRIs with relativistic adversarial supervision achieves the best performance.

#### 4.0.4 DOWNSTREAM TASKS.

To evaluate the clinical utility of the synthesized PET images, we further conduct downstream regression tasks for predicting age, Mini-Mental State Examination (MMSE) and Montreal Cognitive Assessment (MoCA) scores. We compare two input settings: (1) using MRI only, and (2) using MRI combined with synthesized PET (syn-PET) generated from ADNI TAU images. All regression models share the same ResNet-based architecture and training protocol to ensure a fair comparison. The results in Figure 4 demonstrate that incorporating syn-PET consistently improves prediction performance across all three tasks, indicating that the synthesized PET provides complementary functional information beyond structural MRI and is beneficial for clinically relevant cognitive outcome prediction.

## 5 CONCLUSION

This paper presents RelA-Diffusion, a novel relativistic adversarial diffusion framework for synthesizing multi-tracer PET images from multi-sequence MRI. By leveraging T1w and T2-FLAIR inputs, our method captures complementary structural and pathological cues to guide the generation process. The proposed gradient-penalized relativistic adversarial loss enhances anatomical realism and stabilizes training. Extensive experiments demonstrate that RelA-Diffusion achieves superior performance in both visual fidelity and quantitative accuracy compared to existing methods. This highlights its potential to reduce reliance on expensive, high-risk multi-tracer PET imaging in clinical and research settings. In future work, we plan to extend RelA-Diffusion to handle other imaging modalities, such as CT or contrast-enhanced MRI, to support broader clinical applications. We also aim to explore domain adaptation strategies to improve generalizability across diverse scanners.

540 REFERENCES  
541

542 Gunnar Antoni, Mark Lubberink, Sergio Estrada, Jan Axelsson, Kristina Carlson, Lars Lindsjö,  
543 Tanja Kero, Bengt Långström, Sven-Olof Granstam, Sara Rosengren, et al. In vivo visualization  
544 of amyloid deposits in the heart with  $^{11}\text{C}$ -PIB and PET. *Journal of Nuclear Medicine*, 54(2):  
545 213–220, 2013.

546 Farzad Beizaee, Christian Desrosiers, Gregory A Lodygensky, and Jose Dolz. Harmonizing Flows:  
547 Unsupervised MR harmonization based on normalizing flows. In *International Conference on  
548 Information Processing in Medical Imaging*, pp. 347–359. Springer, 2023.

549 M Jorge Cardoso, Wenqi Li, Richard Brown, Nic Ma, Eric Kerfoot, Yiheng Wang, Benjamin Murrey,  
550 Andriy Myronenko, Can Zhao, Dong Yang, et al. MONAI: An open-source framework for deep  
551 learning in healthcare. *arXiv preprint arXiv:2211.02701*, 2022.

552 Alessandro Colasanti, Qi Guo, Nils Muhlert, Paolo Giannetti, Mayca Onega, Rexford D Newbould,  
553 Olga Ciccarelli, Stuart Rison, Charlotte Thomas, Richard Nicholas, et al. In vivo assessment of  
554 brain white matter inflammation in multiple sclerosis with  $^{18}\text{F}$ -PBR111 PET. *Journal of Nuclear  
555 Medicine*, 55(7):1112–1118, 2014.

556 Ian Goodfellow. NIPS 2016 Tutorial: Generative adversarial networks. *arXiv preprint  
557 arXiv:1701.00160*, 2016.

558 Tiankai Hang, Shuyang Gu, Chen Li, Jianmin Bao, Dong Chen, Han Hu, Xin Geng, and Bain-  
559 ing Guo. Efficient diffusion training via min-SNR weighting strategy. In *Proceedings of the  
560 IEEE/CVF International Conference on Computer Vision*, pp. 7441–7451, 2023.

561 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in  
562 Neural Information Processing Systems*, 33:6840–6851, 2020.

563 Shengye Hu, Baiying Lei, Shuqiang Wang, Yong Wang, Zhiguang Feng, and Yanyan Shen. Bi-  
564 directional mapping generative adversarial networks for brain MR to PET synthesis. *IEEE Trans-  
565 actions on Medical Imaging*, 41(1):145–157, 2021.

566 Nick Huang, Aaron Gokaslan, Volodymyr Kuleshov, and James Tompkin. The GAN is dead; Long  
567 live the GAN! A modern gan baseline. *Advances in Neural Information Processing Systems*, 37:  
568 44177–44215, 2024.

569 Clifford R. Jack Jr., Matt A. Bernstein, Nick C. Fox, Paul Thompson, Gene Alexander, Danielle  
570 Harvey, Bret Borowski, Paula J. Britson, Jennifer L. Whitwell, Chadwick Ward, Anders M. Dale,  
571 et al. The Alzheimer’s disease neuroimaging initiative (ADNI): MRI methods. *Journal of Mag-  
572 netic Resonance Imaging: An Official Journal of the International Society for Magnetic Reso-  
573 nance in Medicine*, 27(4):685–691, 2008.

574 Keith A Johnson, Aaron Schultz, Rebecca A Betensky, J Alex Becker, Jorge Sepulcre, Dorene  
575 Rentz, Elizabeth Mormino, Jasmeer Chhatwal, Rebecca Amariglio, Kate Papp, et al. Tau positron  
576 emission tomographic imaging in aging and early Alzheimer disease. *Annals of neurology*, 79(1):  
577 110–119, 2016.

578 A Jolicoeur-Martineau. The relativistic discriminator: A key element missing from standard GAN.  
579 In *International Conference on Learning Representations*, 2019.

580 Vaishnav Nagarajan and J Zico Kolter. Gradient descent GAN optimization is locally stable. *Ad-  
581 vances in neural information processing systems*, 30, 2017.

582 Sebastian Nowozin, Botond Cseke, and Ryota Tomioka. f-GAN: Training generative neural sam-  
583 plers using variational divergence minimization. *Advances in Neural Information Processing  
584 Systems*, 29, 2016.

585 Yongsheng Pan, Mingxia Liu, Yong Xia, and Dinggang Shen. Disease-image-specific learning for  
586 diagnosis-oriented neuroimage synthesis with incomplete multi-modality data. *IEEE Transac-  
587 tions on Pattern Analysis and Machine Intelligence*, 44(10):6839–6853, 2021.

594 Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-  
 595 resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Con-*  
 596 *ference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.

597

598 Adam J Schwarz, Peng Yu, Bradley B Miller, Sergey Shcherbinin, James Dickson, Michael Navit-  
 599 sky, Abhinay D Joshi, Michael D Devous Sr, and Mark S Mintun. Regional profiles of the can-  
 600 didate tau PET ligand 18 F-AV-1451 recapitulate key features of Braak histopathological stages.  
 601 *Brain*, 139(5):1539–1550, 2016.

602

603 Samuel R Walton, Zachary Y Kerr, Rebekah Mannix, Benjamin L Brett, Avinash Chandran,  
 604 Jonathan D DeFreese, Michael A McCrea, Kevin M Guskiewicz, William P Meehan, and Ruben J  
 605 Echemendia. Subjective concerns regarding the effects of sport-related concussion on long-term  
 606 brain health among former NFL players: An NFL-LONG study. *Sports Medicine*, pp. 1–15, 2022.

607

608 Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-  
 609 resolution image synthesis and semantic manipulation with conditional GANs. In *Proceedings of*  
 610 *the IEEE conference on computer vision and pattern recognition*, pp. 8798–8807, 2018.

611

612 Zhendong Wang, Huangjie Zheng, Pengcheng He, Weizhu Chen, and Mingyuan Zhou. Diffusion-  
 613 GAN: Training GANs with diffusion. In *The Eleventh International Conference on Learning*  
 614 *Representations*, 2023.

615

616 Wen Wei, Emilie Poirion, Benedetta Bodini, Matteo Tonietto, Stanley Durrleman, Olivier Colliot,  
 617 Bruno Stankoff, and Nicholas Ayache. Predicting PET-derived myelin content from multise-  
 618 quence MRI for individual longitudinal analysis in multiple sclerosis. *NeuroImage*, 223:117308,  
 619 2020.

620

621 Chun-Fang Xia, Janna Arteaga, Gang Chen, Umesh Gangadharan, Luis F Gomez, Dhanalakshmi  
 622 Kasi, Chung Lam, Qianwa Liang, Changhui Liu, Vani P Mochalra, et al. [18F] T807, a novel tau  
 623 positron emission tomography imaging agent for Alzheimer’s disease. *Alzheimer’s & Dementia*,  
 9(6):666–676, 2013.

624

625 Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with  
 626 denoising diffusion GANs. *arXiv preprint arXiv:2112.07804*, 2021.

627

628 Taofeng Xie, Chentao Cao, Zhuo-xu Cui, Yu Guo, Caiying Wu, Xuemei Wang, Qingneng Li, Zhanli  
 629 Hu, Tao Sun, Ziru Sang, et al. Synthesizing PET images from high-field and ultra-high-field MR  
 630 images using joint diffusion attention model. *Medical Physics*, 51(8):5250–5269, 2024.

631

632 Minhui Yu, Mengqi Wu, Ling Yue, Andrea Bozoki, and Mingxia Liu. Functional imaging con-  
 633 strained diffusion for brain PET synthesis from structural MRI. *arXiv preprint arXiv:2405.02504*,  
 634 2024.

635

636 Jin Zhang, Xiaohai He, Linbo Qing, Feng Gao, and Bin Wang. BPGAN: Brain PET synthesis  
 637 from MRI using generative adversarial network for multi-modal Alzheimer’s disease diagnosis.  
 638 *Computer Methods and Programs in Biomedicine*, 217:106676, 2022.

639

640 Wenjin Zhong, Cong Cong, Ghasem Azemi, Mehnaz Tabassum, Antonio Di Ieva, and Sidong Liu.  
 641 Multi-sequence MRI to multi-tracer PET generation via diffusion model. In *International Sym-*  
 642 *posium on Biomedical Imaging*, pp. 1–4. IEEE, 2025.

643

644 Bo Zhou, Rui Wang, Ming-Kai Chen, Adam P Mecca, Ryan S O’Dell, Christopher H Van Dyck,  
 645 Richard E Carson, James S Duncan, and Chi Liu. Synthesizing multi-tracer PET images for  
 646 Alzheimer’s disease patients using a 3D unified anatomy-aware cyclic adversarial network. In  
 647 *International Conference on Medical Image Computing and Computer Assisted Intervention*, pp.  
 34–43. Springer, 2021.

648  
649

## A APPENDIX

650  
651

## A.1 RELA-DIFFUSION ALGORITHM IMPLEMENTATION

652  
653

## A.1.1 TRAINING PROCEDURE (ALGORITHM 1)

654  
655  
656  
657  
658

The training procedure of RelA-Diffusion is outlined in Algorithm 1. The model is trained on paired multi-sequence MRI ( $x_{T_1}, x_{T_2}$ ) and multi-tracer PET  $x_0$  data. At each iteration, a timestep  $t$  is sampled and Gaussian noise is added to  $x_0$  to obtain a noisy input  $x_t$ . The generator  $G$  takes  $x_t$ , MRI inputs,  $t$ , and tracer label  $c$  to predict the noise  $\hat{\epsilon}$ , from which the clean estimate  $\hat{x}_0$  is reconstructed. The relativistic adversarial loss uses the activation  $f(x) = -\log(1 + e^{-x})$ .

659

**Algorithm 1** Training Procedure for RelA-Diffusion

660

**Input:** Paired data  $\{x_{T_1}, x_{T_2}, x_0, c\}_{i=1}^N$ 

661

**Parameter:** Diffusion steps  $T$ , learning rate  $\eta$ 

662

**Output:** Trained generator  $G$ 

663

1: Initialize generator  $G$ , discriminator  $D$ , and optimizers

664

2: **for** epoch = 1 to  $N_{\text{epochs}}$  **do**

665

3:   **for** each batch  $(x_{T_1}, x_{T_2}, x_0, c)$  **do**

666

4:     Sample timestep  $t \sim \{1, \dots, T\}$  uniformly

667

5:     Sample noise  $\epsilon \sim \mathcal{N}(0, \mathbf{I})$ 

668

6:     Generate noisy PET:  $x_t = \sqrt{\bar{\alpha}_t} x_0 + \sqrt{1 - \bar{\alpha}_t} \epsilon$ 

669

7:     Concatenate  $x_t$ ,  $x_{T_1}$ , and  $x_{T_2}$  along the channel dimension as model input

670

8:     Embed timestep  $t$  and tracer label  $c$ ; inject into  $G$ 

671

9:     Predict noise:  $\hat{\epsilon} = G(x_t, x_{T_1}, x_{T_2}, t, c)$ 

672

10:    Compute  $\hat{x}_0 = \frac{x_t - \sqrt{1 - \bar{\alpha}_t} \hat{\epsilon}}{\sqrt{\bar{\alpha}_t}}$ 

673

11:    **Noise loss:**  $\mathcal{L}_N = \frac{1}{n} \sum_{i=1}^n (\epsilon - \hat{\epsilon})^2$ 

674

12:    **Image loss:**  $\mathcal{L}_I = \frac{1}{n} \sum_{i=1}^n \|x_0 - \hat{x}_0\|$ 

675

13:    **Relativistic adversarial loss:**  $\mathcal{L}_R = f(D(\hat{x}_0) - D(x_0)) + f(D(x_0) - D(\hat{x}_0))$ 

676

14:    **Gradient penalty:**  $\mathcal{L}_G = \|\nabla_{x_0} D(x_0)\|^2 + \|\nabla_{\hat{x}_0} D(\hat{x}_0)\|^2$ 

677

15:    **Total loss:**  $\mathcal{L} = \mathcal{L}_N + \mathcal{L}_I + \mathcal{L}_R + \mathcal{L}_G$ 

678

16:    Update generator  $G$  using  $\nabla \mathcal{L}$ 

679

17:    Update discriminator  $D$  using  $\nabla \mathcal{L}$ 

680

18:   **end for**

681

19: **end for**

682

20: **return** trained generator  $G$ 

683

684

## A.1.2 INFERENCE PROCEDURE (ALGORITHM 2)

685

686  
687  
688  
689  
690

Algorithm 2 describes the inference process for RelA-Diffusion. Given a trained generator  $G$ , multi-sequence MRI inputs  $x_{T_1}$  and  $x_{T_2}$ , and a target tracer label  $c$ , the model synthesizes a PET image via an iterative denoising process. Starting from a pure Gaussian noise image  $x_T \sim \mathcal{N}(0, \mathbf{I})$ , the model runs the reverse diffusion process for  $T$  steps. At each step  $t$ , the generator predicts the noise  $\hat{\epsilon}$ , and uses it to estimate a cleaner image  $x_{t-1}$ . This iterative refinement continues until  $x_0$  is produced.

691

692

693

694

695

696

697

698

699

700

701

702  
 703  
 704  
 705  
 706  
 707  
 708  
 709  
 710  
 711  
 712  
 713  
 714  
 715  
 716  
 717  
 718  
 719  
 720  
 721

---

**Algorithm 2** Inference with Trained ReLA-Diffusion Generator
 

---

**Input:** Multi-sequence MRI inputs  $(x_{T_1}, x_{T_2})$ , tracer label  $c$ , trained generator  $G$

**Parameter:** Diffusion steps  $T$ , noise schedule  $\{\alpha_t, \bar{\alpha}_t, \sigma_t\}_{t=1}^T$

**Output:** Synthesized PET image  $\hat{x}_0$

1: Initialize  $x_T \sim \mathcal{N}(0, \mathbf{I})$  {Start from pure Gaussian noise}  
 2: **for**  $t = T$  to 1 **do**  
 3:   Predict noise:  $\hat{\epsilon} = G(x_t, x_{T_1}, x_{T_2}, t, c)$   
 4:   Sample noise  $z \sim \mathcal{N}(0, \mathbf{I})$  if  $t > 1$ , else  $z = 0$   
 5:   Compute denoised image:

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left( x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \hat{\epsilon} \right) + \sigma_t z$$

6: **end for**  
 7: **return** final image  $\hat{x}_0$

---

731  
 732  
 733  
 734  
 735  
 736  
 737  
 738  
 739  
 740  
 741  
 742  
 743  
 744  
 745  
 746  
 747  
 748  
 749  
 750  
 751  
 752  
 753  
 754  
 755

756

## EXTERNAL TEST SUBJECT IDs

757

758 The following subjects from ADNI were used to test the generalization ability of the proposed  
 759 method in our experiments:

760

761	153_S_6450	033_S_10144	141_S_6872	301_S_6592	033_S_6697	006_S_6770
762	022_S_6797	941_S_6803	067_S_7061	037_S_6216	033_S_10215	305_S_6877
763	153_S_6336	141_S_6240	033_S_6889	941_S_6581	037_S_6141	033_S_10021
764	067_S_6117	301_S_6615	305_S_6498	019_S_6315	070_S_6542	016_S_6773
765	033_S_7114	168_S_6059	003_S_6606	067_S_6957	168_S_6085	022_S_6013

766

767

## A.2 NETWORK ARCHITECTURE.

768

769 (1) The generator is implemented using a U-Net architecture, consisting of 3 downsampling and 3  
 770 upsampling blocks. Each downsampling block contains one convolutional layer for spatial reduction  
 771 and two residual sub-blocks incorporating skip connections. Each residual sub-block consists of two  
 772 convolutional layers, each followed by group normalization and timestep embedding, along with a  
 773 shortcut connection. Positional embeddings for the diffusion timestep  $t$  plus tracer label  $c$  project  
 774 encoded features to match the number of feature channels, enabling effective feature integration.  
 775 The multi-sequence MRI inputs are concatenated along the channel dimension and then fed into  
 776 the U-Net along with the input  $x_t$ . (2) The discriminator is based on a PatchGAN discriminator  
 777 architecture Wang et al. (2018), designed to operate on local image patches rather than the entire  
 778 image. It consists of 3 convolutional layers with decreasing spatial resolution and increasing channel  
 779 depth. Each convolutional layer is followed by LeakyReLU activation (with a negative slope of 0.2)  
 780 and batch normalization, except for the final output layer. The output of the discriminator is a single  
 781 scalar representing the relative realism of the input.

782

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809