Causal PETS: Causality-Informed PET Synthesis from Multi-modal Data

Yujia Li[1,2]   YUJIA.LI@MIRACL.ICT.AC.CN
[1] *Key Laboratory of Intelligent Information Processing of Chinese Academy of Sciences (CAS), Institute of Computing Technology, CAS*
[2] *University of Chinese Academy of Sciences*

Han Li[*3]   [3]*Computer Aided Medical Procedures (CAMP), School of Computation, Information and Technology, Technische Universitaet Muenchen (TUM)*
S Kevin Zhou[1,4,5,6]   SKEVINZHOU@USTC.EDU.CN
[4] *School of Biomedical Engineering, Division of Life Sciences and Medicine, University of Science and Technology of China (USTC), Hefei Anhui, 230026, China*
[5] *Center for Medical Imaging, Robotics, Analytic Computing & Learning (MIRACLE), Suzhou Institute for Advance Research, USTC, Suzhou Jiangsu, 215123, China*
[6] *Key Laboratory of Precision and Intelligent Chemistry, USTC, Hefei Anhui, 230026, China*

Editors: Under Review for MIDL 2025

Abstract

Positron emission tomography (PET) plays a crucial role in diagnosing and monitoring neurological disorders. However, its clinical availability is constrained by high costs, radiation exposure risks, and logistical limitations. In this study, we propose **Causal PETS**, a novel causality-informed multimodal synthesis model for PET image generation. Unlike conventional approaches that rely on a direct transformation from T_1-weighted MRI to PET, our model explicitly captures causal relationships among multimodal data—including MRI, demographic information, and cerebrospinal fluid (CSF) biomarkers—and seamlessly integrates these factors into the PET synthesis process. Through extensive evaluations, we demonstrate that Causal PETS surpasses existing non-causal methods in image clarity and accuracy, particularly in highlighting regions of interest critical for neurological disorders such as Alzheimer's disease (AD). This work underscores the significance of causality in medical image synthesis and highlights the potential of multimodal integration for enhancing clinical decision-making.

Keywords: medical image synthesis, causality, Alzheimer's Disease

1. Introduction

Medical image synthesis offers new solutions to the problem of obtaining certain modality imaging data. For example, Positron Emission Tomography (PET) is pivotal in neuroimaging and is crucial for the diagnosis and monitoring of neurodegenerative diseases such as Alzheimer's Disease (AD) by detecting abnormal molecules (Marcus et al., 2014; Nordberg et al., 2010). However, the widespread use of PET is hindered by several challenges: its reliance on frequent scans for longitudinal studies, the significant expenses of the radiotracers and advanced imaging technology, and the health risks posed by radiation exposure (Nievelstein et al., 2012; Brix et al., 2009). Thus, there is an urgent need to explore alternative approaches for acquiring PET to support diagnostic applications, among which the synthesis of PET from other more available modalities presents a promising solution.

Figure 1: The Causal Graph of PET Synthesis

Recently, deep learning-based medical image synthesis models have shown great potential. (Wang et al., 2021a,b). In the context of PET imaging, most approaches adopt a straightforward architecture for generating target images from source images. Some research synthesize PET images from MRI or CT (Zhang et al., 2022a; Ou et al., 2024b), while others aim at generating high-dose PET images from low-dose PET data (Pan et al., 2024; Shen et al., 2024). While achieving success, this type of image-to-image paradigm, as shown in Fig. 1 (blue path), faces one key challenge: the significant information gap between two image modalities. For instance, generating PET-CT from MRI often suffers from insufficient information representation. The imaging principle of MRI relies on differences in proton relaxation times, which provides structural information. In contrast, PET-CT imaging, based on the positron radiation of radioactive tracers binding to specific molecules, reflects the distribution of these molecules, which MRI inherently lacks.

Multi-modal image synthesis, which integrates complementary information from other modalities, serves as an effective approach to filling the information gap between two image modalities. However, existing multi-modal image synthesis methods typically take all modalities as direct inputs and rely on the network to automatically learn how to utilize them. This approach often suffers from limitations, such as suboptimal exploitation of relationships between different modalities, which in turn leads to inefficient feature fusion and degraded synthesis performance.

To address these limitations, we propose an innovative causality-informed multi-modal synthesis model that explicitly models and leverages causal relationships between multi-modalities to better exploit their complementary information, fill the information gap between image modalities, and improve synthesis performance. Specifically, the novelties and contributions of this paper are as follows: (i) **A novel causality-informed multimodal** medical image synthesis paradigm. Unlike conventional image-to-image translation methods that rely on a single modality, Causal PETS employs a causal graph to explicitly model and leverage causal relationships among multiple modalities, effectively capturing their complementary information for improved PET image synthesis. (ii) Enhanced Performance. Leveraging the causal graph, Causal PETS achieves state-of-the-art (SOTA) reconstruction quality in PET image synthesis. Furthermore, by integrating the synthesized PET images with existing modality data, our approach also attains SOTA classification performance for the early diagnosis and monitoring of Alzheimer's disease (AD). (iii) **High interpretability**. By generating PET images under controlled interventions on specific variables in the model, Causal PETS enables a deeper understanding of the role of each modality, enhancing the interpretability of the synthesis process.
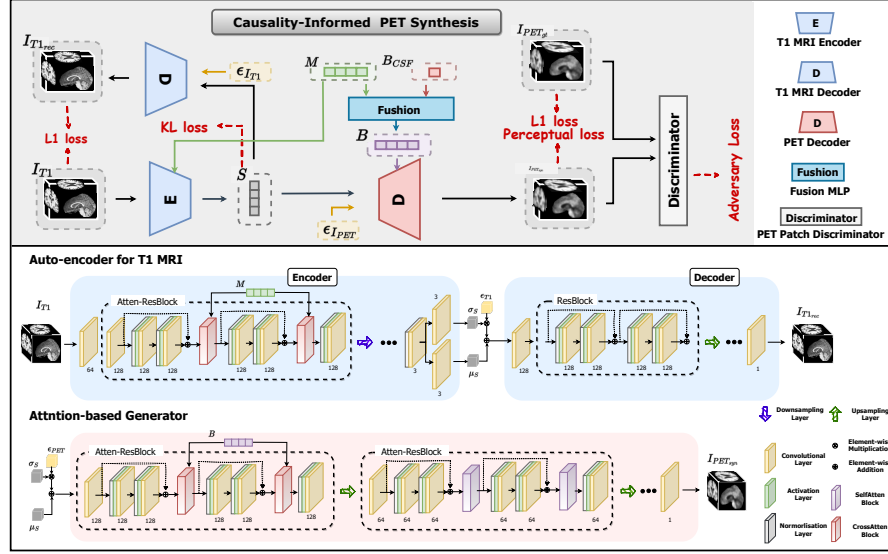
Figure 2: The architecture of the auto-encoder of T_1 MRI and the PET generator.

2. Method: Causal PETS

2.1. Causal Graph & Structural Causal Equations

2.1.1. CAUSAL GRAPH

A Causal Graph (Pearl, 2010) is a graph representing the causal relationships between variables, where nodes represent variables and an edge from node A to node B (A→B) signifies that A has a direct causal effect on B. The variables can be either **Observed Variables**, measured in the dataset or **Latent Variables**, not directly observed.

We construct the causal graph of PET acquisition for AD as Fig.1, which ontains the demographic factor M (Age, Gender, Education levels, APOE4 allele number), brain structure S, biomarker CSF level B_{CSF}, biomarker distribution B, T_1 MRI image I_{T_1}, and PET image I_{PET} , according to clinical research, explained as follows.

$M \rightarrow S$: Demographic factors have a causal impact on brain structure. Brain atrophy occurs as a results of age increasing (Kasantikul et al., 1979), differences in brain volumes are significant between males and females (Gur et al., 2002), and both educational levels and the presence of the APOE4 allele contribute to brain atrophy (Coffey et al., 1999; Lim et al., 2017).

$S \rightarrow I_{T1}$: T_1 MRI relies on the magnetic resonance signals of hydrogen nuclei(Tang et al., 2018) and is directly caused by the brain structure.

$S \rightarrow I_{PET}$: PET imaging (Catafau and Bullich, 2015) relies on the use of radiotracers the uptake of radiotracers may vary depending on the size and shape of brain (S).

$B \rightarrow I_{PET}$: The radiotracers bind to specific molecules and the molecule concentration distribution(B) determines the radiotracer uptake and thus influences I_{PET} (Catafau and Bullich, 2015).

$B \rightarrow B_{CSF}$: CSF biomarkers measure the sampled concentration of molecules in the cerebrospinal fluid (CSF), thus determined by the distribution of specific molecule.

2.1.2. Structural Causal Equations

Structural Causal Equations show how a variable is generate, which can be expressed as

$$Y = f_Y(\mathrm{Pa}(Y), \epsilon_Y), \tag{1}$$

where f_Y is a generative function, $\mathrm{Pa}(Y)$ denotes the set of parent variables of Y, and ϵ_Y is an error term representing all other latent variables affecting Y.

Take the variable I_{PET} as an example. The structural equation is expressed as

$$I_{PET} = f_{I_{PET}}(S, B, \epsilon_{I_{PET}}), \tag{2}$$

where ϵ_Y can be the PET image prototype, the instrumental or the imaging noise, which can affect the PET image but are not included in the model.

For the PET synthesis model, I_{T1} and B_{CSF} can be used to provide the information of their causal parents. In our model, we predict the posterior S and B by

$$S = g_S(I_{T_1}, M), \quad B = g_B(B_{CSF}, M), \tag{3}$$

where g_S is implemented by an encoder and g_B by an Multiple-Layer Perceptron (MLP). Then we use an decoder to implement $f_{I_{PET}}$ in (2).

2.2. Causality-Informed PET Synthesis (Causal PETS)

The Causal PETS Model is shown in Fig. 2, based on the causal graph Fig. 1. We use two decoders as the structural equation for T_1 MRI and PET, described in (2).

When training and infering, the S and B are firstly approximated by

$$S = E(I_{T_1}, M), \quad B = f(B_{CSF}, M), \tag{4}$$

where f denotes the fusion MLP and E denotes the T_1 MRI encoder.

Then $I_{PET_{syn}}$ and $I_{T1_{rec}}$ are generated as the causal structural function,

$$I_{PET_{syn}} = G_P(S, B, \epsilon_{PET}), \quad I_{T1_{rec}} = G_{T_1}(S, \epsilon_{T1}), \tag{5}$$

where G_P and G_{T_1} denotes generator, ϵ_{PET} and ϵ_{T_1} are sampled from normal distribution.

As for the better quality of PET synthesis, we added a discriminator D for adversary training, matching the synthetic data distribution to the target data distribution.

2.2.1. Architectures

In this section we introduce the architectures of networks of the proposed model. Fig. 2 provides an architectural overview. The model architecture details can be found in code[1].

Auto-encoder for $\mathbf{T}_1$ The auto-encoder consists of one Atten-ResBlock, five ResBlocks, three Upsample Layers, and three Downsample Layers. The encoder predict the μ_S and σ_S and the decoder outputs the $I_{T1_{rec}}$. The dimension of feature map is marked in Fig. 2.

---

1. https://github.com/jessyblues/Causality-Informed-PET-Synthesis-from-Multi-modal-Data

Attention-based Generator The attention-based generator consists of three Atten-ResBlocks and three upsampling layers. The first Atten-ResBlock is of cross-attention and the other two are of self-attention.

Fusion MLP and Discriminator The fusion MLP is made up with three linear layers of a latent dimension 128, and the discriminator is chosen as a Patch Discriminator, a discriminator structure based on PatchGAN (Isola et al., 2017a).

2.2.2. LOSS FUNCTION

For the auto-encoder, the commonly used reconstruction loss and the KL loss is used as

$$\mathcal{L}_{\text{AE}} = \mathbb{E}_{x \sim I_{T_1}} \left[\|I_{T_1} - G_{T1}(E(I_{T1}))\|^2 \right] + \text{KL} \left(q_\phi(E(I_{T1})) \| p(S) \right), \tag{6}$$

where $\text{KL}\left(q_\phi(S|I_{T_1})\|p(S)\right)$ is the KL divergence between the approximate posterior and the prior distribution (normal distribution)

For the PET image generator, an L_1 loss and a perceptual loss are incorporated to minimize the absolute pixel-wise difference and the perceptual difference.

$$\mathcal{L}_1(G_P) = \mathbb{E}_{(x,z,m,y) \sim (I_{T_1}, B_{CSF}, M, I_{PET}), \epsilon \sim \mathcal{N}} \|y - G_P(E(x,m), f(z), \epsilon)\|_1, \tag{7}$$

$$\mathcal{L}_{\text{Perceptual}}(G_P) = \mathbb{E}_{(x,z,m,y) \sim (I_{T_1}, B_{CSF}, M, I_{PET}), \epsilon \sim \mathcal{N}} \|V(y) - V(G_P(E(x,m), f(z,m), \epsilon))\|_1. \tag{8}$$

The loss function of the generator G_P and the discriminator D is as follows,

$$\mathcal{L}(D) = \mathbb{E}_{(x,z,m) \sim (I_{T_1}, B_{CSF}, M), \epsilon \sim \mathcal{N}} \left[(D(G_P(E(x,m), f(z,m), \epsilon))^2 \right] + \mathbb{E}_{y \sim I_{PET}} \left[(D(y) - 1)^2 \right], \tag{9}$$

$$\mathcal{L}_{adv}(G_P) = \mathbb{E}_{x \sim I_{T1}, z \sim (B_{CSF}, M), \epsilon \sim \mathcal{N}}[1 - (D(G_P(E(x,m), f(z), \epsilon))^2]. \tag{10}$$

The overall loss function are as follows:

$$\mathcal{L}_{\text{AE}, G_P} = \mathcal{L}_{\text{AE}} + \mathcal{L}_1(G_P) + \lambda_p \mathcal{L}_{Perceptual}(G_P) + \lambda_{adv} \mathcal{L}_{adv}(G_P). \tag{11}$$

where λ_p and λ_{adv} are the hyper-parameters and the details are in appendix.

3. Experiments

3.1. Datasets and Training Details

We train and test on a subset of the Alzheimer's Disease Neuroimaging Initiative database (ADNI) (Petersen et al., 2010), using PET of two different radio-tracers, AV45 and AV1451, with corresponding CSF bio-marker data. The details of dataset are in the appendix.

	dataset AV45				dataset AV1451			
Method	MAE($\times 10^{-1}$) $\downarrow$	SSIM(%) $\uparrow$	PNSR $\uparrow$	$\epsilon_{SUVR} \downarrow$	MAE($\times 10^{-1}$) $\downarrow$	SSIM(%) $\uparrow$	PNSR $\uparrow$	$\epsilon_{SUVR} \downarrow$
CycleGAN (Zhu et al., 2017)	$0.349_{\pm 0.147}$	$92.46_{\pm 1.50}$	$23.192_{\pm 2.558}$	$0.23_{\pm 0.30}$	$0.443_{\pm 0.219}$	$88.26_{\pm 2.10}$	$21.869_{\pm 3.263}$	$0.12_{\pm 0.11}$
Pix2Pix (Isola et al., 2017b)	$0.450_{\pm 0.164}$	$93.78_{\pm 1.46}$	$21.825_{\pm 3.001}$	$0.21_{\pm 0.27}$	$0.251_{\pm 0.081}$	$91.34_{\pm 1.50}$	$25.469_{\pm 1.830}$	$0.12_{\pm 0.12}$
U-Net w/o condition	$0.265_{\pm 0.122}$	$96.73_{\pm 1.33}$	$25.719_{\pm 2.930}$	$0.16_{\pm 0.17}$	$0.230_{\pm 0.112}$	$95.42_{\pm 1.37}$	$26.216_{\pm 2.610}$	$0.11_{\pm 0.10}$
U-Net w/ condition	$0.237_{\pm 0.142}$	$97.19_{\pm 1.28}$	$26.578_{\pm 2.859}$	$0.12_{\pm 0.13}$	$0.229_{\pm 0.119}$	$95.15_{\pm 1.51}$	$26.425_{\pm 2.668}$	$0.09_{\pm 0.09}$
BMGAN (Hu et al., 2021)	$0.320_{\pm 0.116}$	$96.51_{\pm 1.41}$	$24.262_{\pm 2.374}$	$0.18_{\pm 0.22}$	$0.214_{\pm 0.136}$	$94.78_{\pm 1.53}$	$27.644_{\pm 3.691}$	$0.10_{\pm 0.11}$
BPGAN (Zhang et al., 2022b)	$0.321_{\pm 0.114}$	$95.63_{\pm 1.31}$	$24.092_{\pm 2.215}$	$0.18_{\pm 0.24}$	$0.214_{\pm 0.162}$	$93.97_{\pm 1.49}$	$27.688_{\pm 4.285}$	$0.11_{\pm 0.10}$
ControlNet (Zhang et al., 2023)	$0.553_{\pm 0.081}$	$91.90_{\pm 1.78}$	$21.720_{\pm 1.107}$	$0.17_{\pm 0.16}$	$0.941_{\pm 0.167}$	$81.77_{\pm 2.66}$	$19.296_{\pm 1.415}$	$0.10_{\pm 0.09}$
CLDM (Ou et al., 2024a)	$0.329_{\pm 0.128}$	$95.57_{\pm 1.26}$	$23.905_{\pm 2.833}$	$0.23_{\pm 0.15}$	$0.211_{\pm 0.117}$	$97.07_{\pm 2.92}$	$27.896_{\pm 3.215}$	$0.11_{\pm 0.12}$
PASTA (Li et al., 2024)	$0.349_{\pm 0.109}$	$95.33_{\pm 1.20}$	$23.513_{\pm 2.393}$	$0.23_{\pm 0.14}$	$0.394_{\pm 0.171}$	$93.26_{\pm 2.73}$	$23.104_{\pm 3.155}$	$0.12_{\pm 0.09}$
Causal PETS (ours)	$\mathbf{0.224_{\pm 0.104}}$	$\mathbf{97.47_{\pm 1.13}}$	$\mathbf{26.740_{\pm 2.581}}$	$\mathbf{0.10_{\pm 0.13}}$	$\mathbf{0.202_{\pm 0.096}}$	$\mathbf{98.12_{\pm 0.82}}$	$\mathbf{29.687_{\pm 1.905}}$	$\mathbf{0.08_{\pm 0.10}}$

Table 1: Quantitative comparison of PET images synthesised by different methods

Figure 3: The visualizations of synthesized AV45 PET (the first row) and AV1451 PET (the third row) of different methods with the different map against the grountruth.

3.2. PET Image Quality

We evaluate our model's performance in generating PET images, compared against image translation methods Pix2Pix (Isola et al., 2017b), CycleGAN (Zhu et al., 2017), Unet, and MRI-specific networks including BMGAN (Hu et al., 2021), BPGAN (Zhang et al., 2022b) and diffusion-based method including ControlNet (Zhang et al., 2023), CLDM (Ou et al., 2024a), and PASTA (Li et al., 2024). Quantitative results, including mean absolute error (MAE), multi-sacle structural similarity (SSIM) index, and Peak Signal-to-Noise Ratio (PSNR), are detailed in Table 1.

Specifically, we furthermore compute the SUVR (Standardized Uptake Value Ratio) MAE between the synthesized PET and the target real PET. SUVR is a metric to quantify the concentration of radio-tracer uptake in a region of interest (ROI) relative to a reference region. As recommended in clinical research (Schindler et al., 2021; Jack Jr et al., 2018), the cerebral cortex region is set as the ROI with the cerebellar cortex as the reference region. The formula of SUVR computation is provided in the appendix.

Our method achieves the lowest MAE, the highest SSIM and PSNR on both datasets, demonstrating a superior accuracy and the advanced structural similarity. Our method also outperforms other methods in terms of SUVR MAE, demonstrating its effectiveness in synthesizing high-quality PET images in terms of ROI.

Figs. 3 show the slices of the synthesised AV45 and AV1451 PET images respectively. The error map is visualized by subtracting the real PET image from the synthetic one. It shows that our method generates the most authentic PET image and the darkest error map.

3.3. Downstream Tasks

	dataset AV45					dataset AV1451				
Method	F1	AUC	Acc	Prec	Recall	F1	AUC	Acc	Prec	Recall
CycleGAN (Zhu et al., 2017)	0.74367	0.54699	0.75510	0.73312	0.75510	0.86502	0.93009	0.85185	0.88881	0.85185
Pix2Pix (Isola et al., 2017b)	0.76648	0.59096	0.75510	0.78001	0.75510	0.86902	0.92097	0.85185	0.90897	0.85185
Unet w/o condition	0.77533	0.49639	0.80612	0.75409	0.80612	0.90177	0.97872	0.88889	0.94017	0.88889
Unet w/ condition	0.79949	0.78976	0.79082	0.81016	0.79082	0.91000	0.94833	0.90741	0.91355	0.90741
BMGAN (Hu et al., 2021)	0.70168	0.50141	0.68878	0.71553	0.68878	0.89463	0.93921	0.88889	0.90370	0.88889
BPGAN (Zhang et al., 2022b)	0.76482	0.64197	0.76020	0.76973	0.76020	0.88403	0.99696	0.90741	0.91631	0.90741
ControlNet (Zhang et al., 2023)	0.58333	0.72222	0.73333	0.58333	0.58333	0.45455	0.70000	0.83333	0.45455	0.41667
CLDM (Ou et al., 2024a)	0.64000	0.77778	0.80000	0.64000	0.67308	0.91793	0.92358	0.96296	0.91793	0.97917
PASTA (Li et al., 2024)	0.41558	0.50000	0.57143	0.41558	0.32653	0.77778	0.82703	0.83333	0.77778	0.90000
Causal PETS (ours)	**0.82261**	**0.81486**	**0.81122**	**0.83937**	**0.81122**	**0.94260**	**0.96657**	**0.94444**	**0.94213**	**0.94444**
Real Images	0.84894	0.82811	0.84694	0.85117	0.84694	0.96488	0.99392	0.96296	0.97119	0.96296

Table 2: Comparison of pMCI vs sMCI classification results using synthesised PET images.

	dataset AV45					dataset AV1451				
Method	F1	AUC	Acc	Prec	Recall	F1	AUC	Acc	Prec	Recall
PET	0.84894	0.82811	0.84694	0.85117	0.84694	0.96488	0.99392	0.96296	0.97119	0.96296
Tabular	0.77974	0.79096	0.74490	**0.88921**	0.74490	0.88079	0.63526	0.88889	0.87710	0.88889
T_1 and PET	0.81249	0.79257	0.81122	0.81379	0.81122	0.91262	0.95745	0.92593	0.93174	0.92593
T_1 and Tabular	0.78680	0.81767	0.76020	0.83971	0.76020	0.89463	0.91489	0.88889	0.90370	0.88889
PET and Tabular	0.80826	0.81747	0.79082	0.83735	0.79082	0.96296	**0.99696**	0.96296	0.96296	0.96296
T_1 and PET and Tabular	**0.86334**	**0.86827**	**0.85204**	0.88591	**0.85204**	**0.98200**	**0.99696**	**0.98148**	**0.98380**	**0.98148**
PET*	0.82261	0.81486	0.81122	0.83937	0.81122	0.94260	0.96657	0.94444	0.94213	0.94444
PET* and Tabular	0.78625	0.75361	0.79082	0.78204	0.79082	0.92975	0.97264	0.92593	0.93745	0.92593
T_1 and PET* and Tabular	**0.82659**	**0.80341**	**0.81633**	**0.84141**	**0.81633**	**0.94843**	**0.97872**	**0.94444**	**0.96111**	**0.94444**

Table 3: Comparison of pMCI vs sMCI classification results by different modality data.

To further evaluate the synthesized PET images, we employ a downstream task of classification of progressive Mild Cognitive Impairment (pMCI) and stable Mild Cognitive Impairment (sMCI). Mild Cognitive Impairment (MCI) is a stage before dementia. pMCI is likely to progress to AD while sMCI remains relatively stable over time. Distinguishing between pMCI and sMCI is essential for early intervention and treatment planning. Based on the encoder of the T_1 MRI encoder, we train a PET classifier on the real PET images and synthesized PET are only for test. As Table 2 shows, for both the AV45 and AV1451, our method scores the highest in all metrics. Real PET images provide the benchmark and our proposed method closely approximates these results.

Table 3 shows the classification results using different modalities data. PET* denoted the synthesized PET images while PET denoted the real PET images. It can be observed that using real PET images significantly improves the classification performance. When real PET images are unavailable, synthesized PET images (PET*) show considerable improvement over using just T_1 and Tabular data. The combination of T_1, PET*, and Tabular achieves an F1 score of 0.82659 in the AV45 dataset, compared to 0.78680 for the combined T_1 and Tabular. The results indicate the clinical significance of our method.

Figure 4: The SURV comparisons of counterfactual PET

3.4. Interpretability

Generating counterfactual PET images by intervening on variables explains their roles in the model. For example, reducing APOE4 count lowers SUVR in generated PET images ($\Delta n_{APOE4} = -1$), shown in Fig. 4, indicating less amyloid or tau deposition and milder AD, consistent with clinical findings (Lim et al., 2017). This enhances Causal PETS' interpretability. More results are in the appendix.

3.5. Ablation Study

To validate the contribution of each module and loss function component, we conducted ablation experiments, shown in Table 4. Removing I_{T1} reconstruction leads to the highest MAE increase and SSIM drop, while omitting the PET Discriminator also degrades performance, highlighting their roles in image refinement. Additionally, perceptual loss and $\epsilon_{I_{PET}}$ help reduce variance, suggesting that $\epsilon_{I_{PET}}$ improves robustness to image noise.

	dataset AV45				dataset AV1451			
Method	MAE($\times 10^{-1}$) $\downarrow$	SSIM(%) $\uparrow$	PNSR$\uparrow$	ϵ_{SUVR} $\downarrow$	MAE($\times 10^{-1}$) $\downarrow$	SSIM(%) $\uparrow$	PNSR$\uparrow$	ϵ_{SUVR} $\downarrow$
w/o I_{T1} reconstruction	$0.247_{\pm 0.109}$	$97.33_{\pm 1.15}$	$26.25_{\pm 2.53}$	$0.130_{\pm 0.120}$	$0.265_{\pm 0.119}$	$98.11_{\pm 1.08}$	$24.02_{\pm 2.24}$	$0.090_{\pm 0.089}$
w/o PET Discriminator	$0.243_{\pm 0.109}$	$97.29_{\pm 1.10}$	$26.25_{\pm 2.49}$	$0.136_{\pm 0.110}$	$0.235_{\pm 0.102}$	$98.01_{\pm 0.93}$	$26.70_{\pm 1.96}$	$0.091_{\pm 0.094}$
w/o Perceptual loss	$0.233_{\pm 0.111}$	$97.31_{\pm 1.10}$	$26.46_{\pm 2.61}$	$0.124_{\pm 0.120}$	$0.211_{\pm 0.092}$	$97.28_{\pm 0.62}$	$28.26_{\pm 2.07}$	$0.083_{\pm 0.104}$
w/o $\epsilon_{I_{PET}}$	$0.233_{\pm 0.131}$	$97.30_{\pm 1.31}$	$26.36_{\pm 2.85}$	$0.131_{\pm 0.182}$	$0.212_{\pm 0.122}$	$97.28_{\pm 1.49}$	$27.38_{\pm 2.38}$	$0.089_{\pm 0.133}$
Ours	$\mathbf{0.224_{\pm 0.104}}$	$\mathbf{97.47_{\pm 1.13}}$	$\mathbf{26.740_{\pm 2.581}}$	$\mathbf{0.10_{\pm 0.13}}$	$\mathbf{0.202_{\pm 0.096}}$	$\mathbf{98.12_{\pm 0.82}}$	$\mathbf{29.687_{\pm 1.905}}$	$\mathbf{0.08_{\pm 0.10}}$

Table 4: Ablation Study on AV45 and AV1451 Dataset

4. Conclusion

In this work, we propose **Causal PETS**, a novel causality-informed synthesis model for generating PET images from multi-modal data. Our model analyzes the causal relationships between different modalities to generate PET images. Our causality-informed PET synthesis model represents a significant step forward in the integration of multi-modal data for medical imaging. However, our work still has some limitations, e.g., we do not consider the temporal dimension. By addressing the limitations we can enhance the clinical applicability and impact of this approach.

References

Brian B Avants et al. Advanced normalization tools (ants). *Insight j*, 2(365):1–35, 2009.

Gunnar Brix, Elke A Nekolla, Dietmar Nosske, and Jürgen Griebel. Risks and safety aspects related to pet/mr examinations. *European journal of nuclear medicine and molecular imaging*, 36:131–138, 2009.

Ana M Catafau and Santiago Bullich. Amyloid pet imaging: applications beyond alzheimer's disease. *Clinical and translational imaging*, 3:39–55, 2015.

C Edward Coffey, JA Saxton, G Ratcliff, RN Bryan, and JF Lucke. Relation of education to brain size in normal aging: implications for the reserve hypothesis. *Neurology*, 53(1): 189–189, 1999.

Ruben C Gur, Faith Gunning-Dixon, Warren B Bilker, and Raquel E Gur. Sex differences in temporo-limbic and frontal brain volumes of healthy adults. *Cerebral cortex*, 12(9): 998–1003, 2002.

Shengye Hu, Baiying Lei, Shuqiang Wang, Yong Wang, Zhiguang Feng, and Yanyan Shen. Bidirectional mapping generative adversarial networks for brain mr to pet synthesis. *IEEE Transactions on Medical Imaging*, 41(1):145–157, 2021.

Juan Eugenio Iglesias et al. Robust brain extraction across datasets and comparison with publicly available methods. *IEEE Trans. Med. Imag*, 30(9):1617–1634, 2011.

Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017a.

Phillip Isola et al. Image-to-image translation with conditional adversarial networks. In *IEEE/CVF CVPR*, pages 1125–1134, 2017b.

Clifford R Jack Jr, Heather J Wiste, Christopher G Schwarz, Val J Lowe, Matthew L Senjem, Prashanthi Vemuri, Stephen D Weigand, Terry M Therneau, Dave S Knopman, Jeffrey L Gunter, et al. Longitudinal tau pet in ageing and alzheimer's disease. *Brain*, 141(5):1517–1528, 2018.

Keith A Johnson, Reisa A Sperling, Christopher M Gidicsin, Jeremy S Carmasin, Jacqueline E Maye, Ralph E Coleman, Eric M Reiman, Marwan N Sabbagh, Carl H Sadowsky, Adam S Fleisher, et al. Florbetapir (f18-av-45) pet to assess amyloid burden in alzheimer's disease dementia, mild cognitive impairment, and normal aging. *Alzheimer's & Dementia*, 9(5):S72–S83, 2013.

Vira Kasantikul et al. Relation of age and cerebral ventricle size to central canal in man: Morphological analysis. *Journal of neurosurgery*, 51(1):85–93, 1979.

Yitong Li, Igor Yakushev, Dennis M Hedderich, and Christian Wachinger. Pasta: P thology-a ware mri to pet cross-modal translation with diffusion models. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 529–540. Springer, 2024.

Yen Ying Lim, Robert Williamson, Simon M Laws, Victor L Villemagne, Pierrick Bourgeat, Christopher Fowler, Stephanie Rainey-Smith, Olivier Salvado, Ralph N Martins, Christopher C Rowe, et al. Effect of apoe genotype on amyloid deposition, brain volume, and memory in cognitively normal older individuals. *Journal of Alzheimer's Disease*, 58(4): 1293–1302, 2017.

Charles Marcus, Esther Mena, and Rathan M Subramaniam. Brain pet in the diagnosis of alzheimer's disease. *Clinical nuclear medicine*, 39(10):e413–e426, 2014.

Shruti Mishra, Brian A Gordon, Yi Su, Jon Christensen, Karl Friedrichsen, Kelley Jackson, Russ Hornbeck, David A Balota, Nigel J Cairns, John C Morris, et al. Av-1451 pet imaging of tau pathology in preclinical alzheimer disease: defining a summary measure. *Neuroimage*, 161:171–178, 2017.

RAJ Nievelstein, HME Quarles van Ufford, TC Kwee, MB Bierings, I Ludwig, FJA Beek, JMH De Klerk, WP Th M Mali, PW De Bruin, and J Geleijns. Radiation exposure and mortality risk from ct and pet imaging of patients with malignant lymphoma. *European radiology*, 22:1946–1954, 2012.

Agneta Nordberg, Juha O Rinne, Ahmadul Kadir, and Bengt Långström. The use of pet in alzheimer disease. *Nature Reviews Neurology*, 6(2):78–87, 2010.

Zaixin Ou, Yongsheng Pan, Yuanning Li, Fang Xie, Qihao Guo, and Dinggang Shen. Synthesizing abeta-pet via an image and label conditioning latent diffusion model for detecting amyloid status. In *ICASSP 2024 - 2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6610–6614, 2024a. doi: 10.1109/ICASSP48485.2024.10448346.

Zaixin Ou, Yongsheng Pan, Yuanning Li, Fang Xie, Qihao Guo, and Dinggang Shen. Synthesizing aβ-pet via an image and label conditioning latent diffusion model for detecting amyloid status. In *ICASSP 2024-2024 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 6610–6614. IEEE, 2024b.

Shaoyan Pan, Elham Abouei, Junbo Peng, Joshua Qian, Jacob F Wynne, Tonghe Wang, Chih-Wei Chang, Justin Roper, Jonathon A Nye, Hui Mao, et al. Full-dose pet synthesis from low-dose pet using 2d high efficiency denoising diffusion probabilistic model. In *Medical Imaging 2024: Clinical and Biomedical Imaging*, volume 12930, pages 428–435. SPIE, 2024.

Judea Pearl. Causal inference. *Causality: objectives and assessment*, pages 39–58, 2010.

Ronald Carl Petersen, Paul S Aisen, Laurel A Beckett, Michael C Donohue, Anthony Collins Gamst, Danielle J Harvey, CR Jack Jr, William J Jagust, Leslie M Shaw, Arthur W Toga, et al. Alzheimer's disease neuroimaging initiative (adni) clinical characterization. *Neurology*, 74(3):201–209, 2010.

Suzanne E Schindler, Yan Li, Virginia D Buckles, Brian A Gordon, Tammie LS Benzinger, Guoqiao Wang, Dean Coble, William E Klunk, Anne M Fagan, David M Holtzman, et al.

Predicting symptom onset in sporadic alzheimer disease with amyloid pet. *Neurology*, 97 (18):e1823–e1834, 2021.

Chenyu Shen, Changjun Tie, Ziyuan Yang, Na Zhang, and Yi Zhang. Bidirectional condition diffusion probabilistic models for pet image denoising. *IEEE Transactions on Radiation and Plasma Medical Sciences*, 2024.

Xiang Tang, Feng Cai, Dong-Xue Ding, Lu-Lu Zhang, Xiu-Ying Cai, and Qi Fang. Magnetic resonance imaging relaxation time in alzheimer's disease. *Brain research bulletin*, 140: 176–189, 2018.

Tonghe Wang, Yang Lei, Yabo Fu, Jacob F Wynne, Walter J Curran, Tian Liu, and Xiaofeng Yang. A review on medical imaging synthesis using deep learning and its clinical applications. *Journal of applied clinical medical physics*, 22(1):11–36, 2021a.

Tonghe Wang, Yang Lei, Yabo Fu, Jacob F Wynne, Walter J Curran, Tian Liu, and Xiaofeng Yang. A review on medical imaging synthesis using deep learning and its clinical applications. *Journal of applied clinical medical physics*, 22(1):11–36, 2021b.

Jiadong Zhang, Zhiming Cui, Caiwen Jiang, Jingyang Zhang, Fei Gao, and Dinggang Shen. Mapping in cycles: Dual-domain pet-ct synthesis framework with cycle-consistent constraints. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 758–767. Springer, 2022a.

Jin Zhang, Xiaohai He, Linbo Qing, Feng Gao, and Bin Wang. Bpgan: Brain pet synthesis from mri using generative adversarial network for multi-modal alzheimer's disease diagnosis. *Computer Methods and Programs in Biomedicine*, 217:106676, 2022b. ISSN 0169-2607.

Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3836–3847, 2023.

Jun-Yan Zhu et al. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *IEEE/CVF CVPR*, pages 2223–2232, 2017.

Appendix A. Dataset

AV45 PET imaging (Johnson et al., 2013), also known as florbetapir (AV-45) PET imaging, is used to visualize amyloid plaques. Correspondingly, CSF $A\beta$ data is chosen as B_{CSF}, referring to measurements of amyloid-beta peptides in the CSF. AV45 PET images and CSF $A\beta$ data reflect the amyloid pathology in brain. AV1451 PET imaging (Mishra et al., 2017), also known as flortaucipir (AV-1451)PET imaging, is the imaging of tau protein tangles. Meanwhile, CSF tau and p-tau data measure the total tau and phosphorylated tau proteins in the CSF, respectively. They both reflect the tau pathology.

The dataset details for training the synthesis model is as follows. We divided the dataset into training, validation, and test sets in a 4:1:1 ratio, ensuring that all scans from the same individual appear in the same set, thereby preventing data leakage.

AV45 dataset			
Category	CN	MCI	AD
# of subjects	519	476	197
# of sessions	776	954	213
Age	74.84±7.29	73.29±7.44	73.60±7.46
AV1451 dataset			
Category	CN	MCI	AD
# of subjects	187	130	61
# of sessions	314	269	82
Age	72.96±7.32	72.80±7.08	72.73±8.02

Table 5: The Basic Information of ADNI Dataset

For imaging processing, all T_1 MRI are skull-stripped using ROBEX (Iglesias et al., 2011), aligned to the MNI152 space, resampled to 1.5mm isotropic using ANTs (Avants et al., 2009), cropped to dimensions of $96 \times 128 \times 96$, and normalized in voxel values to range [0, 1]. The PET images are skull-stripped using ROBEX (Iglesias et al., 2011), registered to the paired T_1 MRI, normalized to range [0, 1].

To pair I_{PET} with I_{T1} and B_{CSF}, we define a successful pairing as having a measurement interval of less than 6 months. let t_{PET}, t_{T1}, and t_{CSF} represent the time points at which the data were taken. A successful pair is defined as:

$$\max\left(|t_{\text{PET}} - t_{\text{T1}}|, |t_{\text{PET}} - t_{\text{CSF}}|, |t_{\text{T1}} - t_{\text{CSF}}|\right) < 6 \text{ months.}$$

This condition ensures that the measurements I_{PET}, I_{T1}, and B_{CSF} are temporally aligned.

Appendix B. Training Process

We adopt an adversarial training approach, where the generator and discriminator are alternately trained. The ADAM optimizer is used with a learning rate (LR) of 0.0001 for the generator, fusion Network, auto-encoder and a LR of 0.0005 for the discriminator. The training is conducted over 1000 epochs, taking approximately 1.5 days. The batch size is set to 2, and we use 6 NVIDIA TITAN RTX GPUs for parallel training. All code is implemented based on PyTorch. The auto-encoder, generator, and discriminator are built using the basic architectures from MONAI. The $\lambda_{Perceptual}$ is set to 0.02 and λ_{adv} is set to 0.005 in the overall loss function.

Appendix C. SUVR formula

The SUVR value is calculated using the following formula:

$$\text{SUVR} = \frac{\frac{1}{|V_{\text{ROI}}|} \sum_{v \in V_{\text{ROI}}} v}{\frac{1}{|V_{\text{ref}}|} \sum_{v \in V_{\text{ref}}} v}, \tag{1}$$

where V_{ROI} is the set of voxel values in ROI, and V_{ref} is the set of voxel values in the reference region. $|V_{\text{ROI}}|$ and $|V_{\text{ref}}|$ denote the number of voxels in each respective region. As

recommended in clinical research (Schindler et al., 2021; Jack Jr et al., 2018), the cerebral cortex region is set as the ROI with

Appendix D. Interpretability

In this section we generate PET images with intervening on one of the variables in Fig. 1, and the role of each variable within the model can be explained. This counterfactual manipulation allows us to explore the role of variables in our causal PETS model, thereby enhancing the interpretability of the model.

Here we presents the visualisation results of synthesised PET with the intervention on B_{CSF} ($A\beta_{42}$ for AV45 dataset and $p\tau_{181}$ for AV1451 daraset).


Figure 5: The PET images generated by intervening on the $B_{CSF}(A\beta_{42})$, experiments on AV45 dataset. The first row shows the generated PET image and the second row shows the difference map.


Figure 6: The PET images generated by intervening on the $B_{CSF}(p\tau)$, experiments on AV1451 dataset. The first row shows the generated PET image and the second row shows the difference map.

In dataset AV45, B_{CSF} measures the soluble $A\beta_{42}$ concentration in CSF, while PET detects the deposited amyloid in the brain. Lower levels of $A\beta_{42}$ in the CSF are often associated with higher levels of amyloid deposition in the brain because $A\beta_{42}$, which is a form

of amyloid-beta, tends to accumulate in amyloid plaques in the brain, leading to reduced levels in the CSF. Thus, a lower CSF concentration indicates more amyloid deposition in the brain, resulting in a higher-signal in PET image, and vice versa. As Fig. 5 shows, our visualization results reflect the specific locations (mostly cerebral cortex) and patterns of amyloid deposition in the brain as CSF concentration decreases.

In dataset AV1451, B_{CSF} measures $p\tau_{181}$ protein and binds to neurofibrillary tangles, which are aggregates of $p\tau$ associated with AD. The $p\tau_{181}$ protein is primarily generated in the brain and enters the CSF through the blood-brain barrier. Thus, a lower CSF concentration indicates less $p\tau_{181}$ in the brain, resulting in a lower-signal in PET image, and vice versa. As Fig. 6 shows, our visualization results reflect the specific locations and patterns of $p\tau_{181}$ in the brain as CSF concentration increases.