



Article

Enhancing Amyloid PET Quantification: MRI-Guided Super-Resolution Using Latent Diffusion Models

Jay Shah ^{1,2} , Yiming Che ^{1,2}, Javad Sohankar ³, Ji Luo ³, Baoxin Li ^{1,2}, Yi Su ^{3,*} , Teresa Wu ^{1,2}
and for the Alzheimer's Disease Neuroimaging Initiative [†]

¹ School of Computing and Augmented Intelligence, Arizona State University, Tempe, AZ 85281, USA; jgshah1@asu.edu (J.S.); yche14@asu.edu (Y.C.); baoxin.li@asu.edu (B.L.); teresa.wu@asu.edu (T.W.)

² ASU-Mayo Center for Innovative Imaging, Arizona State University, Tempe, AZ 85287, USA

³ Banner Alzheimer's Institute, Banner Health, Phoenix, AZ 85006, USA; j.sohankar@bannerhealth.com (J.S.); ji.luo@bannerhealth.com (J.L.)

* Correspondence: yi.su@bannerhealth.com

[†] Collaborators of the Alzheimer's Disease Neuroimaging Initiative are indicated in the Acknowledgment section.

Abstract: Amyloid PET imaging plays a crucial role in the diagnosis and research of Alzheimer's disease (AD), allowing non-invasive detection of amyloid- β plaques in the brain. However, the low spatial resolution of PET scans limits the accurate quantification of amyloid deposition due to partial volume effects (PVE). In this study, we propose a novel approach to addressing PVE using a latent diffusion model for resolution recovery (LDM-RR) of PET imaging. We leverage a synthetic data generation pipeline to create high-resolution PET digital phantoms for model training. The proposed LDM-RR model incorporates a weighted combination of L_1 , L_2 , and MS-SSIM losses at both noise and image scales to enhance MRI-guided reconstruction. We evaluated the model's performance in improving statistical power for detecting longitudinal changes and enhancing agreement between amyloid PET measurements from different tracers. The results demonstrate that the LDM-RR approach significantly improves PET quantification accuracy, reduces inter-tracer variability, and enhances the detection of subtle changes in amyloid deposition over time. We show that deep learning has the potential to improve PET quantification in AD, effectively contributing to the early detection and monitoring of disease progression.

Keywords: partial volume correction (PVC); positron emission tomography; amyloid; deep learning; diffusion models; medical image super-resolution



Citation: Shah, J.; Che, Y.; Sohankar, J.; Luo, J.; Li, B.; Su, Y.; Wu, T.; for the Alzheimer's Disease Neuroimaging Initiative. Enhancing Amyloid PET Quantification: MRI-Guided Super-Resolution Using Latent Diffusion Models. *Life* **2024**, *14*, 1580. <https://doi.org/10.3390/life14121580>

Academic Editors: Juraj Javor and Ivana Shawkatová

Received: 15 October 2024

Revised: 25 November 2024

Accepted: 26 November 2024

Published: 1 December 2024



Copyright: © 2024 by the authors. Licensee MDPI, Basel, Switzerland. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (<https://creativecommons.org/licenses/by/4.0/>).

1. Introduction

Amyloid imaging is a crucial tool in the diagnosis and research of Alzheimer's disease (AD). It allows for the non-invasive detection of amyloid- β ($A\beta$) plaques in the brain, which is a core neuropathological feature of AD [1]. Detecting $A\beta$ pathology at the earliest stages of AD, before the onset of clinical symptoms, is critical for understanding disease progression, developing intervention techniques, and potentially improving patient outcomes. However, accurate quantification of amyloid using positron emission tomography (PET) imaging is limited due to the low spatial resolution of PET scans [2], which is typically around 5 mm and varies across scanner models and reconstruction algorithms. This causes the partial volume effect (PVE) [3], resulting in a spill-out of signal from target gray matter regions and spill-in from surrounding areas. The magnitude of the PVE depends on the size of the target region relative to the spatial resolution of the scans. In the context of amyloid PET imaging, the size of the target regions varies across subjects and often decreases as the subject ages or with disease progression. Therefore, PVE reduces the accuracy, precision, and statistical power of quantitative amyloid PET measurements. Another well-recognized issue of amyloid PET imaging is harmonizing data acquired using different scanners, tracers, and analytical pipelines. To minimize inter-scanner variabilities, a scanner-specific

harmonization filter is often applied at the cost of further reduced spatial resolution [4]. To minimize the variability of amyloid PET measurements from different analytical pipelines, acquisition protocols, and tracers, the Centiloid scale was defined to linearly transform a particular measurement to this scale [5]. However, this Centiloid approach is designed for standardizing global measures and does not improve the between-measure agreements in terms of their shared variance [6–8]. We hypothesize that effective methods for spatial resolution recovery will improve PET quantification and reduce inter-tracer variabilities in amyloid PET measurements, and in this research, we propose a deep learning approach to achieve the goal.

Several partial volume correction (PVC) methods have been proposed in the literature to mitigate the PVE issue using anatomical information from MRI and CT [2,9–15]. Different from that, Tohka and Reilhac [16] showed that Richardson–Lucy, an iterative deconvolution-based method to recover spatial resolution in PET imaging and an alternative to MR-based approaches, offered comparable accuracy with reduced sensitivity to registration and segmentation errors. However, deconvolution-based correction methods are shown to amplify the image noise [17]. Different correction methods can also produce varying results, making standardization and comparison across studies challenging. Deep learning-based techniques [18,19] have recently been explored to tackle some of these challenges. Deep models can better learn complex patterns of tissue heterogeneity and can perform image denoising, potentially addressing noise amplification issues [19]. Deep models trained on diverse datasets may generalize better to different scanners and acquisition protocols [18], potentially improving the standardization and consensus among multi-center studies.

Instead of focusing on partial volume correction directly, as the PVC methods reviewed above do, an alternative is image super-resolution (SR), which refers to the task of rendering a high-resolution image from its low-resolution counterpart. We contend that PVE may be tackled during the process of rendering high-resolution PET from low-resolution PET. SR is a well-studied research problem in computer vision and image processing [20,21]. Use cases of SR span a broad spectrum, improving existing computer vision tasks [22–24] by improving image spatial resolution and perceptual quality, improving surveillance [25], and enhancing diagnostic accuracy in medical research using different imaging modalities [26–28]. Traditional methods for image SR heavily rely on image statistics [29–31], which has been shown to generate blurry and noisy artifacts in their high-resolution outputs [20]. With the advent of deep learning algorithms, several end-to-end architectures have been proposed where the models learn the mapping of low-resolution to high-resolution images through regression-based learning [32]. However, these methods fail to recover high-frequency details mainly because they learn an average mapping from the training dataset (due to L_1/L_2 loss functions), resulting in overly smooth model outputs and lacking spatial details [21].

To address these limitations, generative models have been explored for SR in recent years. Generative models learn to transform a latent variable z with a tractable prior distribution to a learned data space. Generative adversarial networks (GANs), flow-based methods, and diffusion models are three common generative models used to generate synthetic data. They differ in their core approach: GANs [33] are trained in an adversarial setting with generator and discriminator networks, flow-based methods [34] rely on invertible transformations to manipulate data distributions, while diffusion models [35] iteratively add and then learn to remove noise to generate data. GANs are known to suffer from mode collapse [36], resulting in unstable training and limiting the diversity of generated samples. Flow-based methods can impose topological constraints on the mapping between latent and data spaces, limiting their flexibility in modeling complex data distributions [37]. Except for longer sampling times, diffusion models have shown superior performance in generating high-fidelity medical imaging datasets [38–40]. Unlike GAN priors, diffusion model priors can preserve more information to generate realistic data. Motivated by this, we propose an SR solution based on the diffusion model to improve PET quantification.

The success of any SR methods (e.g., GAN, diffusion) heavily relies on the quantity and quality of the training data. Synthetic data has been substantially useful in medical AI research to alleviate issues such as a lack of datasets, annotations, privacy concerns, and high acquisition costs [41–43]. Data samples are typically artificially generated using domain knowledge or modeling techniques to mimic the characteristics and structure of real data without being directly derived from actual observations. It can be used to train AI models where target data are unavailable or scarce and provides a promising alternative to making AI models generalized to real-world datasets [44–46]. These studies mainly focus on improving detection and segmentation from high-resolution imaging. However, its applicability in enhancing PET quantification remains unexplored.

In this study, (1) we develop a new latent diffusion model for resolution recovery (LDM-RR) in PET imaging. Instead of training the diffusion model to minimize loss on the noise scale, we introduce a composite loss function with three terms: L_1 , L_2 , and MS-SSIM at the noise and image scale to improve MRI-guided reconstruction. (2) We developed a synthetic data generation pipeline to generate PET digital phantoms mimicking high-resolution PET scans for model training. (3) We evaluate the performance of our LDM-RR model in improving the statistical power of detecting longitudinal changes. (4) We evaluate the ability of the LDM-RR model to improve the agreement of amyloid PET measurements acquired using different tracers.

2. Materials and Methods

2.1. Datasets and Simulation Procedure

Imaging data from three different cohorts were used in this study to enable our experiments: (1) the Alzheimer’s Disease Neuroimaging Initiative (ADNI) cohort (adni.loni.usc.edu) [47], (2) the Open Access Series of Imaging Studies-3 (OASIS-3) [48], and (3) the Centiloid Project florbetapir (FBP) calibration dataset [49] (<http://www.gaain.org/centiloid-project>, accessed on 15 October 2023). A subset of the ADNI database containing MRI scans was utilized for data simulation to train the diffusion model, while another subset with FBP scans (Table 1) was employed to evaluate the model’s performance in detecting longitudinal changes. Additionally, paired FBP-PiB imaging data from the OASIS-3 and Centiloid databases (Table 1) were used to further assess the model’s performance in cross-tracer harmonization. Details regarding data selection and simulation are provided in the subsequent sections. The ADNI was launched in 2003 as a public–private partnership led by Principal Investigator Michael W. Weiner, MD. The original goal of ADNI was to test whether serial magnetic resonance imaging (MRI), positron emission tomography (PET), other biological markers, and clinical and neuropsychological assessment can be combined to measure the progression of mild cognitive impairment (MCI) and early Alzheimer’s disease (AD). The current goals include validating biomarkers for clinical trials, improving the generalizability of ADNI data by increasing diversity in the participant cohort and providing data concerning the diagnosis and progression of Alzheimer’s disease to the scientific community.

Table 1. Summary of demographic information of the three cohorts included in this study. * 14 out of 46 were unknown.

Cohort	ADNI	OASIS-3	Centiloid
Sample count	334 (167 baseline-followup FBPs)	113 (FBP-PiB pairs)	46 (FBP-PiB pairs)
Age (SD) years	75.1 (6.9)	68.1 (8.7)	58.4 (21.0)
Education (SD) years	16.1 (2.7)	15.8 (2.6)	NA
Male (%)	182 (54.5%)	48 (42.5%)	27 (58.7%)

Table 1. Cont.

Cohort	ADNI	OASIS-3	Centiloid
Cognitive impairment (%)	236 (70.6%)	5 (4.4%)	24 (52.2%)
APOE4+ (%)	218 (65.3%)	38 (33.6%)	15 (46.9 %*)
PET interval (SD) years	2.0 (0.06)	NA	NA

2.1.1. Data to Train the Diffusion Model

We utilized 3376 MRI scans randomly selected from the ADNI database to generate simulated high-resolution digital phantoms (simDP) and simulated florbetapir (simFBP) using an MR-based procedure as previously described [50] and mimics the distribution of florbetapir (FBP) uptake in participants across a wide range of amyloid burden and clinical status and the noise and spatial resolution characteristics of typical PET images. The specific set of MRIs selected as the basis for simulation does not have a major impact on subsequent experiments and, therefore, was not described in detail. The size of the dataset captures the overall distribution and variability of structural brain differences in the elderly population without losing generalizability. A detailed description of the simulation procedure is discussed in Section 2.2 below. From this simulated dataset, 3038 samples were used to train, and 338 samples were used to validate our LDM-RR model's performance.

2.1.2. Data for Evaluating Longitudinal Power

To evaluate LDM-RR's ability to improve statistical power to detect longitudinal changes in amyloid, we selected 167 ADNI participants with a mean age of 74.1 years (SD = 6.8), who are amyloid positive at baseline using a Centiloid cutoff of 20 [51] and have two-year follow-up (2.0 ± 0.06 years interval) FBP scans. The choice of these participants is to ensure they are on a trajectory to accumulate amyloid during the study period, i.e., having a positive expected rate of amyloid accumulation. Additional demographic information of this cohort is summarized in Table 1.

2.1.3. Data for Evaluating Harmonization Performance

From the OASIS-3 database [48], we identified 113 pairs of FBP-PiB scans with a mean age of 68.1 years (SD = 8.7), and similarly, 46 pairs from the Centiloid project [49] (<http://www.gaain.org/centiloid-project>, accessed on 15 October 2023) with a mean age of 58.4 years (SD = 21.0). Refer to Table 1 for demographic information of these two cohorts. Studies for the cohorts included here were approved by their corresponding institutional review boards and written informed consent was obtained for each participant.

2.2. Image Analysis and Simulation

FreeSurfer v7.3 [52] (Martinos Center for Biomedical Imaging, Charlestown, MA, USA) (<https://surfer.nmr.mgh.harvard.edu/fswiki>, accessed on 15 October 2023) was used to automatically segment T1-weighted MRIs to define the anatomical regions of interests (as defined in the `wmparc.mgz` file). PET images were processed using a FreeSurfer-dependent pipeline that included resolution harmonization filtering, inter-frame motion correction, target frame summation, PET-to-MR registration, and regional and voxel-wise SUVR calculation [50,53]. A mean cortical SUVR (MCSUVR) was calculated as the summary measure of amyloid burden and used to evaluate longitudinal and harmonization performance [53]. For comparison purposes, a Richardson–Lucy algorithm was adopted for resolution recovery (RL-RR) through iterative deconvolution [54,55]. In our experiment, the MATLAB (The Mathworks, Inc., v2021a, Natick, MA, USA) function: `deconvlucy` (<https://www.mathworks.com/help/images/ref/deconvlucy.html>, accessed on 15 October 2023) was called with 20 iterations and an 8 mm full-width-half-max (FWHM) Gaussian

kernel to generate the deconvolved high-resolution PET images and the corresponding MCSUVR estimation.

Similar to previously described by Su et al. [50], the simulation of high-resolution digital phantom (DP) and PET images (simFBP) was performed using segmented MRI as the input. For DP generation, each voxel was assigned a specific intensity value according to tissue type-specific distributions observed from actual FBP SUVR images across the aging and AD spectrum. For non-brain voxels, i.e., those not defined in the `wmparc.mgz` file, the voxel intensity was assigned by randomly scaling the normalized T1-MRI images to simulate moderate non-brain uptake. To generate simFBP images, the DP was smoothed and projected to the sinogram space, adding Poisson noise, and reconstructed back to the image space. We generated the simulation with a range of noise levels as seen in real-world PET scans with a noise equivalent count rate (NECR) of 75 ± 26 kcps [56,57]. The target resolution of the simFBP data is 8 mm in FWHM, approximating the resolution of standardized PET data from ADNI [4]. Figure 1 shows a visual example of a simulated digital phantom (B) and PET image (C) matching a T1-MRI image (A).

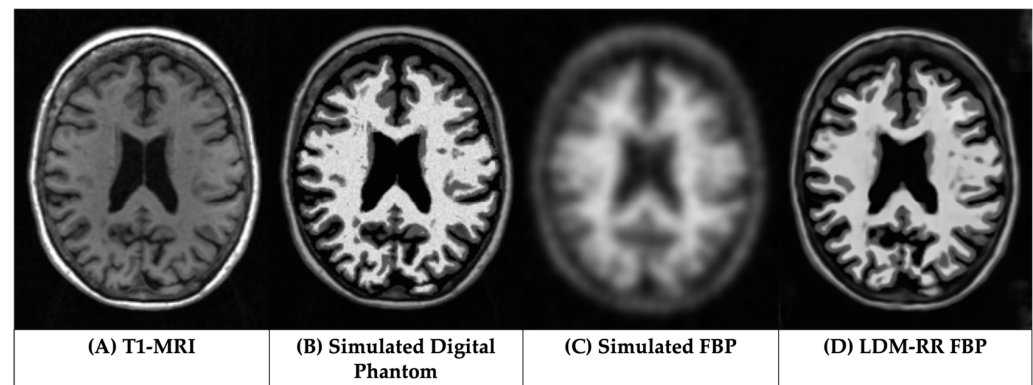


Figure 1. Visualization of a simulated digital phantom (simDP) and simulated FBP (simFBP) from the data simulation pipeline using T1-MRI and the LDM-RR generated synthetic super-resolution FBP.

2.3. LDM-RR: PET Resolution Recovery Framework

We use a latent diffusion model to generate synthetic high-resolution FBP scans given standard low-resolution FBP and matching MRI scans. Figures 2 and 3 give an overview of the training process. Diffusion models have shown impressive results in generating 2D images [58]. However, they are computationally demanding at the training and inference stages. Medical imaging modalities, such as MRI and PET, are more complex as they capture spatial information in 3D. Latent diffusion models operate at a lower-dimensional latent space by compressing useful information from these high-dimensional imaging data.

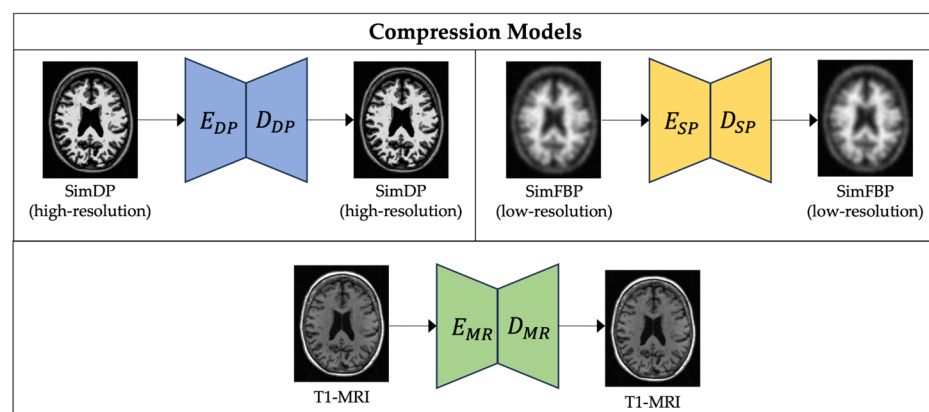


Figure 2. Training three modality-specific autoencoder models to compress high-dimensional simulated DP, simulated FBP and MRI data into a lower-dimensional latent representation.

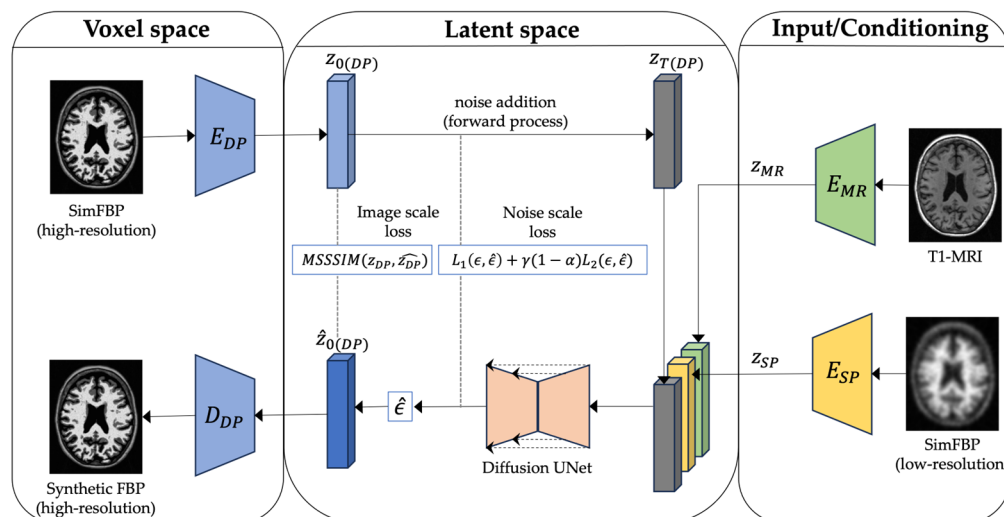


Figure 3. Proposed LDM-RR framework’s training process for PET super-resolution. LDM is conditioned on latent representations of T1-MRI and simFBP and uses a combination of image and noise scale losses to generate corresponding high-resolution simDP.

Our proposed LDM-RR is built upon a state-of-the-art LDM originally proposed to generate 3D brain MRIs [40]. Specifically, it has a 2-stage training process and three different components: an encoder, a diffusion U-Net [59], and a decoder model. The encoder compresses high-dimensional data into a low-dimensional latent representation, diffusion U-Net converts simFBP to simDP in the latent space through a denoising process, and the decoder upsamples the low-dimensional simDP to its original image space. Trained models and implementation code will be made available for reproducibility and further research (<https://github.com/jaygshah/LDM-RR>, accessed on 5 October 2024).

2.3.1. Compression Models

The goal of the compression model is to create a compressed representation of high-dimensional brain images that serve as the foundation for the subsequent diffusion model. We use an autoencoder [40] that compresses the 3D brain images into a lower-dimensional latent representation capturing perceptual representation of original images while preserving essential features to reduce complexity. In the first stage, we train three modality-specific 3D autoencoder models separately for simFBP, simDP, and MRI (see Figure 2), each with a combination of L_1 loss, perceptual loss, a patch-based adversarial objective, and a KL regularization of the latent space [40]. The input to the encoder is a 3D image with dimensions $256 \times 256 \times 256$, and we extract smaller sub-volumes of size $64 \times 64 \times 64$ to fit in GPU memory. The encoder maps these sub-volumes to a latent representation of size $16 \times 16 \times 16$. Once trained, latent representations from these encoders are used as inputs to the diffusion U-Net. See Appendix B for more details on the autoencoder model architectures and hyperparameters used.

2.3.2. Diffusion Model

Diffusion U-Net in LDMs perform denoising by iteratively predicting and removing noise in latent space. Typically, they are trained to minimize the L_2 loss between predicted and actual noise [40,59]. However, for super-resolution, we found that minimizing L_2 loss does not consistently guarantee the recovery of brain structure information in generated outputs, which we have further analyzed in the discussion section. Prior studies have shown that using a mix of image restoration loss can produce high-fidelity images compared to single loss functions [60]. L_2 regularization is sensitive to outliers and can introduce visual artifacts since it penalizes high errors. L_1 , on the other hand, is robust to outliers but suffers from non-differentiability in zero and slow training [61]. Moreover, Zhao et al. [60] showed for image restoration and SR, L_1 , and L_2 penalties fail to capture structure

information and proposed a multi-scale structural similarity index (MS-SSIM) metric. Voxel-level intensity has a high impact on PET quantification [62]. Here, we hypothesize and show through results that existing L_2 loss-based diffusion models do not provide a clinically accurate reconstruction of PET scans. A weighted combination of L_1 , L_2 , and MS-SSIM losses, on the image and noise scales can accurately generate a high-resolution using MRI and simFBP.

To train diffusion models, a small amount of Gaussian noise is progressively added to the data in T steps through a forward noise addition process, forming a Markov chain (Equations (1) and (2)) [35]:

$$q(z_{1:T}|z_0) = \prod_{t=1}^T q(z_t|z_{t-1}) \quad (1)$$

$$q(z_t|z_{t-1}) = N\left(z_t; \sqrt{1 - \beta_t}z_{t-1}, \beta_t I\right) \quad (2)$$

Here, β_t is the fixed variance schedule and z_t follows a pure Gaussian noise distribution after many forward diffusion steps T ($T = 1000$ in our experiments). The diffusion U-Net learns the reverse diffusion process, i.e., denoising z_T to z_0 (Equations (3) and (4)) which can be formulated as [35]:

$$p_\theta(z_{0:T}) = p(z_t) \prod_{t=1}^T p_\theta(z_{t-1}|z_t) \quad (3)$$

$$p_\theta(z_{t-1}|z_t) = N\left(z_{t-1}; \mu_\theta(z_t, t), \sigma_t^2 I\right) \quad (4)$$

$$z_t = \sqrt{\alpha_t}z_0 + \sqrt{1 - \alpha_t}\epsilon \quad (5)$$

where μ_θ represents the denoising neural network (diffusion U-Net) and $\sigma_t^2 = \frac{1 - \bar{\alpha}_{t-1}}{1 - \bar{\alpha}_t} \beta_t$. Traditionally, the diffusion models are trained to predict the added noise in forward diffusion process by minimizing L_2 loss between predicted ($\hat{\epsilon}$) and added noise (ϵ) formulated as [59]:

$$L_\theta = E_{x, \epsilon \sim N(0,1), t} \left[\|\epsilon - \epsilon_\theta(z_t, t)\|_2^2 \right] \quad (6)$$

Furthermore, we can estimate the noise-free latent vector using the predicted noise (ϵ) from the diffusion model using Equation (5) from Ho et al. [35] as:

$$\hat{z}_0 = \frac{z_t - \sqrt{1 - \alpha_t}\hat{\epsilon}}{\sqrt{\alpha_t}} \quad (7)$$

Zhao et al. [60] observed that image reconstruction performance can be improved by adding perceptual image metrics such as MS-SSIM in a network's loss function. It allows capturing structural details at multiple scales while maintaining voxel-level accuracy. While this holds true for current encoder-decoder architectures, to the best of our knowledge, it has not yet been investigated for denoising diffusion networks in latent space. Since our goal is to fuse structure information from T1-MRI to guide the reconstruction, we modify LDM's vanilla loss function (L_θ) on noise scale to a weighted combination of L_2 and MS-SSIM loss on image scale as:

$$loss_1 = (1 - \alpha)L_2(z_0, \hat{z}_0) + \alpha MSSSIM(z_0, \hat{z}_0) \quad (8)$$

Here $\alpha = 0.8$ [60] in Equation (8) is an empirically set hyper-parameter. We explored $\alpha = [0.2, 0.5, 0.8]$. However, $\alpha = 0.8$ resulted in the model with the best performance in reconstructing simDP on the simulated dataset's validation set. While L_2 allows easier optimization in diffusion training due to its convergence properties, it is known to produce an averaging effect, which forces the model to predict values closer to the mean of training data [60]. We argue that using only L_2 loss can help preserve whole image-level properties

but may also produce inaccurate estimates at the voxel level. To this end, we propose L_1 loss at the noise scale to ensure voxel-level details are preserved in the denoising process.

$$loss_2 = L_1(\epsilon, \hat{\epsilon}) = |\epsilon - \hat{\epsilon}| \quad (9)$$

A combined loss function using the two loss terms from Equations (8) and (9),

$$loss_{combined} = L_1(\epsilon, \hat{\epsilon}) + (1 - \alpha)L_2(z_0, \hat{z}_0) + \alpha MSSSIM(z_0, \hat{z}_0) \quad (10)$$

was used to train the LDM-RR model. The combined loss is indeed equivalent to (see Appendix A):

$$loss_{combined} = L_1(\epsilon, \hat{\epsilon}) + \gamma(1 - \alpha)L_2(\epsilon, \hat{\epsilon}) + \alpha MSSSIM(z_0, \hat{z}_0) \quad (11)$$

By minimizing loss on image (z) and noise (ϵ) scales, the LDM-RR model learns to reduce the disparity between the reconstructed high-resolution PET image and the target digital phantom while preserving image-level and voxel-level structure details, and later (voxel-level details) may play an important role for correcting the partial volume effects.

Figure 3 illustrates the second stage of training where only the diffusion U-Net is trained whereas encoder and decoder model parameters are kept frozen. Input to the U-Net is a concatenation of noisy latent representation of simDP ($z_{T(DP)}$) and conditioning of matching MRI (z_{MR}) and simFBP (z_{SP}) latent representations. The model's predicted noise ($\hat{\epsilon}$) can be used to estimate $\widehat{z_{0(DP)}}$ and calculate the combined loss (Equation (11)), which is used to update diffusion U-Net parameters in each training epoch. See Appendix B for more details on model architectures, training, and computational resources.

2.4. Statistical Analysis

2.4.1. Simulated Data Analysis

To evaluate our LDM-RR's actual performance at generating a high-resolution FBP scan, we compare the mean recovery coefficient (RC)—measured as the ratio of synthetic high-resolution PET MCSUVR to that of ground-truth from simulated DP, and variability measured by the standard error (SE). A value closer to 1 indicates higher reconstruction performance. We compare our method to Richardson–Lucy-based resolution recovery (RL-RR), traditional LDM with L_2 noise loss, LDM with L_2 noise and L_1 image loss, and simulated FBP scans without any correction.

2.4.2. Longitudinal Analysis

To evaluate the longitudinal performance of the LDM-RR, for each of the 167 participants, the annualized rate of amyloid accumulation was calculated by dividing the MCSUVR change from baseline to the follow-up visit by the follow-up interval, commonly used in longitudinal PET studies [63]. The imaging data was acquired using standard protocols, and harmonization procedures were performed to reduce variability. The mean and standard deviation of the annualized rate of change were evaluated for each analysis method, i.e., raw measurement, with RL-RR, and with the LDM-RR. A one-sample t -test (one tail) was used to determine whether the annualized rate was significantly greater than zero. A smaller p -value is interpreted as having greater power to detect the longitudinal accumulation of amyloid burden. To further compare the statistical power of different techniques in a longitudinal setting, we estimated the number of participants per arm needed to detect a 25% reduction in amyloid accumulation rate due to treatment with 80% power and a two-tailed type-I error of $p = 0.05$ in hypothetical anti-amyloid treatment trials similar to previous studies [64,65]. A smaller estimated sample size (SS) indicates greater statistical power.

2.4.3. Cross-Tracer Analysis

In the cross-tracer analysis, we evaluated the impact of RL-RR vs. LDM-RR on the agreement of PET-derived global amyloid burden, i.e., MCSUVR, using paired FBP-PIB data from OASIS-3 and the Centiloid project. The imaging data was also collected following standard protocols and underwent harmonization processes to minimize variability, similar to the ADNI study. In our experiment, the raw PIB MCSUVR was used as the reference amyloid burden measurement, and we evaluated whether the corrected FBP MCSUVR is more strongly correlated with PIB MCSUVR using Steiger's test. We also test whether the LDM-RR corrected FBP MCSUVR is more strongly associated with PIB MCSUVR than the RL method.

3. Results

3.1. Qualitative Assessments

Figure 4 showcases corrected FBP scans using RL (Figure 4C) and LDM-RR (Figure 4D) methods in comparison to the real FBP scan without any correction (Figure 4A). The proposed LDM-RR model-generated synthetic FBP image has an improved spatial resolution, with a similar level of anatomical details matching T1-MRI (Figure 4B). A similar example of LDM-RR applied to simulated data is shown in Figure 1D. Although RL-RR does not require an MRI, it generated noisier images and was not able to fully recover the high-resolution details (Figure 4C). The LDM-RR method leverages the high-resolution structural information from MRI to guide the super-resolution process, resulting in PET images with reduced partial volume effects.




Figure 4. Visual comparison of generated FBP scans using RL-RR and our LDM-RR to real FBP and T1- MRI for a sample from OASIS-3 cohort.

3.2. Evaluation on Simulated Data

A visual example of model-generated synthetic FBP from the test set of simulated data is shown in Figure 1D. The mean RC from different diffusion models compared to the RL-RR method and without any correction is shown in Figure 5. Our proposed LDM-RR model was able to better reconstruct target simDP (0.96, SE = 0.004, $p < 0.001$) compared to RL-based correction (0.82, SE = 0.005, $p < 0.001$) and without any corrections (0.76, SE = 0.008, $p < 0.001$). It also performs significantly better compared to a typical LDM architecture [40] for super-resolution (1.32, SE = 0.08, $p < 0.001$) and other combinations of noise and image scale loss (1.58, SE = 0.09, $p < 0.001$). The improvement in recovery coefficient with LDM-RR was also statistically significant ($p < 0.001$) compared to other LDM methods.

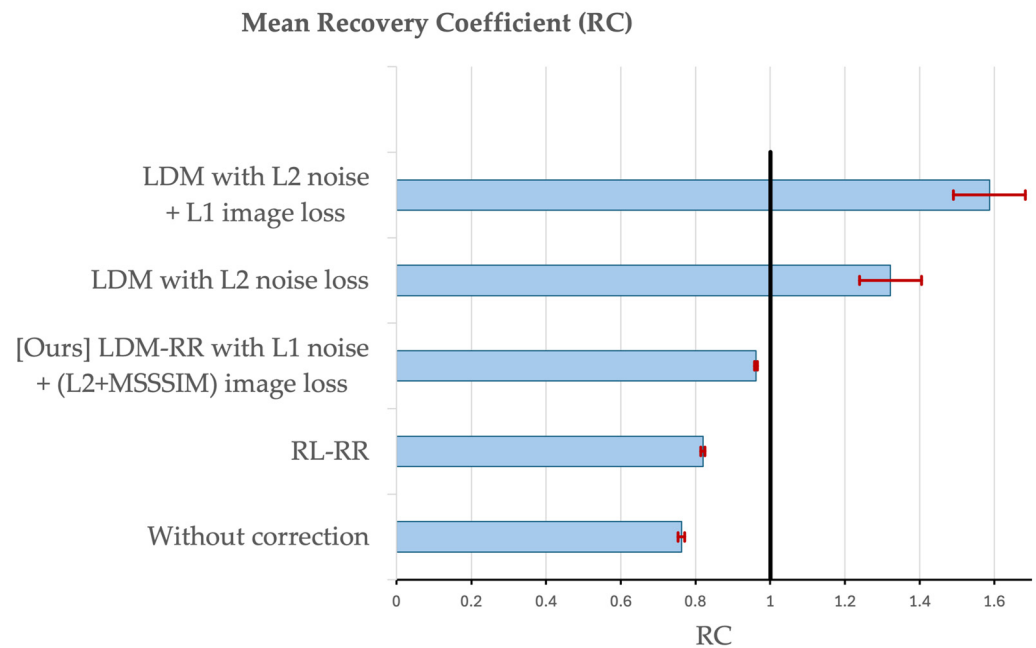


Figure 5. Comparison of mean recovery coefficient (RC) using different methods on a held-out test of 338 samples randomly selected from the simulated dataset. A value closer to 1 indicates high performance.

3.3. Evaluation on Real Longitudinal Amyloid PET Data

Table 2 shows a comparison of statistical power to detect amyloid accumulation in longitudinal studies using the LDM and RL methods for resolution recovery in comparison to measurements from raw FBPs without any correction. The annualized rate of amyloid accumulation was significantly greater than zero for all three methods ($p < 0.0001$), suggesting an increase in brain amyloid burden over time as expected. Notice the annualized rate of amyloid accumulation is the unit of SUVR/year, which is specific to the underlying quantification methods and not directly comparable. Numerically, the p -value was smallest using the LDM-RR and largest without any correction, suggesting our proposed method had the best power in detecting longitudinal changes. Additionally, the LDM-RR required a much smaller sample size estimate to detect a 25% reduction in the amyloid accumulation rate due to treatment in hypothetical anti-amyloid trials. To put the sample size estimation into context, the recently completed TRAILBLAZER-ALZ2 randomized trial of donanemab [66] recruited 860 participants for the treatment arm and 876 for the placebo arm. The donanemab was able to reduce the patient's brain amyloid burden by over 80%. Therefore, our assumed treatment effect is considerably more moderate.

Table 2. Statistical power in detecting longitudinal changes measured by mean, standard deviation, and p -value of an annualized rate of amyloid accumulation and sample size (SS) per arm estimates detecting a 25% reduction in amyloid accumulation rate due to treatment (80% power and a two-tailed type-I error of $p = 0.05$).

Annualized Rate	Raw	RL-RR	LDM-RR
Mean	0.0278	0.0377	0.0459
SD	0.0664	0.0807	0.0881
p -value	1.0×10^{-7}	5.0×10^{-9}	1.3×10^{-10}
SS	1431	1154	926

3.4. Evaluation on Real Cross-Tracer Amyloid PET Data

The performance of LDM-RR and RL-RR methods at harmonizing cross-tracer global amyloid burden measurements is shown in Table 3. Agreement of MCSUVR measurements between tracers significantly improved ($p < 0.001$), as shown by a higher correlation for both LDM and RL-based corrections to the reference measure. The improvements in LDM-RR-based partial volume corrections compared to RL were also statistically significant ($p = 0.042$). The results provide evidence supporting that the inter-tracer variability in PET-derived amyloid burden measurement is at least in part related to the partial volume effect associated with lower spatial resolution and the contaminated signal from the target regions of interest. While the numerical improvement in terms of the Pearson correlation is small, they were statistically significant, suggesting improving image resolution can be one of the strategies for reducing the variability.

Table 3. Comparison of RL and LDM-RR methods in improving the MCSUVR agreement between FBP and PIB tracers shown by Pearson correlation and Steiger’s test.

Method	Pearson Correlation	Steiger’s p -Value
Without correction	0.9163	N/A
RL-RR	0.9308	<0.0001 (RL-RR vs. without correction)
LDM-RR	0.9411	0.0001 (LDM-RR vs. without correction)
		0.0421 (LDM-RR vs. RL-RR)

4. Discussion

It is well recognized that PET imaging has inherently low spatial resolution which leads to PVE, resulting in loss of sensitivity to focal changes and compromised accuracy due to signal contamination [14,67]. Many different techniques have been developed to account for PVE and improve quantitative accuracy [14,16,68,69]. In the context of PET neuroimaging, commonly adopted techniques are often region-based and do not provide high-resolution images [50,70,71]. Voxel-wise approaches do exist [15,16,72]; however, they are known to amplify noise while having limited ability to recover the full spatial resolution or have gone through limited evaluation for targeted applications [73–75]. This study presents a new approach to improving PET quantification leveraging latent diffusion models trained using controlled simulated data. We show that diffusion models have a strong potential to enhance PET quantification through super-resolution. Our LDM-RR model’s performance on longitudinal amyloid and cross-tracer PET data demonstrates that diffusion-based super-resolution (SR) approaches can outperform traditional approaches in tackling the issue of PVE in PET imaging.

We propose an alternative to L_2 loss, which has been a de facto standard in training diffusion models. L_2 penalty pushes the model to reduce large errors, potentially sacrificing high-frequency details (Figure 6B). Moreover, L_2 loss is sensitive to the scale of voxel intensities. In super-resolution tasks, where the goal is to reconstruct fine details (voxel-level), L_2 loss may not be ideal for capturing subtle differences in high-frequency information, whereas L_1 loss may help capture the voxel-level details, which are considered to be crucial to addressing PVE. In addition, it is interesting to observe the added contribution from the multi-scale structural similarity index (MS-SSIM) metric, which confirmed the research findings of Zhao et al. [60]. Visual comparison in Figure 6 shows that using a combined loss at image and noise scales (Figure 6D), the generated high-resolution FBP images have a more accurate representation of brain structure from MRI and voxel-level uptake measurements.




Figure 6. Comparisons of the results from LDM trained using different loss functions.

Simulation data was generated to approximate the PET imaging formation process and the distribution of tracer uptake as observed in real amyloid PET images. The simulated data were used to train the LDM-RR model and evaluate its performance against the ground truth which is otherwise not possible. It should be recognized that simulated data cannot fully replicate the overall distribution and characteristics of real amyloid PET data which may introduce bias to the trained model. More sophisticated simulations can potentially be adopted to minimize this potential bias and improve model performance. Nevertheless, evaluation of model performance in real-world setting is important which we discuss further below.

We selected two commonly encountered scenarios in the investigation of Alzheimer's disease to evaluate the real-world utility and benefits of diffusion-based SR techniques for PET. In the longitudinal analysis, we leveraged data from the ADNI cohort in participants with a baseline visit and a 2-year follow-up to examine the sensitivity and statistical power of different correction methods. The participants were intentionally selected to have moderate to medium-high levels of pathological amyloid burden at baseline to maximize the probability of these participants to accumulate amyloid plaques during the follow-up period, and therefore we expect a positive increase in the MCSUVR measure and deviation from that reflects measurement noise. It is worth noting that both resolution recovery methods (RL-RR and LDM-RR) led to greater numerical values of the rate of amyloid accumulation which reflects the improved recovery coefficient as expected. In the meantime, the standard deviation of the estimated rate also increased numerically which can be a combined effect of the improved recovery coefficient and possible amplification of noise. The net effect of the correction methods is reflected by the p -values of the one-sample t -test applied to the rate data where a smaller p -value indicated a greater statistical power demonstrating a beneficial effect of correcting for PVE. The sample size estimation in hypothetical anti-amyloid treatment trials further confirmed the notion that correcting for PVE improves the longitudinal power. This improvement can lead to reduced experimental costs in longitudinal observational studies and clinical trials which will facilitate treatment development. In a clinical setting this improved power can lead to better patient management by providing more sensitive and accurate monitoring of disease progression once treatment becomes routinely available to patients. While we chose to demonstrate the capability of our proposed technique to improve the quantitative analysis of clinical amyloid PET imaging data, the same principle and method can also be applied to the analysis of preclinical animal PET data. Previous studies have developed advanced algorithms for the analysis of preclinical PET data, e.g., [76]. Super-resolution methods in general and our proposed LDM-RR technique specifically can recover the high-frequency signals lost during the image formation process by leveraging other sources of information such as MR or prior knowledge such as a template and improve the quantitative accuracy of PET-derived measurements in both preclinical and clinical applications.

The second real-world application we tested in this study is the ability of PVE correction to improve agreement between PET-derived measurements from different tracers. Using amyloid PET imaging as an example, currently, there are at least five different PET tracers that are widely used in research studies, clinical trials, and patient management to

measure amyloid burden. It is well recognized that the different tracers behave differently, leading to discrepancies in PET-derived amyloid burden measurements. At least part of this discrepancy is related to the contamination of the target measurement from nuance signals spill-in to our measurements. We demonstrated that both correction methods improved the agreement, and the LDM-RR outperformed the RL method statistically. In the meantime, we recognize the improvement has not reached the level where two tracers can be used interchangeably and may not outperform some of the other techniques we have developed [6,7]. On the other hand, this experiment demonstrated that PVE correction can be one of the strategies we can employ to improve harmonization, and a combination of multiple techniques may be the ultimate solution to fully solve the harmonization problem. We also acknowledge that there are other approaches leveraging deep learning techniques to address issues related to the standardization and harmonization of image-derived measurements, e.g., [77]. Further investigation in this direction is warranted.

This work has a few limitations: One potential limitation of our diffusion model-based framework is the computational complexity to train and validate the model on 3D imaging data. Even with a faster sampling method [78], the inference time is considerably high compared to other generative models (see Appendix B). Due to the sequential nature of the denoising process, this is a known limitation of diffusion models and remains an active area of research. Second, the model was trained on synthetic data rather than real data, which may limit the technique's performance. This is primarily due to the fact that the approach we adopted requires paired data with ground truth high-resolution images for training and validation, which is lacking. Examination of methods that are self-supervised or semi-supervised may allow us to overcome this limitation. We also like to point out that our intended application is narrowly focused on AD-related applications, while the underlying principles can be applied more broadly, although it will be beyond the scope of this paper.

5. Conclusions

We introduced a latent-diffusion-model-based resolution recovery (LDM-RR) method to enhance PET image resolution and mitigate the impact of PVE. Results demonstrate that the LDM-RR method improves spatial resolution while preserving critical amyloid and anatomical information, outperforming traditional methods like Richardson–Lucy (RL) correction. LDM-RR model showed superior performance at reconstructing high-resolution PET images, improved statistical power for detecting longitudinal amyloid accumulation and a strong potential to improve the agreement between measurements obtained from different PET tracers, contributing to better data harmonization across multi-center studies. These findings suggest that diffusion-based super-resolution (SR) techniques offer a promising alternative to conventional PVC methods by overcoming noise amplification issues and achieving better image fidelity.

Author Contributions: Conceptualization, J.S. (Jay Shah), T.W. and Y.S.; methodology, J.S. (Jay Shah), T.W. and Y.S.; software, J.S. (Jay Shah); validation, J.S. (Jay Shah) and Y.S.; formal analysis, J.S. (Jay Shah) and Y.S.; investigation, J.S. (Jay Shah) and Y.S.; data curation, J.S. (Javad Sohankar), J.L. and Y.S.; writing—original draft preparation, J.S. (Jay Shah), T.W. and Y.S.; writing—review and editing, J.S. (Jay Shah), Y.C., T.W. and Y.S.; visualization, J.S. (Jay Shah); supervision, T.W., B.L. and Y.S.; project administration, Y.S.; funding acquisition, Y.S. All authors have read and agreed to the published version of the manuscript.

Funding: This work is supported by the National Institute on Aging (NIA) grant RF1AG073424, P30AG072980, the Arizona Department of Health Services (ADHS) and the state of Arizona (ADHS Grant No. CTR057001).

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: This study did not involve the collection of new human subject data. The research utilized existing datasets (ADNI, OASIS-3, and Centiloid), which were collected and made available by their respective organizations. All participants whose data are included in these

datasets provided informed consent as part of the original studies. Detailed information about the consent process can be found in the respective documentation provided by ADNI (www.adni.loni.usc.edu), OASIS-3 (www.oasis-brains.org), and the Centiloid Project (www.gaain.org/centiloid-project).

Data Availability Statement: This study utilized datasets obtained from existing sources including ADNI, OASIS-3, and the Centiloid Project. ADNI data used in this study is available upon registration and compliance with data use agreements through the ADNI database (www.adni.loni.usc.edu). OASIS-3 data was obtained from the OASIS-3 database, accessible upon registration and approval through the OASIS project website (www.oasis-brains.org). The Centiloid data is available for qualified researchers upon request through the Centiloid Project (www.gaain.org/centiloid-project).

Acknowledgments: Data used in the preparation of this article were obtained from the Alzheimer’s Disease Neuroimaging Initiative (ADNI) database (adni.loni.usc.edu). As such, the investigators within the ADNI contributed to the design and implementation of ADNI and/or provided data but did not participate in the analysis or writing of this report. A complete listing of ADNI investigators can be found at: http://adni.loni.usc.edu/wp-content/uploads/how_to_apply/ADNI_Acknowledgement_List.pdf (accessed on 15 October 2023). Data collection and sharing for the ADNI is funded by the National Institute on Aging (National Institutes of Health Grant U19AG024904). The grantee organization is the Northern California Institute for Research and Education. In the past, ADNI has also received funding from the National Institute of Biomedical Imaging and Bioengineering, the Canadian Institutes of Health Research, and private sector contributions through the Foundation for the National Institutes of Health (FNIH) including generous contributions from the following: AbbVie, Alzheimer’s Association; Alzheimer’s Drug Discovery Foundation; Araclon Biotech; BioClinica, Inc.; Biogen; BristolMyers Squibb Company; CereSpir, Inc.; Cogstate; Eisai Inc.; Elan Pharmaceuticals, Inc.; Eli Lilly and Company; EuroImmun; F. Hoffmann-La Roche Ltd. and its affiliated company Genentech, Inc.; Fujirebio; GE Healthcare; IXICO Ltd.; Janssen Alzheimer Immunotherapy Research & Development, LLC.; Johnson & Johnson Pharmaceutical Research & Development LLC.; Lumosity; Lundbeck; Merck & Co., Inc.; Meso Scale Diagnostics, LLC.; NeuroRx Research; Neurotrack Technologies; Novartis Pharmaceuticals Corporation; Pfizer Inc.; Piramal Imaging; Servier; Takeda Pharmaceutical Company; and Transition Therapeutics. Data were provided in part by OASIS-3: Longitudinal Multimodal Neuroimaging: Principal Investigators: T. Benzinger, D. Marcus, J. Morris; NIH P30 AG066444, P50 AG00561, P30 NS09857781, P01 AG026276, P01 AG003991, R01 AG043434, UL1 TR000448, R01 EB009352. AV-45 doses were provided by Avid Radiopharmaceuticals, a wholly owned subsidiary of Eli Lilly. We would like to thank the study participants for their time and effort in participating in the OASIS and Centiloid projects. We also thank Arizona State University Research Computing [79] for hosting and maintaining our computing resources.

Conflicts of Interest: The authors declare no conflict of interest.

Abbreviations

A β	Amyloid- β
AD	Alzheimer’s disease
ADNI	Alzheimer’s Disease Neuroimaging Initiative
CT	Computed tomography
DP	Digital phantom
FBP	Florbetapir
KL	Kullback–Leibler
L ₁	L1 norm loss function, also known as least absolute errors
L ₂	L2 norm loss function, also known as least squares error
LDM	Latent diffusion model
LDM-RR	Latent diffusion model for resolution recovery
MCSUVR	Mean cortical standard uptake value ratio
MRI	Magnetic resonance imaging
MS-SSIM	Multi-scale structural similarity index
OASIS-3	Open Access Series of Imaging Studies-3
PET	Positron emission tomography
PIB	Pittsburgh compound B

PVC	Partial volume correction
PVE	Partial volume effect
RC	Recovery coefficient
RL	Richardson–Lucy
RL-RR	Richardson–Lucy for resolution recovery
simDP	Simulated digital phantom
simFBP	Simulated florbetapir
SR	Super-resolution
SS	Sample size

Appendix A

Substituting z and \hat{z}_0 from Equations (5) and (7), $loss_1$ can be simplified as:

$$\begin{aligned} loss_1 &= (1 - \alpha) \left\| \frac{z_t - \sqrt{1 - \alpha_t} \epsilon}{\sqrt{\alpha_t}} - \frac{z_t - \sqrt{1 - \alpha_t} \hat{\epsilon}}{\sqrt{\alpha_t}} \right\|_2^2 + \alpha MSSSIM(z_0, \hat{z}_0) \\ &= (1 - \alpha) \left(\frac{z_t - \sqrt{1 - \alpha_t}}{\sqrt{\alpha_t}} \right) \|\epsilon - \hat{\epsilon}\|_2^2 + \alpha MSSSIM(z_0, \hat{z}_0) \\ &= \gamma(1 - \alpha) \|\epsilon - \hat{\epsilon}\|_2^2 + \alpha MSSSIM(z_0, \hat{z}_0) \end{aligned}$$

and substituting $loss_1$ where $\gamma = \frac{z_t - \sqrt{1 - \alpha_t}}{\sqrt{\alpha_t}}$ in the combined loss equation, we obtain:

$$loss_{combined} = L_1(\epsilon, \hat{\epsilon}) + \gamma(1 - \alpha)L_2(\epsilon, \hat{\epsilon}) + \alpha MSSSIM(z_0, \hat{z}_0)$$

Appendix B

Compression Models: Our compression models had a 3-layer 3D AutoencoderKL architecture inspired from Pinaya et al. [40]. We removed all attention layers except at the model's last level. The models had 32 base channels, with a channel multiplier of [1,2,2] and only one residual block per level. Our latent space had a dimensionality of $16 \times 16 \times 16$ with 3 latent channels. We trained our model over 80 epochs with a minibatch of 60, an Adam optimizer, and a base learning rate of 0.0001. We used a patch-based discriminator in our adversarial loss with 32 base channels and a learning rate of 0.0001.

Diffusion Model: Our diffusion model uses the U-net architecture from [1], with 32 base channels, a channel multiplier of [1,2,2], and 1 residual block per level. The input channels are 9 (3 channels each for simFBP, simDP and MRI latents). We used Adam optimizer with a base learning rate of 0.0001. In training, we used a DDPM scheduler with 1000 timesteps, with a linear variance schedule, from 0.0015 to 0.0195. At inference, we use a DDIM scheduler [78] with 250 timesteps, which is a faster sampling method compared to DDPM while maintaining output quality.

Computational resources: All models were trained on a single NVIDIA A100 80 GB GPU. The inference time to generate a synthetic SR FBP was ~10 min on the GPU.

References

- Chapleau, M.; Iaccarino, L.; Soleimani-Meigooni, D.; Rabinovici, G.D. The Role of Amyloid PET in Imaging Neurodegenerative Disorders: A Review. *J. Nucl. Med.* **2022**, *63*, 13S–19S. [[CrossRef](#)] [[PubMed](#)]
- Thomas, B.A.; Erlandsson, K.; Modat, M.; Thurfjell, L.; Vandenberghe, R.; Ourselin, S.; Hutton, B.F. The Importance of Appropriate Partial Volume Correction for PET Quantification in Alzheimer's Disease. *Eur. J. Nucl. Med. Mol. Imaging* **2011**, *38*, 1104–1119. [[CrossRef](#)] [[PubMed](#)]
- Hoffman, E.J.; Huang, S.C.; Phelps, M.E. Quantitation in Positron Emission Computed Tomography: 1. Effect of Object Size. *J. Comput. Assist. Tomogr.* **1979**, *3*, 299–308. [[CrossRef](#)] [[PubMed](#)]
- Joshi, A.; Koeppe, R.A.; Fessler, J.A. Reducing between Scanner Differences in Multi-Center PET Studies. *Neuroimage* **2009**, *46*, 154–159. [[CrossRef](#)]
- Klunk, W.E.; Koeppe, R.A.; Price, J.C.; Benzinger, T.; Devous, M.D.; Jagust, W.; Johnson, K.; Mathis, C.A.; Minhas, D.; Pontecorvo, M.J.; et al. The Centiloid Project: Standardizing Quantitative Amyloid Plaque Estimation by PET. *Alzheimer's Dement.* **2015**, *11*, 1–15.e4. [[CrossRef](#)] [[PubMed](#)]

6. Chen, K.; Ghisays, V.; Luo, J.; Chen, Y.; Lee, W.; Wu, T.; Reiman, E.M.; Su, Y. Harmonizing Florbetapir and PiB PET Measurements of Cortical A β Plaque Burden Using Multiple Regions-of-Interest and Machine Learning Techniques: An Alternative to the Centiloid Approach. *Alzheimer's Dement.* **2024**, *20*, 2165–2172. [[CrossRef](#)]
7. Shah, J.; Gao, F.; Li, B.; Ghisays, V.; Luo, J.; Chen, Y.; Lee, W.; Zhou, Y.; Benzinger, T.L.S.; Reiman, E.M.; et al. Deep Residual Inception Encoder-Decoder Network for Amyloid PET Harmonization. *Alzheimer's Dement.* **2022**, *18*, 2448–2457. [[CrossRef](#)]
8. Shah, J.; Chen, K.; Reiman, E.M.; Li, B.; Wu, T.; Su, Y. Transfer Learning Based Deep Encoder Decoder Network for Amyloid PET Harmonization with Small Datasets. *Alzheimer's Dement.* **2023**, *19*, e062947. [[CrossRef](#)]
9. Alessio, A.M.; Kinahan, P.E. Improved Quantitation for PET/CT Image Reconstruction with System Modeling and Anatomical Priors. *Med. Phys.* **2006**, *33*, 4095–4103. [[CrossRef](#)] [[PubMed](#)]
10. Baete, K.; Nuyts, J.; Laere, K.V.; Van Paesschen, W.; Ceysens, S.; De Ceuninck, L.; Gheysens, O.; Kelles, A.; Van den Eynden, J.; Suetens, P.; et al. Evaluation of Anatomy Based Reconstruction for Partial Volume Correction in Brain FDG-PET. *NeuroImage* **2004**, *23*, 305–317. [[CrossRef](#)]
11. Erlandsson, K.; Dickson, J.; Arridge, S.; Atkinson, D.; Ourselin, S.; Hutton, B.F. MR Imaging-Guided Partial Volume Correction of PET Data in PET/MR Imaging. *PET Clin.* **2016**, *11*, 161–177. [[CrossRef](#)]
12. Meltzer, C.C.; Leal, J.P.; Mayberg, H.S.; Wagner, H.N.; Frost, J.J. Correction of PET Data for Partial Volume Effects in Human Cerebral Cortex by MR Imaging. *J. Comput. Assist. Tomogr.* **1990**, *14*, 561–570. [[CrossRef](#)] [[PubMed](#)]
13. Müller-Gärtner, H.W.; Links, J.M.; Prince, J.L.; Bryan, R.N.; McVeigh, E.; Leal, J.P.; Davatzikos, C.; Frost, J.J. Measurement of Radiotracer Concentration in Brain Gray Matter Using Positron Emission Tomography: MRI-Based Correction for Partial Volume Effects. *J. Cereb. Blood Flow. Metab.* **1992**, *12*, 571–583. [[CrossRef](#)] [[PubMed](#)]
14. Rousset, O.G.; Ma, Y.; Evans, A.C. Correction for Partial Volume Effects in PET: Principle and Validation. *J. Nucl. Med.* **1998**, *39*, 904–911. [[PubMed](#)]
15. Shidahara, M.; Tsoumpas, C.; Hammers, A.; Boussion, N.; Visvikis, D.; Suhara, T.; Kanno, I.; Turkheimer, F.E. Functional and Structural Synergy for Resolution Recovery and Partial Volume Correction in Brain PET. *NeuroImage* **2009**, *44*, 340–348. [[CrossRef](#)] [[PubMed](#)]
16. Tohka, J.; Reilhac, A. Deconvolution-Based Partial Volume Correction in Raclopride-PET and Monte Carlo Comparison to MR-Based Method. *NeuroImage* **2008**, *39*, 1570–1584. [[CrossRef](#)] [[PubMed](#)]
17. Golla, S.S.V.; Lubberink, M.; van Berckel, B.N.M.; Lammertsma, A.A.; Boellaard, R. Partial Volume Correction of Brain PET Studies Using Iterative Deconvolution in Combination with HYPR Denoising. *EJNMMI Res.* **2017**, *7*, 36. [[CrossRef](#)]
18. Matsubara, K.; Ibaraki, M.; Kinoshita, T.; for the Alzheimer's Disease Neuroimaging Initiative. DeepPVC: Prediction of a Partial Volume-Corrected Map for Brain Positron Emission Tomography Studies via a Deep Convolutional Neural Network. *EJNMMI Phys.* **2022**, *9*, 50. [[CrossRef](#)] [[PubMed](#)]
19. Azimi, M.-S.; Kamali-Asl, A.; Ay, M.-R.; Zeraatkar, N.; Hosseini, M.-S.; Sanaat, A.; Dadgar, H.; Arabi, H. Deep Learning-Based Partial Volume Correction in Standard and Low-Dose Positron Emission Tomography-Computed Tomography Imaging. *Quant. Imaging Med. Surg.* **2024**, *14*, 2146–2164. [[CrossRef](#)]
20. Moser, B.B.; Raue, F.; Frolov, S.; Palacio, S.; Hees, J.; Dengel, A. Hitchhiker's Guide to Super-Resolution: Introduction and Recent Advances. *IEEE Trans. Pattern Anal. Mach. Intell.* **2023**, *45*, 9862–9882. [[CrossRef](#)] [[PubMed](#)]
21. Ahmad, W.; Ali, H.; Shah, Z.; Azmat, S. A New Generative Adversarial Network for Medical Images Super Resolution. *Sci. Rep.* **2022**, *12*, 9533. [[CrossRef](#)]
22. Dai, D.; Wang, Y.; Chen, Y.; Van Gool, L. Is Image Super-Resolution Helpful for Other Vision Tasks? In Proceedings of the 2016 IEEE Winter Conference on Applications of Computer Vision (WACV), Lake Placid, NY, USA, 7–10 March 2016; pp. 1–9.
23. Li, H.; Yang, Y.; Chang, M.; Chen, S.; Feng, H.; Xu, Z.; Li, Q.; Chen, Y. SRDiff: Single Image Super-Resolution with Diffusion Probabilistic Models. *Neurocomputing* **2022**, *479*, 47–59. [[CrossRef](#)]
24. Haris, M.; Shakhnarovich, G.; Ukita, N. Task-Driven Super Resolution: Object Detection in Low-Resolution Images. In Proceedings of the Neural Information Processing; Mantoro, T., Lee, M., Ayu, M.A., Wong, K.W., Hidayanto, A.N., Eds.; Springer International Publishing: Cham, Switzerland, 2021; pp. 387–395.
25. Rasti, P.; Uiboupin, T.; Escalera, S.; Anbarjafari, G. Convolutional Neural Network Super Resolution for Face Recognition in Surveillance Monitoring. In Proceedings of the Articulated Motion and Deformable Objects; Perales, F.J., Kittler, J., Eds.; Springer International Publishing: Cham, Switzerland, 2016; pp. 175–184.
26. Li, Y.; Sixou, B.; Peyrin, F. A Review of the Deep Learning Methods for Medical Images Super Resolution Problems. *IRBM* **2021**, *42*, 120–133. [[CrossRef](#)]
27. Greenspan, H. Super-Resolution in Medical Imaging. *Comput. J.* **2009**, *52*, 43–63. [[CrossRef](#)]
28. Isaac, J.S.; Kulkarni, R. Super Resolution Techniques for Medical Image Processing. In Proceedings of the 2015 International Conference on Technologies for Sustainable Development (ICTSD), Mumbai, India, 4–6 February 2015; pp. 1–6.
29. Sun, J.; Xu, Z.; Shum, H.-Y. Image Super-Resolution Using Gradient Profile Prior. In Proceedings of the 2008 IEEE Conference on Computer Vision and Pattern Recognition, Anchorage, AK, USA, 23–28 June 2008; pp. 1–8.
30. Chang, H.; Yeung, D.-Y.; Xiong, Y. Super-Resolution through Neighbor Embedding. In Proceedings of the 2004 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2004, Washington, DC, USA, 27 June–2 July 2004; Volume 1, p. I.

31. Yang, J.; Wright, J.; Huang, T.S.; Ma, Y. Image Super-Resolution Via Sparse Representation. *IEEE Trans. Image Process.* **2010**, *19*, 2861–2873. [[CrossRef](#)] [[PubMed](#)]
32. Wang, Z.; Chen, J.; Hoi, S.C.H. Deep Learning for Image Super-Resolution: A Survey. *IEEE Trans. Pattern Anal. Mach. Intell.* **2021**, *43*, 3365–3387. [[CrossRef](#)] [[PubMed](#)]
33. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; Bengio, Y. Generative Adversarial Networks. *Commun. ACM* **2020**, *63*, 139–144. [[CrossRef](#)]
34. Rezende, D.; Mohamed, S. Variational Inference with Normalizing Flows. In Proceedings of the 32nd International Conference on Machine Learning, PMLR, Lille, France, 6–11 June 2015; pp. 1530–1538.
35. Ho, J.; Jain, A.; Abbeel, P. Denoising Diffusion Probabilistic Models. In Proceedings of the 34th International Conference on Neural Information Processing Systems, Online, 6–12 December 2020; Curran Associates Inc.: Red Hook, NY, USA, 2020; pp. 6840–6851.
36. Li, X.; Jiang, Y.; Rodriguez-Andina, J.J.; Luo, H.; Yin, S.; Kaynak, O. When Medical Images Meet Generative Adversarial Network: Recent Development and Research Opportunities. *Discov. Artif. Intell.* **2021**, *1*, 5. [[CrossRef](#)]
37. Zhang, Q.; Chen, Y. Diffusion Normalizing Flow. In Proceedings of the Advances in Neural Information Processing Systems; Curran Associates, Inc.: Red Hook, NY, USA, 2021; Volume 34, pp. 16280–16291.
38. Khader, F.; Müller-Franzes, G.; Tayebi Arasteh, S.; Han, T.; Haarbuerger, C.; Schulze-Hagen, M.; Schad, P.; Engelhardt, S.; Baeßler, B.; Foersch, S.; et al. Denoising Diffusion Probabilistic Models for 3D Medical Image Generation. *Sci. Rep.* **2023**, *13*, 7303. [[CrossRef](#)] [[PubMed](#)]
39. Müller-Franzes, G.; Niehues, J.M.; Khader, F.; Arasteh, S.T.; Haarbuerger, C.; Kuhl, C.; Wang, T.; Han, T.; Nolte, T.; Nebelung, S.; et al. A Multimodal Comparison of Latent Denoising Diffusion Probabilistic Models and Generative Adversarial Networks for Medical Image Synthesis. *Sci. Rep.* **2023**, *13*, 12098. [[CrossRef](#)] [[PubMed](#)]
40. Pinaya, W.H.L.; Tudosiu, P.-D.; Dafflon, J.; Da Costa, P.F.; Fernandez, V.; Nachev, P.; Ourselin, S.; Cardoso, M.J. Brain Imaging Generation with Latent Diffusion Models. In Proceedings of the Deep Generative Models; Mukhopadhyay, A., Oksuz, I., Engelhardt, S., Zhu, D., Yuan, Y., Eds.; Springer Nature Switzerland: Cham, Switzerland, 2022; pp. 117–126.
41. D’Amico, S.; Dall’Olio, D.; Sala, C.; Dall’Olio, L.; Sauta, E.; Zampini, M.; Asti, G.; Lanino, L.; Maggioni, G.; Campagna, A.; et al. Synthetic Data Generation by Artificial Intelligence to Accelerate Research and Precision Medicine in Hematology. *JCO Clin. Cancer Inform.* **2023**, *7*, e2300021. [[CrossRef](#)] [[PubMed](#)]
42. Rajotte, J.-F.; Bergen, R.; Buckeridge, D.L.; Emam, K.E.; Ng, R.; Strome, E. Synthetic Data as an Enabler for Machine Learning Applications in Medicine. *iScience* **2022**, *25*, 105331. [[CrossRef](#)] [[PubMed](#)]
43. Thambawita, V.; Salehi, P.; Sheshkal, S.A.; Hicks, S.A.; Hammer, H.L.; Parasa, S.; Lange, T.d.; Halvorsen, P.; Riegler, M.A. SinGAN-Seg: Synthetic Training Data Generation for Medical Image Segmentation. *PLoS ONE* **2022**, *17*, e0267976. [[CrossRef](#)] [[PubMed](#)]
44. Gao, C.; Killeen, B.D.; Hu, Y.; Grupp, R.B.; Taylor, R.H.; Armand, M.; Unberath, M. Synthetic Data Accelerates the Development of Generalizable Learning-Based Algorithms for X-Ray Image Analysis. *Nat. Mach. Intell.* **2023**, *5*, 294–308. [[CrossRef](#)]
45. Wang, H.; Zhou, Y.; Zhang, J.; Lei, J.; Sun, D.; Xu, F.; Xu, X. Anomaly Segmentation in Retinal Images with Poisson-Blending Data Augmentation. *Med. Image Anal.* **2022**, *81*, 102534. [[CrossRef](#)]
46. Lyu, F.; Ye, M.; Carlsen, J.F.; Erleben, K.; Darkner, S.; Yuen, P.C. Pseudo-Label Guided Image Synthesis for Semi-Supervised COVID-19 Pneumonia Infection Segmentation. *IEEE Trans. Med. Imaging* **2023**, *42*, 797–809. [[CrossRef](#)] [[PubMed](#)]
47. Weiner, M.W.; Veitch, D.P.; Aisen, P.S.; Beckett, L.A.; Cairns, N.J.; Green, R.C.; Harvey, D.; Jack, C.R., Jr.; Jagust, W.; Morris, J.C.; et al. The Alzheimer’s Disease Neuroimaging Initiative 3: Continued Innovation for Clinical Trial Improvement. *Alzheimer’s Dement.* **2017**, *13*, 561–571. [[CrossRef](#)] [[PubMed](#)]
48. OASIS-3: Longitudinal Neuroimaging, Clinical, and Cognitive Dataset for Normal Aging and Alzheimer Disease | medRxiv. Available online: <https://www.medrxiv.org/content/10.1101/2019.12.13.19014902v1> (accessed on 19 September 2024).
49. Navitsky, M.; Joshi, A.D.; Kennedy, I.; Klunk, W.E.; Rowe, C.C.; Wong, D.F.; Pontecorvo, M.J.; Mintun, M.A.; Devous, M.D., Sr. Standardization of Amyloid Quantitation with Flortbetapir Standardized Uptake Value Ratios to the Centiloid Scale. *Alzheimer’s Dement.* **2018**, *14*, 1565–1571. [[CrossRef](#)] [[PubMed](#)]
50. Su, Y.; Blazey, T.M.; Snyder, A.Z.; Raichle, M.E.; Marcus, D.S.; Ances, B.M.; Bateman, R.J.; Cairns, N.J.; Aldea, P.; Cash, L.; et al. Partial Volume Correction in Quantitative Amyloid Imaging. *NeuroImage* **2015**, *107*, 55–64. [[CrossRef](#)]
51. Royse, S.K.; Minhas, D.S.; Lopresti, B.J.; Murphy, A.; Ward, T.; Koeppe, R.A.; Bullich, S.; DeSanti, S.; Jagust, W.J.; Landau, S.M. Validation of Amyloid PET Positivity Thresholds in Centiloids: A Multisite PET Study Approach. *Alz Res. Ther.* **2021**, *13*, 99. [[CrossRef](#)] [[PubMed](#)]
52. Fischl, B. FreeSurfer. *NeuroImage* **2012**, *62*, 774–781. [[CrossRef](#)]
53. Su, Y.; D’Angelo, G.M.; Vlassenko, A.G.; Zhou, G.; Snyder, A.Z.; Marcus, D.S.; Blazey, T.M.; Christensen, J.J.; Vora, S.; Morris, J.C.; et al. Quantitative Analysis of PiB-PET with FreeSurfer ROIs. *PLoS ONE* **2013**, *8*, e73377. [[CrossRef](#)] [[PubMed](#)]
54. Lucy, L.B. An Iterative Technique for the Rectification of Observed Distributions. *Astron. J.* **1974**, *79*, 745. [[CrossRef](#)]
55. Biggs, D.S.C.; Andrews, M. Acceleration of Iterative Image Restoration Algorithms. *Appl. Opt.* **1997**, *36*, 1766–1775. [[CrossRef](#)]

56. Vandendriessche, D.; Uribe, J.; Bertin, H.; De Geeter, F. Performance Characteristics of Silicon Photomultiplier Based 15-Cm AFOV TOF PET/CT. *EJNMMI Phys.* **2019**, *6*, 8. [[CrossRef](#)] [[PubMed](#)]
57. Reynés-Llompart, G.; Gámez-Cenzano, C.; Romero-Zayas, I.; Rodríguez-Bel, L.; Vercher-Conejero, J.L.; Martí-Climent, J.M. Performance Characteristics of the Whole-Body Discovery IQ PET/CT System. *J. Nucl. Med.* **2017**, *58*, 1155–1161. [[CrossRef](#)] [[PubMed](#)]
58. Saharia, C.; Chan, W.; Saxena, S.; Lit, L.; Whang, J.; Denton, E.; Ghasemipour, S.K.S.; Ayan, B.K.; Mahdavi, S.S.; Gontijo-Lopes, R.; et al. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding. In Proceedings of the 36th International Conference on Neural Information Processing Systems, Orleans, LA, USA, 28 November–9 December 2022; Curran Associates Inc.: Red Hook, NY, USA, 2024; pp. 36479–36494.
59. Rombach, R.; Blattmann, A.; Lorenz, D.; Esser, P.; Ommer, B. High-Resolution Image Synthesis with Latent Diffusion Models. In Proceedings of the 2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR), New Orleans, LA, USA, 18–24 June 2022; pp. 10674–10685.
60. Zhao, H.; Gallo, O.; Frosio, I.; Kautz, J. Loss Functions for Image Restoration With Neural Networks. *IEEE Trans. Comput. Imaging* **2017**, *3*, 47–57. [[CrossRef](#)]
61. Zou, H.; Hastie, T. Regularization and Variable Selection Via the Elastic Net. *J. R. Stat. Soc. Ser. B Stat. Methodol.* **2005**, *67*, 301–320. [[CrossRef](#)]
62. López-González, F.J.; Silva-Rodríguez, J.; Paredes-Pacheco, J.; Niñerola-Baizán, A.; Efthimiou, N.; Martín-Martín, C.; Moscoso, A.; Ruibal, Á.; Roé-Vellvé, N.; Aguiar, P. Intensity Normalization Methods in Brain FDG-PET Quantification. *NeuroImage* **2020**, *222*, 117229. [[CrossRef](#)] [[PubMed](#)]
63. Bollack, A.; Collij, L.E.; García, D.V.; Shekari, M.; Altomare, D.; Payoux, P.; Dubois, B.; Grau-Rivera, O.; Boada, M.; Marquié, M.; et al. Investigating Reliable Amyloid Accumulation in Centiloids: Results from the AMYPAD Prognostic and Natural History Study. *Alzheimer's Dement.* **2024**, *20*, 3429–3441. [[CrossRef](#)] [[PubMed](#)]
64. Su, Y.; Blazey, T.M.; Owen, C.J.; Christensen, J.J.; Friedrichsen, K.; Joseph-Mathurin, N.; Wang, Q.; Hornbeck, R.C.; Ances, B.M.; Snyder, A.Z.; et al. Quantitative Amyloid Imaging in Autosomal Dominant Alzheimer's Disease: Results from the DIAN Study Group. *PLoS ONE* **2016**, *11*, e0152082. [[CrossRef](#)] [[PubMed](#)]
65. Chen, K.; Roontiva, A.; Thiyyagura, P.; Lee, W.; Liu, X.; Ayutyanont, N.; Protas, H.; Luo, J.L.; Bauer, R.; Reschke, C.; et al. Improved Power for Characterizing Longitudinal Amyloid- β PET Changes and Evaluating Amyloid-Modifying Treatments with a Cerebral White Matter Reference Region. *J. Nucl. Med.* **2015**, *56*, 560–566. [[CrossRef](#)] [[PubMed](#)]
66. Sims, J.R.; Zimmer, J.A.; Evans, C.D.; Lu, M.; Ardayfio, P.; Sparks, J.; Wessels, A.M.; Shcherbinin, S.; Wang, H.; Monkul Nery, E.S.; et al. Donanemab in Early Symptomatic Alzheimer Disease: The TRAILBLAZER-ALZ 2 Randomized Clinical Trial. *JAMA* **2023**, *330*, 512–527. [[CrossRef](#)] [[PubMed](#)]
67. Aston, J.A.D.; Cunningham, V.J.; Asselin, M.-C.; Hammers, A.; Evans, A.C.; Gunn, R.N. Positron Emission Tomography Partial Volume Correction: Estimation and Algorithms. *J. Cereb. Blood Flow. Metab.* **2002**, *22*, 1019–1034. [[CrossRef](#)]
68. Meltzer, C.C.; Kinahan, P.E.; Greer, P.J.; Nichols, T.E.; Comtat, C.; Cantwell, M.N.; Lin, M.P.; Price, J.C. Comparative Evaluation of MR-Based Partial-Volume Correction Schemes for PET. *J. Nucl. Med.* **1999**, *40*, 2053–2065. [[PubMed](#)]
69. Teo, B.-K.; Seo, Y.; Bacharach, S.L.; Carrasquillo, J.A.; Libutti, S.K.; Shukla, H.; Hasegawa, B.H.; Hawkins, R.A.; Franc, B.L. Partial-Volume Correction in PET: Validation of an Iterative Postreconstruction Method with Phantom and Patient Data. *J. Nucl. Med.* **2007**, *48*, 802–810. [[CrossRef](#)]
70. Frouin, V.; Comtat, C.; Reilhac, A.; Grégoire, M.-C. Correction of Partial-Volume Effect for PET Striatal Imaging: Fast Implementation and Study of Robustness. *J. Nucl. Med.* **2002**, *43*, 1715–1726. [[PubMed](#)]
71. Rousset, O.G.; Collins, D.L.; Rahmim, A.; Wong, D.F. Design and Implementation of an Automated Partial Volume Correction in PET: Application to Dopamine Receptor Quantification in the Normal Human Striatum. *J. Nucl. Med.* **2008**, *49*, 1097–1106. [[CrossRef](#)] [[PubMed](#)]
72. Boussion, N.; Hatt, M.; Lamare, F.; Bizais, Y.; Turzo, A.; Rest, C.C.-L.; Visvikis, D. A Multiresolution Image Based Approach for Correction of Partial Volume Effects in Emission Tomography. *Phys. Med. Biol.* **2006**, *51*, 1857. [[CrossRef](#)]
73. Thomas, B.A.; Cuplov, V.; Bousse, A.; Mendes, A.; Thielemans, K.; Hutton, B.F.; Erlandsson, K. PETPVC: A Toolbox for Performing Partial Volume Correction Techniques in Positron Emission Tomography. *Phys. Med. Biol.* **2016**, *61*, 7975–7993. [[CrossRef](#)] [[PubMed](#)]
74. Gonzalez-Escamilla, G.; Lange, C.; Teipel, S.; Buchert, R.; Grothe, M.J. PETPVE12: An SPM Toolbox for Partial Volume Effects Correction in Brain PET—Application to Amyloid Imaging with AV45-PET. *NeuroImage* **2017**, *147*, 669–677. [[CrossRef](#)] [[PubMed](#)]
75. Baker, S.L.; Maass, A.; Jagust, W.J. Considerations and Code for Partial Volume Correcting [18F]-AV-1451 Tau PET Data. *Data Brief* **2017**, *15*, 648–657. [[CrossRef](#)] [[PubMed](#)]
76. Giaccone, P.; Benfante, V.; Stefano, A.; Cammarata, F.P.; Russo, G.; Comelli, A. PET Images Atlas-Based Segmentation Performed in Native and in Template Space: A Radiomics Repeatability Study in Mouse Models. In Proceedings of the Image Analysis and Processing. ICIAP 2022 Workshops, Lecce, Italy, 23–27 May 2022; Mazzeo, P.L., Frontoni, E., Sclaroff, S., Distanto, C., Eds.; Springer International Publishing: Cham, Switzerland, 2022; pp. 351–361.
77. Kanaujia, V.K.; Kumar, A.; Yadav, S.P. Advancements in Automatic Kidney Segmentation Using Deep Learning Frameworks and Volumetric Segmentation Techniques for CT Imaging: A Review. *Arch. Comput. Methods Eng.* **2024**, *31*, 3151–3169. [[CrossRef](#)]

78. Song, J.; Meng, C.; Ermon, S. Denoising Diffusion Implicit Models. In Proceedings of the International Conference on Learning Representations, Vienna, Austria, 4 May 2021.
79. Jennewein, D.M.; Lee, J.; Kurtz, C.; Dizon, W.; Shaeffer, I.; Chapman, A.; Chiquete, A.; Burks, J.; Carlson, A.; Mason, N.; et al. The Sol Supercomputer at Arizona State University. In Proceedings of the Practice and Experience in Advanced Research Computing 2023: Computing for the Common Good, Portland, OR, USA, 23–27 July 2023; Association for Computing Machinery: New York, NY, USA, 2023; pp. 296–301.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.