Evaluating Synthetic Diffusion MRI Maps created with Diffusion Denoising Probabilistic Models

Tamoghna Chattopadhyay[a], Saket S. Ozarkar[a], Chirag Jagad[a], Neha Ann Joshy[a], Ketaki Buwa[a], Sophia I. Thomopoulos[a], Julio E. Villalón-Reina[a], and Paul M. Thompson[a]
[a]Imaging Genetics Center, Mark and Mary Stevens Neuroimaging and Informatics Institute, Keck School of Medicine, University of Southern California, Marina del Rey, CA, United States

ABSTRACT

Generative AI models, such as Stable Diffusion, DALL-E, and MidJourney, have recently gained widespread attention as they can generate high-quality synthetic images by learning the distribution of complex, high-dimensional image data. These models are now being applied to medical and neuroimaging data, where AI-based tasks such as diagnostic classification and predictive modeling typically use deep learning methods, such as convolutional neural networks (CNNs) and vision transformers (ViTs), with interpretability enhancements. In our study, we trained latent diffusion models (LDM) and denoising diffusion probabilistic models (DDPM) specifically for generating synthetic Diffusion Tensor Imaging (DTI) maps. We developed models that generate synthetic DTI maps of mean diffusivity by training on real 3D DTI scans, and evaluating the synthetic data's realism and diversity using maximum mean discrepancy (MMD) and multi-scale structural similarity index (MS-SSIM). Additionally, we aim to assess the transfer learning capability of a 3D CNN-based Alzheimer's disease classifier, by pretraining it on sex classification on real and synthetic DTIs. Our approach efficiently produces realistic and diverse synthetic data, potentially helping to create interpretable AI-driven maps for neuroscience research and clinical diagnostics.

Keywords: Generative AI, Denoising Diffusion Model, Alzheimer's Disease, DTI-MD Map

1. INTRODUCTION

In recent years, the intersection of deep learning and neuroimaging has opened new avenues for the generation of synthetic images, which have profound implications for research and clinical practice. Neuroimaging techniques, such as magnetic resonance imaging (MRI) and and diffusion tensor imaging (DTI), can provide detailed insights into the brain's gross anatomy as well as subtle alterations in the brain's white matter microstructure, but acquiring large datasets can be challenging due to the high cost, time constraints, and the need for patient consent for their images to be used in research. Synthetic images generated through deep learning models offer a possible solution to these limitations. By creating realistic neuroimaging data, researchers can augment existing datasets, making it easier to develop and validate new diagnostic tools and treatment strategies. Deep learning has advanced the generation of synthetic neuroimaging data through the use of generative models such as Generative Adversarial Networks (GANs) [1] and Variational Autoencoders (VAEs) [2]. GANs' adversarial training process can produce highly realistic images, but they can be unstable and difficult to train. Common issues with GANs include mode collapse, where the generator produces a limited variety of images, and sensitivity to hyperparameters, which can result in inconsistent output quality. VAEs, by contrast, learn a latent representation of the data by encoding images into a latent space and then decoding them back to their original form. While VAEs are generally more stable and easier to train, they often suffer from blurry reconstructions, which limits their effectiveness in generating high-resolution medical imaging data.

To address these challenges, researchers have increasingly turned to denoising diffusion probabilistic models [3], which offer a robust alternative for generating high-quality synthetic images. Unlike GANs and VAEs, denoising diffusion models gradually transform a simple noise distribution into a complex data distribution through a series of iterative refinements. These models are trained to reverse a noising process, effectively learning to generate data by progressively denoising random noise until a coherent image emerges. This approach can produce more diverse and high-fidelity images, while also being more stable to train compared to GANs. In the context of neuroimaging, denoising diffusion models can generate synthetic images that preserve the fine anatomical details necessary for quantitative analysis but also overcome the common issues of blurriness and mode collapse associated with other generative models. Pinaya et al.

[4] developed latent diffusion models (LDM) to produce realistic brain MRI scans, conditioning the generation process on demographic factors such as age and sex. Additionally, various forms of DDPMs have been employed by other researchers to create synthetic brain MRIs [5,6,7,8]. Diffusion modeling techniques can also be applied to counterfactual image generation [9,10,11], where individual images are modified to align with a reference distribution derived from healthy controls, while preserving the unique anatomical characteristics of the subject. They have also been applied to cross-modal translation [12]. Most of these DDPM models focus on T1-weighted MRI, CT scans or PET scans. In this work, we focus on training diffusion models to generate synthetic diffusion-weighted images, specifically, DTI-MD images - a scalar map commonly generated from diffusion-weighted MRI. We assessed methods for training diffusion models under constraints of limited data, computational time, and memory resources. This involved systematically adjusting the model size, training duration, and incorporating a latent diffusion model to optimize performance. We also aim to evaluate the effectiveness of using synthetic MRIs, generated by our models, as pre-training data to improve the performance of CNN-based Alzheimer's disease (AD) classification.

2. DATA AND PREPROCESSING

Diffusion tensor imaging (DTI) is the most commonly used technique for modeling brain tissue microstructure *in vivo*. A variant of standard MRI, diffusion-weighting is used to sensitize the MRI signal to the rate of water diffusion in a range of spherical directions. The standard diffusion tensor model approximates the local diffusion process using a spatially-varying diffusion tensor, typically represented as a 3D Gaussian approximation. While the tensor model simplifies the diffusion process, making it less complex than higher-order models such as NODDI [13] and MAP-MRI [14], it remains widely used due to its compatibility with single-shell diffusion MRI (dMRI), which is faster and more convenient for clinical applications. DTI is typically summarized using four standard scalar metrics: fractional anisotropy (FA), and mean, axial, and radial diffusivity (MD, AxD, and RD). These metrics describe the shape of the diffusion tensor at each voxel, derived from the three principal eigenvalues of the tensor, which represent the rate of water diffusion in 3 principal directions (eigenvectors) at each point in the tissue. Among these measures, MD is particularly sensitive to changes in cellularity, edema, and necrosis, and is among the most sensitive measures to neurodegenerative diseases such as Alzheimer's disease and Parkinson's disease **you could cite Wendy's paper and Conor's paper** .

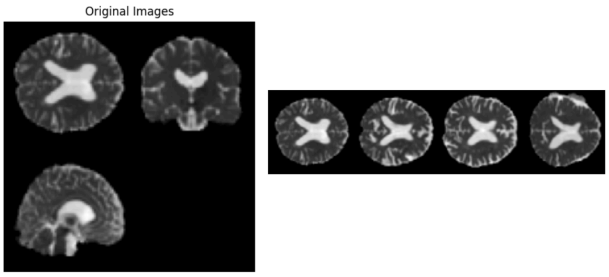

Figure 1. Left - Real DTI-MD (diffusion tensor imaging- mean diffusivity) map from the publicly available Cam-CAN dataset, in 3 planes of section. Right - Real DTI-MD maps used for training where each slice is from a different individual. This shows the diversity of the training images. Note that diffusion MRI maps typically have lower spatial resolution than standard anatomical MRI, but are highly sensitive to pathology.

The primary dataset for our experiments is the Cambridge Centre for Aging and Neuroscience (Cam-CAN) dataset [15], which we selected as it spans a larger proportion of the human lifespan than many other datasets. We used data from a total of 652 healthy control (CN) participants (mean age: 54.29 ±18.59 (SD) years; 322 F/330 M). The secondary dataset we include for downstream experiments is the widely-used, publicly available Alzheimer's Disease Neuroimaging Initiative (ADNI) dataset – a multisite study launched in 2004, designed to improve clinical trials for the prevention and treatment of Alzheimer's disease [16]. The dMRI processing pipeline we used is extensively detailed in [17,18]. All diffusion models were trained on 3D DTI-MD maps from Cam-CAN.

3. METHODS

Diffusion models are generative in nature, approximating the underlying data distribution through a distribution that is learned by iterative approximation, utilizing a trained neural network. During the training phase, noise sampled from a

normal distribution is progressively added to an image according to a noise scheduler, such as the linear scheduler, where noise is introduced at discrete steps. Models such as Diffusion Pixel-Space based [4] DDPM and Latent Diffusion Models (LDM) typically employ a modified U-Net architecture enhanced with a self-attention mechanism. The discrete time steps are also used to condition the model, integrated into the ResNet blocks within the architecture. The primary objective of the diffusion model is to predict the noise added at a random time step, which is then used to minimize a regression loss function, commonly mean squared error (MSE). Once trained, the model can generate synthetic scans through different schedulers, such as the DDPM or its deterministic variant, the Denoising Diffusion Implicit Model (DDIM) [4], with a predefined number of sampling steps. During sampling, the model iteratively transforms pure noise into an image that is drawn from a distribution that closely approximates the training distribution. The choice of sampling scheduler plays a crucial role in determining the quality of the synthetic scans and the time required to generate each scan. The conditional version of the DDPM can be created, adapting the original model to incorporate a cross-attention mechanism, enabling external conditioning through classifier-free guidance. This modification allows the model to generate images based on additional information, such as the sex, age, or diagnosis of the subject. Unlike DDPMs, LDMs utilize autoencoders to compress images into a latent space before applying standard diffusion modeling. This approach helps mitigate the computational challenges associated with training large deep learning models on 3D MRI scans, where memory requirements can be high. The autoencoder in LDMs is trained using a combination of losses, including a reconstruction loss, perceptual loss, KL divergence loss, and an adversarial loss [3,4].We trained one version of each.

As is standard in the evaluation of generative models, we quantitatively assessed our models for realism and diversity of the resulting synthetic images using performance metrics. Metrics were averaged over 500 evaluation pairs, with the Maximum Mean Discrepancy (MMD) [19] and Multi-scale Structural Similarity Index (MS-SSIM) used to measure similarity. Lower MMD values indicate closer distributions (meaning the synthetic data closely resembles the real data distribution), while higher MS-SSIM values denote greater similarity. For realism, MMD and MS-SSIM were used to compare real and synthetic images of the same class. Diversity was evaluated by comparing two unique synthetic images conditioned on the same class, with MS-SSIM also applied to real image pairs as a reference.

4. RESULTS

Table 1. Realism and Diversity Text Metrics for the DDPMs and the LDMs

Model	Image Size	Train Epochs	Scheduler	Realism: MMD (Real vs Synthetic)↓	Realism: MS-SSIM (Real vs Synthetic)↑	Diversity: MS-SSIM (Synthetic vs Synthetic)↓	Diversity: MS-SSIM (Real vs Real)↓
DDPM	32x40x32	1000	DDPM	0.0138 ± 0.0128	0.6031 ± 0.2802	0.5852 ± 0.3140	0.6197 ± 0.2553
DDPM - Conditional	64x64x64	500	DDPM	0.0283 ± 0.0191	0.3509 ± 0.2725	0.4375 ± 0.1577	0.6964 ± 0.2266
LDM	64x64x64	100 (AE) + 1000	DDPM	0.0081 ± 0.0090	0.7205 ± 0.1789	0.8597 ± 0.1413	0.6522 ± 0.2375

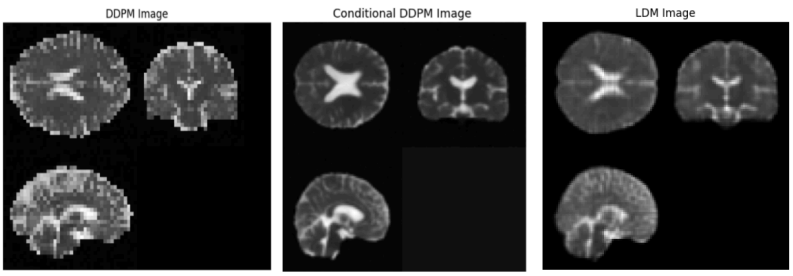

Figure 2. Synthetic DTI-MD maps generated with Left - DDPM, Middle - Conditional DDPM and Right - LDM.

5. FUTURE WORK

By the deadline for full paper submission, we plan to run additional experiments using larger image sizes and use the synthetic images generated for transfer learning. Specifically, we will pretrain three different model combinations—real images, synthetic images, and a combination of both—initially for sex classification. These pretrained models will then be applied to Alzheimer's disease classification on ADNI dataset as the downstream task. We will compare the performance of these three models using balanced accuracy and F1 score as evaluation metrics. These experiments will provide an understanding of how synthetic images may help in boosting performance of diagnostic classification models in medical imaging, by providing more data for pretraining.

ACKNOWLEDGEMENTS

This work was supported by NIH NIA grant U01 AG068057 (‘AI4AD’).

REFERENCES

[1] G. Kwon, et al. “Generation of 3D Brain MRI Using Auto-Encoding Generative Adversarial Networks,” Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics), vol. 11766 LNCS, pp. 118–126, 2019.

[2] A.V. Dalca, et al. "Unsupervised learning of probabilistic diffeomorphic registration for images and surfaces." *Medical Image Analysis* 57 (2019): 226-236.

[3] R. Rombach, et al., “High-Resolution Image Synthesis with Latent Diffusion Models,” in CVPR, 2022, pp. 10674–10685.

[4] W. H. L. Pinaya et al., “Brain Imaging Generation with Latent Diffusion Models,” in MICCAI workshop on Deep Generative Models (DGM4MICCAI), 2022, p. pp 117-126.

[5] A. Ijishakin, et al.“Interpretable Alzheimer’s Disease Classification Via a Contrastive Diffusion Autoencoder,” 2023, [Online]. Available: http://arxiv.org/abs/2306.03022.

[6] W. Peng et al., “Metadata-Conditioned Generative Models to Synthesize Anatomically-Plausible 3D Brain MRIs,” pp. 1–26, 2023, [Online]. Available: http://arxiv.org/abs/2310.04630.

[7] Z. Dorjsembe, et al. “Three-Dimensional Medical Image Synthesis with Denoising Diffusion Probabilistic Models,” in MIDL, 2022, pp. 2–4.

[8] W. Peng, et al. “Generating Realistic 3D Brain MRIs Using a Conditional Diffusion Probabilistic Model,” 2022.

[9] P. Sanchez, et al. “What is Healthy? Generative Counterfactual Diffusion for Lesion Localization,” Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics), vol. 13609 LNCS, pp. 34–44, 2022.

[10] J. Wolleb, et al. “Diffusion Models for Medical Anomaly Detection,” Lect. Notes Comput. Sci. (including Subser. Lect. Notes Artif. Intell. Lect. Notes Bioinformatics), vol. 13438 LNCS, pp. 35–45, 2022.

[11] N.J. Dhinagar, et al. 2024. Counterfactual MRI Generation with Denoising Diffusion Models for Interpretable Alzheimer's Disease Effect Detection. EMBC, 2024.

[12] Y. Li, et al.(2024). PASTA: Pathology-Aware MRI to PET Cross-Modal Translation with Diffusion Models.MICCAI.

[13] F. Deligianni, et al. (2016). NODDI and tensor-based microstructural indices as predictors of functional connectivity. *Plos one*, *11*(4), e0153404.

[14] E. Özarslan, et.al. (2013). Mean apparent propagator (MAP) MRI: a novel diffusion imaging method for mapping tissue microstructure. *NeuroImage*, *78*, 16-32.

[15] J. R. Taylor, et al. (2017). The Cambridge Centre for Ageing and Neuroscience (Cam-CAN) data repository: Structural and functional MRI, MEG, and cognitive data from a cross-sectional adult lifespan sample. *neuroimage*, *144*, 262-269.

[16] D. Veitch, et al., “Understanding disease progression and improving Alzheimer's disease clinical trials: Recent highlights from the Alzheimer's Disease Neuroimaging Initiative,” *Alzheimer’s Dement*. 15(1):106-152 (2019).

[17] S. Thomopoulos, et al., “Diffusion MRI Metrics and their relation to Dementia Severity: Effect of Harmonization Approaches,” (2021) In *17th International Symposium on Medical Information Processing and Analysis* (Vol. 12088, pp. 166-179). SPIE.
[18] T. Chattopadhyay, et al. (2023). Predicting dementia severity by merging anatomical and diffusion MRI with deep 3D convolutional neural networks. In the 18th International Symposium on Medical Information Processing and Analysis (SIPAIM; Vol. 12567, pp. 90-99). SPIE.
[19] A. Gretton, et al. , “A kernel two-sample test,” J. Mach. Learn. Res., vol. 13, pp. 723–773, 2012.