



This MICCAI paper is the Open Access version, provided by the MICCAI Society. It is identical to the accepted version, except for the format and this watermark; the final published version is available on SpringerLink.

Enhancing Spatiotemporal Disease Progression Models via Latent Diffusion and Prior Knowledge

Lemuel Puglisi¹, Daniel C. Alexander², and Daniele Ravi^{2,3,4}

¹ Dept. of Math and Computer Science, University of Catania, Italy

² Centre for Medical Image Computing, University College London, UK

³ School of Physics, Engineering and CS, University of Hertfordshire, UK

⁴ MIFT Department, University of Messina, Italy

lemuel.puglisi@phd.unict.it

Abstract. In this work, we introduce Brain Latent Progression (BrLP), a novel spatiotemporal disease progression model based on latent diffusion. BrLP is designed to predict the evolution of diseases at the individual level on 3D brain MRIs. Existing deep generative models developed for this task are primarily data-driven and face challenges in learning disease progressions. BrLP addresses these challenges by incorporating prior knowledge from disease models to enhance the accuracy of predictions. To implement this, we propose to integrate an auxiliary model that infers volumetric changes in various brain regions. Additionally, we introduce Latent Average Stabilization (LAS), a novel technique to improve spatiotemporal consistency of the predicted progression. BrLP is trained and evaluated on a large dataset comprising 11,730 T1-weighted brain MRIs from 2,805 subjects, collected from three publicly available, longitudinal Alzheimer’s Disease (AD) studies. In our experiments, we compare the MRI scans generated by BrLP with the actual follow-up MRIs available from the subjects, in both cross-sectional and longitudinal settings. BrLP demonstrates significant improvements over existing methods, with an increase of 22% in volumetric accuracy across AD-related brain regions and 43% in image similarity to the ground-truth scans. The ability of BrLP to generate conditioned 3D scans at the subject level, along with the novelty of integrating prior knowledge to enhance accuracy, represents a significant advancement in disease progression modeling, opening new avenues for precision medicine. The code of BrLP is available at the following link: <https://github.com/LemuelPuglisi/BrLP>.

Keywords: Disease Progression · Diffusion Models · Brain MRI

1 Introduction

Neurodegenerative diseases represent a global health challenge, affecting millions of people and leading to extensive morbidity and mortality. The situation is compounded by an increasingly ageing population, putting more strain on healthcare systems and society as a whole. Additionally, the progression of neurodegenerative diseases is characterized by its heterogeneous nature, with a variety of neuropathological patterns arising from different molecular subtypes [20].

In particular, these diseases affect brain regions at varying rates and through distinct mechanisms, highlighting the intricate nature of their pathophysiology [25]. Therefore, we need to develop new methods aimed at better understanding disease development, which will pave the way for more targeted and personalized treatment strategies. Initial approaches used disease progression modeling primarily based on scalar biomarkers [26,10]. Despite the crude representations of these biomarkers, such approaches have been used to enhance our understanding of diseases [4,22]. A natural evolution of these efforts is developing spatiotemporal models, which represent disease progression using rich, high-dimensional imaging biomarkers operating directly on medical scans. Unlike scalar biomarkers, these solutions facilitate the visualization and precise localization of complex patterns of structural changes, thereby offering a more detailed understanding of disease dynamics. Recent approaches have leveraged deep generative techniques, such as Variational Autoencoders (VAEs) [17], Generative Adversarial Networks (GANs) [15,23,7,28,14], and more recently, diffusion models [24], to infer the disease progression at the individual level. In particular, DaniNet [15] is a state-of-the-art model that uses adversarial learning combined with biological constraints to provide individualized predictions of brain MRIs. To mitigate memory requirements, DaniNet generates 2D slices that are then assembled into a 3D volume using a super-resolution module. Another approach is CounterSynth [14], a GAN-based counterfactual synthesis method that can simulate various conditions within a brain MRI, including ageing and disease progression. Lastly, SADM [24] is a diffusion model designed to generate longitudinal scans through autoregressive sampling by using a sequence of prior MRI scans.

The primary challenges of these methods are: 1) improving individualization by conditioning on subject-specific metadata; 2) using longitudinal scans when and if available; 3) enhancing spatiotemporal consistency to achieve a smooth progression across spatial and temporal dimensions; 4) managing the high memory demands imposed by the use of high-resolution 3D medical images [2]. Specifically, DaniNet [15] and CounterSynth [14] are not able to directly use longitudinal data if accessible. SADM [24] is not able to incorporate conditioning on subject-specific metadata and is also memory-intensive. Finally, neither CounterSynth nor SADM offer solutions to enforce spatiotemporal consistency.

In response to these challenges, we introduce BrLP, a novel spatiotemporal model, offering several key contributions: i) we propose to combine a Latent Diffusion Model (LDM) [16] and a ControlNet [27] to generate individualized brain MRIs conditioned on available subject data – addressing challenge 1; ii) we propose to integrate prior knowledge of disease progression by employing an auxiliary model designed to infer volumetric changes in different brain regions, allowing the use of longitudinal data when available – addressing challenge 2; iii) we propose LAS, a technique to improve spatiotemporal consistency in the predicted progression – addressing challenge 3; and iv) we use latent representations of brain MRIs to limit the memory demands for processing 3D scans – addressing challenge 4.

We evaluate BrLP by training it to learn progressive structural changes in the brains of individuals with different cognitive statuses: Cognitively Normal (CN), Mild Cognitive Impairment (MCI), and Alzheimer’s Disease. To do so, we use a large dataset of 11,730 T1-weighted brain MRIs from 2,805 subjects, sourced from three publicly available longitudinal studies on AD. To the best of our knowledge, we are the first to propose a 3D conditional generative model for brain MRI that incorporates prior knowledge of disease progression into the image generation process.

2 Background - Diffusion Models

A Denoising Diffusion Probabilistic Model (DDPM) [5] is a deep generative model with two Markovian processes: forward diffusion and reverse diffusion. In the forward process, Gaussian noise is incrementally added to the original image x_0 over T steps. At each step t , noise is introduced to the current image x_{t-1} by sampling from a Gaussian transition probability defined as $q(x_t | x_{t-1}) := \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t I)$, where β_t follows a variance schedule. If T is sufficiently large, x_T will converge to pure Gaussian noise $x_T \sim \mathcal{N}(0, I)$. The reverse diffusion process aims to revert each diffusion step, allowing the generation of an image from the target distribution starting from pure noise x_T . The reverse transition probability has a Gaussian closed form, $q(x_{t-1} | x_t, x_0) = \mathcal{N}(x_{t-1} | \tilde{\mu}(x_0, x_t), \tilde{\beta}_t)$, conditioned on the real image x_0 . As x_0 is not available during generation, a neural network is trained to approximate $\mu_\theta(x_t, t) \approx \tilde{\mu}(x_0, x_t)$. Following the work proposed in [5], it is possible to reparameterise the mean in terms of x_t and a noise term ϵ , and then use a neural network to predict the noise $\epsilon_\theta(x_t, t) \approx \epsilon$, optimized with the following objective:

$$\mathcal{L}_\epsilon := \mathbb{E}_{t, x_t, \epsilon \sim \mathcal{N}(0, I)} [\|\epsilon - \epsilon_\theta(x_t, t)\|^2]. \quad (1)$$

An LDM [16] extends the DDPM by applying the diffusion process to a latent representation z of the image x , rather than to the image itself. This approach reduces the high memory demand while preserving the quality and flexibility of the models. The latent representation is obtained by training an autoencoder, composed of an encoder \mathcal{E} and a decoder \mathcal{D} , such that the encoder maps the sample x to the latent space $z = \mathcal{E}(x)$, and the decoder recovers it as $x = \mathcal{D}(z)$.

3 Methods - Brain Latent Progression (BrLP)

We now introduce the architecture of BrLP, comprising four key components: an LDM, a ControlNet, an auxiliary model, and a LAS block, each described in successive paragraphs. These four components, summarized in Figure 1, collectively address the challenges outlined in the introduction. In particular, the LDM is designed to generate random 3D brain MRIs that conform to specific covariates, while ControlNet aims to specialize these MRI scans to specific anatomical structures of a subject. Additionally, the auxiliary model leverages prior knowledge of

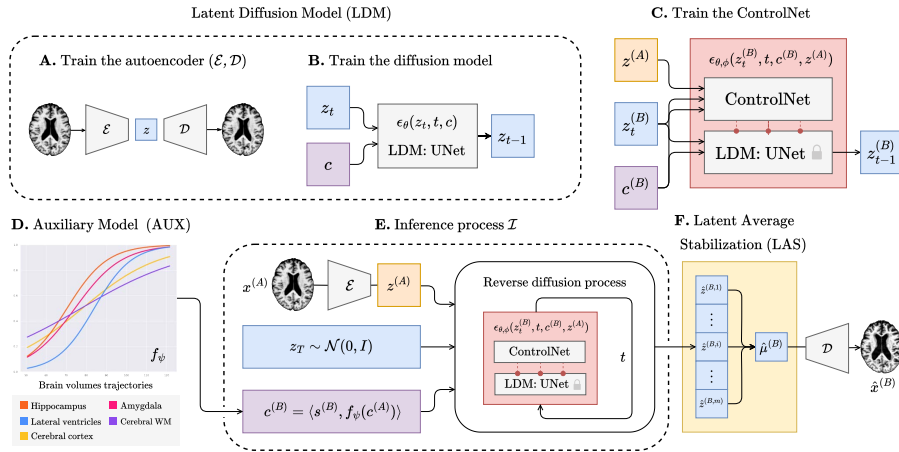


Fig. 1. The overview of BrLP training and inference processes.

disease progression to improve the precision in predicting the volumetric changes of specific brain regions. Finally, the LAS block is used during inference to improve spatiotemporal consistency. Details concerning the training process and hyperparameter settings are provided in Table 1 of the Supplementary Material.

LDM - Learning the Brain MRIs Distribution. Building upon [12], we train an LDM aimed to generate 3D brain MRIs mirroring specific covariates $c = \langle s, v \rangle$, where s includes subject-specific metadata (age, sex, and cognitive status) while v encompasses progression-related metrics such as volumes of brain regions (hippocampus, cerebral cortex, amygdala, cerebral white matter, and lateral ventricles) linked to AD progression [13]. The construction of the LDM is a two-phase process. Initially, we train an autoencoder (\mathcal{E}, \mathcal{D}) (block A in Figure 1) designed to produce a latent representation $z = \mathcal{E}(x)$ for each brain MRI x within our dataset. Subsequently, we train a conditional UNet (block B in Figure 1), represented as ϵ_θ , with network parameters θ , aimed to estimate the noise $\epsilon_\theta(z_t, t, c)$ necessary for reverting from z_t to z_{t-1} , as mentioned in Section 2. We train ϵ_θ by minimizing the loss \mathcal{L}_ϵ (Eq. 1). Covariates c are integrated into the network as conditions using a cross-attention mechanism, in line with [16]. The generation process initiates by sampling random Gaussian noise $z_T \sim \mathcal{N}(0, I)$ and then iteratively reverses each diffusion step $z_t \rightarrow z_{t-1}$ for $t = T, \dots, 1$. Decoding the output z_0 from the final step $t = 1$ yields a synthetic brain MRI $\hat{x} = \mathcal{D}(z_0)$ that follows the specified covariates c .

ControlNet - Conditioning on Subject Brain MRI. The LDM provides only a limited degree of control over the generated brain MRI via the covariates c , and it does not allow for conditioning the model on individual anatomical structures. The purpose of this block is to extend the capabilities of the LDM

to encompass this additional control. To achieve this, we use ControlNet [27], (block C in Figure 1) a neural network designed to work in conjunction with the LDM. We conceptualize ControlNet and LDM as a unified network $\epsilon_{\theta,\phi}$, where θ represents the fixed network’s parameters of the LDM and ϕ denotes the trainable network’s parameters of ControlNet. As in the LDM, $\epsilon_{\theta,\phi}$ is still used to predict the noise $\epsilon_{\theta,\phi}(z_t, t, c, z)$ in the reverse diffusion step $z_t \rightarrow z_{t-1}$, now incorporating $z = \mathcal{E}(x)$ as a condition to encompass the structure of the target brain x during the generation process. To train ControlNet, we use the latent representations $z^{(A)}$ and $z^{(B)}$ from pairs of brain MRIs of the same patient taken at different ages $A < B$. The covariates $c^{(B)}$ associated with $z^{(B)}$ are known and used as target covariates. Each training iteration involves: i) sampling $t \sim U[1, T]$, ii) performing t forward diffusion steps $z^{(B)} \rightarrow z_t^{(B)}$, iii) predicting the noise $\epsilon_{\theta,\phi}(z_t^{(B)}, t, c^{(B)}, z^{(A)})$ to revert $z_t^{(B)} \rightarrow z_{t-1}^{(B)}$, and iv) minimizing the loss \mathcal{L}_ϵ (Eq. 1).

Proposed Auxiliary Model - Leveraging Disease Prior Knowledge. AD-related regions shrink or expand over time and at different rates [13]. Deep-learning-based spatiotemporal models strive to learn these progression rates directly from brain MRIs in a black-box manner, which can be very challenging. To aid this process, we propose incorporating prior knowledge of volumetric changes directly into our pipeline. To do so, we exploit an auxiliary model f_ψ (block D in Figure 1) able to predict how the volumes of AD-related regions change over time and provide this information to the LDM via the progression-related covariates v . The choice of our auxiliary model is tailored to two scenarios, making BrLP flexible for both cross-sectional and longitudinal data. For subjects with a single scan available at age A , we employ a regression model to estimate volumetric changes $\hat{v}^{(B)} = f_\psi(c^{(A)})$ at age B . For subjects with n past visits accessible at ages A_1, \dots, A_n , we predict $\hat{v}^{(B)} = f_\psi(c^{(A_1)}, \dots, c^{(A_n)})$ using Disease Course Mapping (DCM) [18,8], a model specifically designed for disease progression. DCM is intended to provide a more accurate trajectory in alignment with the subject’s history of volumetric changes available. While we employ DCM as a potential solution, any suitable disease progression model can be used in BrLP.

Inference Process. Let $x^{(A)}$ be the input brain MRI from a subject at age A , with known subject-specific metadata $s^{(A)}$ and progression-related volumes $v^{(A)}$ measured from $x^{(A)}$. As summarized in block E from Figure 1, to infer the brain MRI $x^{(B)}$ at age $B > A$, we perform six steps: i) predict the progression-related volumes $\hat{v}^{(B)} = f_\psi(c^{(A)})$ using the auxiliary model; ii) concatenate this information with the subject-specific metadata $s^{(B)}$ to form the target covariates $c^{(B)} = \langle s^{(B)}, \hat{v}^{(B)} \rangle$; iii) compute the latent $z^{(A)} = \mathcal{E}(x^{(A)})$; iv) sample random Gaussian noise $z_T \sim \mathcal{N}(0, I)$; v) run the reverse diffusion process by predicting the noise $\epsilon_{\theta,\phi}(z_t, t, c^{(B)}, z^{(A)})$ to reverse each diffusion step for $t = T, \dots, 1$; and finally vi) employ the decoder \mathcal{D} to reconstruct the predicted brain MRI $\hat{x}^{(B)} = \mathcal{D}(z_0)$ in the imaging domain. This inference process is summarized into a compact notation $\hat{z}^{(B)} = \mathcal{I}(z_T, x^{(A)}, c^{(A)})$ and $\hat{x}^{(B)} = \mathcal{D}(\hat{z}^{(B)})$.

Enhance Inference via Proposed Latent Average Stabilization (LAS).

Variations in the initial value $x_T \sim \mathcal{N}(0, I)$ can lead to slight discrepancies in the results produced by the inference process. These discrepancies are especially noticeable when making predictions over successive timesteps, manifesting as irregular patterns or non-smooth transitions of progression. Therefore, we introduce LAS (block F in Figure 1), a technique to improve spatiotemporal consistency by averaging different results of the inference process. In particular, LAS is based on the assumption that the predictions $\hat{z}^{(B)} = \mathcal{I}(z_T, x^{(A)}, c^{(A)})$ deviate from a theoretical mean $\mu^{(B)} = \mathbb{E}[\hat{z}^{(B)}]$. To estimate the expected value $\mu^{(B)}$, we propose to repeat the inference process m times and average the results:

$$\mu^{(B)} = \mathbb{E}_{z_T \sim \mathcal{N}(0, I)} \left[\mathcal{I}(z_T, x^{(A)}, c^{(A)}) \right] \approx \frac{1}{m} \sum^m \mathcal{I}(z_T, x^{(A)}, c^{(A)}). \quad (2)$$

Similar to before, we decode the predicted scan as $\hat{x}^{(B)} = \mathcal{D}(\mu^{(B)})$. The entire inference process (with $m = 4$) requires ~ 4.8 s per MRI on a consumer GPU.

4 Experiments and Results

Data. We collect a large dataset comprising 11,730 T1-weighted brain MRI scans from 2,805 subjects across various publicly available longitudinal studies: ADNI 1/2/3/GO (1,990 subjects) [11], OASIS-3 (573 subjects) [9], and AIBL (242 subjects) [3]. Each subject has at least two MRIs, and each scan is acquired during a different visit. The average time interval between the initial and follow-up visits is 4.3 years (SD = 3.1), with a maximum span of 16 years. Age, sex, and cognitive status were available from all datasets. The average age is 74 ± 7 years, and 53% of the subjects are male. Based on the final visit, 43.8% of subjects are classified as CN, 25.7% exhibit or develop MCI, and 30.5% exhibit or develop AD. We randomly split data into a training set (80%), a validation set (5%), and a testing set (15%) with no overlapping subjects. The validation set is used for early stopping during training. Each brain MRI is pre-processed using: N4 bias-field correction [21], skull stripping [6], affine registration to the MNI space, intensity normalization [19] and resampling to 1.5 mm^3 . The volumes used as progression-related covariates and for our subsequent evaluation are calculated using SynthSeg 2.0 [1] and are expressed as percentages of the total brain volume to account for individual differences.

Evaluation Metrics. We evaluate BrLP using image-based and volumetric metrics to compare the predicted brain MRI scans with the subjects’ actual follow-up scans. In particular, the Mean Squared Error (MSE) and the Structural Similarity Index (SSIM) are used to assess image similarity between the scans. Volumetric metrics in AD-related regions (hippocampus, amygdala, lateral ventricles, cerebrospinal fluid (CSF), and thalamus) evaluate the model’s accuracy in tracking disease progression. Specifically, the Mean Absolute Error (MAE) between the volumes of actual follow-up scans and the generated brain

Table 1. Results from the ablation study and comparison with baseline methods. MAE (\pm SD) in predicted volumes is expressed as a percentage of total brain volume.

	Method	Config. (AUX)	Image-based metrics		MAE (conditional region volumes)			MAE (unconditional reg. volumes)	
			MSE \downarrow	SSIM \uparrow	Hippocampus \downarrow	Amygdala \downarrow	Lat. Ventricle \downarrow	Thalamus \downarrow	CSF \downarrow
Ablation	Base	-	0.005 \pm 0.003	0.89 \pm 0.03	0.026 \pm 0.023	0.016 \pm 0.015	0.279 \pm 0.347	0.030 \pm 0.023	0.889 \pm 0.681
	Base + AUX	LM	0.005 \pm 0.002	0.90 \pm 0.03	0.024 \pm 0.021	0.015 \pm 0.013	0.279 \pm 0.314	0.030 \pm 0.024	0.851 \pm 0.632
	Base + LAS	-	0.005 \pm 0.002	0.90 \pm 0.03	0.025 \pm 0.022	0.015 \pm 0.014	0.258 \pm 0.330	0.029 \pm 0.022	0.851 \pm 0.659
	Base + LAS + AUX	LM	0.004 \pm 0.002	0.91 \pm 0.03	0.023 \pm 0.021	0.015 \pm 0.014	0.255 \pm 0.303	0.029 \pm 0.023	0.829 \pm 0.624
Single-image (Cross-sectional)									
Comparison Study	DaniNet [15]	-	0.016 \pm 0.007	0.62 \pm 0.16	0.030 \pm 0.030	0.018 \pm 0.017	0.257 \pm 0.222	0.038 \pm 0.030	1.081 \pm 0.814
	CounterSynth [14]	-	0.010 \pm 0.004	0.82 \pm 0.05	0.030 \pm 0.018	0.014 \pm 0.010	0.310 \pm 0.311	0.127 \pm 0.035	0.881 \pm 0.672
	BrLP (Proposed)	LM	0.004 \pm 0.002	0.91 \pm 0.03	0.023 \pm 0.021	0.015 \pm 0.014	0.255 \pm 0.303	0.029 \pm 0.023	0.829 \pm 0.624
	Sequence-aware (Longitudinal)								
	Latent-SADM [24]	-	0.008 \pm 0.002	0.85 \pm 0.02	0.035 \pm 0.027	0.018 \pm 0.015	0.329 \pm 0.328	0.037 \pm 0.028	0.924 \pm 0.705
	BrLP (Proposed)	DCM	0.004 \pm 0.002	0.91 \pm 0.03	0.020 \pm 0.017	0.014 \pm 0.013	0.240 \pm 0.259	0.031 \pm 0.024	0.810 \pm 0.631

MRIs is reported in the results. Notably, CSF and thalamus are excluded from progression-related covariates, enabling the analysis of unconditioned regions in our predictions.

Ablation Study. We conduct an ablation study to assess the contributions of: i) the auxiliary model (AUX) and ii) the proposed technique for spatiotemporal consistency (LAS). The results are presented at the top of Table 1. BrLP without AUX and LAS is referred to as “base”. The experiments demonstrate that both LAS and AUX enhance performance, reducing volumetric errors by 5% and 4%, respectively. An example of improvement achieved with LAS is provided in Figure 2 of the Supplementary Material. Employing both AUX and LAS together offers the optimal setup, achieving an average reduction in volumetric error of 7%. This optimal configuration is used for comparisons against other approaches, with the only variation being the type of auxiliary model used in our pipeline.

Comparison with Baselines. We categorize existing methods into single-image (cross-sectional) and sequence-aware (longitudinal) approaches. Single-image approaches, such as DaniNet [15] and CounterSynth [14], predict progression using just one brain MRI as input. Sequence-aware methods, like SADM [24], leverage a series of prior brain MRIs as input. Due to the large memory demands of SADM, we have re-implemented it using an LDM, allowing the comparisons in our experiments. We refer to it as Latent-SADM. To evaluate all these methods, we conduct two separate experiments. In single-image methods, we predict all subsequent MRIs for a subject based on their initial scan. For sequence-aware methods, we use the first half of a subject’s MRI visits to predict all subsequent MRIs in the latter half. In single-image settings, our approach uses a Linear Model (LM) as the auxiliary model. In contrast, for sequence-aware experiments, we employ the last available MRI in the sequence as the input for BrLP and fit a logistic DCM on the first half of the subject’s visits as the auxiliary model.

Results from our experiments are presented in Table 1. We observe an average decrease of 62% (SD = 10%) in MSE and an average increase of 43% (SD =




Fig. 2. A comparison between the real progression of a 70 y.o. subject with MCI over 15 years and the predictions obtained by BrLP and the baseline methods. Each method shows a predicted MRI (left) and its deviation from the subject’s real brain MRI (right).

18%) in SSIM compared to other baselines. In terms of volumetric measurements across various brain regions, our method shows improvements of 17.55% (SD = 8.79%) over DaniNet, 23.40% (SD = 28.85%) over CounterSynth, and 24.14% (SD = 10.63%) over Latent-SADM. We did not observe any particular differences in the improvement obtained in conditioned and non-conditioned regions. Additionally, paired t-tests ($p < 0.001$) verified the statistical significance of the observed improvements. Finally, Figure 2 presents a visual comparison between the actual progression of a 70-year-old subject over 15 years and the predictions obtained by BrLP and the baseline methods. The results from Latent-SADM and DaniNet exhibit a spatiotemporal mismatch in predicting the lateral ventricles’ enlargement, whereas CounterSynth fails to capture the structural changes observed in the real progression. On the other hand, BrLP shows the most accurate prediction of the brain’s anatomical changes, confirming the previous quantitative findings. It is worth noting that we observed limitations in LAS performance in underrepresented conditions, such as ages over 90 years, resulting in slightly non-monotonic progression (see Case Study 4 in our supplementary video).

5 Conclusion

In this work, we propose BrLP, a 3D spatiotemporal model that accurately captures the progression patterns of neurodegenerative diseases by forecasting the

evolution of 3D brain MRIs at the individual level. While we have showcased the application of our pipeline on brain MRIs, BrLP holds potential for use with other imaging modalities and to model different progressive diseases. Importantly, our framework can be easily extended to integrate additional covariates, such as genetic data, providing further personalized insights into our predictions. BrLP offers promising applications for both pre-trial and post-trial analysis in clinical research. In pre-trial analysis, BrLP can be used for patient stratification, potentially identifying patients at the optimal stage of disease progression when treatment is most likely to be effective. In post-trial analysis, BrLP can be used to create digital twins of participants, simulating untreated disease progression. This enables individualized assessment of treatment effects, potentially reducing the need for control groups and addressing ethical concerns related to participants who do not receive therapeutic benefits.

Acknowledgments. Lemuel Puglisi is enrolled in a PhD program at the University of Catania, fully funded by PNRR (DM 118/2023).

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

References

1. Billot, B., Greve, D.N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A.V., Iglesias, J.E., et al.: Synthseg: Segmentation of brain mri scans of any contrast and resolution without retraining. *Medical image analysis* **86**, 102789 (2023)
2. Blumberg, S.B., Tanno, R., Kokkinos, I., Alexander, D.C.: Deeper image quality transfer: Training low-memory neural networks for 3d images. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2018: 21st International Conference, Granada, Spain, September 16–20, 2018, Proceedings, Part I*. pp. 118–125. Springer (2018)
3. Ellis, K.A., Bush, A.I., Darby, D., De Fazio, D., Foster, J., Hudson, P., Lautenschlager, N.T., Lenzo, N., Martins, R.N., Maruff, P., et al.: The australian imaging, biomarkers and lifestyle (aibl) study of aging: methodology and baseline characteristics of 1112 individuals recruited for a longitudinal study of alzheimer’s disease. *International psychogeriatrics* **21**(4), 672–687 (2009)
4. Eshaghi, A., Young, A.L., Wijeratne, P.A., Prados, F., Arnold, D.L., Narayanan, S., Guttmann, C.R., Barkhof, F., Alexander, D.C., Thompson, A.J., et al.: Identifying multiple sclerosis subtypes using unsupervised machine learning and mri data. *Nature communications* **12**(1), 2078 (2021)
5. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
6. Hoopes, A., Mora, J.S., Dalca, A.V., Fischl, B., Hoffmann, M.: Synthstrip: Skull-stripping for any brain image. *NeuroImage* **260**, 119474 (2022)
7. Jung, E., Luna, M., Park, S.H.: Conditional gan with an attention-based generator and a 3d discriminator for 3d medical image generation. In: *Medical Image Computing and Computer Assisted Intervention—MICCAI 2021: 24th International Conference, Strasbourg, France, September 27–October 1, 2021, Proceedings, Part VI* 24. pp. 318–328. Springer (2021)

8. Koval, I., Bône, A., Louis, M., Lartigue, T., Bottani, S., Marcoux, A., Samper-Gonzalez, J., Burgos, N., Charlier, B., Bertrand, A., et al.: AD course map charts alzheimer’s disease progression. *Scientific Reports* **11**(1), 8020 (2021)
9. LaMontagne, P.J., Benzinger, T.L., Morris, J.C., Keefe, S., Hornbeck, R., Xiong, C., Grant, E., Hassenstab, J., Moulder, K., Vlassenko, A.G., et al.: Oasis-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. *MedRxiv* pp. 2019–12 (2019)
10. Oxtoby, N.P., Alexander, D.C.: Imaging plus x: multimodal models of neurodegenerative disease. *Current opinion in neurology* **30**(4), 371 (2017)
11. Petersen, R.C., Aisen, P.S., Beckett, L.A., Donohue, M.C., Gamst, A.C., Harvey, D.J., Jack, C.R., Jagust, W.J., Shaw, L.M., Toga, A.W., et al.: Alzheimer’s disease neuroimaging initiative (adni): clinical characterization. *Neurology* **74**(3), 201–209 (2010)
12. Pinaya, W.H., 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: *MICCAI Workshop on Deep Generative Models*. pp. 117–126. Springer (2022)
13. Pini, L., Pievani, M., Bocchetta, M., Altomare, D., Bosco, P., Cavedo, E., Galluzzi, S., Marizzoni, M., Frisoni, G.B.: Brain atrophy in alzheimer’s disease and aging. *Ageing research reviews* **30**, 25–48 (2016)
14. Pombo, G., Gray, R., Cardoso, M.J., Ourselin, S., Rees, G., Ashburner, J., Nachev, P.: Equitable modelling of brain imaging by counterfactual augmentation with morphologically constrained 3d deep generative models. *Medical Image Analysis* **84**, 102723 (2023)
15. Ravi, D., Blumberg, S.B., Ingala, S., Barkhof, F., Alexander, D.C., Oxtoby, N.P., Initiative, A.D.N., et al.: Degenerative adversarial neuroimage nets for brain scan simulations: Application in ageing and dementia. *Medical Image Analysis* **75**, 102257 (2022)
16. Rombach, R., Blattmann, A., Lorenz, D., Esser, P., Ommer, B.: High-resolution image synthesis with latent diffusion models. In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. pp. 10684–10695 (2022)
17. Sauty, B., Durrleman, S.: Progression models for imaging data with longitudinal variational auto encoders. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 3–13. Springer (2022)
18. Schiratti, J.B., Allasonnière, S., Colliot, O., Durrleman, S.: A bayesian mixed-effects model to learn trajectories of changes from repeated manifold-valued observations. *The Journal of Machine Learning Research* **18**(1), 4840–4872 (2017)
19. Shinohara, R.T., Sweeney, E.M., Goldsmith, J., Shiee, N., Mateen, F.J., Calabresi, P.A., Jarso, S., Pham, D.L., Reich, D.S., Crainiceanu, C.M., et al.: Statistical normalization techniques for magnetic resonance imaging. *NeuroImage: Clinical* **6**, 9–19 (2014)
20. Tijms, B.M., Vromen, E.M., Mjaavatten, O., Holstege, H., Reus, L.M., van der Lee, S., Wesenhagen, K.E., Lorenzini, L., Vermunt, L., Venkatraghavan, V., et al.: Cerebrospinal fluid proteomics in patients with alzheimer’s disease reveals five molecular subtypes with distinct genetic risk profiles. *Nature Aging* pp. 1–15 (2024)
21. Tustison, N.J., Avants, B.B., Cook, P.A., Zheng, Y., Egan, A., Yushkevich, P.A., Gee, J.C.: N4itk: improved n3 bias correction. *IEEE transactions on medical imaging* **29**(6), 1310–1320 (2010)
22. Vogel, J.W., Young, A.L., Oxtoby, N.P., Smith, R., Ossenkoppele, R., Strandberg, O.T., La Joie, R., Aksman, L.M., Grothe, M.J., Iturria-Medina, Y., et al.:

- Four distinct trajectories of tau deposition identified in alzheimer’s disease. *Nature medicine* **27**(5), 871–881 (2021)
23. Xia, T., Chartsias, A., Wang, C., Tsafaris, S.A., Initiative, A.D.N., et al.: Learning to synthesise the ageing brain without longitudinal data. *Medical Image Analysis* **73**, 102169 (2021)
 24. Yoon, J.S., Zhang, C., Suk, H.I., Guo, J., Li, X.: Sadm: Sequence-aware diffusion model for longitudinal medical image generation. In: *International Conference on Information Processing in Medical Imaging*. pp. 388–400. Springer (2023)
 25. Young, A.L., Marinescu, R.V., Oxtoby, N.P., Bocchetta, M., Yong, K., Firth, N.C., Cash, D.M., Thomas, D.L., Dick, K.M., Cardoso, J., et al.: Uncovering the heterogeneity and temporal complexity of neurodegenerative diseases with subtype and stage inference. *Nature communications* **9**(1), 4273 (2018)
 26. Young, A.L., Oxtoby, N.P., Garbarino, S., Fox, N.C., Barkhof, F., Schott, J.M., Alexander, D.C.: Data-driven modelling of neurodegenerative disease progression: thinking outside the black box. *Nature Reviews Neuroscience* pp. 1–20 (2024)
 27. Zhang, L., Rao, A., Agrawala, M.: Adding conditional control to text-to-image diffusion models. In: *Proceedings of the IEEE/CVF International Conference on Computer Vision*. pp. 3836–3847 (2023)
 28. Zhao, Y., Ma, B., Jiang, P., Zeng, D., Wang, X., Li, S.: Prediction of alzheimer’s disease progression with multi-information generative adversarial network. *IEEE Journal of Biomedical and Health Informatics* **25**(3), 711–719 (2020)