




Contents lists available at ScienceDirect

Computerized Medical Imaging and Graphics

journal homepage: www.elsevier.com/locate/compmedimag

Temporally-aware diffusion model for brain progression modelling with bidirectional temporal regularisation[☆]

Mattia Litrico^{a, ID, *}, Francesco Guarnera^a, Mario Valerio Giuffrida^b, Daniele Ravi^c, Sebastiano Battiato^a

^a University of Catania, Catania, Italy

^b University of Nottingham, Nottingham, United Kingdom

^c University of Messina, Messina, Italy

ARTICLE INFO

Dataset link: <https://github.com/MattiaLitrico/TADM-3D>

Keywords:

Spatial-temporal disease progression
Brain MRI
Diffusion model

ABSTRACT

Generating realistic MRIs to accurately predict future changes in the structure of brain is an invaluable tool for clinicians in assessing clinical outcomes and analysing the disease progression at the patient level. However, current existing methods present some limitations: (i) some approaches fail to explicitly capture the relationship between structural changes and time intervals, especially when trained on age-imbalanced datasets; (ii) others rely only on scan interpolation, which lack clinical utility, as they generate intermediate images between timepoints rather than future pathological progression; and (iii) most approaches rely on 2D slice-based architectures, thereby disregarding full 3D anatomical context, which is essential for accurate longitudinal predictions. We propose a 3D Temporally-Aware Diffusion Model (TADM-3D), which accurately predicts brain progression on MRI volumes. To better model the relationship between time interval and brain changes, TADM-3D uses a pre-trained *Brain-Age Estimator* (BAE) that guides the diffusion model in the generation of MRIs that accurately reflect the expected age difference between baseline and generated follow-up scans. Additionally, to further improve the temporal awareness of TADM-3D, we propose the *Back-In-Time Regularisation* (BITR), by training TADM-3D to predict bidirectionally from the baseline to follow-up (forward), as well as from the follow-up to baseline (backward). Although predicting past scans has limited clinical applications, this regularisation helps the model generate temporally more accurate scans. We train and evaluate TADM-3D on the OASIS-3 dataset, and we validate the generalisation performance on an external test set from the NACC dataset. The code is available at <https://github.com/MattiaLitrico/TADM-3D>.

1. Introduction

Predicting how brain structures evolve over time in MRI scans is a key challenge in medical imaging. The modelling of temporal brain trajectories has proven useful in multiple applications, including recovering missing scans in longitudinal data (Ravi et al., 2019), acting as a virtual placebo (Ravi et al., 2022), aiding in patient stratification (Pombo et al., 2023; Ravi et al., 2019, 2022; Young et al., 2024), and supporting both diagnosis and prognosis of Alzheimer's disease (AD) (Bowles et al., 2018). Although AD diagnosis is often based on neuropsychological and behavioural evaluations, imaging data plays a crucial role in revealing structural brain changes associated with the disease, even at early stages (Jack et al., 2018).

Recent advancements in AI have led to the emergence of novel spatio-temporal technologies designed to model disease progression

(Liu et al., 2006), enabling more accurate predictions of structural brain changes. Among these, generative approaches have gained attention for their ability to accurately simulate future MRI scans based on earlier scans. Early methods used Generative Adversarial Networks (GANs) to generate future MRIs (Pombo et al., 2023; Ravi et al., 2022). Recently, Denoising Diffusion Probabilistic Models (DDPMs) have become the state-of-the-art for generative models, demonstrating superior performance also in this domain. Sequence-Aware Diffusion Model (SADM) (Yoon et al., 2022) combines a diffusion model with a transformer to predict follow-up MRIs. Similarly, Diffusion Deformable Model (DDM) (Kim and Ye, 2022) and DiffuseMorph (Kim et al., 2022) use diffusion models to estimate the deformation field between scans. Differently, BrLP (Puglisi et al., 2024) proposes the training of a latent diffusion model in conjunction with an autoencoder and an auxiliary model to generate individual MRI scans based on longitudinal data.

[☆] This article is part of a Special issue entitled: 'GenAI for Medical Imaging' published in Computerized Medical Imaging and Graphics.

* Corresponding author.

E-mail address: mattia.litrico@phd.unict.it (M. Litrico).

<https://doi.org/10.1016/j.compmedimag.2025.102688>

Received 18 August 2025; Received in revised form 15 December 2025; Accepted 15 December 2025

Available online 16 December 2025

0895-6111/© 2025 The Authors. Published by Elsevier Ltd. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

However, these methods often show poor performance in accurately modelling the temporal dynamics of morphological changes under neurodegenerative diseases. Some approaches (Pombo et al., 2023; Ravi et al., 2022; Xia et al., 2019a, 2021) do not explicitly capture the relationship between structural MRI changes and time intervals, while others are limited in their ability to generate future scans (Kim and Ye, 2022; Kim et al., 2022) or require longitudinal data at inference time (Yoon et al., 2022).

To address these challenges, we introduce TADM-3D, a novel 3D diffusion-based framework specifically designed to predict future brain MRI scans. Our method learns the statistical distribution of brain changes over specified time intervals, capturing the complex, non-linear patterns of neuroanatomical trajectories directly on MRI data. Rather than directly generating follow-up MRI scans, which can be prone to artifacts, TADM-3D focuses on predicting the voxel-wise intensity differences between baseline and follow-up images, reducing the complexity of the task and mitigating generation errors.

In contrast to other approaches, our model is conditioned on the age difference between the input and output scans, rather than on the patient's absolute age. This encourages the model to learn the relation between structural brain changes and the time interval between baseline and follow-up scans, rather than with specific age values. Since identical age gaps can occur at various absolute ages, this approach alleviates the need to include samples from all age groups, which is particularly advantageous when certain age ranges are underrepresented in the training data. For instance, a 5-year gap could apply both to a patient scanned at ages 60 and 65, and to another scanned at 80 and 85. While the output ages differ (65 vs. 85), the age gap remains the same (5 years).

To further improve the temporal awareness, we use a Brain-Age Estimator (BAE) to estimate the age difference between the baseline and the generated scans. During training, these predicted age differences are incorporated into the loss function, encouraging the generation of scans that accurately reflect the expected temporal interval between input and prediction.

Lastly, we introduce a *Back-In-Time Regularisation* (BITR) strategy at training time, to improve the temporal awareness of the model. Based on the intuition that a temporal-aware model should be able to predict both forward and backwards in time, we randomly swap the roles of the baseline and follow-up scans, and we train the model bidirectionally to predict future or past scans, alternately. This simple regularisation strategy encourages the model to learn how the brain anatomically changes going forward and backwards in time.

The method is trained and tested on the OASIS-3 dataset (LaMontagne et al., 2019). Moreover, we also use an external test set from the NACC dataset to evaluate the generalisation performance of TADM-3D on out-of-distribution data. As evaluation metrics, we use image-based similarity scores and brain region volumes between real and predicted follow-up MRIs. TADM-3D overcomes previous methods on both image-based and volumetric metrics, demonstrating its effectiveness on modelling the temporal brain evolution.

Additionally, our qualitative analysis demonstrates that our model more effectively reproduces the temporal progression of brains.

This work extends our previous MICCAI 2024 conference paper (Litrico et al., 2024) in several ways:

- The method was extended to operate on 3D MRI scans, rather than single 2D slices, allowing the model to capture richer spatial context and more complex anatomical relationships across all three dimensions.
- We introduced BITR, a technique aimed to improve the temporal awareness. This approach trains the model by considering not only forward progression but also backwards changes over time.
- The impact of the cognitive status as a conditioning variable was evaluated, allowing the model to account for clinically relevant differences. This conditioning helps capture variations in brain structure associated with different stages of cognitive decline (Cole et al., 2018).

- The comparisons were extended to include the most recent 3D methods, such as CounterSynth (Pombo et al., 2023) and BrLP (Puglisi et al., 2024) and we evaluate TADM-3D on an external test set to assess its generalisability.

2. Related works

The study of neurodegenerative diseases through MRIs has bloomed in the last years (Rondinella et al., 2023). Most existing approaches propose simulators that model temporal changes in brain structure. Simulators receive high-dimensional data as input, such as 3D MRIs of a given subject, and predict the changes in the brain MRIs over a specified period and under specific subject's conditions, such as neurodegenerative diseases. Most of the existing simulators are based on recent deep generative methods, including Generative Adversarial Networks (GANs) (Pombo et al., 2023; Ravi et al., 2022; Xia et al., 2019b), Variational Autoencoders (VAEs) (He et al., 2024), Normalising Flows (NFs) (Wilms et al., 2022) and Diffusion Models (DMs) (Puglisi et al., 2024; Yoon et al., 2022).

Early methods (Ravi et al., 2022; Xia et al., 2019b) used GANs to simulate subject-specific brain changes conditioned on the presence of a neurodegenerative disease. For instance, DANINet (Ravi et al., 2022) employs adversarial training in conjunction with biological constraints to improve the generation process. Nevertheless, these approaches produce synthetic 2D slices, without making use of full 3D volumetric information. More recently, CounterSynth (Pombo et al., 2023) introduced a 3D GAN-based framework that models 3D deformations, instead of directly manipulating image pixels. By leveraging morphology constraints, deformations are applied to the input MRI to reflect brain changes over time. However, CounterSynth only predicts structural changes without modelling temporal evolutions. Other related approaches proposed methodologies based on VAEs (He et al., 2024) and NFs (Wilms et al., 2022), but they produce low resolution scans and rely on morphologically constrained transformations, limiting their applicability.

More recently, the use of diffusion models has been introduced in medical image generation. DDM (Kim and Ye, 2022) proposed to combine a diffusion and a deformation module to learn how to interpolate between two MRIs. This combination enables the modelling of smooth anatomical transitions by simulating plausible intermediate brain states. Similarly, DiffuseMorph (Kim et al., 2022) leveraged a diffusion model to predict deformation fields between two MRIs. This field is then applied to warp one image into another, effectively enabling the synthesis of interpolated images along a continuous trajectory. Both DDM and DiffuseMorph are effective in learning realistic anatomical transformations, but they are limited to interpolation tasks. As a result, their applicability in clinical settings is constrained, since they are unable to generate predictive scans.

SADM (Yoon et al., 2022) introduced a sequence-aware diffusion model, incorporating a transformer-based architecture to better handle temporal dependencies in longitudinal data. The transformer generates a sequence of prior MRIs to extract a latent representation that encodes the subject's disease trajectory. This representation is then used to condition the generation process of the diffusion model, allowing for the synthesis of future brain MRIs. It necessarily requires a sequence of input MRI scans, which are rarely available in a real-world context, especially at an early stage of the patient's treatment. Moreover, it does not use subject metadata, which is important to better model the disease progression. BrLP (Puglisi et al., 2024) introduced a methodology that leverages an autoencoder, a latent diffusion model, and an auxiliary model (ControlNet) to generate individual MRI scans. However, the quality of the produced scans is highly correlated to the effectiveness of the autoencoder features.

TADM-3D overcomes these limitations as follows: (i) unlike CounterSynth, our model learns to model the temporal progression of brain changes, making it capable of forecasting future scans; (ii) in contrast

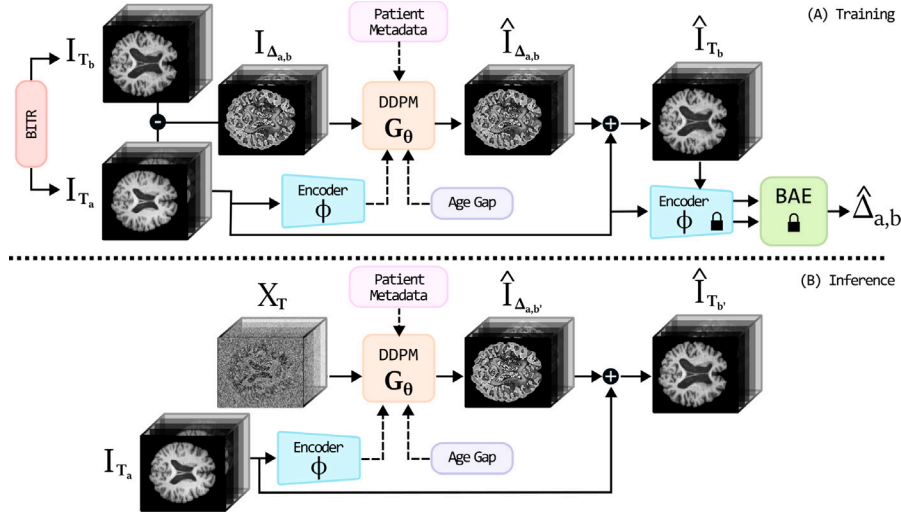


Fig. 1. TADM-3D. (A) Given a baseline scan I_{T_a} and a follow-up scan I_{T_b} , a residual image $I_{\Delta_{a,b}}$ is computed by subtracting the two scans. This residual is then corrupted with noise and denoised by a Denoising Diffusion Probabilistic Model (DDPM), producing the denoised residual $\hat{I}_{\Delta_{a,b}}$. The denoised residual image and the scan I_{T_a} are added together to estimate the scan \hat{I}_{T_b} at time T_b (see Section 3.2). The baseline scan is encoded with Φ to produce a latent representation z_a , which, along with patient metadata, is used to condition the DDPM (see Section 3.3). The estimated follow-up scan \hat{I}_{T_b} is also encoded to extract z_b . Representations z_a and z_b are then fed into a BAE to predict the time interval $\hat{\Delta}_{a,b}$ between the scans (see Section 3.4). To promote time awareness, the model is also trained to predict past scans with a probability $p = 0.5$ by swapping the roles of baseline and follow-up scans (see Section 3.5). The PADLOCK indicates frozen parameters. Dashed lines indicate model conditioning. (B) At inference time, given the baseline scan I_{T_a} and random noise X_T , the DDPM predicts the residual $\hat{I}_{\Delta_{a,b'}}$ that summed to the baseline I_{T_a} produces the predicted follow-up $\hat{I}_{T_b'}$.

to DDM and DiffuseMorph, TADM-3D is conditioned with age differences, allowing it to predict scans at arbitrary future time points; (iii) unlike SADM, TADM-3D does not need a sequence of input MRI scans, improving its applicability in clinical context. Moreover, our model incorporates subject-specific metadata to improve the accuracy and personalisation of predictions; (iv) differently than BrLP, we train the diffusion model to directly predict MRI scans, without the need for any auxiliary additional model to post-process the output, thereby allowing an end-to-end training process and improving the prediction quality.

3. Proposed method

3.1. Background

DDPMs (Ho et al., 2020) are generative models that learn how to gradually convert a Gaussian data distribution into another distribution by applying a Markov chain process. During the forward process (diffusion), Gaussian noise is progressively and incrementally added to the input data x_0 from the given data distribution $q(x_0)$. This occurs over a fixed number of steps, gradually transforming the sample into a latent variable distribution $q(x_t)$, as following:

$$q(x_1, x_2, \dots, x_T | x_0) = \prod_{t=1}^T q(x_t | x_{t-1}) \quad (1)$$

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

where $t \in 1, \dots, T$ indicates the diffusion step, \mathcal{N} is the Gaussian distribution, β_t is the variance of the noise, and \mathbf{I} is the identity matrix.

During the reverse process (denoising), the model is trained to learn to invert the diffusion process, by progressively turning the noise latent variable distribution $p_\theta(x_t)$ into the data distribution $p_\theta(x_0)$, parameterised by θ . During this process, the model is trained to learn the Gaussian transformations, as follows:

$$p_\theta(x_0, x_1, \dots, x_{T-1} | x_T) = \prod_{t=1}^T p_\theta(x_{t-1} | x_t) \quad (3)$$

$$p_\theta(x_{t-1} | x_t) = \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \sigma_\theta^2(x_t, t) \mathbf{I}) \quad (4)$$

$$p(x_t) = \mathcal{N}(x_t; \mathbf{0}, \mathbf{I}) \quad (5)$$

where $\sigma_\theta^2(x_t, t)$ is the variance at the step t and $\mu_\theta(x_t, t)$ is the mean of the Gaussian distribution. Once trained, the model is able to denoise random noise, generating high-quality images.

3.2. Temporally-aware diffusion model

Our proposed training pipeline, depicted in Fig. 1, comprises three main components: (i) a Denoising Diffusion Probabilistic Model (DDPM), responsible for modelling the brain progression and generating high quality scans; (ii) an Encoder, which extracts latent representation from MRI scans; and (iii) a Brain-Age Estimator (BAE), which is leveraged to encourage temporal consistency by estimating the time interval between scans.

During training, we use MRI pairs, i.e. I_{T_a} and I_{T_b} , corresponding to scans acquired from the same subject at two distinct timepoints T_a and T_b , respectively. From these images, we compute a residual image $I_{\Delta_{a,b}} = I_{T_b} - I_{T_a}$, which captures the voxel-wise intensity changes occurring over the time interval $\Delta_{a,b} = T_b - T_a$. This residual is used to train the DDPM at predicting the residual $\hat{I}_{\Delta_{a,b}}$. To generate the output follow-up scans \hat{I}_{T_b} , we add the predicted residual to the baseline scan, such that $\hat{I}_{T_b} = I_{T_a} + \hat{I}_{\Delta_{a,b}}$. Additionally, we leverage the BAE model to estimate the time interval $\hat{\Delta}_{a,b}$ between I_{T_a} and the estimated \hat{I}_{T_b} . The estimated time interval is then integrated in the loss function to regularise the DDPM generation process. At inference, we utilise the scan I_{T_a} acquired at time T_a along with the target time interval $\Delta_{a,b'}$ to produce a future scan $I_{T_b'}$, at time T_b' . To achieve patient individualisation during the generation process, the DDPM is conditioned on a latent embedding obtained by applying an encoder Φ to the baseline scan I_{T_a} , along with auxiliary patient-specific metadata (e.g., cognitive status and age).

3.3. Conditioning strategies

To generate residual images, we condition the DDPM using: (i) baseline latent representation $\Phi(I_{T_a})$ extracted by the encoder Φ on the

baseline scan I_{T_a} ; (ii) the time interval $\Delta_{a,b}$ between T_a and T_b ; (iii) patient's specific metadata.

Baseline Scan Encoding. Our goal is to predict patient-specific brain trajectories. To achieve this, the DDPM is conditioned with the encoded representation z_a derived from the baseline scan I_{T_a} , using an encoder Φ . This latent features capture the patient-specific anatomy, enabling the generation of personalised outputs.

Age Difference. Using age directly to model progression fails to explicitly capture the temporal evolution of structural changes in brain MRIs, and it demands age-balanced datasets, which are often unavailable in real-world settings. To address this issue, we propose conditioning the model on the age difference between scans, denoted as $\Delta_{a,b}$. Since the same age gaps can arise from scan pairs acquired at different ages, this strategy reduces the need of uniformly sampling all age groups during training. This is especially beneficial when certain age ranges are underrepresented in the dataset. In our approach, we encode $\Delta_{a,b}$ using positional encoding (Vaswani et al., 2017) before integrating it into the model.

Patient-Specific Metadata. Together with the embedding of the age gap, we concatenate embedding vectors obtained from patient-specific metadata. Specifically, we use both the age (A) of the patient at time T_a , and cognitive status (D) to further condition our model. Baseline age provides essential context to the model, as neurodegenerative diseases progress differently depending on the age of the patient. Note that, unlike previous approaches that condition using the age of the patient w.r.t. the follow-up scan, we use the age at baseline to accurately ground the starting point of the disease trajectory.

3.4. Estimating brain age for improving the temporal awareness

To guide the model in generating scans that accurately reflect the expected age difference between the input and predicted scans, we incorporate a Brain Age Estimator (BAE) (Jónsson et al., 2019) into our training pipeline. During DDPM training, the BAE evaluates whether the generated follow-up scan exhibits brain changes consistent with the expected time interval from the baseline. By including the predicted age difference as part of the loss function, this mechanism encourages the diffusion model to generate outputs that not only appear realistic but also accurately capture the desired temporal progression. This approach provides feedback to the model, thereby enhancing the temporal consistency of the predictions. The BAE model is pre-trained on the training set, and its parameters remain fixed during the DDPM training.

Given a baseline scan I_{T_a} and a generated follow-up scan \hat{I}_{T_b} , we compute the estimated age difference as $\hat{\Delta}_{a,b} = \Psi(\Phi(\hat{I}_{T_b})) - \Psi(\Phi(I_{T_a}))$, where Φ is the encoder and Ψ is the BAE model. When the DDPM prediction \hat{I}_{T_b} accurately matches the ground-truth image I_{T_b} , the estimated age difference $\hat{\Delta}_{a,b}$ aligns with the actual age gap $\Delta_{a,b}$, ensuring that the generated output reflects the correct temporal progression. Any differences between the predicted and true age gaps are backpropagated during training, serving as a regularisation. This feedback loop iteratively refines the generation process, enhancing the model's ability to produce scans that accurately reflect the temporal progression.

3.5. Back-In-Time Regularisation

TADM-3D is trained using pairs of MRI scans I_{T_a} and I_{T_b} acquired from the same patient at two distinct time points. Here I_{T_a} indicates the baseline scan acquired at age A , and I_{T_b} is the follow-up scan acquired after a time interval $\Delta_{a,b}$. In this way, the diffusion model learns to predict the brain structural changes that occur due to the combined effect of ageing and disease progression over the time interval. However, a temporal-aware model should also be able to predict past scans, e.g. generating a baseline scan I_{T_a} from a follow-up scan I_{T_b} .

To this aim, we propose to include a *Back-In-Time Regularisation* (BITR) strategy. Specifically, at each training step, we randomly swap the roles of the two scans I_{T_a} and I_{T_b} with a probability $p = 0.5$.

Algorithm 1 Pseudocode of TADM-3D's training process.

Input: Pairs of MRI scans (I_{T_a}, I_{T_b}), time interval $\Delta_{a,b}$, patient metadata (age at baseline A , cognitive status D).

Output: Trained DDPM model G_θ for generating future MRI scans

Training:

for each training step **do**

 Sample a pair (I_{T_a}, I_{T_b}) from the training set

 Compute residual image $I_{\Delta_{a,b}} = I_{T_b} - I_{T_a}$

 Extract latent representation

$z_a = \Phi(I_{T_a})$ using the encoder Φ

 Sample $p \sim \mathcal{B}(0.5)$

if $p = 1$ **then**

 Flip sign in $\Delta_{a,b}$

end if

 Sample noise $\epsilon \sim \mathcal{N}(0, 1)$

 Compute noisy residual $\tilde{I}_{\Delta_{a,b}}$ at diffusion step t

 Predict noise using DDPM:

$\hat{\epsilon} = G_\theta(\tilde{I}_{\Delta_{a,b}}, t; z_a, \Delta_{a,b}, A, D)$

 Compute DDPM loss: $\mathcal{L}^{DDL} = \|\hat{\epsilon} - \epsilon\|_2^2$

 Generate predicted residual $\hat{I}_{\Delta_{a,b}}$ from the DDPM

 Compute predicted follow-up scan $\hat{I}_{T_b} = I_{T_a} + \hat{I}_{\Delta_{a,b}}$

 Extract latent representation $z_b = \Phi(\hat{I}_{T_b})$

 Predict age difference using BAE: $\hat{\Delta}_{a,b} = \Psi(z_b) - \Psi(z_a)$

 Compute BAE loss: $\mathcal{L}^{BAE} = (\hat{\Delta}_{a,b} - |\Delta_{a,b}|)^2$

 Compute total loss: $\mathcal{L}^{Tot} = \mathcal{L}^{DDL} + \mathcal{L}^{BAE}$

 Update DDPM parameters θ using gradients of \mathcal{L}^{Tot}

end for

Consequently, the diffusion model is trained to predict the follow-up from the baseline scan for approximately half of the training iterations and the baseline from the follow-up in the other half of the training steps. Note that this swapping is exclusively applied in the training phase. Although simple, this strategy encourages the model to learn a better relationship between anatomical changes in the brain and time interval, leading to an improved temporal awareness.

3.6. Overall framework

Training. A complete overview of the training process is described in Algorithm 1. During the diffusion process, the DDPM learns to estimate the noise ϵ incorporated into $I_{\Delta_{a,b}}$. This is achieved by minimising the following loss function:

$$\mathcal{L}^{DDL} = \mathbb{E}_{\epsilon \sim \mathcal{N}(0,1), \tilde{I}_{\Delta_{a,b}}, t} \left[\|G_\theta(\tilde{I}_{\Delta_{a,b}}, t; z_a, \Delta_{a,b}, A, D) - \epsilon\|_2^2 \right], \quad (6)$$

where, G_θ represents the DDPM with parameters θ , and t denotes the diffusion timestep. Furthermore, as discussed in Section 3.4, the prediction of BAE is incorporated as an extra component in the loss function. Specifically, we define the loss on the expected brain age difference as follows:

$$\mathcal{L}^{BAE} = (\hat{\Delta}_{a,b} - |\Delta_{a,b}|)^2. \quad (7)$$

The BAE loss is computed using the absolute value of $\Delta_{a,b}$ to account for iterations where the sign of $\Delta_{a,b}$ is flipped due to the BITR. The BAE loss helps the model in generating scans matching the actual age gap, independently whether the generated scan is forward or backward in time.

Finally, the overall loss is obtained by combining Eqs. (6) and (7):

$$\mathcal{L}^{Tot} = \mathcal{L}^{DDL} + \mathcal{L}^{BAE} \quad (8)$$

Inference. Given a baseline MRI I_{T_a} , the model generates a future MRI \hat{I}_{T_b} at any desired time interval $\Delta_{a,b'}$ after the baseline. The generation process starts with a random Gaussian noise input X_T , which is iteratively refined by the network $G_\theta(X_T, t; z_a, \Delta_{a,b'}, A, D)$. The generated residual image $\hat{I}_{\Delta_{a,b'}}$ is finally combined with the baseline scan to yield the follow-up prediction.

Table 1

Results on the internal test set of OASIS-3. Performance are evaluated in terms of image-based and region volumes errors with respect to other methods. The MAE in region volumes is expressed as a percentage of total brain volume.

	Method	MSE ↓	SSIM ↑	Region volumes error (%) ↓				
				Hippocampus	Amygdala	Lat. Ventricle	Thalamus	CSF
All patients	DaniNet (Ravi et al., 2022)	0.016 ± 0.007	0.623 ± 0.162	0.030 ± 0.030	0.018 ± 0.017	0.257 ± 0.222	0.038 ± 0.030	1.081 ± 0.814
	CounterSynth (Pombo et al., 2023)	0.008 ± 0.004	0.861 ± 0.052	0.030 ± 0.018	0.016 ± 0.010	0.273 ± 0.311	0.041 ± 0.035	0.881 ± 0.672
	BrLP (Puglisi et al., 2024)	0.005 ± 0.002	0.887 ± 0.017	0.028 ± 0.018	0.017 ± 0.009	0.264 ± 0.271	0.039 ± 0.021	0.882 ± 0.645
	TADM-3D	0.004 ± 0.001	0.902 ± 0.014	0.017 ± 0.019	0.015 ± 0.014	0.228 ± 0.187	0.027 ± 0.015	0.642 ± 0.412
CN patients	DaniNet (Ravi et al., 2022)	0.018 ± 0.009	0.628 ± 0.167	0.031 ± 0.031	0.016 ± 0.018	0.261 ± 0.214	0.036 ± 0.032	1.091 ± 0.818
	CounterSynth (Pombo et al., 2023)	0.009 ± 0.004	0.884 ± 0.056	0.034 ± 0.019	0.017 ± 0.011	0.274 ± 0.315	0.040 ± 0.037	0.871 ± 0.675
	BrLP (Puglisi et al., 2024)	0.005 ± 0.003	0.885 ± 0.018	0.027 ± 0.017	0.017 ± 0.010	0.261 ± 0.270	0.037 ± 0.020	0.884 ± 0.638
	TADM-3D	0.005 ± 0.001	0.904 ± 0.015	0.014 ± 0.018	0.015 ± 0.013	0.225 ± 0.184	0.026 ± 0.016	0.638 ± 0.419
MCI patients	DaniNet (Ravi et al., 2022)	0.015 ± 0.007	0.623 ± 0.162	0.030 ± 0.030	0.017 ± 0.017	0.258 ± 0.221	0.038 ± 0.030	1.082 ± 0.812
	CounterSynth (Pombo et al., 2023)	0.008 ± 0.004	0.860 ± 0.052	0.029 ± 0.018	0.015 ± 0.010	0.273 ± 0.311	0.042 ± 0.035	0.880 ± 0.674
	BrLP (Puglisi et al., 2024)	0.004 ± 0.002	0.887 ± 0.017	0.029 ± 0.018	0.017 ± 0.009	0.264 ± 0.271	0.038 ± 0.021	0.881 ± 0.645
	TADM-3D	0.004 ± 0.002	0.901 ± 0.011	0.016 ± 0.020	0.015 ± 0.015	0.229 ± 0.183	0.028 ± 0.01	0.645 ± 0.414
AD patients	DaniNet (Ravi et al., 2022)	0.015 ± 0.009	0.618 ± 0.167	0.029 ± 0.029	0.021 ± 0.018	0.252 ± 0.214	0.040 ± 0.032	1.070 ± 0.816
	CounterSynth (Pombo et al., 2023)	0.007 ± 0.004	0.839 ± 0.056	0.027 ± 0.019	0.016 ± 0.011	0.272 ± 0.314	0.041 ± 0.037	0.892 ± 0.676
	BrLP (Puglisi et al., 2024)	0.006 ± 0.003	0.889 ± 0.018	0.028 ± 0.017	0.017 ± 0.010	0.267 ± 0.269	0.042 ± 0.020	0.881 ± 0.639
	TADM-3D	0.003 ± 0.002	0.901 ± 0.016	0.021 ± 0.020	0.015 ± 0.014	0.230 ± 0.182	0.027 ± 0.009	0.643 ± 0.413

Table 2

Results on the external test set of NACC evaluating TADM-3D's generalisation performance in terms of image-based and region volumes errors in comparison to other methods. The MAE in region volumes is expressed as a percentage of total brain volume.

	Method	MSE ↓	SSIM ↑	Region volumes error (%) ↓				
				Hippocampus	Amygdala	Lat. Ventricle	Thalamus	CSF
All patients	DaniNet (Ravi et al., 2022)	0.017 ± 0.007	0.611 ± 0.181	0.032 ± 0.031	0.018 ± 0.016	0.232 ± 0.210	0.039 ± 0.032	1.154 ± 0.871
	CounterSynth (Pombo et al., 2023)	0.011 ± 0.003	0.813 ± 0.042	0.030 ± 0.020	0.014 ± 0.010	0.283 ± 0.314	0.111 ± 0.034	1.173 ± 0.731
	BrLP (Puglisi et al., 2024)	0.005 ± 0.002	0.909 ± 0.023	0.024 ± 0.023	0.014 ± 0.013	0.213 ± 0.350	0.030 ± 0.024	1.044 ± 0.788
	TADM-3D	0.004 ± 0.002	0.902 ± 0.017	0.020 ± 0.020	0.013 ± 0.011	0.235 ± 0.200	0.029 ± 0.021	0.833 ± 0.543
CN patients	DaniNet (Ravi et al., 2022)	0.016 ± 0.006	0.591 ± 0.163	0.030 ± 0.028	0.017 ± 0.014	0.220 ± 0.190	0.037 ± 0.029	1.134 ± 0.851
	CounterSynth (Pombo et al., 2023)	0.010 ± 0.002	0.793 ± 0.038	0.028 ± 0.018	0.013 ± 0.009	0.269 ± 0.294	0.105 ± 0.031	1.153 ± 0.711
	BrLP (Puglisi et al., 2024)	0.004 ± 0.001	0.889 ± 0.021	0.023 ± 0.021	0.013 ± 0.012	0.202 ± 0.330	0.028 ± 0.022	1.024 ± 0.768
	TADM-3D	0.003 ± 0.001	0.882 ± 0.015	0.019 ± 0.018	0.012 ± 0.018	0.223 ± 0.180	0.023 ± 0.018	0.813 ± 0.523
MCI patients	DaniNet (Ravi et al., 2022)	0.017 ± 0.005	0.615 ± 0.184	0.032 ± 0.030	0.018 ± 0.014	0.230 ± 0.220	0.038 ± 0.038	1.151 ± 0.870
	CounterSynth (Pombo et al., 2023)	0.012 ± 0.002	0.817 ± 0.047	0.031 ± 0.027	0.014 ± 0.011	0.287 ± 0.311	0.115 ± 0.033	1.170 ± 0.737
	BrLP (Puglisi et al., 2024)	0.006 ± 0.001	0.911 ± 0.026	0.022 ± 0.022	0.014 ± 0.015	0.212 ± 0.352	0.031 ± 0.025	1.042 ± 0.786
	TADM-3D	0.003 ± 0.003	0.905 ± 0.013	0.019 ± 0.023	0.013 ± 0.010	0.236 ± 0.210	0.029 ± 0.024	0.832 ± 0.547
AD patients	DaniNet (Ravi et al., 2022)	0.018 ± 0.008	0.631 ± 0.199	0.034 ± 0.034	0.019 ± 0.018	0.244 ± 0.230	0.041 ± 0.035	1.174 ± 0.891
	CounterSynth (Pombo et al., 2023)	0.012 ± 0.004	0.833 ± 0.046	0.032 ± 0.022	0.015 ± 0.011	0.297 ± 0.334	0.117 ± 0.037	1.193 ± 0.751
	BrLP (Puglisi et al., 2024)	0.006 ± 0.003	0.929 ± 0.025	0.025 ± 0.025	0.015 ± 0.014	0.224 ± 0.370	0.032 ± 0.026	1.064 ± 0.808
	TADM-3D	0.005 ± 0.003	0.922 ± 0.019	0.021 ± 0.022	0.014 ± 0.012	0.247 ± 0.220	0.035 ± 0.022	0.853 ± 0.563

4. Experimental results

4.1. Datasets

We train TADM-3D on 2535 T1-weighted (T1w) brain MRI scans from 634 subjects from the OASIS-3 dataset (LaMontagne et al., 2019). Scans span a longitudinal interval of approximately 15 years, with an average interval of 3.2 years between baseline and follow-up, capturing a broad spectrum of ageing-related changes. The dataset includes participants aged between 42 and 95 years classified as cognitively normal (CN), mild cognitive impairment (MCI), and Alzheimer's disease (AD). We evaluate TADM-3D on both internal and external test sets, to also assess generalisation performance with out-of-distribution data. For the external dataset, we used data from the NACC dataset (Beekly et al., 2007), including 2257 T1w MRIs from 962 subjects. Scan time intervals range from 4 months to 13 years between initial and follow-up MRI, with an average of 3.8 years. Given that the minimum interval is less than 1 year, we express the age difference in terms of months. In both datasets, we build pairs of baseline and follow-up scans with MRI acquired from the same patient, with an average number of 3 pairs per patient. Data from the external dataset is used solely for evaluation. To note that all MRI scans are linearly registered to the MNI152 template to standardise orientation and size, but baseline and follow-up scans are not directly co-registered. This leads to some intrinsic misalignment between pairs, which may introduce artifacts affecting prediction accuracy. However, our model demonstrates robust

performance against these minor inconsistencies, as supported by our quantitative and qualitative results.

4.2. Implementation details

All MRI volumes were linearly registered using the MNI152 template to normalise spatial orientation and scale, and skull stripping using the FSL library (Jenkinson et al., 2012). Baseline and follow-up scans are not directly co-registered to avoid potential distortions from enforcing alignment across structurally changing regions. The dataset is divided into a training set (70%), a validation set (10%) and a test set (20%). For both the training of diffusion and BAE models, we use the AdamW optimiser with a learning rate of 0.0001 and a weight decay of 0.001, a batch size of 16, and a cosine-based learning rate scheduler. Following Li et al. (2022), we employ the U-Net architecture as backbone of the diffusion model G_θ . For the encoder Φ , we employ the same backbone as the one used in the diffusion model encoder and it is trained in an end-to-end manner. Results for all state-of-the-art methods were obtained using their publicly available implementations.

4.3. Evaluation metrics

To evaluate the performance of our method, we employed both image-based similarity metrics and region-specific volumetric analyses in anatomically relevant brain areas. For the image-based evaluation, we use the *Structural Similarity Index Measure* (SSIM) and *Mean*




Fig. 2. Temporal progression on a 62-year-old subject with AD from the internal test set, generated by our approach against BrLP (Puglisi et al., 2024), CounterSynth (Pombo et al., 2023) and DaniNet (Ravi et al., 2022). We show predicted MRIs on the centre slice on the left and the corresponding heatmap of prediction error on the right.

Squared Error (MSE), measuring the similarity between the generated and ground-truth MRI scans in terms of structural fidelity and pixel-wise reconstruction accuracy, respectively. On the other side, volumetric metrics in AD-related regions (lateral ventricles, cerebrospinal fluid as CSF, hippocampus, amygdala, and thalamus) are computed to evaluate the modelling of disease progression in TADM-3D. We use *SynthSeg* 2.0 (Billot et al., 2023) to segment the brain and compute the region volumes, which are expressed as percentages of the entire brain to account for personal variations.

The accuracy of predicted regional volumes is then assessed by computing the Mean Absolute Error (MAE) with the ground-truth volumes.

4.4. Comparative analysis

Table 1 shows the quantitative results on the internal test set obtained by TADM-3D compared to other state-of-the-art 3D approaches (Pombo et al., 2023; Puglisi et al., 2024; Ravi et al., 2022; Yoon et al., 2022). Overall, TADM-3D outperforms them across the board, reducing the MSE and increasing the SSIM by +0.001 and +0.15, respectively. Results on the region’s volumes show that TADM-3D achieves the lowest error in different brain regions. In particular, we reduce the error on Hippocampus, Amygdala and Ventricles by approximately 40%, 12%, and 13%, respectively. On Thalamus and CSF, TADM-3D reduces the error by 30% and 27%, respectively. These results show that TADM-3D generates accurate follow-up images with respect to existing approaches, achieving better performance in both image-based and volumetric metrics. On the external test set (see Table 2), TADM-3D obtains similar performance w.r.t. the internal evaluation, showing robustness to generalisation challenges. TADM-3D achieves the best performance on the MSE and on volumetric errors of 4 out of 5 analysed brain regions, while BrLP unexpectedly improves its results w.r.t. the internal test set on SSIM and on volumetric error of the Lateral Ventricle, obtaining the lowest error. To evaluate model performance on both healthy and diseased patients, we report quantitative results for each cognitive group in the internal (Table 1) and external (Table 2) evaluations. All methods show slightly higher performance on CN than on MCI and AD, which is expected since disease-related brain changes introduce greater variability and are more difficult to model. Nonetheless, TADM-3D maintains consistently strong performance across all groups. This suggests that it effectively models disease progression and remains robust to varying levels of cognitive impairment.

Furthermore, Fig. 2 presents a sample of qualitative results illustrating the prediction of MRI scans at various ages for a 62-year-old patient diagnosed with AD. Results show that TADM-3D provides a better approximation of the temporal evolution of the brain compared

to other methods. Our model reliably captures ventricular expansion with reasonable accuracy and maintains consistently competitive errors across age differences. Compared to CounterSynth (Pombo et al., 2023) and DaniNet (Ravi et al., 2022), TADM-3D generally demonstrates superior performance in the ventricular region. Although BrLP (Puglisi et al., 2024) shows results closer to ours, TADM-3D still exhibits a more balanced performance across all brain regions, with only minor peaks on the cortex. This underscores the robustness of our solution in modelling brain changes over time.

4.5. Ablation studies

In this section, we present an ablation study to assess the contribution of the key components of TADM-3D. Specifically, we evaluate the effect of removing conditioning on age differences, patient-specific conditioning variables, and excluding the BAE module and BTR from training. Furthermore, we compare the performance of our 3D generative architecture against a 2D alternative, where slices are generated independently to reconstruct full 3D volumes (named TADM-2.5D).

Table 3 (first row) presents the performance of TADM-3D without using patient metadata. Results show that, removing the conditioning with such metadata, we observe a minimal performance loss. In the second row, we observe that removing the conditioning of the model on age gaps highly drops performance, demonstrating the effectiveness of our strategy over the current state-of-the-art. Another case of performance drop is when BAE is not used (third row), indicating its usefulness in supporting the generation process. In the fourth row, we show the impact of removing the BTR with a significant drop of performance, demonstrating its contribution in improving the training. Finally, in the last row, we assess the performance using a 2D diffusion model (as our previous work Litrico et al., 2024) rather than a 3D one. To reconstruct full 3D volumes, we train a 2D model to generate each slice independently, using the slice index as a conditioning variable to preserve spatial coherence across the volume. Results show a drastically reduction in performance when using the 2D model, demonstrating that a 3D model captures better structural information during training.

4.6. Is TADM-3D modelling the disease progression?

We perform an experiment to investigate the effect of conditioning on TADM-3D predictions. Specifically, we aim to assess whether the model’s outputs are appropriately generated with respect to the patient’s cognitive status. To do so, we use our trained model with scans obtained from patients with AD, but we deliberately condition the model, specifying that the patient is Cognitively Normal (CN). We

Table 3

Ablation studies: showing the results of the absence of the different components of TADM-3D. Finally, we also show the impact of using a 2D model to individually generate scans rather than a native 3D architecture.

Method	MSE ↓	SSIM ↑	Region volumes error (%) ↓				
			Hippocampus	Amygdala	Lat. Ventricle	Thalamus	CSF
TADM-3D w/o patient metadata	0.006 ± 0.001	0.899 ± 0.017	0.019 ± 0.011	0.015 ± 0.018	0.235 ± 0.201	0.030 ± 0.019	0.648 ± 0.557
TADM-3D w/o age gaps cond.	0.011 ± 0.006	0.812 ± 0.016	0.027 ± 0.013	0.017 ± 0.021	0.251 ± 0.231	0.037 ± 0.029	0.747 ± 0.555
TADM-3D w/o BAE	0.009 ± 0.005	0.872 ± 0.024	0.025 ± 0.014	0.016 ± 0.025	0.249 ± 0.212	0.036 ± 0.026	0.691 ± 0.512
TADM-3D w/o BITR	0.008 ± 0.003	0.874 ± 0.021	0.023 ± 0.016	0.016 ± 0.022	0.241 ± 0.200	0.031 ± 0.021	0.667 ± 0.491
TADM-2.5D	0.014 ± 0.006	0.779 ± 0.089	0.030 ± 0.027	0.016 ± 0.015	0.299 ± 0.277	0.060 ± 0.041	0.899 ± 0.643
TADM-3D	0.004 ± 0.001	0.902 ± 0.014	0.017 ± 0.019	0.015 ± 0.014	0.228 ± 0.187	0.027 ± 0.015	0.642 ± 0.412

Table 4

Evaluating the impact of incorrect conditioning on cognitive status in TADM-3D predictions.

Method	MSE ↓	SSIM ↑	Region volumes error (%) ↓				
			Hippocampus	Amygdala	Lat. Ventricle	Thalamus	CSF
TADM-3D w/ wrong cond.	0.006 ± 0.005	0.895 ± 0.012	0.030 ± 0.013	0.017 ± 0.021	0.234 ± 0.187	0.030 ± 0.026	0.679 ± 0.417
TADM-3D	0.004 ± 0.001	0.902 ± 0.014	0.017 ± 0.019	0.015 ± 0.014	0.228 ± 0.187	0.027 ± 0.015	0.642 ± 0.412




Fig. 3. Example of limitation of TADM-3D in predicting long-time trajectory for a 65-year-old patient. The brain's evolution in predictions after 6.5 and 9 years is inconsistent, showing limitations in capturing long-time dependencies.

then generate predicted future scans under both the incorrect (CN) and correct (AD) cognitive conditions, and compare the resulting volumetric predictions. This setup allows us to isolate the influence of cognitive status on the generative behaviour of the model. Results in [Table 4](#) show that prediction errors generally increase when incorrect cognitive conditioning is applied, particularly in the hippocampal region, a structure notably affected in AD. This supports the hypothesis that the model is influenced by the conditioning and is capable of capturing distinct anatomical progressions associated with neurodegeneration.

5. Discussions

5.1. Limitations

Despite the promising results achieved by TADM-3D, some limitations still warrant further investigation. Specifically, [Fig. 3](#) illustrates an example of the challenges faced by TADM-3D in predicting long-term trajectories for a 65-year-old patient. The model tends to lose accuracy when predicting brain evolution over long time intervals, exhibiting increased errors, particularly for intervals spanning several years between scans. To assess TADM-3D ability in capturing long-time progressions, [Table 5](#) provides a detailed quantitative analysis across different age gap intervals (0–12 months, 12–48 months, and >48 months) and cognitive status (cognitively normal, MCI, and AD). Results show that TADM-3D reduces performance for age gaps exceeding 48 months, while still maintaining competitive results overall. This observation highlights the challenge of effectively modelling long-term temporal dependencies and predicting anatomical changes in the long future. Moreover, in AD patients, these challenges are slightly amplified, as their brain structures can significantly change within a long time interval. Another limitation is that, although TADM-3D is conditioned on age differences and some patient metadata, it does not yet incorporate other important clinical factors such as genetic information, comorbidities, or medication effects, that could significantly influence disease trajectories.

5.2. Applications

Disease progression modelling plays a crucial role in many aspects of clinical practice and holds significant promise for improving patient care. By simulating the anatomical changes that occur over time, TADM-3D enables earlier diagnosis and more accurate prognostic assessments, which are essential for timely and targeted interventions. Predicting personalised progression trajectories allows for individualised treatment planning, thereby optimising care strategies. In the context of clinical trials, our model can generate virtual follow-up scans or synthetic control arms, which can enhance trial design and reduce patient burden. TADM-3D also helps overcome limitations caused by missing or irregularly acquired longitudinal scans by effectively filling gaps and supporting robust longitudinal analyses. From a research perspective, modelling disease trajectories provides insights into the underlying pathophysiological mechanisms and temporal dynamics, contributing to novel biomarker discovery. Moreover, our model supports educational efforts by delivering visual representations of expected progression, improving communication among clinicians, patients, and their families. Finally, TADM-3D may help address data scarcity and biases in medical research by synthesising data for underrepresented populations and costly imaging modalities, including uncommon and expensive scans such as PET and CT, where generating synthetic samples remains particularly challenging.

In summary, these applications highlight the importance of our model for precision medicine, enabling earlier, more personalised, and more effective care for progressive diseases such as Alzheimer's.

5.3. Future works

Building upon our pipeline, multiple directions could be followed to further advance TADM-3D. One direction is to enhance the model's ability to capture longer-term temporal dependencies. Another one involves integrating multiple data modalities beyond structural MRI, such as other medical scans, genomic information and clinical records.

Table 5
Quantitatively results on different age gaps, evaluating the ability of TADM-3D to perform long-time predictions.

	Age Gap	MSE ↓	SSIM ↑	Region volumes error (%) ↓				
				Hippocampus	Amygdala	Lat. Ventricle	Thalamus	CSF
All patients	0–12 months	0.003 ± 0.001	0.904 ± 0.016	0.017 ± 0.025	0.015 ± 0.018	0.225 ± 0.192	0.027 ± 0.018	0.638 ± 0.445
	12–48 months	0.005 ± 0.001	0.896 ± 0.016	0.018 ± 0.021	0.016 ± 0.016	0.230 ± 0.195	0.028 ± 0.017	0.647 ± 0.430
	> 48 months	0.006 ± 0.002	0.890 ± 0.020	0.020 ± 0.023	0.018 ± 0.018	0.233 ± 0.205	0.030 ± 0.019	0.654 ± 0.440
CN patients	0–12 months	0.003 ± 0.001	0.905 ± 0.013	0.013 ± 0.018	0.014 ± 0.013	0.223 ± 0.180	0.026 ± 0.014	0.622 ± 0.402
	12–48 months	0.004 ± 0.001	0.899 ± 0.015	0.017 ± 0.020	0.015 ± 0.015	0.228 ± 0.194	0.027 ± 0.016	0.643 ± 0.429
	> 48 months	0.005 ± 0.001	0.893 ± 0.019	0.019 ± 0.022	0.017 ± 0.017	0.231 ± 0.204	0.029 ± 0.018	0.650 ± 0.439
MCI patients	0–12 months	0.004 ± 0.001	0.903 ± 0.019	0.016 ± 0.024	0.015 ± 0.016	0.226 ± 0.145	0.025 ± 0.017	0.636 ± 0.481
	12–48 months	0.005 ± 0.001	0.896 ± 0.016	0.018 ± 0.021	0.016 ± 0.016	0.230 ± 0.195	0.028 ± 0.017	0.647 ± 0.430
	> 48 months	0.006 ± 0.002	0.890 ± 0.020	0.020 ± 0.023	0.018 ± 0.018	0.233 ± 0.205	0.030 ± 0.019	0.654 ± 0.440
AD patients	0–12 months	0.002 ± 0.001	0.903 ± 0.015	0.018 ± 0.020	0.013 ± 0.015	0.227 ± 0.194	0.025 ± 0.016	0.637 ± 0.422
	12–48 months	0.005 ± 0.002	0.893 ± 0.017	0.019 ± 0.022	0.017 ± 0.017	0.232 ± 0.196	0.029 ± 0.018	0.651 ± 0.431
	> 48 months	0.007 ± 0.002	0.887 ± 0.021	0.021 ± 0.024	0.019 ± 0.019	0.235 ± 0.206	0.031 ± 0.020	0.658 ± 0.441

This multi-modal integration could provide more comprehensive and personalised information on the neurodegenerative progression. Additionally, improving the model's robustness and adaptability to diverse clinical and imaging settings could be beneficial. Indeed, approaches based on federated learning frameworks could facilitate model deployment across institutions without requiring data sharing. Finally, extending the framework to other organs would further broaden its clinical impact.

6. Conclusion

In this work, we introduced TADM-3D, a diffusion-based approach for 3D brain progression modelling that directly predicts the intensity difference between baseline and follow-up MRI scans, effectively capturing structural changes over time. Our method addresses several key limitations of existing techniques by conditioning predictions on age differences rather than on target ages, allowing for more accurate temporal modelling without requiring age-balanced datasets. Furthermore, we proposed a Back-In-Time Regularisation strategy, enhancing the model's temporal awareness by training the model to predict both forward and backwards in time.

We extensively evaluated TADM-3D on the OASIS-3 dataset, achieving superior results in both similarity metrics and region volume estimation. Moreover, to assess the generalisability of our method, we also tested TADM-3D on an external test set from the NACC dataset, achieving comparable results w.r.t. the internal evaluation. Qualitative analyses further highlighted the model's capacity to better capture the progression of brain structures over time, particularly in regions highly susceptible to neurodegenerative diseases.

CRedit authorship contribution statement

Mattia Litrico: Writing – original draft, Validation, Methodology, Conceptualization. **Francesco Guarnera:** Writing – review & editing. **Mario Valerio Giuffrida:** Writing – review & editing, Methodology. **Daniele Ravi:** Writing – review & editing, Supervision. **Sebastiano Battiato:** Writing – review & editing, Supervision.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

The NACC database is funded by NIA/NIH Grant U24 AG072122. NACC data are contributed by the NIA-funded ADCRCs: P30 AG062429 (PI James Brewer, MD, PhD), P30 AG066468 (PI Oscar Lopez, MD), P30 AG062421 (PI Bradley Hyman, MD, PhD), P30 AG066509 (PI Thomas Grabowski, MD), P30 AG066514 (PI Mary Sano, PhD), P30 AG066530 (PI Helena Chui, MD), P30 AG066507 (PI Marilyn Albert, PhD), P30 AG066444 (PI David Holtzman, MD), P30 AG066518 (PI Lisa Silbert, MD, MCR), P30 AG066512 (PI Thomas Wisniewski, MD), P30 AG066462 (PI Scott Small, MD), P30 AG072979 (PI David Wolk, MD), P30 AG072972 (PI Charles DeCarli, MD), P30 AG072976 (PI Andrew Saykin, PsyD), P30 AG072975 (PI Julie A. Schneider, MD, MS), P30 AG072978 (PI Ann McKee, MD), P30 AG072977 (PI Robert Vassar, PhD), P30 AG066519 (PI Frank LaFerla, PhD), P30 AG062677 (PI Ronald Petersen, MD, PhD), P30 AG079280 (PI Jessica Langbaum, PhD), P30 AG062422 (PI Gil Rabinovici, MD), P30 AG066511 (PI Allan Levey, MD, PhD), P30 AG072946 (PI Linda Van Eldik, PhD), P30 AG062715 (PI Sanjay Asthana, MD, FRCP), P30 AG072973 (PI Russell Swerdlow, MD), P30 AG066506 (PI Glenn Smith, PhD, ABPP), P30 AG066508 (PI Stephen Strittmatter, MD, PhD), P30 AG066515 (PI Victor Henderson, MD, MS), P30 AG072947 (PI Suzanne Craft, PhD), P30 AG072931 (PI Henry Paulson, MD, PhD), P30 AG066546 (PI Sudha Seshadri, MD), P30 AG086401 (PI Erik Roberson, MD, PhD), P30 AG086404 (PI Gary Rosenberg, MD), P20 AG068082 (PI Angela Jefferson, PhD), P30 AG072958 (PI Heather Whitson, MD), P30 AG072959 (PI James Leverenz, MD). The work of Francesco Guarnera has been supported by MUR in the framework of PNRR PE0000013, under project "Future Artificial Intelligence Research FAIR".

Data availability

The code is available at <https://github.com/MattiaLitrico/TADM-3D>.

References

- Beekly, D.L., Ramos, E.M., Lee, W.W., Deitrich, W.D., Jacka, M.E., Wu, J., Hubbard, J.L., Koepsell, T.D., Morris, J.C., Kukull, W.A., NIA Alzheimer's Disease Centers, 2007. The National Alzheimer's Coordinating Center (NACC) database: the Uniform Data Set. *Alzheimer Dis. Assoc. Disord.* 21 (3), 249–258.
- Billot, B., Greve, D.N., Puonti, O., Thielscher, A., Van Leemput, K., Fischl, B., Dalca, A.V., Iglesias, J.E., 2023. SynthSeg: Segmentation of brain MRI scans of any contrast and resolution without retraining. *Med. Image Anal.* 86, 102789. <http://dx.doi.org/10.1016/j.media.2023.102789>.

- Bowles, C., Gunn, R., Hammers, A., Rueckert, D., 2018. Modelling the progression of Alzheimer's disease in MRI using generative adversarial networks. *Med. Imaging* 2018: Image Process. 10574, 105741K.
- Cole, J.H., Ritchie, S.J., Bastin, M.E., Hernández, V., Muñoz Maniega, S., Royle, N., Corley, J., Pattie, A., Harris, S.E., Zhang, Q., et al., 2018. Brain age predicts mortality. *Mol. Psychiatry* 23 (5), 1385–1392.
- He, R., Ang, G., Tward, D., 2024. Individualized multi-horizon MRI trajectory prediction for Alzheimer's disease. arXiv:2408.02018.
- Ho, J., Jain, A., Abbeel, P., 2020. Denoising diffusion probabilistic models. *Adv. Neural Inf. Process. Syst.* 33, 6840–6851.
- Jack, Jr., C.R., Bennett, D.A., Blennow, K., Carrillo, M.C., Dunn, B., Haeberlein, S.B., Holtzman, D.M., Jagust, W., Jessen, F., Karlawish, J., et al., 2018. NIA-AA research framework: toward a biological definition of Alzheimer's disease. *Alzheimer's & Dement.* 14 (4), 535–562.
- Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M., 2012. FSL. *Neuroimage* 62 (2), 782–790.
- Jónsson, B.A., Björnsdóttir, G., Thorgeirsson, T., Ellingsen, L.M., Walters, G.B., Gudbjartsson, D., Stefansson, H., Stefansson, K., Ulfarsson, M., 2019. Brain age prediction using deep learning uncovers associated sequence variants. *Nat. Commun.* 10 (1), 5409.
- Kim, B., Han, I., Ye, J.C., 2022. DiffuseMorph: Unsupervised deformable image registration using diffusion model. In: *Computer Vision – ECCV 2022: 17th European Conference, Tel Aviv, Israel, October 23–27, 2022, Proceedings, Part XXXI*. Springer-Verlag, Berlin, Heidelberg, pp. 347–364.
- Kim, B., Ye, J.C., 2022. Diffusion deformable model for 4D temporal medical image generation. In: *Medical Image Computing and Computer Assisted Intervention – MICCAI 2022: 25th International Conference, Singapore, September 18–22, 2022, Proceedings, Part I*. Springer-Verlag, Berlin, Heidelberg, pp. 539–548.
- 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., 2019. OASIS-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and Alzheimer disease. MedRxiv 2019-2012, Cold Spring Harbor Laboratory Press.
- Li, H., Yang, Y., Chang, M., Chen, S., Feng, H., Xu, Z., Li, Q., Chen, Y., 2022. Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing* 479, 47–59.
- Litrico, M., Guarnera, F., Giuffrida, M.V., Ravi, D., Battiato, S., 2024. TADM: Temporally-Aware Diffusion Model for Neurodegenerative Progression on Brain MRI. In: *Proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024, vol. LNCS 15002*, Springer Nature Switzerland.
- Liu, D., Kelly, M., Gong, P., 2006. A spatial-temporal approach to monitoring forest disease spread using multi-temporal high spatial resolution imagery. *Remote Sens. Environ.* 101 (2), 167–180.
- Pombo, G., Gray, R., Cardoso, M.J., Ourselin, S., Rees, G., Ashburner, J., Nachev, P., 2023. Equitable modelling of brain imaging by counterfactual augmentation with morphologically constrained 3D deep generative models. *Med. Image Anal.* 84, 102723. <http://dx.doi.org/10.1016/j.media.2022.102723>, URL <https://www.sciencedirect.com/science/article/pii/S1361841522003516>.
- Puglisi, L., Alexander, D.C., Ravi, D., 2024. Enhancing spatiotemporal disease progression models via latent diffusion and prior knowledge. In: *Linguraru, M.G., Dou, Q., Feragen, A., Giannarou, S., Glocker, B., Lekadir, K., Schnabel, J.A. (Eds.), Medical Image Computing and Computer Assisted Intervention. MICCAI 2024*, Springer Nature Switzerland, Cham, pp. 173–183.
- Ravi, D., Alexander, D.C., Oxtoby, N.P., Alzheimer's Disease Neuroimaging Initiative, 2019. Degenerative adversarial neuroimage nets: generating images that mimic disease progression. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. Springer, pp. 164–172.
- Ravi, D., Blumberg, S.B., Ingala, S., Barkhof, F., Alexander, D.C., Oxtoby, N.P., Alzheimer's Disease Neuroimaging Initiative, et al., 2022. Degenerative adversarial neuroimage nets for brain scan simulations: Application in ageing and dementia. *Med. Image Anal.* 75, 102257.
- Rondinella, A., Guarnera, F., Giudice, O., Ortis, A., Russo, G., Crispino, E., Pappalardo, F., Battiato, S., 2023. Enhancing multiple sclerosis lesion segmentation in multimodal MRI scans with diffusion models. In: *2023 IEEE International Conference on Bioinformatics and Biomedicine. BIBM, IEEE*, pp. 3733–3740.
- Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, Ł., Polosukhin, I., 2017. Attention is all you need. In: *Guyon, I., Luxburg, U.V., Bengio, S., Wallach, H., Fergus, R., Vishwanathan, S., Garnett, R. (Eds.), Advances in Neural Information Processing Systems, Vol. 30*. Curran Associates, Inc.
- Wilms, M., Bannister, J.J., Mouches, P., MacDonald, M.E., Rajashekar, D., Langner, S., Forkert, N.D., 2022. Invertible modeling of bidirectional relationships in neuroimaging with normalizing flows: Application to brain aging. *IEEE Trans. Med. Imaging* 41 (9), 2331–2347.
- Xia, T., Chartsias, A., Tsafaris, S.A., 2019a. Consistent brain ageing synthesis. In: *Shen, D., Liu, T., Peters, T.M., Staib, L.H., Essert, C., Zhou, S., Yap, P.-T., Khan, A. (Eds.), Medical Image Computing and Computer Assisted Intervention. MICCAI 2019*, Springer International Publishing, Cham, pp. 750–758.
- Xia, T., Chartsias, A., Wang, C., Tsafaris, S.A., 2019b. Learning to synthesise the ageing brain without longitudinal data. *Med. Image Anal.* 73, 102169, URL <https://api.semanticscholar.org/CorpusID:208637395>.
- Xia, T., Chartsias, A., Wang, C., Tsafaris, S.A., 2021. Learning to synthesise the ageing brain without longitudinal data. *Med. Image Anal.* 73, 102169. <http://dx.doi.org/10.1016/j.media.2021.102169>, URL <https://www.sciencedirect.com/science/article/pii/S1361841521002152>.
- Yoon, J.S., Zhang, C., Suk, H.-I., Guo, J., Li, X., 2022. SADM: Sequence-aware diffusion model for longitudinal medical image generation. In: *Information Processing in Medical Imaging*. URL <https://api.semanticscholar.org/CorpusID:254823541>.
- Young, A.L., Oxtoby, N.P., Garbarino, S., Fox, N.C., Barkhof, F., Schott, J.M., Alexander, D.C., 2024. Data-driven modelling of neurodegenerative disease progression: thinking outside the black box. *Nature Rev. Neurosci.* 25 (2), 111–130. <http://dx.doi.org/10.1038/s41583-023-00779-6>.