



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.

TADM: Temporally-Aware Diffusion Model for Neurodegenerative Progression on Brain MRI

Mattia Litrico¹, Francesco Guarnera¹, Mario Valerio Giuffrida², Daniele Ravi^{3,4,5}, and Sebastiano Battiato¹

¹ University of Catania, Catania, Italy
`mattia.litrico@phd.unict.it`

² University of Nottingham, Nottingham, United Kingdom

³ Centre for Medical Image Computing, University College London, UK

⁴ MIFT Department, University of Messina, Italy

⁵ University of Hertfordshire, Hatfield, United Kingdom

Abstract. Generating realistic images to accurately predict changes in the structure of brain MRI can be a crucial tool for clinicians. Such applications can help assess patients' outcomes and analyze how diseases progress at the individual level. However, existing methods developed for this task present some limitations. Some approaches attempt to model the distribution of MRI scans directly by conditioning the model on patients' ages, but they fail to explicitly capture the relationship between structural changes in the brain and time intervals, especially on age-unbalanced datasets. Other approaches simply rely on interpolation between scans, which limits their clinical application as they do not predict future MRIs. To address these challenges, we propose a Temporally-Aware Diffusion Model (TADM), which introduces a novel approach to accurately infer progression in brain MRIs. TADM learns the distribution of structural changes in terms of intensity differences between scans and combines the prediction of these changes with the initial baseline scans to generate future MRIs. Furthermore, during training, we propose to leverage a pre-trained Brain-Age Estimator (BAE) to refine the model's training process, enhancing its ability to produce accurate MRIs that match the expected age gap between baseline and generated scans. Our assessment, conducted on 634 subjects from the OASIS-3 dataset, uses similarity metrics and region sizes computed by comparing predicted and real follow-up scans on 3 relevant brain regions. TADM achieves large improvements over existing approaches, with an average decrease of 24% in region size error and an improvement of 4% in similarity metrics. These evaluations demonstrate the improvement of our model in mimicking temporal brain neurodegenerative progression compared to existing methods. We believe that our approach will significantly benefit clinical applications, such as predicting patient outcomes or improving treatments for patients. Our code is publicly available at <https://github.com/MattiaLitrico/TADM-Temporally-Aware-Diffusion-Model-for-Neurodegenerative-Progression-on-Brain-MRI>

Keywords: Spatial-temporal Disease Progression · Brain MRI · Diffusion Model

1 Introduction

The capability to forecast structural changes in brain MRIs over time is critical in medical imaging. The prediction of temporal brain trajectories has proven its usefulness in several applications, such as recovering missing images in longitudinal data, as a potential virtual placebo or for patient stratification [12–14, 20]; or for diagnosis and prognosis of Alzheimer’s disease [1]. While Alzheimer’s disease (AD) diagnosis commonly relies on neuropsychological and behavioural assessments, imaging data significantly aid in identifying characteristic disease effects on the brain, even in its early stages [4]. Recent advancements in Artificial Intelligence (AI) have driven the development of sophisticated spatial-temporal disease progression models [11], empowering accurate prediction of brain structural progression. In particular, generative models have been proposed to simulate future MRI scans starting from past MRI scans used as inputs. One of the AI solutions employed in this context involves the training of Generative Adversarial Networks (GANs). For example, 4D-DANINet [14] utilizes adversarial training alongside a series of biologically informed constraints to enhance the generation process. Another approach proposed in [12] uses a conditional 3D GAN with morphology constraints to predict deformations, instead of directly manipulating image pixels. More recently, approaches based on Denoising Diffusion Probabilistic Models (DDPMs) have demonstrated exceptional performance in this domain. For instance, in [19], a diffusion model is combined with a transformer network. The transformer generates a latent representation from a sequence of input MRIs, which is used to condition the generation process of the diffusion model. Another notable solution is Diffusion Deformable Model (DDM) [8], which introduces a methodology to combine a diffusion and a deformation module to generate images that interpolate between two MRI scans. Similarly, DiffuseMorph [7] trains a diffusion model to estimate a deformation field between two scans. This deformation field facilitates the translation of one input image into another, thereby enabling the generation of interpolated scans.

Due to the complexity of age-related changes in brain morphology during disease progression [2], all these approaches often fail to accurately capture the corresponding temporal evolution, facing the following limitations: (i) approaches that operate conditioning on patient’s age, such as [12, 14, 17, 18], do not explicitly capture the relationship between structural changes in brain MRIs and the time interval over which these changes occur; additionally, they require age-balanced datasets, which are often lacking in real-world applications; (ii) approaches based on interpolation, including DDM [8] and DiffuseMorph [7], have limitations in their capability to generate MRIs beyond the two input scans, reducing their relevance in clinical applications that require predicting future scans; and (iii) other approaches, like SADM [19], require longitudinal data at inference time, limiting again their application in real-world contexts where a series of scans for the patient are not available.

To address these issues, we propose TADM, a novel diffusion-based approach for brain progression modelling which operates directly on T1-weighted MRIs. Our model is trained to learn the distribution of brain changes within a specified

time interval. To achieve this, we employ a three-fold strategy. Firstly, TADM learns to predict the intensity difference between baseline and follow-up MRIs. This avoids the need to generate entirely new scans, reduces the complexity of the problem, and mitigates generation errors. Secondly, we condition the model on the age gap between the input and output scans rather than directly on the output age, aiming to better learn the relationship between observed differences and the time interval. Given that the same age gap can arise between scans acquired at different ages, conditioning on age gap avoids the necessity of including samples from every age group in the training set. This is particularly beneficial when the dataset has limited samples in some age groups. Lastly, we propose to leverage a Brain-Age Estimator (BAE) to predict age differences between two scans. During training, these predicted age gaps are used in the loss function of our model, allowing the generation of images that accurately reflect the expected age gaps between the inputs and the predictions.

We evaluate our method on the OASIS-3 dataset [9]. Specifically, we compute similarity metrics and region size in three areas of the brain to estimate the difference between real and predicted follow-up brain images. TADM improves similarity metrics by 4% and obtains the best performance on estimating region size reducing the error by 24%. Additionally, our qualitative analysis shows visual improvements in our approach in terms of better mimicking the temporal progression of brains.

In conclusion, the contributions of this work are: (i) introducing TADM, a diffusion-based approach for modelling brain progression trained on T1-weighted MRIs; (ii) learning the distribution of intensity differences between MRI scans to reduce the complexity of the generation process; (iii) conditioning on age gap to better capture the relationship between brain changes and time intervals; and (iv) proposing to leverage BAE to allow the generation of images that accurately reflect the expected age gap.

2 Background

Denoising Diffusion Probabilistic Models (DDPMs) [3] are generative models that learn a Markov chain process to convert a Gaussian distribution into data distribution. During the diffusion process, Gaussian noise is added in successive steps to a sample x_0 from the given data distribution $q(x_0)$, to convert it into a latent variable distribution $q(x_t)$, as follows:

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

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

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

Then, during the reverse process, the latent variable distribution $p_\theta(x_t)$ is transformed progressively into the data distribution $p_\theta(x_0)$, which is parameterized by θ , by training the model to learn the following Gaussian transformations:

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

$$p_\theta(x_{t-1} | x_t) = \mathcal{N}(x_{t-1}; \mu_0(x_t, t), \sigma_0(x_t, t)^2 \mathbf{I}) \quad (4)$$

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

where $\mu_0(x_t, t)$ represents the mean of the Gaussian distribution, and $\sigma_0(x_t, t)^2$ denotes the variance, at the t reverse step.

3 Proposed Method

Here we provide the details of TADM. Our pipeline is depicted in Figure 1 and consists of the following blocks: (a) a DDPM; (b) an Encoder; and (c) BAE.

During training, we use pairs of MRI scans denoted as I_{T_a} and I_{T_b} , obtained from the same patient at two time points T_a and T_b . These scans are used to compute a residual image $I_{\Delta_{a,b}} = I_{T_b} - I_{T_a}$ that we leverage to train the DDPM aimed at predicting residuals $\hat{I}_{\Delta_{a,b}}$. To generate the output scan \hat{I}_{T_b} at time T_b we then add the predicted residual $\hat{I}_{\Delta_{a,b}}$ to the baseline scan I_{T_a} at time T_a . Additionally, to achieve patient individualization, the DDPM is conditioned with a latent representation extracted by the encoder ϕ from I_{T_a} and other patient-specific data (*i.e.* cognitive status and age). Finally, we leverage BAE to predict the time interval $\Delta_{a,b} = T_b - T_a$ between I_{T_a} and the estimated \hat{I}_{T_b} to aid the generation process of the DDPM. During inference, we use the scan I_{T_a} at time T_a together with the desired time interval $\Delta_{a,b'}$ as an input to generate a future unseen scan $I_{T_{b'}}$ at time $T_{b'}$ of the same patient.

3.1 Conditioning the Diffusion Model

We condition the DDPM to generate residual images using the following information: (i) the image representation $\phi(I_{T_a})$ extracted by the encoder ϕ on the baseline; (ii) the time interval $\Delta_{a,b}$; (iii) other patient’s specific data.

Image Representation. To obtain individualization at the subject level, we condition the model using a latent representation z_a of the baseline scan I_{T_a} . In particular, the latent representation is obtained leveraging a pretrained encoder based on Residual-in-Residual Dense Blocks (RRDB) [16].

Time Interval (Age Gap). Conditioning the progression directly on age does not explicitly capture the relationship between structural changes in brain MRIs and the time interval over which these changes occur. Moreover, this strategy necessitates age-balanced datasets, which are difficult to observe in real-world scenarios. To tackle this limitation, we propose to condition the model using




Fig. 1. TADM comprises three main parts (surrounded by coloured dashed boxes). **RED:** The DDPM takes a residual image, calculated as the difference between two scans acquired at two the time points T_a and T_b , on which random noise ϵ is applied. The DDPM is trained to denoise the residual image through several diffusion steps. The denoised residual image $\hat{I}_{\Delta_{a,b}}$ and the scan I_{T_a} are summed together to estimate the scan \hat{I}_{T_b} at time T_b . **BLUE:** Here, we encode the scan I_{T_a} to extract a representation z_a used to condition the DDPM, in conjunction with other patient-specific data. **GREEN:** The estimated \hat{I}_{T_b} is provided to the encoder to extract the features z_b that, together with the previously extracted features z_a , are provided as inputs to a BAE to predict the time interval $\hat{\Delta}_{a,b}$. The PADLOCK indicates a model with frozen parameters.

the age gap between scans $\Delta_{a,b}$. Since the same age gap can occur between scans acquired at different ages, conditioning on the age gap eliminates the need for including samples from every age group in the training set. This is particularly advantageous when the dataset has limited samples in certain age groups. In our implementation, we encode $\Delta_{a,b}$ using positional encoding [15] before incorporating it into the model.

Patient-Specific Data. We also condition the model using the patient’s cognitive status (D) and age at baseline (A). Indeed, age at baseline is crucial information as diseases progress at different rates over the course of ageing. However, by only using $\Delta_{a,b}$, our model would not capture such age-related progression information. Note that this is different from previous work on age conditioning, as they use age at the prediction rather than age at baseline.

3.2 Leveraging BAE to improve the temporal awareness

To encourage the model to generate images that accurately reflect the expected age gap between the input and the prediction, we propose to leverage a BAE model [6] to predict the age gap between two MRI scans. BAE is trained offline on our training set and it is not further fine-tuned during the DDPM training. Specifically, given the baseline I_{T_a} and the generated \hat{I}_{T_b} scans, the predicted age gap is computed as $\hat{\Delta}_{a,b} = \Psi(\Phi(\hat{I}_{T_b})) - \Psi(\Phi(I_{T_a}))$, where Ψ is the BAE model. This information will be used later to train the DDPM and improve the generation. In particular, if the DDPM generates a scan \hat{I}_{T_b} that closely approximates the ground truth I_{T_b} , the predicted age gap $\hat{\Delta}_{a,b}$ should closely match the actual age gap $\Delta_{a,b}$. Any error in the estimation of the age gap

Table 1. Comparison study: Results showing the performance in terms of image-base and region size in comparison to other methods.

Method			Region Size Error (%) ↓			
	SSIM ↑	PSNR ↑	Gray Matter	White Matter	Cerebrospinal Fluid	Total Brain
DiffuseMorph [7]	0.68	19.67	10.40 ± 3.45	3.49 ± 2.58	4.65 ± 2.80	46.30 ± 7.51
4D-DaniNet [14]	0.65	16.99	2.21 ± 1.08	2.57 ± 1.98	3.12 ± 3.65	9.31 ± 8.72
DDM [8]	0.69	19.59	2.44 ± 1.35	3.05 ± 2.74	4.37 ± 3.12	10.85 ± 11.64
TADM (Proposed)	0.72	20.51	1.69 ± 1.54	1.85 ± 2.20	2.70 ± 2.29	6.84 ± 5.00

will be corrected through backpropagation in the diffusion model to refine the generation process.

3.3 Overall Framework

In this section, we will provide a complete overview of the framework.

Training. During the diffusion process, the DDPM is trained to predict the noise ϵ added to the input $I_{\Delta_{a,b}}$. This is obtained by minimising the following objective:

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

where G_{θ} is the DDPM parametrized by θ and t is the diffusion timestep.

Additionally, as mentioned in Section 3.2, we incorporate the output from BAE as an additional term in the loss function of the DDPM. Specifically, we define the loss on the expected brain age gap as follows:

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

Note that the gradient of this loss updates only the DDPM parameters θ .

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

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

Inference. The model takes as input a baseline MRI \mathbf{I}_{T_a} and predicts a follow-up MRI $\mathbf{I}_{T_{b'}}$ with an arbitrarily time interval $\Delta_{a,b'}$ with respect to the baseline. The reverse process starts from a Gaussian noise variable X_T that is progressively denoised through $G_{\theta}(X_t, t; z_a, \Delta_{a,b'}, A, D)$. The predicted residual image $\hat{\mathbf{I}}_{\Delta_{a,b'}}$ is then added to the baseline MRI to generate the predicted follow-up.

4 Experimental Results

Implementation Details. To evaluate our approach, we use 2,535 T1-weighted (T1w) brain MRIs from 634 subjects from the OASIS-3 dataset [9]. Scans were acquired over a period of 15 years and subjects are between 42 and 95 years old,




Fig. 2. Comparison of the temporal progression on a 66-year-old subject with AD, obtained by our approach against other state-of-the-art methods. We show predicted slice-MRIs on the left and the corresponding error with the subject’s real brain MRI on the right.

classified as cognitively normal, Mild Cognitive Impairment (MCI), and AD. We apply linear registration through the MNI152 template and skull removal using the FSL library [5] to all the MRIs. The dataset is divided into training set (70%), validation set (10%) and test set (20%). We used the validation set to optimize the hyperparameters of BAE.

Following [10], we adopt the U-Net as architecture of the diffusion model G_θ and the same hyperparameters. Results of all baseline methods are obtained using their publicly available codes.

Evaluation Metrics. We compute image-based metrics and region size in relevant brain areas to assess the performance of our method. Specifically, for the image-based metrics we use the Structural Similarity Index Measure (SSIM) and Peak Signal-to-Noise Ratio (PSNR) between the generated and the actual MRIs. On the other side, region sizes are used to evaluate the accuracy of disease progression. The regions considered in our experiment are: (i) Gray Matter, (ii) White Matter, and (iii) Cerebrospinal Fluid (CSF). We use the FMRIB’s Automated Segmentation Tool [5] to compute the area of the regions which are expressed as percentages of the Total Brain to account for individual differences. The error is calculated as the mean absolute error between the area on the region from the predicted scan and ground truth MRIs.

Comparison Results. In Table 1, we present quantitative results obtained by our method compared to other state-of-the-art approaches [7, 8, 14]. The results demonstrate that we outperform the state-of-the-art in SSIM and PSNR by +0.03 and +0.84, respectively. Results on the size of the region show that our method achieves the lowest error in all the considered brain regions. In partic-

Table 2. Additional analysis: results showing the contribution of the different components of TADM, including BAE, the use of patient’s specific data. Finally, we also show the impact of using the conditioning on age rather than age gap as we proposed.

Method	SSIM \uparrow PSNR \uparrow		Region Size Error (%) \downarrow			
			Gray Matter	White Matter	Cerebrospinal Fluid	Total Brain
TADM w/o patient’s data	0.71	20.32	1.78 ± 1.44	1.97 ± 2.14	2.72 ± 1.98	7.85 ± 5.17
TADM w/o BAE	0.69	20.08	2.44 ± 2.12	2.02 ± 2.13	3.85 ± 3.67	9.77 ± 8.23
TADM w/ age cond.	0.68	19.71	4.12 ± 3.48	4.98 ± 2.45	5.65 ± 3.32	11.95 ± 7.34
TADM	0.72	20.51	1.69 ± 1.54	1.85 ± 2.20	2.70 ± 2.29	6.84 ± 5.00

ular, we reduce the error on grey and white matter regions by approximately 30% and 8%, respectively. Regarding the CSF and Total Brain, the reduction of error is nearly 29%, demonstrating that our model generates high-quality follow-up scans compared to the current state-of-the-art. Additionally, in Figure 2, we show one example of qualitative results in predicting MRI scans at different time points. The figure shows that TADM offers a better approximation of the brain’s temporal evolution compared to other methods. Specifically, our predictions depict a notably accurate alignment of the ventricular expansion over time. Finally, TADM exhibits fewer minor disparities in the brain cortex compared to other methods.

Additional Analysis. In this section, we explore the impact of integrating BAE into our pipeline and conditioning the model on patient’s patient-specific data. Additionally, we evaluate the model’s performance when it is conditioned on either the age or the age gap as proposed in our method. In the first row in Table 2, we show the results of our method without incorporating the patient’s specific data. This outcome highlights that the absence of patient-specific data results in a minimal reduction in performance, indicating that the patient-specific data contributes minimally to the pipeline. When BAE is not used in our pipeline (second row in Table 2), we notice an evident decrease in performance. This indicates that leveraging BAE is essential to support the generation process. Lastly, in the third row, we observe that conditioning the model on age rather than age gap drastically reduces performance by an average of approximately 60%, demonstrating the effectiveness of our idea of conditioning on age gaps.

5 Conclusion and Future Works

In this paper, we propose TADM, a novel approach designed to accurately mimic brain neurodegenerative progression in MRIs. We evaluated TADM on the OASIS-3 dataset, demonstrating superior performance compared to existing approaches. TADM focuses on 2D scans, as we strive to develop a new data-driven pipeline capable of improving the accuracy of current methods. Nonetheless, our method can be easily extended to 3D scans. Furthermore, our pipeline presents exciting prospects for data augmentation of underrepresented samples

in medical imaging datasets. This feature holds significant promise, especially in the context of less common and more expensive modalities, *e.g.* PET and CT scans, where the generation of synthetic samples remains a critical challenge.

6 Acknowledgment

Francesco Guarnera is funded by MUR within the framework of PNRR PE0000013, under project “Future Artificial Intelligence Research – FAIR”. Data were provided by OASIS-3 [9].

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

References

1. Bowles, C., Gunn, R., Hammers, A., Rueckert, D.: Modelling the progression of alzheimer’s disease in mri using generative adversarial networks. *Medical Imaging 2018: Image Processing* **10574**, 105741K (2018)
2. Dickstein, D.L., Kabaso, D., Rocher, A.B., Luebke, J.I., Wearne, S.L., Hof, P.R.: Changes in the structural complexity of the aged brain. *Aging cell* **6**(3), 275–284 (2007)
3. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in neural information processing systems* **33**, 6840–6851 (2020)
4. 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.: NIA-AA research framework: toward a biological definition of alzheimer’s disease. *Alzheimer’s & Dementia* **14**(4), 535–562 (2018)
5. Jenkinson, M., Beckmann, C.F., Behrens, T.E., Woolrich, M.W., Smith, S.M.: FSL. *Neuroimage* **62**(2), 782–790 (2012)
6. Jónsson, B.A., Bjornsdottir, G., Thorgeirsson, T., Ellingsen, L.M., Walters, G.B., Gudbjartsson, D., Stefansson, H., Stefansson, K., Ulfarsson, M.: Brain age prediction using deep learning uncovers associated sequence variants. *Nature communications* **10**(1), 5409 (2019)
7. Kim, B., Han, I., Ye, J.C.: 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*. p. 347–364. Springer-Verlag, Berlin, Heidelberg (2022)
8. Kim, B., Ye, J.C.: 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*. p. 539–548. Springer-Verlag, Berlin, Heidelberg (2022)
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. Li, H., Yang, Y., Chang, M., Chen, S., Feng, H., Xu, Z., Li, Q., Chen, Y.: Srdiff: Single image super-resolution with diffusion probabilistic models. *Neurocomputing* **479**, 47–59 (2022)

11. Liu, D., Kelly, M., Gong, P.: A spatial–temporal approach to monitoring forest disease spread using multi-temporal high spatial resolution imagery. *Remote sensing of environment* **101**(2), 167–180 (2006)
12. 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). <https://doi.org/https://doi.org/10.1016/j.media.2022.102723>, <https://www.sciencedirect.com/science/article/pii/S1361841522003516>
13. Ravi, D., Alexander, D.C., Oxtoby, N.P., Initiative, A.D.N.: Degenerative adversarial neuroimage nets: generating images that mimic disease progression. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention*. pp. 164–172. Springer (2019)
14. 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)
15. Vaswani, A., Shazeer, N., Parmar, N., Uszkoreit, J., Jones, L., Gomez, A.N., Kaiser, L.u., Polosukhin, I.: 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. (2017)
16. Wang, X., Yu, K., Wu, S., Gu, J., Liu, Y., Dong, C., Qiao, Y., Loy, C.C.: Esrgan: Enhanced super-resolution generative adversarial networks. In: Leal-Taixé, L., Roth, S. (eds.) *Computer Vision – ECCV 2018 Workshops*. pp. 63–79. Springer International Publishing, Cham (2019)
17. Xia, T., Chartsias, A., Tsiftaris, S.A.: 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*. pp. 750–758. Springer International Publishing, Cham (2019)
18. Xia, T., Chartsias, A., Wang, C., Tsiftaris, S.A.: Learning to synthesise the ageing brain without longitudinal data. *Medical Image Analysis* **73**, 102169 (2021). <https://doi.org/https://doi.org/10.1016/j.media.2021.102169>, <https://www.sciencedirect.com/science/article/pii/S1361841521002152>
19. Yoon, J.S., Zhang, C., Suk, H.I., Guo, J., Li, X.: Sadm: Sequence-aware diffusion model for longitudinal medical image generation. In: *Information Processing in Medical Imaging* (2022), <https://api.semanticscholar.org/CorpusID:254823541>
20. 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* **25**(2), 111–130 (Jan 2024). <https://doi.org/10.1038/s41583-023-00779-6>, <http://dx.doi.org/10.1038/s41583-023-00779-6>