

LEARNING PATIENT-SPECIFIC DISEASE DYNAMICS WITH LATENT FLOW MATCHING FOR LONGITUDINAL IMAGING GENERATION

Hao Chen^{1*} **Rui Yin**^{2*} **Yifan Chen**¹ **Qi Chen**³ **Chao Li**^{1,4}

¹University of Cambridge, ²Nanjing First Hospital, Nanjing Medical University

³Johns Hopkins University ⁴University of Dundee

ABSTRACT

Understanding disease progression is a central clinical challenge with direct implications for early diagnosis and personalized treatment. While recent generative approaches have attempted to model progression, key mismatches remain: disease dynamics are inherently continuous and monotonic, yet latent representations are often scattered, lacking semantic structure, and diffusion-based models disrupt continuity with a random denoising process. In this work, we propose to treat the disease dynamics as a velocity field and leverage Flow Matching (FM) to align the temporal evolution of patient data. Unlike prior methods, it captures the intrinsic dynamics of disease, making the progression more interpretable. However, a key challenge remains: in latent space, autoencoders (AEs) do not guarantee alignment across patients or correlation with clinical-severity indicators (*e.g.*, age and disease conditions). To address this, we propose to learn patient-specific latent alignment, which enforces patient trajectories to lie along a specific axis, with magnitude increasing monotonically with disease severity. This leads to a consistent and semantically meaningful latent space. Together, we present Δ -LFM, a framework for modeling patient-specific latent progression with flow matching. Across three longitudinal MRI benchmarks, Δ -LFM demonstrates strong empirical performance and, more importantly, offers a new framework for interpreting and visualizing disease dynamics.

1 INTRODUCTION

Modeling disease progression is a fundamental problem in healthcare (Suk & Shen, 2013; Wang et al., 2020), with profound implications for early detection and individualized treatment. Disease does not evolve uniformly: differences in anatomy, genetics, and environment give rise to heterogeneous trajectories, which are an intrinsic property of disease dynamics (Real & Biek, 2007; Parratt et al., 2016; Li et al., 2025). Accounting for this heterogeneity is essential for accurately characterizing temporal changes, identifying preclinical biomarkers, and designing personalized therapies that can improve clinical outcomes (Frisoni et al., 2010). However, most progression models capture only population-level trends, overlooking individualized variation (Rokuss et al., 2025; Yang et al., 2025b). This limitation obscures early-stage signals and misaligns disease severity with underlying anatomy, ultimately reducing clinical utility. Capturing patient-specific dynamics is therefore critical, but remains an underexplored challenge (Lai et al., 2025).

Recent advances in generative modeling have opened new opportunities for studying disease progression. Traditional biomarker-based approaches often reduce complexity to coarse measures (*e.g.*, volume changes (Jack Jr et al., 2003; Tabrizi et al., 2011) or disease rating (O’Bryant et al., 2008; Young et al., 2020)), which can obscure heterogeneous disease dynamics. In contrast, longitudinal medical image generation offers a principled framework for visualizing symptomatic transitions and individualized patterns of anatomical change. Rather than limiting analysis to numeric outcomes, synthesized images yield interpretable and more informative visual representations that

*Equal contribution. Correspondence to: hc666@cam.ac.uk.

support structured analyses and offer more actionable insights for clinical practice (Wu et al., 2024; Chen et al., 2024; 2025).

Several methods have attempted to model disease progression. For instance, DiffuseMorph (Kim et al., 2022) introduces a morph field to describe pixel-wise displacements, and TADM (Litrico et al., 2024) incorporates temporal awareness by predicting residual images conditioned on age gaps. Despite these advances, generating temporally consistent and individualized disease trajectories remains an open problem. Beyond high-fidelity image synthesis, it is crucial for generative models to preserve fine-grained anatomical structures and accurately capture patient-specific dynamics. Although BrLP (Puglisi et al., 2024) takes a step toward personalization by modeling volume changes, such guidance is coarse and provides only limited control over individualized trajectories.

Contributions: Motivated by these gaps, we propose a personalized generative framework for modeling disease progression based on flow matching. Our approach, termed Progression Latent Flow Matching (Δ -LFM), is designed to generate trajectories at arbitrary future time points within a patient-specific latent space, ensuring temporal consistency and clinical interpretability. The key contributions are summarized as follows:

- We reformulate conventional flow matching, which typically models dynamics from $t = 0$ to $t = 1$, into a formulation where the time variable t explicitly encodes the future time gap. This extension transforms the horizon from a normalized range $[0, 1]$ to a meaningful temporal range $[0, T]$, enabling prediction at arbitrary future time points T with consistent temporal semantics.
- We introduce a patient-specific latent space, regularized by a new *ArcRank* loss, which enforces chronology-aware alignment and captures individualized disease dynamics.
- We show that the learned latent space supports interpretable visualization of patient trajectories. Notably, although disease severity is not used for training, the latent representations naturally reflect severity levels, providing clinically meaningful structure.
- We further validate Δ -LFM on longitudinal MRI benchmarks and propose a progression-specific evaluation metric tailored to disease modeling called Δ -RMAE. Results demonstrate improved imaging fidelity and more accurate alignment with actual disease progression.

2 RELATED WORK

Generative models for disease progression. Generative modeling has emerged as a powerful strategy for simulating longitudinal disease changes. Early efforts relied on generative adversarial networks (GANs), either by introducing morphological priors to simulate aging effects (Ravi et al., 2019) or by predicting deformation fields with 3D conditional GANs instead of direct intensity values (Ravi et al., 2022).

More recently, diffusion models (Ho et al., 2020) have become the dominant paradigm due to their superior fidelity and stability. Extensions of this framework introduce temporal awareness and anatomical constraints. Sequence-Aware Diffusion Model (SADM) leverages multiple scans to autoregressively generate future MRIs (Yoon et al., 2023), while DiffuseMorph (Kim et al., 2022) produces smooth voxel-wise deformation fields across timepoints, thereby preserving topology. Progressive Image Editing (PIE) (Liang et al., 2023) edits disease-related regions in medical images using text-guided diffusion models to reflect how a patient’s condition evolves over time. Kyung et al. (Kyung et al., 2024) adopt electronic health records (EHRs), including medical history, to model temporal changes and enable medically meaningful generation. Temporally-Aware Diffusion Model (TADM) conditions on age maps and adopts a residual prediction strategy, modeling progression as incremental changes rather than absolute reconstructions (Litrico et al., 2024).

Conditioning strategies further enhance subject specificity and anatomical fidelity. Brain Latent Progression (BrLP) employs ControlNet with volumetric ratio conditioning to synthesize individualized disease trajectories while integrating population priors (Puglisi et al., 2024; 2025). Similarly, BrainMRDiff fuses multiple structural masks (e.g., white matter, gray matter, ventricles) into a unified conditioning representation, substantially improving anatomical realism (Bhattacharya et al., 2025). ImageFlowNet (Liu et al., 2025a) learns multiscale joint patient embeddings and position-parameterized neural flow fields to forecast full image-level trajectories.

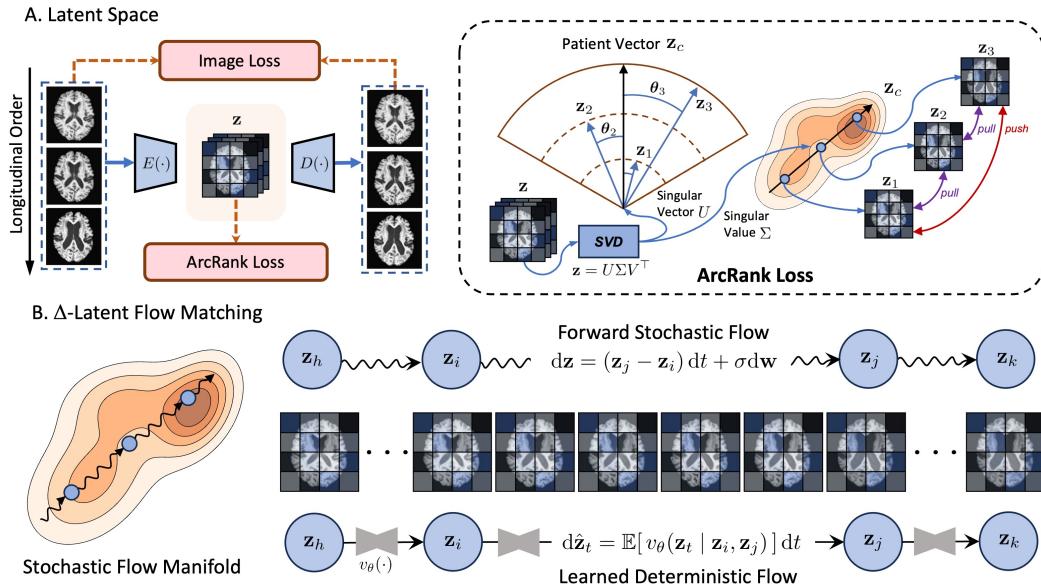


Figure 1: Δ -LFM overview. Δ -LFM operates in the latent space. In the first stage, an autoencoder constructs the latent space with ArcRank Loss to capture patient-specific disease trajectories. In the second stage, flow matching predicts disease progression along the trajectory over the total time horizon T/dt , where, for example, $T = j - i$, to adapt the actual prediction span rather than relying on a fixed interval.

In contrast to these approaches, which primarily focus on realism and population-informed priors, our method targets patient-specific disease trajectories while explicitly modeling heterogeneity through a flow-matching generative framework.

Subject-level consistency and biomedical plausibility. Ensuring subject-specific consistency and biomedical plausibility is a central challenge in modeling synthetic disease progression. Early approaches such as the Longitudinal Variational Autoencoder (Ramchandran et al., 2021) and Longitudinal Self-Supervised Learning (LNE) (Zhao et al., 2021) focus on learning patient-temporal dynamics by constructing temporally coherent latent spaces, thereby producing individualized representations for downstream tasks.

Subsequent works introduced explicit mechanisms to regularize longitudinal consistency. Ren *et al.* (Ren et al., 2023) enforce subject-level coherence through local and multi-scale spatio-temporal features within a self-supervised framework. In a complementary direction, Treatment-Aware Diffusion (TaDiff) (Liu et al., 2025b) incorporates prior scans and treatment information to jointly generate future MRIs and tumor segmentations, thereby improving biological plausibility by conditioning on clinically relevant variables.

3 METHODOLOGY

Our proposed approach consists of two key components: (i) patient-specific latent learning via *Ar**cRank* (Sec. 3.2), and (ii) arbitrary-time progression modeling via Δ -LFM (Sec. 3.3).

Before delving into these components, we first present the preliminaries in Sec. 3.1.

3.1 PRELIMINARIES

Flow matching (Lipman et al., 2023; Tong et al., 2024) trains continuous normalizing flows by directly aligning vector fields instead of relying on stochastic denoising objectives as in diffusion models. Let π_0 and π_1 denote the source and target distributions, respectively. Flow matching defines a family of intermediate distributions $\{\pi_t\}_{t \in [0,1]}$ along a continuous trajectory connecting

π_0 and π_1 . The goal is to learn a time-dependent velocity field $v_\theta(\mathbf{x}, t)$ such that the solution of the ordinary differential equation (ODE) (Lu et al., 2022) transports π_0 exactly to π_1 at $t = 1$:

$$\frac{d\mathbf{x}_t}{dt} = v_\theta(\mathbf{x}_t, t), \quad (1)$$

where the velocity field $v_\theta(\mathbf{x}, t)$ is trained to ensure that samples from π_0 flow continuously into π_1 .

Unlike score-based diffusion models, which approximate the score function $\nabla_{\mathbf{x}} \log \pi_t(\mathbf{x})$ via noisy denoising objectives, flow matching bypasses score estimation by directly regressing the neural vector field $v_\theta(\mathbf{x}, t)$ onto analytically known target fields. For example, under conditional flow matching (Lipman et al., 2023), one defines a reference path $\mathbf{x}_t = (1-t)\mathbf{x}_0 + t\mathbf{x}_1$ between samples $\mathbf{x}_0 \sim \pi_0$ and $\mathbf{x}_1 \sim \pi_1$, yielding a target velocity

$$v^*(\mathbf{x}_t, t) = \mathbf{x}_1 - \mathbf{x}_0, \quad (2)$$

which is constant along the trajectory. The training objective then minimizes

$$\mathcal{L}_{\text{FM}}(\theta) = \mathbb{E}_{\mathbf{x}_0 \sim \pi_0, \mathbf{x}_1 \sim \pi_1, t \sim \mathcal{U}[0, 1]} \|v_\theta(\mathbf{x}_t, t) - v^*(\mathbf{x}_t, t)\|^2, \quad (3)$$

which yields stable training and ensures the learned flow recovers the correct marginal π_1 at $t = 1$.

To further incorporate stochasticity (Albergo et al., 2023), the deterministic flow can be extended to a stochastic differential equation of the form

$$d\mathbf{z}_t = v^*(\mathbf{z}_t, t) dt + \sigma d\mathbf{w}_t, \quad (4)$$

where \mathbf{w}_t is a Wiener process and σ controls the noise scale. This extension allows flow matching to model both deterministic transport and stochastic variability, providing greater flexibility in capturing complex data distributions.

3.2 ARC RANK: LONGITUDINAL ALIGNMENT OF PATIENT LATENT

To enhance sensitivity to patient-specific dynamics, we introduce a chronologically-aware latent representation learned via a new contrastive objective, termed *ArcRank Loss*. This objective aligns representations across time for the same individual by enforcing both angular consistency and temporal ordering in the embedding space. Let \mathbf{z} denote the latent representation of an input \mathbf{x} , extracted from a Variational Autoencoder (VAE) (Kingma & Welling, 2013) encoder $E(\cdot)$. The loss encourages

$$\angle(\mathbf{z}_i) \approx \angle(\mathbf{z}_j) \quad \wedge \quad \|\mathbf{z}_i\| \prec \|\mathbf{z}_j\| \quad \text{if } t_i < t_j, \quad \mathbf{z}_i, \mathbf{z}_j \in \mathcal{P}^k, \quad (5)$$

where \mathcal{P}^k denotes the latent feature set of patient k , and t here represents the imaging capture time. Here, $\angle(\cdot)$ represents the direction of a latent vector, and $\|\cdot\|$ is its magnitude. Intuitively, samples from the same patient are encouraged to align along a consistent latent direction, while the trajectory respects chronological order (further along the trajectory corresponds to later disease stages).

In practice, we extract these components via Singular Value Decomposition (SVD) (Klema & Laub, 1980) and write

$$\angle \mathbf{z} = U, \quad \|\mathbf{z}\| = \Sigma, \quad \text{where } U\Sigma V^\top = \text{SVD}(\mathbf{z}). \quad (6)$$

where U captures the orientation (angle) and Σ encodes the scaling (norm).

ArcRank Loss Computation. For a given patient k , let $\mathcal{P}^k = \{\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_n\}$ denote latent features at times $t_1 < t_2 < \dots < t_n$. We decompose each feature via SVD:

$$\angle(\mathbf{z}_t) = U_t, \quad \|\mathbf{z}_t\| = \Sigma_t. \quad (7)$$

The angular consistency loss encourages stability of feature orientations across time, while the temporal ranking loss enforces monotonic growth of feature magnitudes:

$$\mathcal{L}_{\text{Arc}} = \sum_{\substack{i < j \\ \text{same patient}}} |U_i - U_j|, \quad \mathcal{L}_{\text{Rank}} = \sum_{\substack{i < j \\ \text{same patient}}} \max(0, m - (\Sigma_j - \Sigma_i)), \quad t_i < t_j, \quad (8)$$

where $m > 0$ is a margin hyperparameter. The final objective combines the two components:

$$\mathcal{L}_{\text{ArcRank}} = \lambda_{\text{arc}} \mathcal{L}_{\text{Arc}} + \lambda_{\text{rank}} \mathcal{L}_{\text{Rank}}, \quad (9)$$

where $\lambda_{\text{arc}}, \lambda_{\text{rank}} > 0$ balance the contributions.

In practice, the ranking term $\mathcal{L}_{\text{Rank}}$ tends to push timepoints for the same patient increasingly far apart, and the margin m is shared across all pairs (i, j) , i.e., it is not adaptively optimized for different temporal gaps. To mitigate this effect, we introduce an additional pull term that penalizes excessive separation in magnitude between temporally adjacent features:

$$\mathcal{L}_{\text{Pull}} = |\Sigma_j - \Sigma_i|. \quad (10)$$

We augment the ranking objective as

$$\tilde{\mathcal{L}}_{\text{Rank}} = \mathcal{L}_{\text{Rank}} + \mathcal{L}_{\text{Pull}}, \quad (11)$$

using the same overall weight as $\mathcal{L}_{\text{Rank}}$ in the total loss. This pull term softly encourages consecutive timepoints to remain close in latent space, while still allowing the hinge-based push term to enforce the desired temporal ordering.

To stabilize training, for each pair (i, j) used in the ArcRank loss, we apply the stop-gradient operator $\text{sg}(\cdot)$ to the representation of i .

3.3 Δ -LFM: FLOW MATCHING ALONG PATIENT TRAJECTORIES

While ArcRank enforces patient-specific chronological alignment, it does not explicitly capture how latent states evolve over time. To address this limitation, we introduce Δ -LFM, a flow-matching objective that learns smooth mappings across arbitrary time intervals within a patient’s trajectory. Unlike standard flow matching, which interpolates between fixed source and target distributions (e.g., from $0 \rightarrow 1$), Δ -LFM models temporal transitions on disease-specific time scales T , where T denotes the future time relative to the baseline. For clarity, we refer to the conventional $[0, 1]$ interval as *Physical Sampling*, and to our proposed $[0, T]$ formulation as *Temporal Sampling*.

Formally, given a patient k , let $\langle \mathbf{z}_i, \mathbf{z}_j \rangle$ denote paired latent states at times $t_i < t_j$. We define the latent transition field $v_\theta(\cdot)$, parameterized by a neural network, which approximates the continuous-time velocity of latent dynamics. The goal is to match this learned velocity field to the empirical difference between latent features across time:

$$v^*(i, j) \approx \frac{\mathbf{z}_j - \mathbf{z}_i}{t_j - t_i}, \quad t_j > t_i, \quad (12)$$

The associated flow-matching loss is then given by

$$\mathcal{L}_{\text{LFM}}(\theta) = \sum_{i < j} \left| v_\theta(i, j) - v^*(i, j) \right|^2. \quad (13)$$

Inference. Given an input state \mathbf{x}_i observed at time t_i , the encoder maps it to a latent representation \mathbf{z}_i . To predict the target state \mathbf{x}_j at a later time t_j , Δ -LFM integrates the learned velocity field forward in time. Specifically, we fix a discretization step size $dt > 0$ (e.g., $dt = 0.01$), so that the total number of integration steps is

$$N = \frac{t_j - t_i}{dt}. \quad (14)$$

At each step, the latent state is updated according to the learned dynamics

$$\mathbf{z}_{i+dt} = \mathbf{z}_i + dt \cdot v_\theta(\mathbf{z}_i, t_i), \quad (15)$$

with time advanced as $t_{i+dt} = t_i + dt$. Repeated application of this update yields an approximation of the latent state at t_j , which is then decoded to produce the predicted observation $\hat{\mathbf{x}}_j$.

3.4 DISCUSSION

The use of Flow Matching. We use flow matching because it naturally models continuous trajectories while keeping intermediate states interpretable, which is crucial in our setting. It learns a continuous velocity field between observed endpoints and allows direct sampling of intermediate states, without relying on long numerical integrations. Compared to latent ODEs, this avoids sensitivity to ODE solver choices and reduces computational cost. While rectified flow methods are effective, they are typically built on diffusion trajectories; in contrast, flow matching gives a more direct and tractable formulation of the probability path that is better aligned with our goal of faithful temporal interpolation.

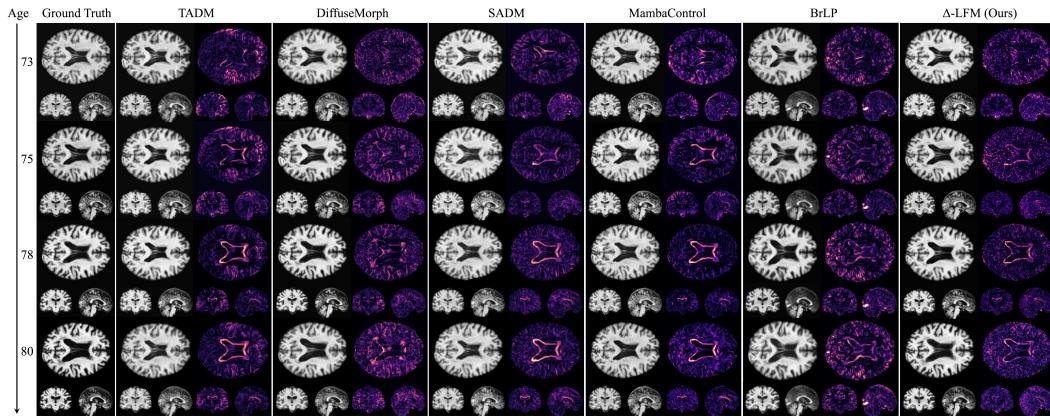


Figure 2: **Method comparison.** The first column shows the ground truth. All predictions are initialized from the same baseline scan at age 71. Odd columns show predicted MRIs; even columns show residual maps relative to the ground truth. Residuals use the `magma` colormap, scaled to the 1st to 99th percentiles of the residual distribution, enhancing error visibility around the lateral ventricles (central butterfly-shaped region) and cortical gray matter (outer surface).

Latent-space trajectories reasoning. One central assumption in our framework is that disease progression can be represented as a straight line in latent space, structured along a linear trajectory where the magnitude encodes disease stage and the direction preserves patient identity. While this design simplifies prediction, it raises an important question: can a straight line adequately capture both the shared and divergent aspects of disease dynamics across patients? We argue that this concern is valid. In our work, however, the linear trajectory is not intended to reproduce the full biological complexity of disease, but rather to provide a stable and interpretable scaffold that highlights the primary mode of variation corresponding to disease advancement.

Temporal uneven progression. We formulate flow matching (FM) over the interval $[0, T]$, interpreting the index as a progression coordinate rather than physical clock time. However, although this formulation accounts for time differences, it implicitly assumes that progression is evenly distributed along the disease trajectory. In practice, progression is often uneven: extended periods of stability may be punctuated by sudden changes, or vice versa. The constant-velocity assumption in FM fails to capture this heterogeneity, effectively treating patient scans as uniformly spaced along the trajectory. To address this limitation, we introduce conditioning on the target time T together with patient-specific attributes (e.g., sex, age, and clinical status) to enable the model to capture heterogeneous progression patterns that vary substantially over time.

Design details of ArcRank loss. The proposed ArcRank loss uses SVD to jointly capture the direction and magnitude of latent trajectories. Alternatives such as cosine similarity for angles and absolute values for magnitudes treat these aspects separately and can become unstable when latent features are noisy or vary in scale. In contrast, SVD provides an orthogonal basis in which singular vectors naturally describe trajectory directions and singular values represent their magnitudes. This unified representation is compact and numerically stable, allowing ArcRank loss to align trajectories along meaningful progression directions while penalizing inconsistent scaling.

4 EXPERIMENTS

4.1 EXPERIMENT SETTINGS

We build our study on three major Alzheimer’s disease (AD) cohorts: ADNI (Petersen et al., 2010), OASIS-3 (LaMontagne et al., 2019), and AIBL (Ellis et al., 2009). Every MRI undergoes the same preprocessing pipeline: N4 bias correction (Tustison et al., 2010), skull stripping (Hoopes et al., 2022), resampling to 1.5 mm^3 voxels. We then normalize tissue contrasts and register all scans to a refined standard brain template. We use a random 80/5/15 split for train/validation/test.

Table 1: Quantitative comparison of methods in terms of image quality metrics, including PSNR (dB) and SSIM ($\times 1.0E2$). Results are reported as mean \pm standard deviation. *Type* denotes the prediction strategy: “Direct” generates outputs directly, “Deform” estimates a transformation field, “Noise” applies diffusion-based denoising, “ Δ Noise” predicts and denoises progression residuals, and “ Δ Flow” models progression through gradual residual prediction.

Method	Type	ADNI		AIBL		OASIS	
		PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow	PSNR \uparrow	SSIM \uparrow
CardiacAging (Campello et al., 2022)	Direct	27.78 \pm 1.49	92.04 \pm 0.99	28.41 \pm 1.41	90.30 \pm 0.90	26.23 \pm 1.45	86.17 \pm 2.06
CounterSynth (Pombo et al., 2023)	Direct	28.24 \pm 1.31	92.84 \pm 0.95	28.96 \pm 1.25	92.10 \pm 0.91	26.79 \pm 1.42	86.45 \pm 1.98
DiffuseMorph (Kim et al., 2022)	Deform	29.56 \pm 1.63	93.57 \pm 0.93	29.17 \pm 1.54	92.70 \pm 1.02	28.13 \pm 1.36	88.16 \pm 1.69
SADM (Yoon et al., 2023)	Noise	26.94 \pm 2.28	85.15 \pm 2.72	27.97 \pm 1.91	89.40 \pm 1.78	26.74 \pm 1.80	86.36 \pm 2.37
TADM (Litrico et al., 2024)	Δ Noise	27.89 \pm 1.91	90.78 \pm 1.54	28.58 \pm 1.77	91.70 \pm 1.01	26.29 \pm 1.52	85.51 \pm 1.93
ImageFlowNet (Liu et al., 2025a)	Flow	28.37 \pm 1.18	92.96 \pm 0.88	29.08 \pm 1.12	92.23 \pm 0.84	27.92 \pm 1.29	87.63 \pm 1.82
BrLP (Puglisi et al., 2024)	Noise	28.51 \pm 1.77	91.52 \pm 1.31	28.96 \pm 1.36	92.21 \pm 0.97	27.98 \pm 1.40	87.93 \pm 1.44
MambaControl (Yang et al., 2025a)	Noise	29.72 \pm 1.04	93.60 \pm 0.96	29.86 \pm 1.21	93.17 \pm 0.98	28.24 \pm 1.28	88.30 \pm 1.52
Δ -LFM (Ours)	Δ Flow	30.59 \pm 0.89	94.62 \pm 0.85	30.52 \pm 1.03	93.92 \pm 0.88	29.01 \pm 1.19	89.36 \pm 1.29

We benchmark against a broad spectrum of recent approaches. Direct-prediction methods include CardiacAging (Campello et al., 2022) and CounterSynth (Pombo et al., 2023). Deformation-based modeling represented by DiffuseMorph (Kim et al., 2022). Diffusion-based methods span multiple variants: SADM (Yoon et al., 2023), which denoises full images; TADM (Litrico et al., 2024), which denoises residuals; and two ControlNet-based (Zhang et al., 2023) models, BrLP (Puglisi et al., 2024) and MambaControl (Yang et al., 2025a).

Metrics. We evaluate reconstruction quality with two standard image-level metrics: Peak Signal-to-Noise Ratio (PSNR) and Structural Similarity Index Measure (SSIM). To capture anatomical fidelity, we followed BrLP (Puglisi et al., 2024) to report region-level mean absolute error (MAE) over clinically relevant structures (hippocampus, amygdala, lateral ventricles, cerebrospinal fluid (CSF), and thalamus).

However, measuring image similarity alone is insufficient in the longitudinal setting. Scans from the same patient naturally exhibit high similarity because they share the same biological structure. The subtle deformations caused by disease are often small and easily overshadowed by the stable, unaffected anatomy. As a result, conventional metrics such as PSNR and SSIM tend to report inflated scores, failing to reflect clinically meaningful progression.

The true signal of interest lies in the *residual differences* between baseline and follow-up scans. These residuals often encode the disease trajectory. To capture this, we introduce a residual-based metric. Specifically, let

$$\Delta_{\text{gt}} = \mathbf{x}_T - \mathbf{x}_0, \quad \Delta_{\text{gen}} = \hat{\mathbf{x}}_T - \mathbf{x}_0, \quad (16)$$

where Δ denotes the residual image, \mathbf{x} and $\hat{\mathbf{x}}$ denote ground-truth and generated MRIs, respectively. The metric is defined as the Residual-based Relative Mean Absolute Error (Δ -RMAE):

$$\Delta\text{-RMAE} = \frac{|\Delta_{\text{gt}} - \Delta_{\text{gen}}|}{\frac{1}{2}(|\Delta_{\text{gt}}| + |\Delta_{\text{gen}}|)}. \quad (17)$$

By definition, $\Delta\text{-RMAE} \in [0, 2]$. A smaller value indicates that the predicted residual aligns with the ground-truth residual, *i.e.*, the model has correctly captured disease-relevant evolution. If the model predicts almost no change (essentially copying the baseline image), then $\Delta_{\text{gen}} \approx 0$. In this case, both numerator and denominator are dominated by $|\Delta_{\text{gt}}|$, and the score approaches 2. Likewise, if the prediction deviates strongly, $|\Delta_{\text{gen}}|$ dominates, again driving the score toward 2.

4.2 BENCHMARK COMPARISONS

We present qualitative results in Figure 2. In this subject, disease progression primarily affects the lateral ventricles (the central butterfly-shaped region) and the cortical gray matter (outer surface). SADM, TADM, and MambaControl fail to capture changes in the lateral ventricles, resulting in pronounced butterfly-shaped error in the residual image. DiffuseMorph and BrLP better capture ventricular changes, but leave substantial residuals along the cortical surface. Our method also

Table 2: Quantitative comparison results on clinical structure faithfulness metrics are reported in region-based MAE and Δ -RMAE (Residual-based Relative Mean Absolute Error). Results are reported as mean \pm standard deviation.

Method	Type	ADNI		AIBL		OASIS	
		Region MAE \downarrow	Δ -RMAE \downarrow	Region MAE \downarrow	Δ -RMAE \downarrow	Region MAE \downarrow	Δ -RMAE \downarrow
CardiacAging (Campello et al., 2022)	Direct	0.289 \pm 0.33	0.771 \pm 0.12	0.281 \pm 0.35	0.698 \pm 0.13	0.345 \pm 0.35	0.739 \pm 0.13
CounterSynth (Pombo et al., 2023)	Direct	0.272 \pm 0.32	0.704 \pm 0.11	0.256 \pm 0.30	0.672 \pm 0.12	0.298 \pm 0.34	0.690 \pm 0.12
DiffuseMorph (Kim et al., 2022)	Deform	0.230 \pm 0.28	0.516 \pm 0.10	0.226 \pm 0.27	0.482 \pm 0.10	0.279 \pm 0.29	0.503 \pm 0.11
SADM (Yoon et al., 2023)	Noise	0.283 \pm 0.34	0.746 \pm 0.11	0.278 \pm 0.30	0.693 \pm 0.12	0.365 \pm 0.39	0.728 \pm 0.12
TADM (Litrico et al., 2024)	Δ Noise	0.263 \pm 0.32	0.697 \pm 0.11	0.243 \pm 0.30	0.654 \pm 0.11	0.291 \pm 0.34	0.685 \pm 0.11
ImageFlowNet (Liu et al., 2025a)	Flow	0.259 \pm 0.30	0.589 \pm 0.11	0.240 \pm 0.29	0.561 \pm 0.12	0.283 \pm 0.32	0.574 \pm 0.12
BrLP (Puglisi et al., 2024)	Noise	0.257 \pm 0.32	0.630 \pm 0.10	0.278 \pm 0.31	0.594 \pm 0.11	0.335 \pm 0.36	0.622 \pm 0.11
MambaControl (Yang et al., 2025a)	Noise	0.225 \pm 0.30	0.554 \pm 0.09	0.249 \pm 0.30	0.525 \pm 0.10	0.299 \pm 0.32	0.561 \pm 0.10
Δ -LFM (Ours)	Δ Flow	0.210 \pm 0.28	0.436 \pm 0.08	0.226 \pm 0.29	0.417 \pm 0.10	0.262 \pm 0.30	0.473 \pm 0.08

Table 3: Ablation study results (mean across datasets). We evaluate the effect of different auto-encoder training strategies (Arc Loss, Rank Loss, ArcRank Loss) and flow-matching manifold sampling strategies ($[0, 1]$ vs. $[0, T]$). The Cond. represent using condition or not. The row with gray background indicated the proposed method.

Model	Auto-Encoder	FM Sampling	Cond.	Image Quality		Structure Fidelity	
				PSNR \uparrow	SSIM \uparrow	Region MAE \downarrow	Δ -RMAE \downarrow
LFM (Baseline)	-	$[0, 1]$	\times	27.59 \pm 1.78	88.62 \pm 1.63	0.357 \pm 0.342	0.552 \pm 0.146
Conditional LFM	-	$[0, 1]$	\checkmark	28.46 \pm 1.22	91.20 \pm 1.15	0.299 \pm 0.312	0.486 \pm 0.101
Δ -LFM	-	$[0, T]$	\checkmark	28.78 \pm 1.15	90.97 \pm 1.07	0.286 \pm 0.295	0.472 \pm 0.091
	w/ Arc Loss	$[0, 1]$	\checkmark	29.52 \pm 0.98	92.38 \pm 0.95	0.251 \pm 0.282	0.457 \pm 0.086
	w/ Rank Loss	$[0, 1]$	\checkmark	28.36 \pm 1.21	91.15 \pm 1.11	0.272 \pm 0.285	0.474 \pm 0.088
	w/ ArcRank Loss	$[0, 1]$	\checkmark	29.83 \pm 0.95	92.74 \pm 0.93	0.243 \pm 0.277	0.454 \pm 0.088
	w/ ArcRank Loss	$[0, T]$	\checkmark	30.04 \pm 1.04	92.63 \pm 1.03	0.233 \pm 0.282	0.442 \pm 0.087

shows error in the ventricles; however, these are less intense, with thinner boundaries and smaller cortical-surface errors, demonstrating improved overall accuracy.

We report PSNR and SSIM in Table 1, and Region MAE together with Δ -RMAE in Table 2. Across all three datasets, our proposed Δ -LFM consistently achieves high performance.

In Table 1, Δ -LFM improves PSNR by a clear margin over the strongest baseline. On ADNI, it reaches 30.59 dB, exceeding the second-best method (MambaControl, 29.72 dB) by +0.87 dB. Similar gains are observed on AIBL (+0.66 dB) and OASIS (+0.77 dB). The SSIM results mirror this trend: Δ -LFM achieves 94.62, 93.92, and 89.36, improving upon the best baselines by roughly +1.0, +0.7, and +1.1 points, respectively. These consistent improvements indicate that Δ -LFM produces reconstructions with both higher fidelity and sharper structural alignment.

Table 2 further validates these observations with volumetric error metrics. Δ -LFM yields the lowest Region MAE on ADNI (0.210) and OASIS (0.262), and matches the best-performing method (DiffuseMorph) on AIBL (0.226). More importantly, it achieves the lowest Δ -RMAE across all datasets: 0.436 on ADNI, 0.417 on AIBL, and 0.473 on OASIS. Compared to the strongest baseline (MambaControl), this represents relative error reductions of about 21%, 21%, and 16%, respectively.

4.3 ABLATION STUDY

We provide visualization results of the learned latent space in Figure 3 on test cases in ADNI dataset, where t-SNE is used to project the representations into two dimensions. As shown in the left panel, the latent trajectories of individual patients are well aligned. Interestingly, the right panel reveals that the latent space also organizes according to diagnostic status, despite the fact that this information was never used during training.

To assess the contribution of each component in our framework, we conducted an ablation study across three datasets, as summarized in Table 3. The experiments examine the effect of (i) introducing conditional information, (ii) different auto-encoder training strategies (Arc Loss, Rank Loss,

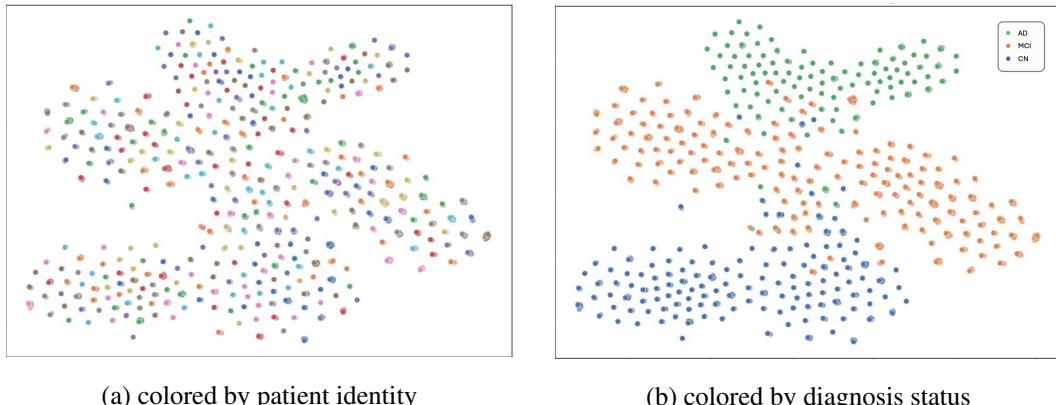


Figure 3: t-SNE projection of the learned longitudinal latent space into two dimensions. (a) Coloring by subject ID shows that scans from the same patient cluster together. (b) Coloring by disease status shows that scans with the same diagnosis are grouped closely. CN: Cognitively Normal. MCI: Mild Cognitive Impairment. AD: Alzheimer’s Disease (Dementia Stage).

Table 4: Ablation on the proposed components.

Angular loss	Ranking loss	Sampling	PSNR	SSIM
Arc Loss	–	[0, T]	29.88	92.31
–	Rank Loss	[0, T]	28.85	91.10
Cosine Similarity	–	[0, 1]	29.13	91.10
Cosine Similarity	Simple Rank Loss	[0, 1]	28.32	89.92
Arc Loss	Simple Rank Loss	[0, 1]	28.44	89.87

ArcRank Loss), and (iii) different flow-matching (FM) manifold sampling strategies, namely physical sampling $[0, 1]$ versus temporal progression modeling $[0, T]$.

The baseline Latent Flow Matching (LFM) (Dao et al., 2023) model without conditioning performs the worst across all metrics, showing notably low SSIM/PSNR and high error values. Incorporating conditional information (Conditional LFM) significantly improves image quality (SSIM from 88.62 to 91.20; PSNR from 27.59 to 28.46) and reduces both Region MAE and Δ -RMAE.

Within the Δ -LFM ablation, we observe two consistent trends. First, the choice of loss function has a direct impact: using Arc Loss alone yields clear improvements, whereas Rank Loss alone slightly undermines performance. This is because Arc Loss encourages patient trajectories to follow a meaningful direction, while Rank Loss without Arc imposes only a weak magnitude constraint. In contrast, their combination (ArcRank Loss) achieves the best balance, suggesting that structural alignment (Arc) and temporal ordering (Rank) are complementary. Second, extending the FM sampling strategy from $[0, 1]$ to $[0, T]$ further improves results by promoting gradual residual modeling over meaningful temporal intervals.

Finally, we visualize qualitative progression predictions in Figure 4 using cases from three datasets, demonstrating the ability of Δ -LFM to predict future timepoints at arbitrary intervals. The results show disease-related deformations primarily affecting the ventricles and surrounding cortical areas, which is consistent with clinical observations.

We present qualitative visualizations of predicted disease progression in Figure 4. Each row illustrates disease trajectories over nine years generated by Δ -LFM. The intermediate steps provide clinically interpretable insights into the progression process, capturing characteristic neurodegenerative patterns such as ventricular enlargement and cortical thinning in the temporal and parietal regions. These deformations align with established clinical observations, underscoring the biological plausibility of our model.

To further verify each proposed component, we provide additional ablations in Table 4. We evaluate cosine similarity $\mathcal{L}_{\text{cosine}}$ as a surrogate for Arc loss and a simple ranking loss $\mathcal{L}_{\text{simple}}$ as a surrogate

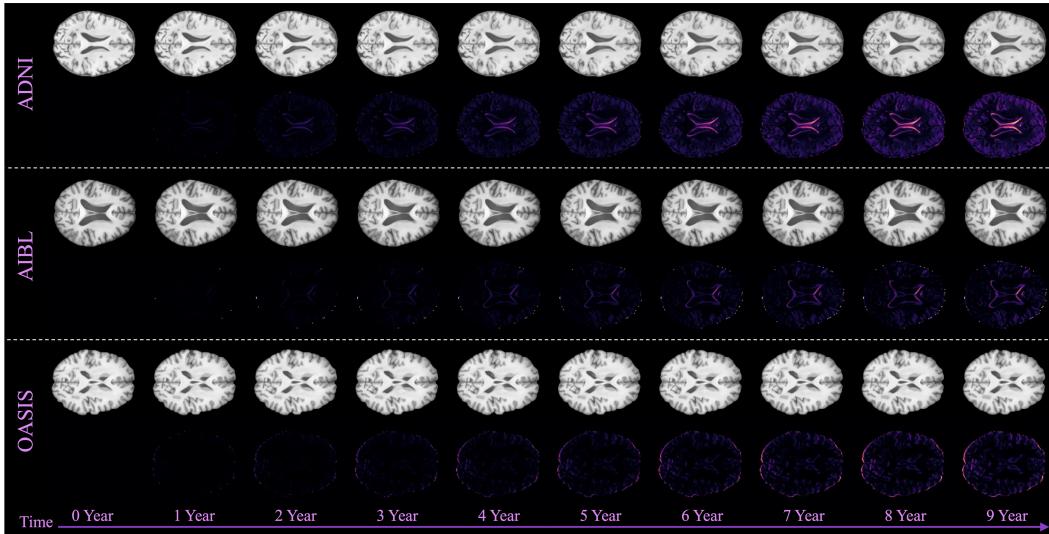


Figure 4: Continuous disease trajectory predicted by Δ -LFM at one-year intervals. Odd rows show the baseline MRI (year 0) and predicted MRIs from years 1–9; even rows show residual maps relative to baseline. Progression highlights structural abnormalities, primarily in atrophy-related regions (e.g., hippocampus, ventricles, surrounding cortex).

for Rank loss, defined as follows:

$$\mathcal{L}_{\text{cosine}}(i, j) = 1 - \frac{\langle \mathbf{z}_i, \mathbf{z}_j \rangle}{\|\mathbf{z}_i\| \|\mathbf{z}_j\|}, \quad \mathcal{L}_{\text{simple}}(i, j) = \max(0, d_j - d_i). \quad (18)$$

Cosine similarity performs slightly worse than Arc loss, which is expected since it only enforces a global angular alignment and does not capture the fine-grained latent directional structure modeled by Arc loss. In contrast, the simple rank loss without a margin often leads to latent mode collapse, explaining its poor performance when combined with either Arc loss or cosine similarity.

Loss Weight. λ_{arc} is relatively sensitive. Values $\lambda_{\text{arc}} > 0.01$ consistently cause collapse, while 0.0025–0.0075 are stable across runs, with $\lambda_{\text{arc}} = 0.0075$ noticeably degrading AE reconstruction. We therefore set $\lambda_{\text{arc}} = 0.005$ as a compromise, trading a small decrease in AE reconstruction for improved flow-matching performance (+0.4 PSNR over 0.0025).

For λ_{rank} , values above 0.05 hinder convergence, whereas the range 0.005–0.02 yields similar behavior. We choose $\lambda_{\text{rank}} = 0.01$, which gives a slight PSNR gain (+0.05).

Speed Analysis ArcRank involves an SVD operation, which is computational expensive. In practice, we use `torch.linalg.svd(z, full_matrices=False)`, which gives a $6\times$ speedup compared to the full matrices variant (0.0552s \rightarrow 0.0092s on average over 50 runs). Overall, ArcRank increases AE training time by $\sim 40\%$ (1.4 \times), which we regard as a reasonable overhead given the resulting performance gains.

5 CONCLUSION

We introduce **Progression Latent Flow Matching (Δ -LFM)**, a framework that unifies the proposed ArcRank loss with temporal flow matching to capture patient-specific disease trajectories. To evaluate progression in isolation, we propose a new metric, Δ -RMAE. On longitudinal MRI benchmarks, Δ -LFM consistently achieves higher fidelity, lower error, and tighter alignment with ground-truth progression than prior methods. Beyond accuracy, it yields temporally coherent and clinically interpretable trajectories. These results show that Δ -LFM provides a new direction for progression modeling, with the potential to generalize across other temporal generation tasks.

REFERENCES

Michael S Albergo, Nicholas M Boffi, and Eric Vanden-Eijnden. Stochastic interpolants: A unifying framework for flows and diffusions. *arXiv preprint arXiv:2303.08797*, 2023. 4

Moinak Bhattacharya, Saumya Gupta, Annie Singh, Chao Chen, Gagandeep Singh, and Prateek Prasanna. Brainmrddiff: A diffusion model for anatomically consistent brain mri synthesis. *arXiv preprint arXiv:2504.04532*, 2025. 2

Víctor M Campello, Tian Xia, Xiao Liu, Pedro Sanchez, Carlos Martín-Isla, Steffen E Petersen, Santi Seguí, Sotirios A Tsaftaris, and Karim Lekadir. Cardiac aging synthesis from cross-sectional data with conditional generative adversarial networks. *Frontiers in Cardiovascular Medicine*, 9:983091, 2022. 7, 8

Qi Chen, Yuxiang Lai, Xiaoxi Chen, Qixin Hu, Alan Yuille, and Zongwei Zhou. Analyzing tumors by synthesis. *Generative Machine Learning Models in Medical Image Computing*, pp. 85–110, 2024. 2

Qi Chen, Xinze Zhou, Chen Liu, Hao Chen, Wenxuan Li, Zekun Jiang, Ziyan Huang, Yuxuan Zhao, Dexin Yu, Junjun He, et al. Scaling tumor segmentation: Best lessons from real and synthetic data. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 24001–24013, 2025. 2

Quan Dao, Hao Phung, Binh Nguyen, and Anh Tran. Flow matching in latent space. *arXiv preprint arXiv:2307.08698*, 2023. 9

Kathryn A Ellis, Ashley I Bush, David Darby, Daniela De Fazio, Jonathan Foster, Peter Hudson, Nicola T Lautenschlager, Nat Lenzo, Ralph N Martins, Paul Maruff, 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. 6

Giovanni B Frisoni, Nick C Fox, Clifford R Jack Jr, Philip Scheltens, and Paul M Thompson. The clinical use of structural mri in alzheimer disease. *Nature reviews neurology*, 6(2):67–77, 2010. 1

Pengfei Guo, Can Zhao, Dong Yang, Ziyue Xu, Vishwesh Nath, Yucheng Tang, Benjamin Simon, Mason Belue, Stephanie Harmon, Baris Turkbey, et al. Maisi: Medical ai for synthetic imaging. In *2025 IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pp. 4430–4441. IEEE, 2025. 14

Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 2

Andrew Hoopes, Jocelyn S Mora, Adrian V Dalca, Bruce Fischl, and Malte Hoffmann. Synthstrip: Skull-stripping for any brain image. *NeuroImage*, 260:119474, 2022. 6

CR Jack Jr, M Slomkowski, S Gracon, TM Hoover, JP Felmlee, K Stewart, Y Xu, M Shiung, PC O'brien, R Cha, et al. Mri as a biomarker of disease progression in a therapeutic trial of milameline for ad. *Neurology*, 60(2):253–260, 2003. 1

Boah Kim, Inhwa Han, and Jong Chul Ye. Diffusemorph: Unsupervised deformable image registration using diffusion model. In *European conference on computer vision*, pp. 347–364. Springer, 2022. 2, 7, 8

Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 4

Virginia Klema and Alan Laub. The singular value decomposition: Its computation and some applications. *IEEE Transactions on automatic control*, 25(2):164–176, 1980. 4

Daeun Kyung, Junu Kim, Tackeun Kim, and Edward Choi. Towards predicting temporal changes in a patient's chest x-ray images based on electronic health records. *arXiv preprint arXiv:2409.07012*, 2024. 2

Yuxiang Lai, Jike Zhong, Vanessa Su, and Xiaofeng Yang. Patient-specific autoregressive models for organ motion prediction in radiotherapy. *arXiv preprint arXiv:2505.11832*, 2025. 1

Pamela J LaMontagne, Tammie LS Benzinger, John C Morris, Sarah Keefe, Russ Hornbeck, Chengjie Xiong, Elizabeth Grant, Jason Hassenstab, Krista Moulder, Andrei G Vlassenko, et al. Oasis-3: longitudinal neuroimaging, clinical, and cognitive dataset for normal aging and alzheimer disease. *MedRxiv*, pp. 2019–12, 2019. 6

Wenxuan Li, Xinze Zhou, Qi Chen, Tianyu Lin, Pedro RAS Bassi, Szymon Plotka, Jaroslaw B Cwikla, Xiaoxi Chen, Chen Ye, Zheren Zhu, et al. Pants: The pancreatic tumor segmentation dataset. *arXiv preprint arXiv:2507.01291*, 2025. 1

Kaizhao Liang, Xu Cao, Kuei-Da Liao, Tianren Gao, Wenqian Ye, Zhengyu Chen, Jianguo Cao, Tejas Nama, and Jimeng Sun. Pie: Simulating disease progression via progressive image editing. 2023. 2

Yaron Lipman, Ricky T. Q. Chen, Heli Ben-Hamu, Maximilian Nickel, and Matt Le. Flow matching for generative modeling, 2023. URL <https://arxiv.org/abs/2210.02747>. 3, 4

Mattia Litrico, Francesco Guarnera, Mario Valerio Giuffrida, Daniele Ravì, and Sebastiano Battiato. Tadm: Temporally-aware diffusion model for neurodegenerative progression on brain mri. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 444–453. Springer, 2024. 2, 7, 8, 14

Chen Liu, Ke Xu, Liangbo L Shen, Guillaume Huguet, Zilong Wang, Alexander Tong, Danilo Bzdok, Jay Stewart, Jay C Wang, Lucian V Del Priore, et al. Imageflownet: Forecasting multiscale image-level trajectories of disease progression with irregularly-sampled longitudinal medical images. In *ICASSP 2025-2025 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2025a. 2, 7, 8

Q. Liu, E. Fuster-Garcia, I. Thokle Hovden, B. J. MacIntosh, E. O. S. Grødem, P. Brandal, C. Lopez-Mateu, D. Sederevicius, K. Skogen, T. Schellhorn, A. Bjørnerud, and K. Eeg Emblem. Treatment-aware diffusion probabilistic model for longitudinal MRI generation and diffuse glioma growth prediction. *IEEE Transactions on Medical Imaging*, 44(6):2449–2462, June 2025b. doi: 10.1109/TMI.2025.3533038. 3

Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017. 14

Cheng Lu, Yuhao Zhou, Fan Bao, Jianfei Chen, Chongxuan Li, and Jun Zhu. Dpm-solver: A fast ode solver for diffusion probabilistic model sampling in around 10 steps. *Advances in neural information processing systems*, 35:5775–5787, 2022. 4

OpenAI. Chatgpt [large language model]. <https://chat.openai.com/chat>, 2023. 17

Sid E O'Bryant, Stephen C Waring, C Munro Cullum, James Hall, Laura Lacritz, Paul J Massman, Philip J Lupo, Joan S Reisch, Rachelle Doody, Texas Alzheimer's Research Consortium, et al. Staging dementia using clinical dementia rating scale sum of boxes scores: a texas alzheimer's research consortium study. *Archives of neurology*, 65(8):1091–1095, 2008. 1

Steven R Parratt, Elina Numminen, and Anna-Liisa Laine. Infectious disease dynamics in heterogeneous landscapes. *Annual Review of Ecology, Evolution, and Systematics*, 47(1):283–306, 2016. 1

Ronald Carl Petersen, Paul S Aisen, Laurel A Beckett, Michael C Donohue, Anthony Collins Gamst, Danielle J Harvey, Clifford R Jack, William J Jagust, Leslie M Shaw, Arthur W Toga, et al. Alzheimer's disease neuroimaging initiative (adni): clinical characterization. *Neurology*, 74(3):201–209, 2010. 6

Guilherme Pombo, Robert Gray, M Jorge Cardoso, Sébastien Ourselin, Geraint Rees, John Ashburner, and Parashkev Nachev. Equitable modelling of brain imaging by counterfactual augmentation with morphologically constrained 3d deep generative models. *Medical Image Analysis*, 84:102723, 2023. 7, 8

Lemuel Puglisi, Daniel C Alexander, and Daniele Ravì. Enhancing spatiotemporal disease progression models via latent diffusion and prior knowledge. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pp. 173–183. Springer, 2024. 2, 7, 8, 14

Lemuel Puglisi, Daniel C Alexander, and Daniele Ravì. Brain latent progression: Individual-based spatiotemporal disease progression on 3d brain mrис via latent diffusion. *arXiv preprint arXiv:2502.08560*, 2025. 2

Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *Journal of machine learning research*, 21(140):1–67, 2020. 14

Siddharth Ramchandran, Gleb Tikhonov, Kalle Kujanpää, Miika Koskinen, and Harri Lähdesmäki. Longitudinal variational autoencoder, 2021. URL <https://arxiv.org/abs/2006.09763>. 3

Daniele Ravi, Daniel C. Alexander, and Neil P. Oxtoby. Degenerative Adversarial NeuroImage Nets: Generating Images that Mimic Disease Progression. In Dinggang Shen, Tianming Liu, Terry M. Peters, Lawrence H. Staib, Caroline Essert, Sean Zhou, Pew-Thian Yap, and Ali Khan (eds.), *Medical Image Computing and Computer Assisted Intervention – MICCAI 2019*, Lecture Notes in Computer Science, pp. 164–172, Cham, 2019. Springer International Publishing. ISBN 978-3-030-32248-9. doi: 10.1007/978-3-030-32248-9_19. 2

Daniele Ravi, Stefano B Blumberg, Silvia Ingala, Frederik Barkhof, Daniel C Alexander, Neil P Oxtoby, Alzheimer's Disease Neuroimaging Initiative, et al. Degenerative adversarial neuroimage nets for brain scan simulations: Application in ageing and dementia. *Medical Image Analysis*, 75:102257, 2022. 2

Leslie A Real and Roman Biek. Spatial dynamics and genetics of infectious diseases on heterogeneous landscapes. *Journal of the Royal Society Interface*, 4(16):935–948, 2007. 1

Mengwei Ren, Neel Dey, Martin A. Styner, Kelly Botteron, and Guido Gerig. Local spatiotemporal representation learning for longitudinally-consistent neuroimage analysis, 2023. URL <https://arxiv.org/abs/2206.04281>. 3

Maximilian Rokuss, Yannick Kirchhoff, Seval Akbal, Balint Kovacs, Saikat Roy, Constantin Ulrich, Tassilo Wald, Lukas T Rotkopf, Heinz-Peter Schlemmer, and Klaus Maier-Hein. Lesionlocator: Zero-shot universal tumor segmentation and tracking in 3d whole-body imaging. In *Proceedings of the Computer Vision and Pattern Recognition Conference*, pp. 30872–30885, 2025. 1

Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pp. 234–241. Springer, 2015. 14

Heung-Il Suk and Dinggang Shen. Deep learning-based feature representation for ad/mci classification. In *International conference on medical image computing and computer-assisted intervention*, pp. 583–590. Springer, 2013. 1

Sarah J Tabrizi, Rachael I Scahill, Alexandra Durr, Raymund AC Roos, Blair R Leavitt, Rebecca Jones, G Bernhard Landwehrmeyer, Nick C Fox, Hans Johnson, Stephen L Hicks, et al. Biological and clinical changes in premanifest and early stage huntington’s disease in the track-hd study: the 12-month longitudinal analysis. *The Lancet Neurology*, 10(1):31–42, 2011. 1

Gemini Team, Rohan Anil, Sebastian Borgeaud, Jean-Baptiste Alayrac, Jiahui Yu, Radu Soricu, Johan Schalkwyk, Andrew M Dai, Anja Hauth, Katie Millican, et al. Gemini: a family of highly capable multimodal models. *arXiv preprint arXiv:2312.11805*, 2023. 17

Alexander Tong, Kilian Fatras, Nikolay Malkin, Guillaume Huguet, Yanlei Zhang, Jarrid Rector-Brooks, Guy Wolf, and Yoshua Bengio. Improving and generalizing flow-based generative models with minibatch optimal transport, 2024. URL <https://arxiv.org/abs/2302.00482>. 3

Nicholas J Tustison, Brian B Avants, Philip A Cook, Yuanjie Zheng, Alexander Egan, Paul A Yushkevich, and James C Gee. N4itk: improved n3 bias correction. *IEEE transactions on medical imaging*, 29(6):1310–1320, 2010. 6

Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. *Advances in neural information processing systems*, 30, 2017. 14

Ge Wang, Jong Chul Ye, and Bruno De Man. Deep learning for tomographic image reconstruction. *Nature machine intelligence*, 2(12):737–748, 2020. 1

Linshan Wu, Jiaxin Zhuang, Xuefeng Ni, and Hao Chen. Freetumor: Advance tumor segmentation via large-scale tumor synthesis. *arXiv preprint arXiv:2406.01264*, 2024. 2

Hao Yang, Tao Tan, Shuai Tan, Weiqin Yang, Kunyan Cai, Calvin Chen, and Yue Sun. Mambacontrol: Anatomy graph-enhanced mamba controlnet with fourier refinement for diffusion-based disease trajectory prediction. *arXiv preprint arXiv:2505.09965*, 2025a. 7, 8

Yijun Yang, Zhao-Yang Wang, Qiuping Liu, Shuwen Sun, Kang Wang, Rama Chellappa, Zongwei Zhou, Alan Yuille, Lei Zhu, Yu-Dong Zhang, et al. Medical world model: Generative simulation of tumor evolution for treatment planning. *arXiv preprint arXiv:2506.02327*, 2025b. 1

Jee Seok Yoon, Chenghao Zhang, Heung-Il Suk, Jia Guo, and Xiaoxiao Li. Sadm: Sequence-aware diffusion model for longitudinal medical image generation. In *International Conference on Information Processing in Medical Imaging*, pp. 388–400. Springer, 2023. 2, 7, 8

Peter NE Young, Mar Estrella, Emma Coomans, Meera Srikrishna, Helen Beaumont, Anne Maass, Ashwin V Venkataraman, Rikki Lissaman, Daniel Jiménez, Matthew J Betts, et al. Imaging biomarkers in neurodegeneration: current and future practices. *Alzheimer’s research & therapy*, 12(1):49, 2020. 1

Lvmin Zhang, Anyi Rao, and Maneesh Agrawala. Adding conditional control to text-to-image diffusion models. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pp. 3836–3847, 2023. 7

Qiang Zhao, Zhiqiang Liu, Ehsan Adeli, and Kilian M. Pohl. Longitudinal self-supervised learning. *Medical Image Analysis*, 71:102051, 2021. doi: 10.1016/j.media.2021.102051. 3

Appendix

A TRAINING SETTING

We adopt the AdamW optimizer (Loshchilov & Hutter, 2017) for all experiments. For the autoencoder (AE) and 3D U-Net (Ronneberger et al., 2015), we trained with a learning rate of 1×10^{-3} and a batch size of 2 for 300 epochs. For the flow matching model based on the 3D U-Net, we used a reduced learning rate of 3×10^{-5} with a larger batch size of 4 for 200 epochs.

B CONDITIONAL GENERATION

Following Litrico *et al.* (Litrico et al., 2024), we encode the current sample time in flow matching into continuous embeddings using sinusoidal positional encodings. These time embeddings act as an additive bias to the input, ensuring that the generative process is aware of its temporal position.

Building on this, and inspired by BrLP (Puglisi et al., 2024), we introduce *explicit control signals* that combine patient-specific attributes with temporal queries. In particular, we condition the model not only on the start time, the current sample time, and the end query time, but also on the patient attributes. This design enables the model to incorporate both subject-specific and task-specific information when shaping the generative trajectory.

To inject these conditional signals into the U-Net backbone, we explored three strategies: (1) Additive biasing, (2) Cross-attention (Vaswani et al., 2017) and (3) Adaptive Layer Normalization (AdaLN) (Raffel et al., 2020).

Among these approaches, our experiments show that AdaLN is most effective with our temporal sampling over the interval $[0, T]$. It adapts the network representation in a temporally aware and context-sensitive manner, allowing the generative process to be steered smoothly by patient attributes and target query times.

The implementation of AdaLN is realized through an initial one-layer MLP that processes the condition signals, followed by additional MLPs at each U-Net decoder layer to generate the corresponding AdaLN modulation parameters.

C AUTO-ENCODER ABLATION

Table 5 summarizes the effect of AE capacity and training crop size. For the capacity ablation, we vary the encoder channels while fixing the crop size to 64^3 . For the crop-size ablation, we fix the AE to $[64, 128, 256]$ and compare 48^3 vs. 64^3 crops.

Table 5: Ablation on AE capacity and training crop size.

Channels	AE capacity (crop 64^3)		Crop size (AE [64, 128, 256])			
	Crop	PSNR	SSIM	Crop	PSNR	SSIM
$[64, 128, 256]$	64^3	30.04	92.63	64^3	30.04	92.63
$[64, 128, 128]$	48^3	29.89	92.55	48^3	29.51	91.71
$[64, 128]$	48^3	29.62	92.10			

We do not use deeper or wider AEs, as these configurations lead to out-of-memory (OOM) errors on our RTX A6000 GPU (48 GB). Based on the above results, we use channels $[64, 128, 256]$ and crop size 64^3 as our default configuration. Larger crops (e.g., 72^3) also result in OOM on our hardware. To keep the comparison with baselines fair, we do not use a pretrained AE, although initializing with MAISI (Guo et al., 2025) can further improve PSNR (from 30.04 to 31.97 for $[64, 128, 256]$).

Table 6: Quantitative results across different prediction horizons (years into the future). Measurements are reported using PSNR (dB), SSIM ($\times 1.0\text{E}2$), Region MAE, and Δ -RMAE.

Year	PSNR \uparrow	SSIM \uparrow	Region MAE \downarrow	Δ -RMAE \downarrow
1	31.61 ± 0.87	93.08 ± 0.75	0.209 ± 0.107	0.386 ± 0.107
2	31.79 ± 0.92	93.79 ± 0.77	0.217 ± 0.106	0.369 ± 0.097
3	31.09 ± 0.95	93.51 ± 0.82	0.206 ± 0.118	0.388 ± 0.089
4	32.48 ± 0.93	93.95 ± 0.89	0.201 ± 0.132	0.395 ± 0.084
5	32.14 ± 0.98	93.59 ± 0.88	0.215 ± 0.148	0.405 ± 0.080
6	31.37 ± 1.02	93.07 ± 0.95	0.224 ± 0.163	0.414 ± 0.081
7	30.12 ± 0.98	92.58 ± 1.07	0.241 ± 0.176	0.436 ± 0.082
8	29.83 ± 1.05	92.18 ± 1.15	0.261 ± 0.191	0.459 ± 0.089
9	28.09 ± 1.07	90.65 ± 1.20	0.279 ± 0.207	0.487 ± 0.085
10	28.61 ± 1.23	91.41 ± 1.37	0.302 ± 0.246	0.552 ± 0.104
11	28.71 ± 1.13	91.32 ± 1.22	0.294 ± 0.218	0.506 ± 0.094
13	26.98 ± 1.29	89.96 ± 1.44	0.339 ± 0.259	0.579 ± 0.109

D LONGITUDINAL PROGRESSION QUANTITATIVE RESULTS

Table 6 summarizes predictive performance across different future horizons. For short- to mid-term forecasts (1–5 years), reconstruction metrics remain consistently high: PSNR hovers around 31–32 dB and SSIM above 93%, indicating strong fidelity in reproducing structural details. Region MAE is also relatively stable in this range (~ 0.20), suggesting reliable local accuracy, while Δ -RMAE values remain below 0.41, reflecting well-preserved progression dynamics.

Beyond 6 years, however, all metrics exhibit gradual degradation. PSNR and SSIM decline steadily, dropping to 28 dB and below 91% by year 10, and further to 26.98 dB / 89.96% at year 13. Concurrently, Region MAE and Δ -RMAE increase monotonically, indicating that both absolute errors in regional values and their temporal progression deviate more strongly with longer horizons. Notably, the sharp rise in Δ -RMAE after year 10 (> 0.50) highlights a growing difficulty in capturing long-term progression trends.

We also observe that the prediction trend is not strictly linear as the prediction horizon increases. This is mainly because different time horizons are estimated from different subsets of subjects. In particular, fewer subjects have follow-up data at 10–13 years compared to 1–5 years, which introduces higher variability and heterogeneity, leading to slightly inconsistent trends.

E ON THE STABILITY OF Δ -RMAE

Residual-based metrics such as Δ -RMAE can in principle be affected by imperfect registration and intensity inhomogeneity, and our preprocessing cannot guarantee perfectly aligned, bias-free images. Here we analyze how such imperfections influence the metric.

In our setting, Δ -RMAE is defined voxel-wise as

$$\Delta\text{-RMAE} = \frac{|\Delta_{\text{gt}} - \Delta_{\text{gen}}|}{\frac{1}{2}(|\Delta_{\text{gt}}| + |\Delta_{\text{gen}}|)}, \quad (19)$$

so that $\Delta\text{-RMAE} \in [0, 2]$, with values close to 2 indicating large disagreement between the two residuals. We model misregistration and residual bias-field effects as an additive perturbation Δ_{mis} on the ground-truth residual, yielding

$$\Delta\text{-RMAE}_{\text{noise}} = \frac{|\Delta_{\text{mis}} + \Delta_{\text{gt}} - \Delta_{\text{gen}}|}{\frac{1}{2}(|\Delta_{\text{mis}} + \Delta_{\text{gt}}| + |\Delta_{\text{gen}}|)}. \quad (20)$$

When Δ_{mis} is small relative to Δ_{gt} and Δ_{gen} , numerator and denominator are perturbed in a comparable way, so $\Delta\text{-RMAE}_{\text{noise}}$ remains close to the ideal Δ -RMAE. As Δ_{mis} increases, the main effect is a mild compression of the dynamic range (the effective upper bound becomes slightly below 2), rather than a systematic bias toward any particular model.

To quantify this effect, we perform a simple sensitivity analysis by adding Gaussian noise with standard deviation σ to the residuals and recomputing Δ -RMAE. Table 7 reports the mean, standard

Table 7: Sensitivity of Δ -RMAE to additive Gaussian noise on the residuals.

σ (noise)	Mean Δ -RMAE	Standard Deviation	Bias vs. True
0.00	2.000000	0.000000	0.000000
0.05	1.999842	0.002315	-0.000158
0.10	1.999375	0.005982	-0.000625
0.15	1.998214	0.011430	-0.001786
0.20	1.996102	0.019752	-0.003898
0.25	1.992994	0.031245	-0.007006
0.30	1.989124	0.046521	-0.010876
0.35	1.983612	0.064210	-0.016388
0.40	1.976348	0.085402	-0.023652
0.45	1.968740	0.107895	-0.031260
0.50	1.960255	0.133005	-0.039745
0.55	1.950492	0.160884	-0.049508
0.60	1.939101	0.191156	-0.060899
0.65	1.926345	0.222974	-0.073655
0.70	1.911890	0.255780	-0.088110
0.75	1.896232	0.288901	-0.103768
0.80	1.879110	0.322410	-0.120890
0.85	1.861775	0.355122	-0.138225
0.90	1.843569	0.386230	-0.156431
0.95	1.834021	0.415104	-0.165979
1.00	1.826570	0.462351	-0.173430

deviation, and bias relative to the noiseless value (2.0). Even for relatively large noise levels (σ up to 1.0), the induced bias remains small compared to the full range [0, 2], supporting the robustness of Δ -RMAE to moderate misregistration.

F CLINICAL CONDITIONING: SETUP AND VERIFICATION

We conduct a series of analyses to assess the robustness of the model to perturbations in the conditioning metadata.

Age. Age serves as the temporal anchor for longitudinal prediction, as future timepoints are defined by (*baseline age* + *time span*). Injecting noise into age would therefore shift the implied prediction timepoint, so that the generated image no longer corresponds to a meaningful clinical follow-up. Such an ablation would be difficult to interpret, and for this reason we do not treat age as a noise-robust attribute.

Sex. To evaluate robustness with respect to sex, we intentionally flipped the sex attribute and re-evaluated the model. On ADNI, performance changed only marginally, from PSNR/SSIM of 30.59/94.62 to 30.12/94.47, indicating limited sensitivity to this source of label noise.

Clinical status. In clinical practice, diagnostic labels are inherently noisy. To reflect this, we inject Gaussian noise into the clinical status during training to model diagnostic uncertainty. Specifically, we use zero-mean Gaussian noise with standard deviation equal to one-third of the status class interval (interval = 3).

Under perturbed status labels, performance remains reasonable, with PSNR/SSIM degrading only slightly from 30.59/94.62 to 29.81/93.58, demonstrating robustness to clinically realistic label noise.

G LIMITATIONS

Our work, while demonstrating the potential of latent flow matching for longitudinal imaging generation, has several limitations that we wish to acknowledge.

Task-specific limitation. In the current study, we restrict our focus to AD, which provides one of the most widely available longitudinal imaging datasets in the medical domain. This restriction is

pragmatic, as AD offers abundant imaging data suitable for model training and evaluation. However, it inevitably narrows the generalizability of our findings. In future work, we aim to extend our framework to other progressive diseases such as brain tumors, where longitudinal imaging carries unique challenges, including rapid progression patterns, heterogeneous lesion appearance, and treatment-induced alterations. Moreover, we encourage exploration of our method beyond medical imaging, where modeling temporal dynamics may shed light on broader scientific and engineering problems.

Dataset processing limitation. Although we followed standard MRI preprocessing protocols, variations across scanners and acquisition protocols introduce substantial heterogeneity in the data. Such differences may affect the stability of model training and the reliability of generated outputs. Furthermore, we identified occasional metadata inconsistencies (*e.g.*, mislabeling of imaging orientation), which led to registration failures in certain cases. While these issues were infrequent, they highlight the importance of rigorous data curation and quality control in longitudinal imaging research. Future work may benefit from harmonization techniques or scanner-invariant representations to further mitigate such variability.

H FUTURE WORK

In future research, we plan to enrich our framework in several important directions.

First, the incorporation of explicit anatomical information could enable more fine-grained modeling of disease progression. By constraining generation with structure-aware priors (*e.g.*, hippocampus, ventricles, cortical regions), the model may better capture local pathological changes while preserving biological plausibility.

Second, beyond modeling each patient in isolation, future work may explore leveraging inter-patient disease trend similarity and dissimilarity. Identifying cohorts of patients with comparable progression patterns may enable transfer of knowledge across individuals, while simultaneously capturing unique patient-specific deviations. In addition, modeling nonlinear progression dynamics remains an important direction for future work.

Addressing dataset heterogeneity remains another critical challenge. Future works may investigate harmonization techniques and scanner-invariant representations to mitigate variability introduced by different acquisition protocols and hardware. Such advances will be essential for scaling the framework to large, multi-center cohorts and for ensuring robust generalization in real-world clinical applications.

I DISCUSSION OF LLM USAGE

We used large language models (LLMs), including ChatGPT ([OpenAI, 2023](#)) and Gemini ([Team et al., 2023](#)), as writing assistants for grammar correction and style refinement. All technical content, methodology, and experiments were developed by the authors.