



Physics-Grounded Disentangled Flow Modeling for Brain Disease Progression Trajectory

Jun Wang and Peirong Liu*

Department of Electrical and Computer Engineering,
Data Science and AI Institute,
Johns Hopkins University
{jwang674, pliu53}@jh.edu

Abstract. Forecasting longitudinal brain lesion evolution is critical for disease monitoring and treatment planning. Existing approaches typically learn a direct mapping from a baseline image to a future observation, without explicitly modeling the physical mechanisms underlying the lesion progression. Such an entangled modeling of structural deformation and image intensity variation limits physical plausibility, model generalization, and interpretability. To address this, we propose **PDF**, a **P**hysics-grounded **D**isentangled **F**low matching framework for longitudinal brain disease forecasting. We explicitly decompose the longitudinal modeling of lesion growth into two processes, each learned by a dedicated flow matching network: morphology evolution, which captures lesion growth and structural deformation; and intensity evolution, which models signal changes driven by variations in lesion concentration. To enforce physics-grounded constraints, we introduce a PDE-regularized loss based on lesion growth dynamics, that enforces a diffusion-reaction-advection formulation for morphological evolution. Experiments on three public longitudinal datasets spanning diverse brain diseases demonstrate state-of-the-art performance, validating the effectiveness of the disentangled modeling framework and physics-grounded learning design. Code is publicly available at <https://github.com/jhuldr/PDF>.

Keywords: Physics-Informed Learning; Longitudinal Disease Modeling

1 Introduction

Accurate prediction of longitudinal brain disease progression is clinically important, as anticipating lesion evolution supports treatment planning, therapy adjustment, and risk assessment [8, 16, 38, 45]. Advances in imaging technologies have enabled routine longitudinal magnetic resonance imaging (MRI) acquisition, making it possible to model disease trajectories and forecast future structural changes [19, 22, 29]. Despite these advances, longitudinal prediction of

* Corresponding author.




Fig. 1: Disentangled modeling of brain disease progression. **Top:** irregularly sampled longitudinal MRI observations illustrating disease progression trajectory over time. **Bottom:** overview of the proposed disentangled framework, where morphology evolution models lesion growth and associated structural deformation, while intensity evolution captures signal changes induced by lesion concentration variations.

disease progression remains challenging due to limited data, irregular time intervals between scans, and imaging or protocol variability across time points, which together hinder reliable and generalizable progression modeling [11, 30, 35].

Recent advances in continuous-time and generative modeling have substantially improved longitudinal disease progression modeling, particularly through diffusion- and flow-based approaches [5, 15, 22, 23, 50]. I²SB [23] introduces a Schrödinger Bridge formulation within diffusion models to capture transitions between complex image distributions. ImageFlowNet [22] models disease evolution as a vector field learned through a position-parameterized neural ordinary differential equation (ODE), while trajectory flow matching [50] trains neural stochastic differential equations in a simulation-free manner to model continuous trajectories. Despite their flexibility, these methods often model progression as direct image-to-image translation from baseline. Such formulation implicitly entangles structural deformation and intensity variation without modeling the underlying physical mechanisms of lesion growth. As a result, they risk generating longitudinal predictions that lack physical consistency and reliability.

To address these limitations, we propose **PDF**, a **P**hysics-grounded **D**isentangled **F**low matching framework for longitudinal prediction of brain disease progression. Instead of directly generating a future scan, we decompose the brain disease progression into two complementary processes: morphological evolution and intensity evolution, as shown in Fig. 1. The former captures lesion growth and its associated structural deformation, while the latter models signal changes driven by variations in lesion concentration. Each process is learned

by a dedicated flow matching network, enabling structured and interpretable modeling of disease progression trajectories.

To enforce physically consistent lesion growth and improve the reliability of longitudinal prediction, we incorporate prior knowledge from lesion growth dynamics into morphological evolution modeling. Inspired by the lesion growth modeling through Fisher-KPP partial differential equation (PDE) [26, 27], we propose a PDE-regularized loss that constrains the predicted velocity fields to satisfy the diffusion-reaction-advection formulation. We demonstrated that such physics-grounded constraints encourage biologically plausible structural changes and improve longitudinal consistency. Our main contributions are:

- We propose a disentangled framework that explicitly decomposes brain disease progression into morphology evolution and intensity evolution, enabling interpretable and physically consistent modeling of disease dynamics.
- We introduce a physics-grounded regularization based on the Fisher-KPP PDE, constraining the predicted flow fields to follow a diffusion-reaction-advection formulation, and thereby enforcing biologically plausible structural evolution.
- We conduct experiments on three longitudinal brain MRI datasets with diverse modalities and diseases. Our method achieves state-of-the-art performance while producing more physically consistent predictions.

2 Related Work

2.1 Generative Modeling for Longitudinal Disease Progression

Generative and continuous-time models have recently become a dominant paradigm for modeling longitudinal aging analysis and disease progression [1, 17, 25]. Recent advances in diffusion and flow-based models enable flexible modeling of complex high-dimensional distributions over time [5, 15, 22, 23, 31, 46, 50].

Diffusion-based generative models [13, 37] have demonstrated remarkable performance in image synthesis and have been widely adopted in medical imaging tasks. To better model transitions between distributions, I²SB [23] formulates image-to-image translation as learning a Schrödinger bridge between endpoint distributions, allowing more direct modeling of temporal transitions beyond standard Gaussian-to-data diffusion. Such bridge-based perspectives provide a principled way to connect observations across time.

Flow-based continuous-time formulations offer a natural alternative for trajectory modeling. Neural ODEs [6] and neural SDEs [20] parameterize temporal dynamics through differential equations, accommodating irregular sampling. Continuous normalizing flows [12] further enable flexible density transport via learned vector fields. Flow Matching [21] further offers an efficient objective for learning probability paths without simulation. Building upon these advances, ImageFlowNet [22] learns multiscale image representations and evolves them using position-parameterized neural ODE/SDEs to forecast disease trajectories from baseline scans. Trajectory Flow Matching (TFM) [50] leverages flow matching to train neural SDEs in a simulation-free manner, improving scalability and

stability for stochastic time series modeling. IMMFM [15] further proposes interpolative multi-marginal flow matching to jointly learn stochastic dynamics consistent with multiple observed time points. Δ -LFM [5] models disease progression as a velocity field in a semantically aligned latent space, enabling interpretable patient-specific trajectory generation. While these methods demonstrate strong performance in longitudinal forecasting, they typically learn progression end-to-end, without explicitly disentangling structural deformation from appearance variation, or incorporating domain-specific priors governing lesion growth.

2.2 Physics-Based Modeling of Brain Lesions

Physics-based models characterize lesion growth through mechanistic descriptions of underlying biological processes. A common formulation relies on reaction-diffusion PDEs, which model lesion cell density evolution via proliferation and diffusive spread. Grounded in classical population dynamics, these models have been widely used to describe tumor and lesion progression [7, 32, 41, 43, 44]. A canonical example is the Fisher-KPP equation, which combines diffusion and logistic growth to model wave-like lesion expansion [33]. Analytical and numerical studies of such dynamics provide insights into propagation speed, spatial spread, and parameter influence [28, 42].

A key challenge is the inverse problem: estimating patient-specific tumor cell distributions or model parameters from imaging data [14, 18]. Prior work commonly formulates tumor personalization as a PDE-constrained optimization problem, which requires repeated forward PDE solves and is therefore computationally expensive [24, 36]. To address scalability, recent works integrate deep learning into physics-based inference. Learn-Morph-Infer [9] proposes amortized inference to directly predict patient-specific tumor cell distributions from MRI, significantly reducing personalization time while remaining compatible with reaction-diffusion models. GliODIL [3] softly assimilates imaging data and Fisher-KPP constraints through a discrete optimization framework, improving individualized radiotherapy planning. Similarly, [47] incorporates a learnable deep-learning prior to initialize physics-based evolutionary sampling, accelerating convergence in inverse tumor modeling. Physics-regularized multi-modal image assimilation [2] further integrates PDE growth and biomechanical constraints as regularization terms within a learnable discrete framework. Despite improved efficiency and hybrid modeling strategies, these approaches primarily focus on parameter estimation or hidden cell reconstruction. Their physics constraints are applied during inverse inference, rather than embedded within the forward generative trajectory for longitudinal image forecasting.

3 Method

Sec. 3.1 introduces our physics-grounded lesion evolution model, including PDE-based lesion growth dynamics (Sec. 3.1.1) and PDE-regularized progression learning (Sec. 3.1.2). Sec. 3.2 presents the disentangled flow matching framework, covering morphology evolution, intensity evolution, and the inference procedure.




Fig. 2: Physics-simulated lesion progression and projection of real progression onto the physics span. Top: lesion cell concentration evolution simulated under the diffusion-reaction-advection formulation. **Bottom:** projection of real lesion concentration changes onto the same diffusion-reaction-advection span, illustrating how real disease progression can be approximated within a physics-grounded subspace.

3.1 Physics-Grounded Lesion Evolution Modeling

3.1.1 PDE-based Lesion Growth Dynamics

A commonly used formulation for modeling lesion evolution is the Fisher-Kolmogorov-Petrovsky-Piskunov (Fisher-KPP) type diffusion-reaction PDE, which describes the spatiotemporal dynamics of lesion cell concentration [27, 48].

Let $c(\mathbf{x}, t) \in [0, 1]$ denote the normalized lesion cell concentration at spatial location $\mathbf{x}$ and time t . We have the Fisher-KPP dynamics as

$$\frac{\partial c}{\partial t} = D \Delta c + \rho c(1 - c), \quad (1)$$

where $D = D(t)$ is a diffusion coefficient which controls the spatial infiltration, capturing wave-like cell expansion, random movement, and spreading; $\rho = \rho(t)$ is the proliferation (i.e., reaction) rate governing local growth such as cell division and multiplication. To further account for the transport effects induced by larger structural deformation or mass effect [39], Eq. (1) is often augmented with an advection term and becomes

$$\frac{\partial c}{\partial t} = D \Delta c + \rho c(1 - c) - \mathbf{u} \cdot \nabla c, \quad (2)$$

$\mathbf{u} = \mathbf{u}(\mathbf{x}, t)$ is a spatiotemporally varying advection field. These three components correspond to diffusion, reaction, and advection processes, respectively.

3.1.2 PDE-Regularized Progression Learning

In practice, the underlying patient-specific PDE parameters $(D, \rho, \mathbf{u})$ in Eq. (2) are unknown. Instead of explicitly solving the full inverse problem, we design a physics-grounded formulation that embeds the PDE into learning.

Physics-Simulated Progression. Given two longitudinal observations, c_{t_1} at t_1 and c_{t_2} at t_2 , we first simulate a progression based on Eq. (2), as illustrated in Fig. 2 (top). We begin with the diffusion-reaction components and estimate the optimal (D, ρ) via least squares fitting:

$$(D, \rho) = \arg \min_{D, \rho} \left\| \frac{c_{t_2} - c_{t_1}}{t_2 - t_1} - D \Delta c_{t_1} - \rho c_{t_1} (1 - c_{t_1}) \right\|_2^2. \quad (3)$$

Given the estimated (D, ρ) , and the corresponding diffusion-reaction prediction $\tilde{c}_{t_2}$ from c_{t_1} , the remaining discrepancy between $\tilde{c}_{t_2}$ and the observed c_{t_2} is attributed to transport effects.

Then we approximate the residual motion for $\mathbf{u}$ in the advection term through image-based transport estimation. Specifically, we estimate a dense displacement field $\mathbf{d}(\mathbf{x})$ that aligns $\tilde{c}_{t_2}$ to c_{t_2} :

$$\mathbf{d} = \arg \min_{\mathbf{d}} \mathcal{E}(\tilde{c}_{t_2}(\mathbf{x}), c_{t_2}(\mathbf{x} + \mathbf{d}(\mathbf{x}))), \quad (4)$$

where $\mathcal{E}(\cdot)$ denotes a variational optical-flow objective. In practice, we employ a TV-L1 optical flow formulation [49] to obtain $\mathbf{d}$. The advection field is then

$$\mathbf{u}(\mathbf{x}) = -\frac{\mathbf{d}(\mathbf{x})}{t_2 - t_1}, \quad (5)$$

which converts displacement into instantaneous transport velocity. The negative sign follows from the semi-Lagrangian formulation used in forward simulation.

As shown in Fig. 2 (top), the full diffusion-reaction-advection in Eq. (2) is obtained with $(D, \rho, \mathbf{u})$ to form a physics-driven progression that closely approximates the observed longitudinal evolution.

Projection of Real Progression onto Physics Span. Given two observations c_{t_1} and c_{t_2} , the evolution of lesion cell concentration is approximated as

$$v_{\text{real}} = \frac{c_{t_2} - c_{t_1}}{t_2 - t_1}. \quad (6)$$

With the estimated fields $(D, \rho, \mathbf{u})$, and following the formulation in Eq. (2), v_{real} can be approximated via a combination of diffusion-, reaction-, and advection-induced components, evaluated at t_1 :

$$v_{\text{real}}(\mathbf{x}) \approx D \Delta c_{t_1}(\mathbf{x}) + \rho c_{t_1}(\mathbf{x})(1 - c_{t_1}(\mathbf{x})) - \mathbf{u}(\mathbf{x}) \cdot \nabla c_{t_1}(\mathbf{x}). \quad (7)$$

Accordingly, we define the physics span formed by the diffusion, reaction, and advection components as

$$\mathcal{S}(c_{t_1}; \mathbf{u}) = \text{span} \{ \Delta c_{t_1}, c_{t_1}(1 - c_{t_1}), -\mathbf{u} \cdot \nabla c_{t_1} \}. \quad (8)$$

Therefore, projecting v_{real} onto $\mathcal{S}(c_{t_1}; \mathbf{u})$ reduces to solving a least-squares problem over the scalar coefficients (α, β, γ) :

$$(\alpha, \beta, \gamma) = \arg \min_{\alpha, \beta, \gamma} \| v_{\text{real}} - \alpha \Delta c_{t_1} - \beta c_{t_1}(1 - c_{t_1}) - \gamma (-\mathbf{u} \cdot \nabla c_{t_1}) \|_2^2, \quad (9)$$




Fig. 3: Overview of the proposed PDF framework. **Top-left:** morphology flow matching network, which describes lesion growth and structural deformation. The proposed PDE-regularized loss minimizes the residual of the predicted velocity after projection onto the physics span. **Top-right:** intensity flow matching network, which models signal evolution conditioned on the image and lesion region. **Bottom:** inference procedure, where morphological evolution is first predicted to update lesion-induced structural changes, followed by intensity evolution to generate the final prediction.

and the projected velocity is $\Pi_S(v_{\text{real}}) = \alpha \Delta c_{t_1} + \beta c_{t_1}(1 - c_{t_1}) + \gamma (-\mathbf{u} \cdot \nabla c_{t_1})$.

To quantify the degree to which real longitudinal evolution conforms to this PDE-induced structure, we report the normalized projection error

$$\epsilon = \frac{\|v_{\text{real}} - \Pi_S(v_{\text{real}})\|_2}{\|v_{\text{real}}\|_2}, \quad (10)$$

which is visualized as a residual map in Fig. 2 (bottom).

PDE-Regularized Loss. To incorporate this physics prior into learning, we constrain the predicted velocity field v_θ (i.e., evolution of lesion cell concentration in Eq. (6)) to remain close to the physics span $\mathcal{S}(c_{t_1}; \mathbf{u})$. Specifically, for a given v_θ , we estimate its projection onto $\mathcal{S}(c_{t_1}; \mathbf{u})$ by solving

$$(\alpha, \beta, \gamma) = \arg \min_{\alpha, \beta, \gamma} \|v_\theta - \alpha \Delta c_{t_1} - \beta c_{t_1}(1 - c_{t_1}) - \gamma (-\mathbf{u} \cdot \nabla c_{t_1})\|_2^2. \quad (11)$$

Denoting the projected velocity as

$$\Pi_S(v_\theta) = \alpha \Delta c_{t_1} + \beta c_{t_1}(1 - c_{t_1}) + \gamma (-\mathbf{u} \cdot \nabla c_{t_1}), \quad (12)$$

we define the PDE regularization term as the squared projection residual:

$$\mathcal{L}_{\text{PDE}} = \|v_\theta - \Pi_S(v_\theta)\|_2^2. \quad (13)$$

This regularizer penalizes deviations of the estimated lesion cell concentration evolution from the diffusion-reaction-advection subspace, thereby embedding physics-grounded constraints into learning without explicitly solving a full PDE inverse problem.

3.2 Disentangled Flow Matching Framework

As illustrated in Fig. 3, instead of directly predicting a future scan from a baseline image, we decompose longitudinal image prediction of lesion progression into two complementary components: *morphology evolution*, capturing lesion growth and structural deformation, and *intensity evolution*, modeling the overall image appearance changes conditioned on the predicted lesion evolution morphology. Each component is modeled with a corresponding flow matching module.

Morphology Flow Matching. Let M_{t_1} and M_{t_2} denote the binary lesion masks at time t_1 and t_2 , respectively, with time interval as $t_2 - t_1$. To obtain a continuous field for PDE modeling, each binary mask is converted into a smooth signed distance function and mapped to a normalized lesion cell concentration $c_t(\mathbf{x}) \in [0, 1]$, where larger values indicate higher lesion cell density. The resulting concentrations c_{t_1} and c_{t_2} are then used to estimate the PDE parameters, $(D, \rho, \mathbf{u})$, in Sec. 3.1.2.

We learn a velocity field that describes the morphological evolution associated with lesion growth,

$$v_\theta = f_\theta(M_{t_1}; I_{t_1}, t_2 - t_1), \quad (14)$$

where I_{t_1} , the image at t_1 , serves as a structural context, $t_2 - t_1$ encodes temporal conditioning. The network f_θ predicts the instantaneous morphological velocity. In flow matching, the empirical velocity target is approximated by

$$v_{\text{morph}} = \frac{M_{t_2} - M_{t_1}}{t_2 - t_1}, \quad (15)$$

and the standard flow matching objective is defined as,

$$\mathcal{L}_{\text{FM}} = \|v_\theta - v_{\text{morph}}\|_2^2. \quad (16)$$

Incorporating the PDE regularization defined in Eq. (13) to enforce physics consistency, we have

$$\mathcal{L}_{\text{morph}} = \mathcal{L}_{\text{FM}} + \lambda \mathcal{L}_{\text{PDE}}, \quad (17)$$

where λ controls the strength of the physics-grounded constraint.

Intensity Flow Matching. While morphology flow accounts for structural deformation, intensity evolution captures signal variations conditioned on lesion evolution. During training, we adopt a flow-matching scheme in image space. Specifically, a noisy sample z_τ is generated along a predefined interpolation trajectory in image space, parameterized by $\tau \in [0, 1]$. The intensity evolution velocity is defined as

$$v_\phi = g_\phi(z_\tau; M), \quad (18)$$

Following the standard flow matching formulation, the target velocity $v_{\text{intensity}}$ is defined by the chosen interpolation path in image space. The image intensity learning objective is therefore

$$\mathcal{L}_{\text{intensity}} = \|v_\phi - v_{\text{intensity}}\|_2^2. \quad (19)$$

This formulation accounts for changes in image intensity, in addition to structural deformations resulting from lesion morphology.

Inference. At inference time, the predicted morphological velocity v_θ is first integrated over the time interval $(t_2 - t_1)$ to obtain the evolved lesion mask $\hat{M}_{t_2}$. To obtain the dense structural deformation, we estimate a displacement field ($\mathbf{f}$) between $\hat{M}_{t_2}$ and M_{t_1} with a variational optical flow formulation. Specifically,

$$\mathbf{f}(\mathbf{x}) = \arg \min_{\mathbf{f}} \mathcal{E}(\hat{M}_{t_2}(\mathbf{x}), M_{t_1}(\mathbf{x} + \mathbf{f}(\mathbf{x}))), \quad (20)$$

where $\mathcal{E}$ denotes a TV-L1 optical flow objective.

Given the estimated dense morph field $\mathbf{f}(\mathbf{x})$, the baseline image I_{t_1} at time t_1 is warped using a semi-Lagrangian formulation:

$$\hat{I}^{\text{morph}}(\mathbf{x}) = I_{t_1}(\mathbf{x} + \mathbf{f}(\mathbf{x})), \quad (21)$$

where bilinear interpolation is used for sampling, and the sampling coordinates are clipped to remain within structural boundaries. This step models structural deformation induced by lesion growth.

While $\hat{I}^{\text{morph}}$ captures structural displacement, it does not account for the image intensity changes induced by the estimated lesion cell concentration variations. To model appearance evolution, we apply the intensity flow matching module to the warped image, $\hat{I}^{\text{morph}}$. Specifically, $\hat{I}^{\text{morph}}$ serves as the structural initialization, and is refined by integrating the learned intensity velocity field conditioned on the baseline lesion mask M_{t_1} . To preserve lesion-specific texture patterns, high-frequency components extracted from I_{t_1} are incorporated as additional guidance during refinement.

The final prediction $\hat{I}_{t_2}$ is obtained after intensity flow integration, producing appearance changes that are consistent with both the predicted morphology and the intrinsic lesion texture.

4 Experiments

4.1 Experimental Settings

Datasets. We evaluate our framework on three publicly available longitudinal brain disease cohorts, covering diverse modalities and pathology types. In each dataset, images are longitudinally co-registered per subject during preprocessing.

- **UCSF.** The UCSF-ALPTDG dataset [10] provides longitudinal MRI scans of 302 adult patients with *brain glioma tumors*. Each patient has two time points

with expert-annotated lesion masks for precise evaluation of tumor progression. In this work, we use the T1-weighted MRI modality for longitudinal modeling.

- **LUMIERE.** The LUMIERE dataset [40] contains longitudinal MRI scans of 91 patients with *glioblastoma tumors*, along with corresponding lesion masks. Each patient underwent one pre-operative scan up to five years, resulting in 795 longitudinal series with 2-18 time points per case. Following prior work, we retain only post-operative scans to model natural tissue evolution after surgery, and use contrast-enhanced T1-weighted MRI scans in our experiments.

- **LMSLS.** The LMSLS dataset [4] contains longitudinal FLAIR MRI scans from 79 data points with *multiple sclerosis (MS) lesions*. Each subject has an average of 4.4 visits collected over approximately five years with expert-annotated lesion masks. We adopt FLAIR-weighted MRI for longitudinal lesion modeling.

Lesion Cell Concentration Map Construction. For each binary lesion mask M , we first convert it into a continuous cell concentration map, via a signed distance function (SDF) that measures the distance from each pixel to the lesion boundary (∂M). The signed distance function $d(\mathbf{x})$ is defined as

$$d(\mathbf{x}) = \begin{cases} -\min_{\mathbf{y} \in \partial M} |\mathbf{x} - \mathbf{y}|, & \mathbf{x} \in M, \\ \min_{\mathbf{y} \in \partial M} |\mathbf{x} - \mathbf{y}|, & \mathbf{x} \notin M, \end{cases} \quad (22)$$

where negative (positive) values indicate locations inside (outside) the lesion. We further normalize $d(\mathbf{x})$ with sigmoid transformation $c(\mathbf{x}) = 1 / \left(1 + \exp\left(\frac{d(\mathbf{x})}{\tau}\right) \right)$, where τ controls the smoothness of the transition around the lesion boundary. This transformation produces a continuous lesion cell concentration field $c(\mathbf{x}) \in [0, 1]$, where larger values correspond to higher lesion cell density.

Evaluation Metrics. We evaluate longitudinal prediction performance from both image fidelity and disease region consistency perspectives. For lesion-level assessment, we compute the Dice-Sørensen coefficient (DSC) and Hausdorff distance (HD) between predicted and ground-truth lesion masks. DSC measures volumetric overlap between regions, while HD quantifies the maximum boundary discrepancy. Lesion regions are obtained using dataset-specific image segmentation networks following [22]. For image-level assessment, we report average L1 error, peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM) for reconstruction quality and structural consistency of the predicted images.

Implementation Details. All time points in each longitudinal series were rigidly registered to ensure spatial alignment across visits. We partition the data into training, validation, and test sets using a 6:2:2 ratio. To prevent data leakage, splitting was performed at the longitudinal series level, ensuring that all scans from the same patient were assigned to the same subset. All flow matching models are implemented using a U-Net backbone with a shared encoder-decoder structure. Experiments were conducted on an NVIDIA RTX PRO 6000 GPU with CUDA version 13.0 and PyTorch 2.9.1.


































	Ground-truth (t_2)	Input (t_1)	T-UNet	I ² SB	TFM	ImageFlowNet	PDF (Ours)
	PSNR↑/SSIM↑	28.4/0.88	28.1/0.88	28.4/0.88	28.4/0.88	26.4/0.81	28.8/0.89
Image (Tumor)							
	DSC↑/HD↓	0.53/24.2	0.53/20.4	0.54/24.2	0.54/24.0	0.48/24.6	0.74/18.2
w/ Lesion Region							
	PSNR↑/SSIM↑	28.4/0.88	27.1/0.89	26.4/0.89	26.4/0.88	26.9/0.75	26.8/0.88
Image (Tumor)							
	DSC↑/HD↓	0.52/21.6	0.52/21.6	0.51/21.6	0.51/21.6	0.32/19.7	0.72/15.6
w/ Lesion Region							
	PSNR↑/SSIM↑	31.6/0.91	31.9/0.92	31.6/0.91	31.6/0.91	31.5/0.90	31.6/0.92
Image (MS)							
	DSC↑/HD↓	0.75/11.4	0.78/12.7	0.77/11.4	0.75/11.4	0.74/11.5	0.82/9.49
w/ Lesion Region							

Fig. 4: Comparison of longitudinal modeling methods on two tumor cases and one MS lesion case. Each case is presented in two rows: the first row shows the predicted images with PSNR/SSIM, the second row shows images overlaid with corresponding lesion regions, along with DSC/HD. Compared to the competing methods, **PDF** produces more accurate lesion boundaries and improved structural consistency.

4.2 Performance Comparison

We compare our method with four representative longitudinal disease modeling approaches across all datasets: Time-conditional UNet (T-UNet) [34], I²SB [23], Trajectory Flow Matching (TFM) [50], ImageFlowNet [22]. T-UNet integrates temporal information by injecting time embeddings into intermediate UNet representations. I²SB is a diffusion-based image-to-image Schrödinger Bridge model. TFM formulates temporal evolution as flow matching in image space. ImageFlowNet models longitudinal dynamics via neural-ODE-based flow prediction.

Tab. 1 reports quantitative results on UCSF, LUMIERE, and LMSLS. On UCSF, our method achieves the highest DSC of 0.752 and the lowest HD of 25.952, outperforming T-UNet with a DSC of 0.728 and an HD of 27.322. Al-

Method	UCSF (Tumor)					LUMIERE (Tumor)					LMSLS (MS)				
	DSC $\uparrow$	HD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	L1 $\downarrow$	DSC $\uparrow$	HD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	L1 $\downarrow$	DSC $\uparrow$	HD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	L1 $\downarrow$
T-UNet [34]	0.728 ± 0.004	27.322 ± 0.812	31.739 ± 0.103	0.920 ± 0.002	0.085 ± 0.002	0.520 ± 0.020	38.135 ± 1.925	36.753 ± 0.185	0.845 ± 0.004	0.097 ± 0.003	0.673 ± 0.005	26.354 ± 0.932	29.248 ± 0.118	0.805 ± 0.003	0.083 ± 0.002
I ² SB [23]	0.725 ± 0.005	28.839 ± 0.941	31.393 ± 0.097	0.920 ± 0.002	0.093 ± 0.002	0.522 ± 0.035	36.105 ± 1.720	37.206 ± 0.348	0.953 ± 0.003	0.084 ± 0.003	0.682 ± 0.006	23.617 ± 0.784	29.312 ± 0.094	0.893 ± 0.002	0.068 ± 0.003
TFM [50]	0.696 ± 0.006	31.701 ± 1.284	30.659 ± 0.125	0.898 ± 0.003	0.091 ± 0.002	0.492 ± 0.015	39.211 ± 1.530	36.317 ± 0.406	0.894 ± 0.001	0.088 ± 0.002	0.670 ± 0.007	24.438 ± 0.851	29.441 ± 0.102	0.889 ± 0.003	0.062 ± 0.004
ImageFlowNet [22]	0.724 ± 0.004	28.402 ± 0.893	31.675 ± 0.108	0.909 ± 0.003	0.085 ± 0.002	0.493 ± 0.033	41.994 ± 1.653	37.178 ± 0.346	0.916 ± 0.003	0.078 ± 0.003	0.687 ± 0.005	24.296 ± 0.741	28.972 ± 0.113	0.890 ± 0.003	0.072 ± 0.002
PDF (Ours)	0.752 ± 0.003	25.952 ± 0.654	31.364 ± 0.092	0.920 ± 0.001	0.072 ± 0.002	0.534 ± 0.013	33.672 ± 1.229	37.216 ± 0.349	0.957 ± 0.002	0.071 ± 0.002	0.696 ± 0.004	20.867 ± 0.623	29.314 ± 0.087	0.900 ± 0.002	0.058 ± 0.002

Table 1: Quantitative comparison across three longitudinal datasets.

though the PSNR of our method is 31.364, slightly lower than 31.739 from T-UNet, the improved lesion-level metrics indicate more accurate boundary alignment. On the LUMIERE dataset, our method achieves the best DSC of 0.534 and the lowest HD of 33.672, surpassing both T-UNet and I²SB in segmentation accuracy. PSNR and SSIM are 37.216 and 0.957, respectively, demonstrating stable intensity reconstruction under post-operative tissue evolution. On LMSLS, which involves subtle and gradual lesion progression, our method achieves the highest DSC of 0.696 and the lowest HD of 20.867, compared with 0.673 DSC and 26.354 HD from T-UNet. These results suggest improved sensitivity to small-scale lesion evolution.

Fig. 4 presents qualitative comparisons. The first two rows show a glioma case with a notable tumor region and structural deformation. Compared to the ground-truth follow-up, the initial input yields a DSC of 0.53 and an HD of 24.2, while our method achieves 0.74 and 18.2 and produces boundaries that closely match the ground truth. All competing models tend to underestimate the lesion morphological changes or generate less coherent lesion expansion. This demonstrates the advantage of our *Morphology Evolution* module. The third and fourth rows present a glioma case with prominent intensity changes. Our method achieves a DSC of 0.72 and an HD of 15.6, compared with 0.52 and 21.6 from the input, while better preserving appearance consistency. This highlights the contribution of our *Intensity Evolution* module. The last two rows illustrate an MS case, with subtle lesion progression. Our model successfully captures the lesion evolution in the bottom-right region of the image, while competing methods struggle to do so. Overall, the qualitative results are consistent with the quantitative improvements observed across diverse disease dynamics.

Notes on evaluation metrics and performance interpretation. We observe that image-level metrics (PSNR, SSIM, L1) may not fully reflect the lesion progression performance, as lesion changes are typically small compared to the whole brain region. Therefore, lesion-level metrics (DSC and HD) are more representative, as they focus on lesion regions in particular. As shown in Fig. 5, although all models achieve satisfactory image-level performance, our method demonstrates much stronger progression modeling and more accurately captures aggressive tumor progression (DSC = 0.85), including the ventricle distortion, whereas other models struggle to model such pronounced structural changes.




Fig. 5: An edge case of aggressive lesion progression. This example illustrates a case in which aggressive lesion progression causes *structural distortion* of the originally normal ventricle. Only PDF successfully captures this severe progression pattern.

Component		UCSF (Tumor)				LUMIERE (Tumor)				LMSLS (MS)			
Morph	Intensity	DSC ↑	HD ↓	PSNR ↑	SSIM ↑	DSC ↑	HD ↓	PSNR ↑	SSIM ↑	DSC ↑	HD ↓	PSNR ↑	SSIM ↑
✓		0.750	26.038	31.343	0.920	0.532	34.066	37.216	0.958	0.694	21.855	29.359	0.900
	✓	0.749	25.806	31.383	0.920	0.529	35.143	37.025	0.957	0.690	21.749	29.307	0.900
✓	✓	0.752	25.952	31.364	0.920	0.534	33.672	37.075	0.957	0.696	20.867	29.314	0.900

Table 2: Ablation of disentangled morphology and intensity evolutions across three longitudinal datasets. We evaluate morphology-only, intensity-only, and joint models to demonstrate the complementary strengths of the two components.

4.3 Ablation Studies

Disentangled Components. We perform ablation studies on the disentangled morphology and intensity evolution modeling by enabling each separately and comparing them with the jointly trained full model. Since the intensity branch requires a predicted lesion mask, we use the mask from the morphology branch while disabling structural deformation.

Tab. 2 reports the quantitative results. The morphology branch alone improves lesion-level metrics, particularly DSC, highlighting the importance of structural deformation modeling. The intensity branch alone yields competitive PSNR and SSIM, reflecting strong appearance modeling. When combined, the full model achieves the best or near-best performance across datasets. In particular, on UCSF and LMSLS, the joint model improves DSC and reduces HD compared to either single component, indicating that morphology and intensity modeling provide complementary benefits.

Fig. 6 provides qualitative visualizations of the disentangled components. Applying the morphology branch first produces a structurally aligned prediction that yields structurally consistent evolution, improving DSC from 0.84 to 0.89. The subsequent intensity evolution refines lesion boundaries and internal textures, further increasing DSC to 0.91. The structural alignment corrects lesion-induced spatial shifts, while intensity modeling captures residual appearance changes, together producing more accurate longitudinal lesion trajectories.

PDE-Regularized Loss. We also evaluate the effect of the proposed PDE-regularized loss by comparing models trained with and without $\mathcal{L}_{\text{PDE}}$ (Eq. (13)).



Fig. 6: Qualitative visualization of disentangled prediction. From left to right: ground-truth and input images with lesions; original lesion at t_1 and predicted lesion mask, estimated morph field, morph-warped image, intensity-evolved final prediction.

Loss	UCSF (Tumor)				LUMIERE (Tumor)				LMSLS (MS)			
	DSC $\uparrow$	HD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	DSC $\uparrow$	HD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	DSC $\uparrow$	HD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$
w/o $\mathcal{L}_{\text{PDE}}$	0.745	26.216	31.359	0.920	0.524	34.211	37.137	0.953	0.690	22.337	29.312	0.898
w/ $\mathcal{L}_{\text{PDE}}$	0.752	25.952	31.364	0.920	0.534	33.672	37.216	0.957	0.696	20.867	29.314	0.900

Table 3: Ablation of PDE-regularized loss across three longitudinal datasets. We compare models trained with and without $\mathcal{L}_{\text{PDE}}$. The model with $\mathcal{L}_{\text{PDE}}$ consistently improves lesion-level metrics while maintaining image-level performance.

Tab. 3 summarizes the quantitative results. Across all three datasets, incorporating $\mathcal{L}_{\text{PDE}}$ consistently improves lesion-level metrics. On UCSF, DSC increases from 0.745 to 0.752 while HD decreases from 26.216 to 25.952. On LUMIERE, both DSC and HD show improvements, and on LMSLS the gain is more pronounced, with DSC increasing from 0.690 to 0.696 and HD decreasing from 22.337 to 20.867. These results suggest that PDE regularization improves structural consistency in longitudinal prediction.

Fig. 7 further illustrates the effect of $\mathcal{L}_{\text{PDE}}$ on a representative case. The first two columns show the ground-truth and input images, with their lesion regions outlined. The third column visualizes the predicted mask-evolution velocity with and without PDE regularization. The fourth column presents the energy decomposition of the velocity projected onto the predefined PDE basis and the corresponding residual ratio. Without $\mathcal{L}_{\text{PDE}}$, the residual energy is 0.20, indicating that a substantial portion of the predicted dynamics lies outside the physics-based span. With PDE regularization, the residual decreases to 0.12, showing improved alignment with the underlying physical basis. Consistently, the predicted lesion DSC increases from 0.69 to 0.73, with better shape agreement with the ground truth. These results suggest that $\mathcal{L}_{\text{PDE}}$ promotes physically consistent evolution and improves structural accuracy.

Key Parameter. We further investigate the influence of the PDE regularization weight λ in the morphological loss (Eq. (17)), where λ controls the balance between data-driven flow matching and physics-grounded PDE regularization.

As shown in Tab. 4, moderate regularization ($\lambda = 0.01$) consistently achieves the best overall lesion-level performance across datasets. Larger values ($\lambda = 0.1$ and $\lambda = 1$) slightly degrade DSC and HD, suggesting that overly strong




Fig. 7: Effect of PDE-regularized loss on longitudinal prediction. Left to right: ground-truth and input images; predicted morphology evolution velocity; energy projection within physics span and its residual; predicted images and lesion regions with and without $\mathcal{L}_{PDE}$. $\mathcal{L}_{PDE}$ reduces residual energy and improves lesion evolution.

Weight	UCSF				LUMIERE				LMSLS			
	DSC $\uparrow$	HD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	DSC $\uparrow$	HD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$	DSC $\uparrow$	HD $\downarrow$	PSNR $\uparrow$	SSIM $\uparrow$
$\lambda = 1$	0.745	26.224	31.482	0.920	0.526	35.241	37.134	0.925	0.691	21.634	29.307	0.898
$\lambda = 0.1$	0.746	25.635	31.321	0.920	0.529	34.128	37.163	0.936	0.693	21.402	29.322	0.898
$\lambda = 0.01$	0.752	25.952	31.364	0.920	0.534	33.672	37.216	0.957	0.696	20.867	29.314	0.900

Table 4: Ablation of key parameter across three longitudinal brain MRI datasets. We evaluate different values of λ across longitudinal brain MRI datasets.

physical constraints may hinder accurate deformation modeling. These results indicate that appropriately balancing data-driven learning and physics-based regularization is crucial for stable and effective longitudinal forecasting.

5 Conclusion

We presented **PDF**, a physics-grounded disentangled flow matching framework for longitudinal brain disease progression. By disentangling disease morphological evolution from image intensity evolution, our approach enables structured and interpretable modeling of disease progression. Incorporating a PDE-regularized loss grounded in lesion growth dynamics further constrains the predicted trajectories within a biologically meaningful diffusion-reaction-advection formulation, improving physical plausibility and longitudinal consistency. Experiments across three longitudinal datasets demonstrate that **PDF** achieves state-of-the-art performance while producing more coherent lesion evolution patterns. These results highlight the importance of integrating disentangled modeling and physics-grounded constraints for reliable longitudinal disease forecasting.

Acknowledgements. We gratefully acknowledge the support of NVIDIA Corporation for providing the NVIDIA RTX PRO 6000 GPU used in this research.

References

1. Bai, H., Hong, Y.: Noder: Image sequence regression based on neural ordinary differential equations. In: International Conference on Medical Image Computing and Computer-Assisted Intervention (2024)
2. Balcerak, M., Amiranashvili, T., Wagner, A., Weidner, J., Karnakov, P., Paetzold, J.C., et al.: Physics-regularized multi-modal image assimilation for brain tumor localization. *Advances in Neural Information Processing Systems* (2024)
3. Balcerak, M., Weidner, J., Karnakov, P., Ezhov, I., Litvinov, S., Koumoutsakos, P., et al.: Individualizing glioma radiotherapy planning by optimization of a data and physics-informed discrete loss. *Nature Communications* (2025)
4. Carass, A., Roy, S., Jog, A., Cuzzocreo, J.L., Magrath, E., Gherman, A., et al.: Longitudinal multiple sclerosis lesion segmentation: resource and challenge. *NeuroImage* (2017)
5. Chen, H., Yin, R., Chen, Y., Chen, Q., Li, C.: Learning patient-specific disease dynamics with latent flow matching for longitudinal imaging generation. In: International Conference on Learning Representations (2026)
6. Chen, R.T., Rubanova, Y., Bettencourt, J., Duvenaud, D.K.: Neural ordinary differential equations. In: *Advances in Neural Information Processing Systems* (2018)
7. Clatz, O., Sermesant, M., Bondiau, P.Y., Delingette, H., Warfield, S.K., Malandain, G., Ayache, N.: Realistic simulation of the 3-d growth of brain tumors in mr images coupling diffusion with biomechanical deformation. *IEEE Transactions on Medical Imaging* (2005)
8. Durso-Finley, J., Falet, J.P., Nichyporuk, B., Douglas, A., Arbel, T.: Personalized prediction of future lesion activity and treatment effect in multiple sclerosis from baseline mri. In: International Conference on Medical Imaging with Deep Learning (2022)
9. Ezhov, I., Scibilia, K., Frantza, K., Steinbauer, F., Shit, S., Zimmer, L., et al.: Learn-morph-infer: a new way of solving the inverse problem for brain tumor modeling. *Medical Image Analysis* (2023)
10. Fields, B.K., Calabrese, E., Mongan, J., Cha, S., Hess, C.P., Sugrue, L.P., et al.: The university of california san francisco adult longitudinal post-treatment diffuse glioma mri dataset. *Radiology: Artificial Intelligence* (2024)
11. Fonteijn, H.M., Clarkson, M.J., Modat, M., Barnes, J., Lehmann, M., Ourselin, S., Fox, N.C., Alexander, D.C.: An event-based disease progression model and its application to familial alzheimer’s disease. In: Biennial International Conference on Information Processing in Medical Imaging (2011)
12. Grathwohl, W., Chen, R.T.Q., Bettencourt, J., Sutskever, I., Duvenaud, D.: Ffjord: Free-form continuous dynamics for scalable reversible generative models. In: International Conference on Learning Representations (2019)
13. Ho, J., Jain, A., Abbeel, P.: Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems* (2020)
14. Hogue, C., Davatzikos, C., Biros, G.: An image-driven parameter estimation problem for a reaction–diffusion glioma growth model with mass effects. *Journal of Mathematical Biology* (2008)
15. Islam, M.M., Kuipers, T.P., Vadgama, S., de Vente, C., Khan, A., Sánchez, C.I., Bekkers, E.J.: Longitudinal flow matching for trajectory modeling. In: International Conference on Artificial Intelligence and Statistics (2026)

16. Jackson, P.R., Juliano, J., Hawkins-Daarud, A., Rockne, R.C., Swanson, K.R.: Patient-specific mathematical neuro-oncology: using a simple proliferation and invasion tumor model to inform clinical practice. *Bulletin of mathematical biology* (2015)
17. Jian, B., Pan, J., Li, Y., Bongratz, F., Li, R., Rueckert, D., Wiestler, B., Wachinger, C.: Temporal conditioning for longitudinal brain mri registration and aging analysis. *IEEE Transactions on Medical Imaging* (2026)
18. Konukoglu, E., Clatz, O., Bondiau, P.Y., Delingette, H., Ayache, N.: Extrapolating glioma invasion margin in brain magnetic resonance images: suggesting new irradiation margins. *Medical Image Analysis* (2010)
19. Lachinov, D., Chakravarty, A., Grechenig, C., Schmidt-Erfurth, U., Bogunović, H.: Learning spatio-temporal model of disease progression with neuralodes from longitudinal volumetric data. *IEEE Transactions on Medical Imaging* (2023)
20. Li, X., Wong, T.K.L., Chen, R.T., Duvenaud, D.: Scalable gradients for stochastic differential equations. In: *International Conference on Artificial Intelligence and Statistics*. PMLR (2020)
21. Lipman, Y., Chen, R.T.Q., Ben-Hamu, H., Nickel, M., Le, M.: Flow matching for generative modeling. In: *International Conference on Learning Representations* (2023)
22. Liu, C., Xu, K., Shen, L.L., Huguet, G., Wang, Z., Tong, A., et al.: Image-flownet: Forecasting multiscale image-level trajectories of disease progression with irregularly-sampled longitudinal medical images. In: *IEEE International Conference on Acoustics, Speech and Signal Processing* (2025)
23. Liu, G.H., Vahdat, A., Huang, D.A., Theodorou, E.A., Nie, W., Anandkumar, A.: $\mathbb{I}^2\text{sb}$: Image-to-image schrödinger bridge. In: *International Conference on Machine Learning* (2023)
24. Mang, A., Gholami, A., Davatzikos, C., Biros, G.: Pde-constrained optimization in medical image analysis. *Optimization and Engineering* (2018)
25. Marcus, D.S., Fotenos, A.F., Csernansky, J.G., Morris, J.C., Buckner, R.L.: Open access series of imaging studies: longitudinal mri data in nondemented and demented older adults. *Journal of Cognitive Neuroscience* (2010)
26. Martens, C., Rovai, A., Bonatto, D., Metens, T., Debeir, O., Decaestecker, C., Goldman, S., Van Simaey, G.: Deep learning for reaction-diffusion glioma growth modeling: Towards a fully personalized model? *Cancers* (2022)
27. Metzcar, J., Jutzeler, C.R., Macklin, P., Köhn-Luque, A., Brüningk, S.C.: A review of mechanistic learning in mathematical oncology. *Frontiers in Immunology* (2024)
28. Murray, J.D.: *Mathematical biology i: An introduction*. Interdisciplinary applied mathematics (2002)
29. Nabrawi, E., Alanazi, A.T., Al Alkhaibari, E.: Imaging in healthcare: A glance at the present and a glimpse into the future. *Cureus* (2023)
30. Nguyen, H.H., Blaschko, M.B., Saarakkala, S., Tiulpin, A.: Clinically-inspired multi-agent transformers for disease trajectory forecasting from multimodal data. *IEEE Transactions on Medical Imaging* (2023)
31. Puglisi, L., Alexander, D.C., Ravi, D., Initiative, A.D.N., et al.: Brain latent progression: Individual-based spatiotemporal disease progression on 3d brain mris via latent diffusion. *Medical Image Analysis* (2025)
32. Raissi, M., Perdikaris, P., Karniadakis, G.E.: Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics* (2019)

33. Rockne, R., Rockhill, J., Mrugala, M., Spence, A., Kalet, I., Hendrickson, K., et al.: Predicting the efficacy of radiotherapy in individual glioblastoma patients in vivo: a mathematical modeling approach. *Physics in Medicine & Biology* (2010)
34. Ronneberger, O., Fischer, P., Brox, T.: U-net: Convolutional networks for biomedical image segmentation. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2015)
35. Ruan, C., Huang, C., Yang, Y.: Comprehensive evaluation of multimodal ai models in medical imaging diagnosis: From data augmentation to preference-based comparison. In: *International Conference on Bioinformatics and Computational Biology* (2025)
36. Scheufele, K., Subramanian, S., Mang, A., Biros, G., Mehl, M.: Image-driven biophysical tumor growth model calibration. *SIAM journal on scientific computing: a publication of the Society for Industrial and Applied Mathematics* (2020)
37. Song, Y., Sohl-Dickstein, J., Kingma, D.P., Kumar, A., Ermon, S., Poole, B.: Score-based generative modeling through stochastic differential equations. In: *International Conference on Learning Representations* (2021)
38. Stupp, R., Brada, M., Van Den Bent, M., Tonn, J.C., Pentheroudakis, G., Group, E.G.W., et al.: High-grade glioma: Esmo clinical practice guidelines for diagnosis, treatment and follow-up. *Annals of oncology* (2014)
39. Subramanian, S., Scheufele, K., Himthani, N., Biros, G.: Multiatlas calibration of biophysical brain tumor growth models with mass effect. In: *International Conference on Medical Image Computing and Computer-Assisted Intervention* (2020)
40. Suter, Y., Knecht, U., Valenzuela, W., Notter, M., Hwer, E., Schuch, P., et al.: The lumiere dataset: Longitudinal glioblastoma mri with expert rano evaluation. *Scientific Data* (2022)
41. Swanson, K.R., Alvord Jr, E.C., Murray, J.D.: A quantitative model for differential motility of gliomas in grey and white matter. *Cell proliferation* (2000)
42. Swanson, K.R., Bridge, C., Murray, J., Alvord Jr, E.C.: Virtual and real brain tumors: using mathematical modeling to quantify glioma growth and invasion. *Journal of the Neurological Sciences* (2003)
43. Tracqui, P., Cruywagen, G., Woodward, D., Bartoo, G., Murray, J., Alvord Jr, E.: A mathematical model of glioma growth: the effect of chemotherapy on spatio-temporal growth. *Cell proliferation* (1995)
44. Tunç, B., Hormuth II, D.A., Biros, G., Yankeelov, T.E.: Modeling of glioma growth with mass effect by longitudinal magnetic resonance imaging. *IEEE Transactions on Biomedical Engineering* (2021)
45. Wang, G., Ye, J.C., De Man, B.: Deep learning for tomographic image reconstruction. *Nature Machine Intelligence* (2020)
46. Wang, J., Liu, P.: Usb: Unified synthetic brain framework for bidirectional pathology-healthy generation and editing. *arXiv preprint arXiv:2512.00269* (2025)
47. Weidner, J., Ezhov, I., Balcerak, M., Metz, M.C., Litvinov, S., Kaltenbach, S., et al.: A learnable prior improves inverse tumor growth modeling. *IEEE Transactions on Medical Imaging* (2024)
48. Yin, A., Moes, D.J.A., van Hasselt, J.G., Swen, J.J., Guchelaar, H.J.: A review of mathematical models for tumor dynamics and treatment resistance evolution of solid tumors. *CPT: pharmacometrics & systems pharmacology* (2019)
49. Zach, C., Pock, T., Bischof, H.: A duality based approach for realtime tv-l 1 optical flow. In: *Joint pattern recognition symposium* (2007)
50. Zhang, X.N., Pu, Y., Kawamura, Y., Loza, A., Bengio, Y., Shung, D., Tong, A.: Trajectory flow matching with applications to clinical time series modelling. In: *Advances in Neural Information Processing Systems* (2024)