

ImageFlowNet: Forecasting Multiscale Image-Level Trajectories of Disease Progression with Irregularly-Sampled Longitudinal Medical Images

Chen Liu¹, Ke Xu¹, Liangbo L Shen², Guillaume Huguet^{3,4}, Zilong Wang^{3,5}, Alexander Tong^{3,4}, Danilo Bzdok^{3,5}, Jay Stewart², Jay C Wang^{1,2,6}, Lucian V Del Priore¹, and Smita Krishnaswamy¹

¹Yale University

²University of California

³Mila -Quebec AI Institute

⁴Université de Montréal

⁵McGill University

⁶Northern California Retina Vitreous Associates

October 20, 2024

ImageFlowNet: Forecasting Multiscale Image-Level Trajectories of Disease Progression with Irregularly-Sampled Longitudinal Medical Images

Chen Liu^{1*} Ke Xu^{1*} Liangbo L. Shen² Guillaume Huguet^{3,4} Zilong Wang^{3,5}
Alexander Tong^{3,4§} Danilo Bzdok^{3,5§} Jay Stewart^{2§} Jay C. Wang^{1,2,6§}
Lucian V. Del Priore^{1§} Smita Krishnaswamy^{1§}

¹Yale University ²University of California, San Francisco ³Mila - Quebec AI Institute
⁴Université de Montréal ⁵McGill University ⁶Northern California Retina Vitreous Associates

* These authors are joint first authors: {chen.liu.c12482, k.xu}@yale.edu. § Senior authors.
Please direct correspondence to: smita.krishnaswamy@yale.edu or lucian.delpriore@yale.edu.

Abstract

Advances in medical imaging technologies have enabled the collection of longitudinal images, which involve repeated scanning of the same patients over time, to monitor disease progression. However, predictive modeling of such data remains challenging due to high dimensionality, irregular sampling, and data sparsity. To address these issues, we propose ImageFlowNet, a novel model designed to forecast disease trajectories from initial images while preserving spatial details. ImageFlowNet first learns multiscale joint representation spaces across patients and time points, then optimizes deterministic or stochastic flow fields within these spaces using a position-parameterized neural ODE/SDE framework. The model leverages a UNet architecture to create robust multiscale representations and mitigates data scarcity by combining knowledge from all patients. We provide theoretical insights that support our formulation of ODEs, and motivate our regularizations involving high-level visual features, latent space organization, and trajectory smoothness. We validate ImageFlowNet on three longitudinal medical image datasets depicting progression in geographic atrophy, multiple sclerosis, and glioblastoma, demonstrating its ability to effectively forecast disease progression and outperform existing methods. Our contributions include the development of ImageFlowNet, its theoretical underpinnings, and empirical validation on real-world datasets. The official implementation is available at <https://github.com/KrishnaswamyLab/ImageFlowNet>.

1 Introduction

Advances in medical imaging technologies such as X-ray, computed tomography (CT), optical coherence tomography (OCT), and magnetic resonance imaging (MRI) combined with improved storage capacity and practices have enabled the collection of longitudinal medical images that track disease progression [1–3]. However, predictive modeling using such data is challenging due to the high dimensionality of images, irregular time intervals between samples, and the sparsity of data in most patients (see Appendix B for more background). These challenges often lead to methods that undermine the spatial-temporal nature of the data and instead treat them as time series of hand-crafted features, losing rich spatial information within the images [4–14] (see Fig. 1, top panel).

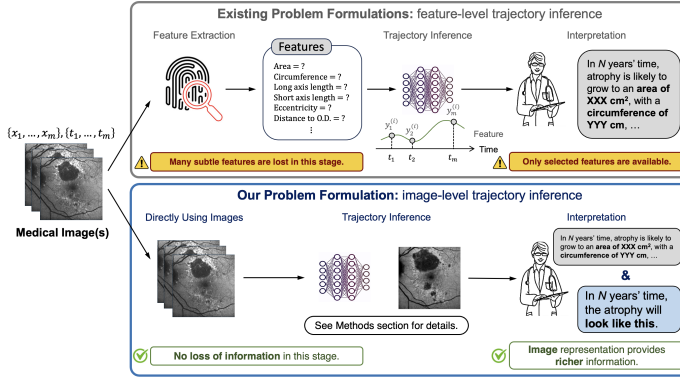


Figure 1: Advantages of image-level trajectory inference.

To address all these issues, we introduce ImageFlowNet, a model designed to forecast disease progression from initial images while addressing the aforementioned challenges and preserving spatial detail (see Fig. 1, bottom panel). ImageFlowNet learns multiscale joint representation spaces from all patients’ images and optimizes deterministic or stochastic flow fields over these spaces using a modified position-parameterized neural ODE/SDE framework.

ImageFlowNet learns multiscale representations with a UNet [15]

backbone (implementation adopted from [16]) while enforcing robustness to variations such as scaling, rotation, and contrast through extensive augmentation. Next, a vector field representing flows in each representation space is learned with a position-parameterized neural ODE framework. Unlike the standard parameterization by time $f_\theta(z_t, t)$, we parameterize the derivative by each vector’s position in the embedding space $f_\theta(z_t)$, which ensures that the space is shared across all patients at all times. Due to the nature of such joint embedding spaces, this approach mitigates patient-level data scarcity. We also present a stochastic alternative that generates non-deterministic trajectories.

In addition to presenting these networks, we theoretically establish the equivalent expressive power of the ODE and demonstrate the connections of ImageFlowNet to dynamic optimal transport. We then present empirical results on one longitudinal retinal image dataset with geographic atrophy and two longitudinal brain image datasets with multiple sclerosis and glioblastoma. Our main contributions are as follows.

1. Proposing ImageFlowNet to forecast trajectories of disease progression in the image domain.
2. Learning multiscale joint patient representation spaces that integrate the knowledge from all observed trajectories and remedy the data scarcity issue of any single patient.
3. Designing a multiscale position-parameterized ODE/SDE and providing theoretical rationales.
4. Showcasing results on three medical image datasets with sparse longitudinal progression data.

2 Preliminaries and Background

Problem Formulation and Notation We consider a set of N longitudinal image series $\{X^{(m)}\}_{m=1}^N$, where the m -th series $X^{(m)}$ contains $n_m \geq 2$ images and all images have the same dimension $\mathbb{R}^{H \times W \times C}$. These images $X^{(m)} = \{x_1^{(m)}, x_2^{(m)}, \dots, x_{n_m}^{(m)}\}$ are acquired at time points $T^{(m)} = \{t_1^{(m)}, t_2^{(m)}, \dots, t_{n_m}^{(m)}\}$ where $t_1^{(m)} < t_2^{(m)} < \dots < t_{n_m}^{(m)} \in \mathbb{R}$. Note that we do not assume any specific sampling schedule, such as uniform sampling over time. In addition, the time points for different series are not necessarily the same. This represents a very common scenario of a medical record containing N patients with multiple visits, where visit schedules can be irregular over time and heterogeneous among patients. For simplicity, we will omit the script m when considering the same series. Our task is to predict any image x_k with $1 < k \leq n$ given any subset of the earlier images $\tilde{X}_{<k}$ and their corresponding time points $\tilde{T}_{<k}$, where $\emptyset \subset \tilde{X}_{<k} \subseteq \{x_1, x_2, \dots, x_{k-1}\}$.

Neural Ordinary Differential Equations (Neural ODEs) Neural ODEs [17] model the evolution of a variable over time by considering the ODE in Eqn (1a), where f_θ is parameterized by a neural network. Since the gradient field f_θ is defined at every time point, future states can be modeled deterministically from an earlier state by integration, as shown in Eqn (1b).

$$\frac{dy(\tau)}{d\tau} = f_{\theta}(y(\tau), \tau) \quad (1a) \quad y(t_1) = y(t_0) + \int_{t_0}^{t_1} f_{\theta}(y(\tau), \tau) d\tau \quad (1b)$$

In practice, the integration step is performed by an ODE solver: $y(t_0) + \int_{t_0}^{t_1} f_{\theta}(y(\tau), \tau) d\tau = \text{ODESolve}(y(t_0), f_{\theta}, t_0, t_1)$. The gradient field f_{θ} can be optimized by any loss function $L(\cdot)$ that takes the result from the solver as input, in the form of $L(\text{ODESolve}(y(t_0), f_{\theta}, t_0, t_1), *args)$.

Neural Stochastic Differential Equations (Neural SDEs) Neural SDEs [18] inject stochasticity into deterministic neural ODEs by additionally considering Brownian motion $\{W_t\}_{t \geq 0}$ in the equation (Eqn (2)). f_{θ} and σ_{ϕ} respectively model the drift and diffusion components.

$$dy(\tau) = f_{\theta}(y(\tau), \tau) d\tau + \sigma_{\phi}(y(\tau), \tau) dW_{\tau} \quad (2)$$

UNet UNet [15] is a convolutional neural network architecture originally designed for biomedical image segmentation, but has later been found competent in many tasks such as image-to-image translation [19], style transfer [20], and image generation as in diffusion models [16]. It has a distinctive U-shaped structure with a contraction path that extracts multiscale features and an expansion path that recovers spatial resolution, along with skip connections that achieve residual learning [21].

3 ImageFlowNet

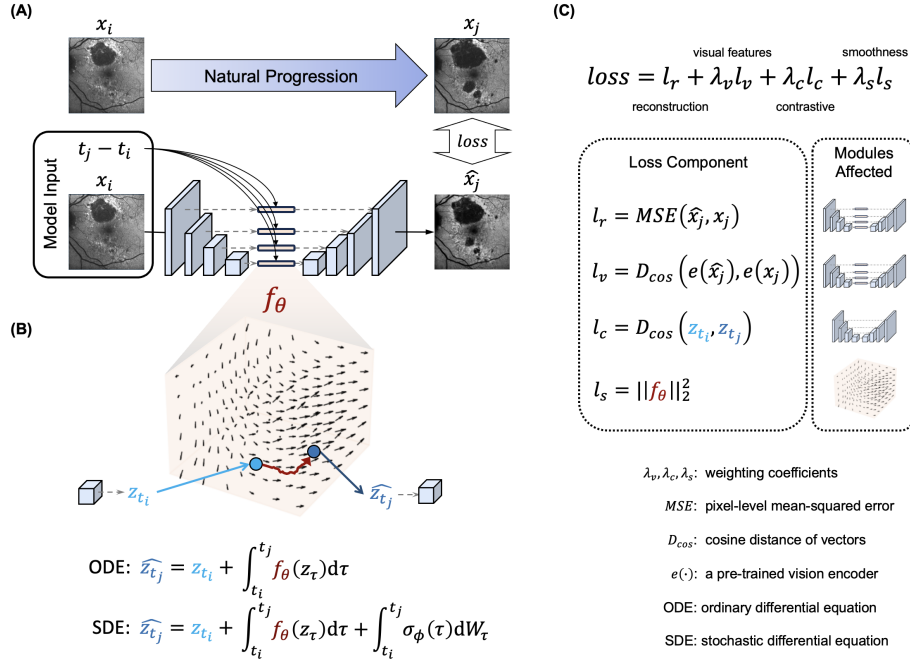


Figure 2: Overview of the proposed ImageFlowNet. (A) The model uses an earlier image x_i at time t_i as well as the change in time $t_j - t_i$ to forecast the future image x_j at time t_j . (B) For each hidden layer, a separate flow field f_{θ} is used to model the joint patient embedding space. Trajectory inference can be performed by integration along this flow field. It should be noted that the change in time $t_j - t_i$ is sufficient for integration in practice, while the exact time values t_i and t_j are included in the integral merely for mathematical clarity. (C) The learning objective has four components. The loss function and modules affected by each component are illustrated.

ImageFlowNet models the spatial-temporal dynamics of longitudinal images by first establishing a joint patient representation space (Section 3.1), and then flowing the representations of earlier time points to later time points (Section 3.2). This approach incorporates all images at all times during training, and thus addresses the issue of data scarcity at the individual patient level without

impairing the inference capabilities. From an engineering perspective, ImageFlowNet extracts latent representations at various resolutions and reassembles them to an image (i.e., the *spatial* aspect) after evolving these latent representations over time (i.e., the *temporal* aspect) along a learned flow field.

3.1 Learning Multiscale Spaces of Joint Patient Representations

As illustrated in Fig. 2 (A), we first learn joint embedding spaces for training samples in the hidden layers within the contraction path of a UNet. Multiscale representations are extracted from an input image x_i acquired at time t_i to produce B representations, one for each hidden layer, at R resolutions, which we denote $\{z_{t_i}^{(b)}\}_{b=1}^B$, with $B \geq R$.

During training, images are augmented with transformations that may naturally occur during acquisition, including reflection, rotation, shifting, rescaling, random brightness and contrast, and additive noise. Augmented versions effectively enlarge the sample variety and better populate the joint embedding spaces. This increases the chance that, during inference, a new image is embedded close to images seen in the training set and can leverage the learned dynamics around that local cohort.

3.2 Learning Multiscale Flow Fields on Joint Patient Representations

As depicted in Fig. 2 (B), at each hidden layer that spans different granularities of the image, a flow field is learned to evolve the joint patient representations at that scale. The flow field $f_\theta^{(b)}$ parameterizes the flow gradient as described in Eqn (3a), so that, given an initial position and a time duration, the trajectory can be computed through integration. $f_\theta^{(b)}$ is implemented as a 2-layer convolutional neural network whose input and output dimensions match the dimension of $z_{t_i}^{(b)}$. The latent representation $z_{t_j}^{(b)}$ corresponding to the future time t_j can be inferred using Eqn (3b), using Eqn (4) or similar variants. Finally, these multiscale representations meet at the expansion path to compose an output image. From now on, we will omit the superscript $\cdot^{(b)}$ if the specific hidden layer is not emphasized.

$$\frac{dz_\tau^{(b)}}{d\tau} = f_\theta^{(b)}(z_\tau^{(b)}) \quad \text{for } b \in [1, B] \quad (3a) \quad z_{t_j}^{(b)} = z_{t_i}^{(b)} + \int_{t_i}^{t_j} f_\theta^{(b)}(z_\tau^{(b)}) d\tau \quad \text{for } b \in [1, B] \quad (3b)$$

$$\begin{aligned} \hat{x}_j &= \text{ResBlock}(\text{Concat}(\tilde{z}_{t_j}^{(2)}, z_{t_j}^{(1)})), \text{ where} \\ \tilde{z}_{t_j}^{(b)} &= \text{Upsample}(\text{ResBlock}(\text{Concat}(\tilde{z}_{t_j}^{(b+1)}, z_{t_j}^{(b)}))) \text{ for } b \in [2, B-1], \text{ with } \tilde{z}_{t_j}^{(B)} = z_{t_j}^{(B)} \end{aligned} \quad (4)$$

3.3 Training and Inference

Training Objectives As in other neural differential frameworks, we use an ODE solver to compute the integral. Any loss function on the inferred latent representation z_{t_j} can be backpropagated through the ODE solver. Since the inferred image \hat{x}_j only depends on the expansion path of ImageFlowNet and the inferred latent representations $\{z_{t_i}^{(b)}\}_{b=1}^B$, the same principle applies to the loss functions on \hat{x}_j as well.

As shown in Eqn (5), our loss function contains four components, which we will explain below. During training, the learnable parameters in the entire ImageFlowNet are penalized by the first two components, while contrastive regularization only affects the UNet backbone and smoothness regularization only affects the flow field, as described in Fig. 2 (C).

$$\begin{aligned} \text{loss} &= \overbrace{\frac{1}{HWC} \sum_{h \in H} \sum_{w \in W} \sum_{c \in C} \|\hat{x}_j[h, w, c] - x_j[h, w, c]\|_2^2}^{\text{1 } l_r: \text{ reconstruction}} + \overbrace{\lambda_v \left(-\frac{e(\hat{x}_j)^\top e(x_j)}{\|e(\hat{x}_j)\|_2 \|e(x_j)\|_2} \right)}^{\text{2 } l_v: \text{ visual feature}} \\ &+ \overbrace{\lambda_c \left(-\frac{p_d(p_j(z_{t_i}))^\top p_j(z_{t_j})}{2\|p_d(p_j(z_{t_i}))\|_2 \|p_j(z_{t_j})\|_2} - \frac{p_d(p_j(z_{t_j}))^\top p_j(z_{t_i})}{2\|p_d(p_j(z_{t_j}))\|_2 \|p_j(z_{t_i})\|_2} \right)}^{\text{3 } l_c: \text{ contrastive learning (SimSiam)}} + \overbrace{\lambda_s \|f_\theta\|_2^2}^{\text{4 } l_s: \text{ trajectory smoothness}} \end{aligned} \quad (5)$$

- ① *Image reconstruction* is achieved by MSE, attending to low-level features on the pixel level.
- ② *Visual feature regularization* guides the network to produce images that resemble the ground truth on high-level features judged by a ConvNeXT [22] encoder pretrained on ImageNet [23].
- ③ *Contrastive learning regularization* organizes a well-structured ImageFlowNet latent space, by encouraging proximity of representations from images within the same longitudinal series, following the SimSiam formulation [24].
- ④ *Trajectory smoothness regularization* leverages a theorem in convex optimization (Lemma 2.2 in [25]) to enforce smoothness of trajectories by regularizing the norm of the field. *Notably, this achieves Lipschitz continuity, satisfying a crucial assumption for our theoretical results.*

Full Trajectories Are Used for Training Although training is performed over pairs of observations, since all samples within the longitudinal series are embedded in the same flow field f_θ and trajectories between all pairs are used for optimization, the model is indeed trained by the full trajectories.

Inferring Trajectories During inference, for a new patient with one or more previous observations, we can use any observation to serve as the starting point and obtain the future prediction using Eqn (3b). This approach has a low demand on access to patient history. In case we really want to take advantage of the entire patient history, we can perform test-time optimization to fine-tune f_θ using the measured time series and infer the patient trajectory afterwards (see Section 5.5).

ImageFlowNet_{SDE} as a Stochastic Variant We formulate a Langevin-type SDE as an alternative to our ODE, given by $dz_\tau^{(b)} = f_\theta^{(b)}(z_\tau^{(b)})d\tau + \sigma_\phi^{(b)}(\tau)dW_\tau$ for $b \in [1, B]$. This is motivated by stochasticity in patient trajectories and the need to model alternative outcomes from the same starting point, which is not possible in a deterministic ODE. In this construction, we decompose the dynamics into a deterministic drift function f_θ and a stochastic diffusion term σ_ϕ . This SDE is guaranteed to have a unique strong solution under mild assumptions [26].

An important characteristic unique to SDEs is their ability to model alternative trajectories due to the stochastic diffusion term. This feature enables SDEs to estimate uncertainty by generating multiple trajectories to infer the distribution of potential outcomes. By analyzing the variability and frequency of these outcomes, we can quantify the likelihood of various progression trends.

4 Theoretical Results

A Position-Parameterized ODE Unlike the original version, our ODE is parameterized by time *indirectly*: $f_\theta(z_t)$ instead of $f_\theta(z_t, t)$. This altered formulation learns a more compact vector field which alleviates the data scarcity issue, and allows the space to model disease states regardless of the time from disease onset. Importantly, we demonstrate that this formulation is **theoretically equivalent in expressiveness** to the original formulation (Proposition 4.1) and achieves **better empirical performance** (Table 3 in the Empirical Results section). The proof is shown in Appendix A.

Proposition 4.1. *Let f_θ be a continuous function that satisfies the Lipschitz continuity and linear growth conditions. Also, let the initial state $y(t_0) = y_0$ satisfy the finite second moment requirement ($\mathbb{E}[\|y(t_0)\|^2] < \infty$). Suppose $z(t_0)$ is the latent representation learned by ImageFlowNet in the initial state corresponding to t_0 . Then, our neural ODEs (Eqn (3a)) are at least as expressive as the original neural ODEs (Eqn (1a)), and their solutions capture the same dynamics.*

Connections to Dynamic Optimal Transport Dynamic Optimal Transport [27] aims to find the optimal transport plan to achieve the minimal transport cost between the original and target distribution. Here, we show that ImageFlowNet falls into the framework of dynamic optimal transport of images with ground distance based on the UNet representation. The proof is shown in Appendix A.

Proposition 4.2. *If we consider an image as a distribution over a 2D grid, ImageFlowNet is equivalently solving a dynamic optimal transport problem, as it meets three essential criteria: (1) matching the density, (2) smoothing the dynamics, and (3) minimizing the transport cost, where the ground distance is the Euclidean distance in the latent joint embedding space.*

5 Empirical Results

5.1 Preprocessing: Aligning Longitudinal Images with Image Registration

Longitudinal image datasets often face the problem of spatial misalignment among images acquired at different time points. This phenomenon is nearly inevitable, as minor adjustments in position or angle can disrupt the exact alignment between pixels. To address this problem, we spatially align all images in a longitudinal series during the preprocessing stage using keypoint detection [28] along with a perspective transform for retinal images, and affine registration [29] for brain scans. More details are described in Appendix D.

5.2 Baseline Methods

Image extrapolation methods is the most straightforward method for inferring future images. We included linear extrapolation [30] and cubic spline extrapolation [31] for comparison.

Time-conditional UNet (T-UNet) integrates time by adding a time-embedding tensor to representations throughout UNet and is a key component of diffusion models [16, 32]. The sinusoidal waveform is commonly used for time embedding, similar to the position encoding in transformers [33].

Time-aware diffusion model (T-Diffusion) is a modification to existing diffusion models by changing the diffusion time schedule. We introduce time awareness by considering the diffusion steps as equally spaced time intervals and dynamically adjusting the number of diffusion steps to match the time gap. Furthermore, our implementation is based on a specific diffusion model called image-to-image Schrödinger bridge (I2SB) [34], which directly maps from the input image to the output image without a noising and denoising process as in many others [16, 32, 35]. This allows it to produce high-quality images at any arbitrary diffusion time step, which is critical to our use case.

5.3 Datasets

Retinal Images We used longitudinal retinal images from the METforMIN study [36, 37] that monitored patients across 12 clinical centers with geographic atrophy, an advanced form of age-related macular degeneration that is slowly progressive and can lead to loss of vision. The dataset contains fundus autofluorescence images of 132 eyes over 2-5 visits at irregular intervals for a duration of up to 24 months.

Brain Multiple Sclerosis Images We used longitudinal FLAIR-weighted MRI scans from the LMSLS dataset [38] monitoring patients with Multiple Sclerosis (MS) over an average of 4.4 time points for approximately 5 years. We obtained 79 longitudinal image series from this dataset.

Brain Glioblastoma Images We used longitudinal contrastive T1-weighted MRI scans from the LUMIERE dataset [39] that tracked 91 glioblastoma (GBM) patients who underwent a pre-operative scan and repeated post-operative scans for up to 5 years. We obtained a set of 795 longitudinal image series each with 2-18 time points. Only post-operative images are kept in each series to model the natural change of tissues after surgery.

For all datasets, we took caution to avoid data leakage: data were partitioned on the level of longitudinal series into train/validation/test sets, and images from the same patient would only go to the same set. Disease regions were delineated by standalone image segmentation networks, one for each dataset. We quantified image similarity using peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM), residual map similarity using mean absolute error (MAE) and mean squared error (MSE), and disease region overlap using Dice-Sørensen coefficient (DSC) and Hausdorff distance (HD). Disease regions were delineated by 3 standalone image segmentation networks, one for each dataset. Mean values and standard deviations are reported from 3 independent runs with different random seeds. Training and implementation details, as well as information about the metrics, can be found in Appendices E and F.

5.4 ImageFlowNet Forecasts Images and Preserves Visual Traits of Disease Progression

As shown in Figure 3, the classic extrapolation methods and the time-conditional UNet often struggle to capture the changes in the disease-affected regions, implying that these methods are not quite

	Ground Truth x_j	Input x_i	Linear \hat{x}_j	Cubic Spline \hat{x}_j	T-UNet \hat{x}_j	T-Diffusion \hat{x}_j	ImageFlowNet _{ode} \hat{x}_j	ImageFlowNet _{ode++} \hat{x}_j	ImageFlowNet _{sde} \hat{x}_j	ImageFlowNet _{sde++} \hat{x}_j	
Retinal Geographic Atrophy	PSNR / SSIM		20.09 / 0.572	18.72 / 0.540	18.72 / 0.540	19.15 / 0.550	19.76 / 0.543	20.27 / 0.539	21.56 / 0.587	21.27 / 0.558	21.98 / 0.597
	MAE / MSE		0.156 / 0.039	0.177 / 0.054	0.177 / 0.054	0.168 / 0.049	0.130 / 0.026	0.137 / 0.029	0.131 / 0.028	0.139 / 0.030	0.125 / 0.025
	DSC / HD		0.810 / 7.28	0.848 / 7.00	0.848 / 7.00	0.850 / 6.94	0.841 / 6.35	0.902 / 5.39	0.876 / 6.71	0.884 / 5.83	0.925 / 6.32
Multiple Sclerosis	PSNR / SSIM		30.98 / 0.914	30.38 / 0.911	25.19 / 0.876	31.34 / 0.913	30.94 / 0.914	31.61 / 0.913	33.81 / 0.913	31.29 / 0.909	33.93 / 0.915
	MAE / MSE		0.024 / 0.003	0.028 / 0.003	0.057 / 0.012	0.025 / 0.003	0.025 / 0.003	0.023 / 0.003	0.019 / 0.002	0.025 / 0.003	0.019 / 0.002
	DSC / HD		0.803 / 8.00	0.494 / 28.86	0.311 / 35.36	0.823 / 8.00	0.803 / 8.00	0.846 / 7.00	0.844 / 3.16	0.813 / 6.00	0.854 / 3.16
Glioblastoma	PSNR / SSIM		30.98 / 0.914	29.20 / 0.865	32.60 / 0.899	31.34 / 0.909	29.72 / 0.889	31.88 / 0.914	33.80 / 0.909	32.81 / 0.911	31.37 / 0.903
	MAE / MSE		0.026 / 0.004	0.031 / 0.005	0.023 / 0.003	0.022 / 0.003	0.026 / 0.004	0.021 / 0.003	0.019 / 0.002	0.021 / 0.002	0.022 / 0.003
	DSC / HD		0.136 / 45.54	0.006 / 47.64	0.000 / 63.19	0.127 / 45.88	0.138 / 45.54	0.078 / 44.60	0.150 / 16.12	0.093 / 42.95	0.103 / 51.97

Figure 3: Qualitative comparison of image forecasting results on the retinal geographic atrophy, multiple sclerosis and glioblastoma datasets. “++” means using the 3 regularizations in Eqn (5).

capable of modeling the complex time evolution of the underlying disease processes. This is especially obvious when inspecting the residual maps. Extrapolation methods work significantly better on the glioblastoma dataset, where more historical data is available during inference. The diffusion model can represent atrophy growth in the retinal image, but not in the other cases. On the other hand, our proposed ImageFlowNet variants better represent disease-related changes.

As summarized in Table 1, our proposed methods achieve improved quantitative results, as demonstrated by higher image similarity, smaller residual maps, and better prediction of atrophy. The final ranking indicates that our SDE formulation using visual feature, contrastive learning, and trajectory smoothness regularizations is the best, while our ODE formulation with the three regularizations comes second, followed by the same two models penalized by the image reconstruction loss only.

Another notable phenomenon arises when we break down the retinal images into subsets. Compared to the other alternatives, our proposed methods show similar atrophy prediction performance (DSC, HD) for eyes with “minor atrophy growth”, but significantly better for eyes with “major atrophy growth”. Major/minor growth is defined by whether the ground truth masks differ by more than 0.1 in DSC. This implies that while the other methods may be on par with us in performing image reconstruction, our method is better at modeling the actual disease progression dynamics.

In Figure 4, we visualize the joint patient representations with and without contrastive learning regularization. The latent space of the ImageFlowNet bottleneck layer is visualized after projection into the 2D/3D PHATE space [41]. Indeed, the contrastive loss helped organize better structures in the latent space, as is evident in fewer global-range connections and smoother transitions over time.

Table 1: Image forecasting performance: metric(x_j, \hat{x}_j). $\hat{x}_j = \mathcal{F}(x_i, t_i, t_j), \forall i < j$.

†Extrapolation methods use the entire history. “+” means using the 3 regularizations in Eqn (5).

Dataset	Metric	Linear† [30]	Cubic Spline† [31]	T-UNet [40]	T-Diffusion [34]	ImageFlowNet _{ODE} (ours)	ImageFlowNet _{ODE++} (ours)	ImageFlowNet _{ODE} (ours)	ImageFlowNet _{ODE++} (ours)
Retinal Images <i>all cases</i>	PSNR ↑	20.22±0.00	19.79±0.00	22.06±0.33	22.29±0.33	22.63±0.26	22.74±0.25	22.32±0.29	22.89 ±0.28
	SSIM ↑	0.535±0.000	0.505±0.000	0.635±0.015	0.624±0.016	0.646±0.012	0.647±0.013	0.651 ±0.015	0.651 ±0.012
	MAE ↓	0.163±0.000	0.177±0.000	0.126±0.005	0.122±0.004	0.119±0.004	0.118±0.004	0.124±0.005	0.115 ±0.004
	MSE ↓	0.050±0.000	0.060±0.000	0.029±0.002	0.027±0.002	0.024±0.001	0.024±0.001	0.027±0.002	0.023 ±0.001
	DSC ↑	0.833±0.000	0.756±0.000	0.872±0.012	0.867±0.014	0.874±0.012	0.873±0.011	0.885 ±0.011	0.883 ±0.012
1	HD ↓	51.64±0.00	54.30±0.00	44.59±4.66	44.41±4.74	42.68 ±4.82	47.10±4.89	48.14±4.87	45.14±4.89
<i>minor atrophy</i>	PSNR ↑	21.36±0.00	21.08±0.00	22.56±0.55	22.99±0.55	23.23±0.34	23.44±0.33	23.28±0.36	23.63 ±0.43
	SSIM ↑	0.599±0.000	0.586±0.000	0.662±0.023	0.657±0.024	0.682±0.018	0.685±0.018	0.693 ±0.018	0.687 ±0.019
	MAE ↓	0.141±0.000	0.147±0.000	0.121±0.007	0.114±0.007	0.110±0.005	0.108±0.004	0.109±0.005	0.106 ±0.005
	MSE ↓	0.038±0.000	0.042±0.000	0.027±0.003	0.024±0.002	0.021±0.002	0.020 ±0.002	0.021±0.002	0.020 ±0.002
	DSC ↑	0.900±0.000	0.874±0.000	0.949 ±0.004	0.949 ±0.004	0.936±0.009	0.939±0.007	0.948±0.005	0.948±0.006
2	HD ↓	38.15±0.00	41.67±0.00	35.74±5.67	29.40 ±4.77	34.59±6.20	39.86±6.40	31.66 ±5.21	36.98±6.04
<i>major atrophy</i>	PSNR ↑	19.02±0.00	18.41±0.00	21.40±0.33	21.68±0.32	21.94±0.34	22.01±0.33	22.01±0.30	22.10 ±0.31
	SSIM ↑	0.468±0.000	0.420±0.000	0.607±0.017	0.588±0.017	0.607±0.014	0.606±0.014	0.607±0.014	0.613 ±0.013
	MAE ↓	0.186±0.000	0.210±0.000	0.135±0.006	0.131±0.006	0.129±0.006	0.129±0.006	0.128±0.005	0.126 ±0.005
	MSE ↓	0.063±0.000	0.080±0.000	0.032±0.003	0.030±0.002	0.028±0.002	0.028±0.002	0.027±0.002	0.027 ±0.002
	DSC ↑	0.762±0.000	0.631±0.000	0.784±0.016	0.779±0.019	0.807±0.014	0.803±0.012	0.817 ±0.016	0.814 ±0.017
3	HD ↓	65.97±0.00	67.73±0.00	61.43±7.26	60.36±7.37	51.28 ±7.13	54.79±7.19	65.65±7.17	53.81 ±7.49
Brain MS Images	PSNR ↑	30.07±0.00	29.56±0.00	31.55±0.20	31.57±0.23	32.01±0.19	32.34±0.20	32.40±0.20	32.41 ±0.20
	SSIM ↑	0.895±0.000	0.888±0.000	0.909±0.003	0.907±0.003	0.914±0.002	0.915 ±0.002	0.913±0.002	0.915 ±0.002
	MAE ↓	0.028±0.000	0.030±0.000	0.024±0.001	0.024±0.001	0.023±0.000	0.021 ±0.000	0.021±0.000	0.021 ±0.000
	MSE ↓	0.004±0.000	0.005±0.000	0.004±0.000	0.004±0.000	0.003 ±0.000	0.003 ±0.000	0.003 ±0.000	0.003 ±0.000
	DSC ↑	0.739±0.000	0.682±0.000	0.774±0.007	0.771±0.007	0.775±0.007	0.777 ±0.007	0.777±0.007	0.774±0.007
4	HD ↓	22.73±0.00	26.23±0.00	22.00±1.30	20.91 ±1.23	22.38±1.28	21.72±1.16	22.21±1.27	21.28 ±1.27
Brain GBM Images	PSNR ↑	35.32±0.00	33.60±0.00	35.73±0.13	35.49±0.17	35.86±0.12	35.90±0.14	35.77±0.12	35.79±0.15
	SSIM ↑	0.929±0.000	0.895±0.000	0.935±0.001	0.940±0.001	0.940±0.001	0.943 ±0.001	0.937±0.001	0.939±0.001
	MAE ↓	0.017±0.000	0.024±0.000	0.015±0.000	0.014 ±0.000	0.014 ±0.000	0.014 ±0.000	0.015±0.000	0.015±0.000
	MSE ↓	0.002±0.000	0.005±0.000	0.001±0.000	0.002±0.000	0.001 ±0.000	0.001 ±0.000	0.001 ±0.000	0.001 ±0.000
	DSC ↑	0.300±0.000	0.287±0.000	0.258±0.018	0.253±0.017	0.302 ±0.019	0.266±0.018	0.286±0.019	0.287±0.017
5	HD ↓	170.44±0.00	165.62±0.00	195.52±7.69	189.61±7.64	198.19±7.78	185.14±7.69	196.37±7.74	181.66±7.66
1, 4, 5	Rank ↓	6.3±1.6	7.3±2.0	4.9±1.4	4.6±1.9	2.9±1.9	2.3±1.6	3.4±2.0	2.1 ±1.3
1, 2, 3, 4, 5	Rank ↓	6.5±1.3	7.6±1.5	4.9±1.5	4.5±1.8	3.1±1.6	2.7±1.7	3.0±1.8	2.0 ±1.2




Figure 4: Joint representation space and the effect of contrastive learning regularization. Red dots are the observed disease states and arrows connect adjacent transitions. Normalized time is color coded. (A) Without regularization ($\lambda_c = 0$). (B) With contrastive learning regularization ($\lambda_c = 0.01$).

5.5 Test-Time Optimization Improves Prediction Leveraging the Entire Patient History

While it only requires a single observation to infer the trajectory of a new patient, we can further improve ImageFlowNet performance if we use the entire patient history to perform test-time optimization. More specifically, we could take the trained ImageFlowNet and fine-tune the flow field f_θ with the previous measurements $\{x_1, x_2, \dots, x_{n-1}\}$, and then predict x_n with the fine-tuned model.

We investigated the effect of test-time optimization using longitudinal series with at least 3 images from the retinal image dataset. The results for ImageFlowNet_{ODE} are summarized in Table 2, and similar trends are observed in other model variants. This indicates the possibility of trading computations for performance when the patient’s history is accessible.

Table 2: Effect of test-time optimization.

Iterations	Learning Rate	PSNR↑	SSIM↑	MAE↓	MSE↓	DSC↑	HD↓
N/A	N/A	22.31	0.643	0.123	0.027	0.827	51.07
1	10 ⁻⁴	22.52	0.646	0.120	0.025	0.829	48.97
1	10 ⁻⁵	22.36	0.643	0.122	0.027	0.827	51.02
1	10 ⁻⁶	22.31	0.643	0.123	0.027	0.827	51.07
10	10 ⁻⁴	20.63	0.605	0.157	0.042	0.749	64.79
10	10 ⁻⁵	22.59	0.646	0.119	0.025	0.829	49.92
10	10 ⁻⁶	22.36	0.644	0.122	0.027	0.827	51.01
100	10 ⁻⁴	19.63	0.571	0.177	0.056	0.726	70.12
100	10 ⁻⁵	20.92	0.614	0.152	0.040	0.759	58.76
100	10 ⁻⁶	22.61	0.646	0.119	0.025	0.829	49.74

5.6 Modeling Alternative Trajectories from the Same Starting Point with ImageFlowNet_{SDE}




Figure 5: ImageFlowNet_{SDE} alternative trajectories. Multiple inferences predict non-identical disease progressions, and their vectors in the joint representations space indeed follow different trajectories.

ries during training, which might be an interesting direction for future research.

5.7 Ablations Studies

We performed ablations on (1) time vs. position parameterization of ODE, (2) single-scale vs. multiscale ODE, and (3) effects of 3 regularizations.

Flow Field Formulation Previously, we decided on the $f_\theta(z_t)$ formulation analytically, and here we support our decision with empirical evidence (Table 3).

Single-scale vs Multiscale ODEs The UNet architecture uses hierarchical hidden layers to extract multiscale representations. Starting at the image resolution and ending at the bottleneck layer (bottom of the “U”), the model produces increasingly higher-level and more global representations. In this study, we analyze the advantages of multiscale ODEs. Moreover, there might be multiple hidden layers at the same resolution. On which representations should we perform trajectory inference?

To study this, we explored the following settings: (1) infer a single-scale z_t from the bottleneck layer, (2) infer multiscale $\{z_t\}$ at all layers and use distinct f_θ for each resolution, but all hidden layers of the same resolution share the same f_θ , and (3) infer multiscale $\{z_t\}$ at all layers and use distinct f_θ for each hidden layer. The empirical results as shown in Table 4 indicate that modeling all representations separately would lead to the best performance.

Note: To avoid confusion, all of these hidden layers produce outputs that are bridged by skip connections from the contraction path to the expansion path.

Effects of Regularizations We ablated visual feature regularization (Table 5), contrastive learning regularization (Table 6), and trajectory smoothness regularization (Table 7).

Table 3: Flow field formulation.

	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	MSE \downarrow	DSC \uparrow	HD \downarrow
$f_\theta(z_t, t)$	22.42	0.643	0.123	0.027	0.872	48.38
$f_\theta(z_t)$	22.63	0.646	0.119	0.024	0.874	42.68

Table 4: Latent representation.

	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	MSE \downarrow	DSC \uparrow	HD \downarrow
bottleneck only	22.33	0.639	0.122	0.026	0.850	48.13
all unique resolutions	22.49	0.643	0.122	0.025	0.859	43.39
all unique layers	22.63	0.646	0.119	0.024	0.874	42.68

Table 5: Visual feature regularization.

λ_v	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	MSE \downarrow	DSC \uparrow	HD \downarrow
0	22.63	0.646	0.119	0.024	0.874	42.68
0.001	22.65	0.658	0.118	0.024	0.872	44.27
0.01	22.64	0.650	0.120	0.025	0.872	45.89
0.1	22.57	0.647	0.120	0.025	0.869	50.69
1	22.54	0.634	0.124	0.027	0.867	48.13

Table 6: Contrastive regularization.

λ_c	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	MSE \downarrow	DSC \uparrow	HD \downarrow
0	22.63	0.646	0.119	0.024	0.874	42.68
0.001	22.63	0.646	0.119	0.025	0.872	46.23
0.01	22.65	0.652	0.118	0.024	0.875	42.18
0.1	22.38	0.651	0.121	0.025	0.871	45.30
1	22.25	0.644	0.121	0.025	0.868	46.85

Table 7: Smoothness regularization.

λ_s	PSNR \uparrow	SSIM \uparrow	MAE \downarrow	MSE \downarrow	DSC \uparrow	HD \downarrow
0	22.63	0.646	0.119	0.024	0.874	42.68
0.001	22.38	0.649	0.123	0.027	0.870	46.91
0.01	22.65	0.648	0.119	0.024	0.870	45.71
0.1	22.70	0.657	0.118	0.024	0.878	47.44
1	22.69	0.655	0.118	0.024	0.875	45.16

6 Related Works

Disease Progression Modeling in Longitudinal Medical Images Most existing methods for modeling disease progression operate in the vector space of hand-selected features. The event-based model (EBM) [4, 5] and discriminative EBM [6] use bivariate mixture modeling and univariate Gaussian modeling, respectively. Both methods analyze disease progression at the group level and do not predict the outcome for individuals, the need of which led to sequence modeling techniques. Liu et al. [10] used XGBoost [42] to model the progression of breast cancer. Lei et al. [11] used a polynomial network to describe selected image statistics of patients with Alzheimer’s disease. LSTM [43] and Transformer [33] have also been recruited to predict the progression of diseases such as Alzheimer’s [14, 13] and COVID-19 [12] in longitudinal medical images.

Disease Progression Modeling in the Image Space Due to the challenges discussed in the Introduction, very few works tackle the disease progression problem in the image space. STLSTM [44] and LDDMM [45] use LSTMs to model time evolution, which is a sequential model that does not handle irregular sampling over time — a limitation commonly seen in video prediction models as well [46–49]. ManifoldExtrap [50] projects images to a StyleGAN latent space and performs a linear walk whose direction and distance are determined by the nearest neighbor found in this latent space. This method only models the transition between the baseline visit and the follow-up visit and does not model the continuous-time evolution. Upgrading the latent space navigation (e.g. using ODEs to model continuous time) might be a good solution, which we leave to future investigations. Lachinov and his collaborators present an approach [51, 52] most similar to ours, as they also aim to model disease progression at the image level. However, their method is limited to segmentation output and does not explore multiscale trajectories, our innovative differential equation formulations, or our comprehensive suite of regularizations. These key elements distinguish our work from theirs.

Neural ODEs for Disease Progression Modeling To better accommodate irregular sampling, a common situation in longitudinal medical images, researchers later investigated differential equation models [53], especially neural ODEs [17]. Neural ODEs have been successfully applied to predict the dynamics of individual patients with Alzheimer’s [54] and COVID-19 [55], but exclusively on biomarkers and/or attributes extracted from images rather than on the medical images themselves.

Applying Neural ODEs to Images At the birth of neural ODEs [17], the possibility of applying them to images was discussed, but only as a drop-in replacement for residual blocks in a ResNet model, such as for image classification [56]. UNode [57] and follow-ups [58, 59] adapted it to work on an image-to-image task for the first time on image segmentation, but they treated the ODE as additional trainable parameters beyond convolution blocks and *did not exploit its ability to model time*. Among similar endeavors, we are the first to use neural ODEs in the natural and intuitive manner by actually modeling how latent representations evolve over time in spatial-temporal data.

7 Conclusion

We introduced ImageFlowNet, a deep learning framework which uses joint representation spaces and multiscale position-parameterized differential equations to infer trajectories in irregularly-sampled longitudinal images. We provided theoretical evidence to support its soundness and demonstrated its empirical effectiveness across three longitudinal medical image datasets. We believe that our method offers a promising approach to image-level trajectory analysis that can model progression in medical image datasets, a relatively underexplored yet highly promising area of research.

8 Limitations and Broader Impacts

Limitations We have not yet studied the capabilities of using ImageFlowNet in a clinical context for patient diagnosis, which we plan to cover in a follow-up study.

Broader Impacts Our work can help us understand spatial-temporal systems, state transitions in longitudinal images, and in particular disease progression in longitudinal medical images. To the best of our knowledge, our work has no negative societal impact.

Acknowledgements

This work was supported in part by the National Science Foundation (NSF Career Grant 2047856) and the National Institute of Health (NIH 1R01GM130847-01A1, NIH 1R01GM135929-01).

References

- [1] Eman Nabrawi, Abdullah T Alanazi, and Eman Al Alkhaibari. Imaging in healthcare: A glance at the present and a glimpse into the future. *Cureus*, 15(3), 2023.
- [2] Karl Ricanek and Tamirat Tesafaye. Morph: A longitudinal image database of normal adult age-progression. In *7th international conference on automatic face and gesture recognition (FGRO6)*, pages 341–345. IEEE, 2006.
- [3] 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*, pages 2019–12, 2019.
- [4] Hubert M Fonteijn, Matthew J Clarkson, Marc Modat, Josephine Barnes, Manja Lehmann, Sebastien Ourselin, Nick C Fox, and Daniel C Alexander. An event-based disease progression model and its application to familial alzheimer’s disease. In *Information Processing in Medical Imaging: 22nd International Conference, IPMI 2011, Kloster Irsee, Germany, July 3-8, 2011. Proceedings 22*, pages 748–759. Springer, 2011.
- [5] Hubert M Fonteijn, Marc Modat, Matthew J Clarkson, Josephine Barnes, Manja Lehmann, Nicola Z Hobbs, Rachael I Scahill, Sarah J Tabrizi, Sebastien Ourselin, Nick C Fox, et al. An event-based model for disease progression and its application in familial alzheimer’s disease and huntington’s disease. *NeuroImage*, 60(3):1880–1889, 2012.
- [6] Vikram Venkatraghavan, Esther E Bron, Wiro J Niessen, and Stefan Klein. A discriminative event based model for alzheimer’s disease progression modeling. In *Information Processing in Medical Imaging: 25th International Conference, IPMI 2017, Boone, NC, USA, June 25-30, 2017, Proceedings 25*, pages 121–133. Springer, 2017.
- [7] Jiecheng Lu, Xu Han, Yan Sun, and Shihao Yang. Cats: Enhancing multivariate time series forecasting by constructing auxiliary time series as exogenous variables. *arXiv preprint arXiv:2403.01673*, 2024.
- [8] Ziyou Guo, Yan Sun, and Tieru Wu. Weits: A wavelet-enhanced residual framework for interpretable time series forecasting. *arXiv preprint arXiv:2405.10877*, 2024.
- [9] Yan Sun and Shihao Yang. Manifold-constrained gaussian process inference for time-varying parameters in dynamic systems. *Statistics and Computing*, 33(6):142, 2023.
- [10] Pei Liu, Bo Fu, Simon X Yang, Ling Deng, Xiaorong Zhong, and Hong Zheng. Optimizing survival analysis of xgboost for ties to predict disease progression of breast cancer. *IEEE Transactions on Biomedical Engineering*, 68(1):148–160, 2020.
- [11] Baiying Lei, Mengya Yang, Peng Yang, Feng Zhou, Wen Hou, Wenbin Zou, Xia Li, Tianfu Wang, Xiaohua Xiao, and Shuqiang Wang. Deep and joint learning of longitudinal data for alzheimer’s disease prediction. *Pattern Recognition*, 102:107247, 2020.
- [12] Jamil Ahmad, Abdul Khader Jilani Saudagar, Khalid Mahmood Malik, Waseem Ahmad, Muhammad Badruddin Khan, Mozaherul Hoque Abul Hasanat, Abdullah AlTameem, Mohammed AlKhathami, and Muhammad Sajjad. Disease progression detection via deep sequence learning of successive radiographic scans. *International journal of environmental research and public health*, 19(1):480, 2022.
- [13] Huy Hoang Nguyen, Matthew B Blaschko, Simo Saarakkala, and Aleksei Tiulpin. Clinically-inspired multi-agent transformers for disease trajectory forecasting from multimodal data. *IEEE transactions on medical imaging*, 2023.

- [14] Anza Aqeel, Ali Hassan, Muhammad Attique Khan, Saad Rehman, Usman Tariq, Seifedine Kadry, Arnab Majumdar, and Orawit Thinnukool. A long short-term memory biomarker-based prediction framework for alzheimer’s disease. *Sensors*, 22(4):1475, 2022.
- [15] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *Medical Image Computing and Computer-Assisted Intervention–MICCAI 2015: 18th International Conference, Munich, Germany, October 5-9, 2015, Proceedings, Part III 18*, pages 234–241. Springer, 2015.
- [16] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33:6840–6851, 2020.
- [17] Ricky TQ Chen, Yulia Rubanova, Jesse Bettencourt, and David K Duvenaud. Neural ordinary differential equations. *Advances in neural information processing systems*, 31, 2018.
- [18] Patrick Kidger, James Foster, Xuechen Li, and Terry J Lyons. Neural sdes as infinite-dimensional gans. In *International conference on machine learning*, pages 5453–5463. PMLR, 2021.
- [19] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [20] Lvmin Zhang, Yi Ji, Xin Lin, and Chunping Liu. Style transfer for anime sketches with enhanced residual u-net and auxiliary classifier gan. In *2017 4th IAPR Asian conference on pattern recognition (ACPR)*, pages 506–511. IEEE, 2017.
- [21] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [22] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 11976–11986, 2022.
- [23] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE conference on computer vision and pattern recognition*, pages 248–255. Ieee, 2009.
- [24] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 15750–15758, 2021.
- [25] Stephen J Wright and Benjamin Recht. *Optimization for data analysis*. Cambridge University Press, 2022.
- [26] YongKyung Oh, Dongyoung Lim, and Sungil Kim. Stable neural stochastic differential equations in analyzing irregular time series data. *arXiv preprint arXiv:2402.14989*, 2024.
- [27] Jean-David Benamou and Yann Brenier. A computational fluid mechanics solution to the monge-kantorovich mass transfer problem. *Numerische Mathematik*, 84(3):375–393, 2000.
- [28] Jiazhen Liu, Xirong Li, Qijie Wei, Jie Xu, and Dayong Ding. Semi-supervised keypoint detector and descriptor for retinal image matching. In *European Conference on Computer Vision*, pages 593–609. Springer, 2022.
- [29] Brian B Avants, Nick Tustison, Gang Song, et al. Advanced normalization tools (ants). *Insight j*, 2(365):1–35, 2009.
- [30] Thierry Blu, Philippe Thévenaz, and Michael Unser. Linear interpolation revitalized. *IEEE Transactions on Image Processing*, 13(5):710–719, 2004.
- [31] Sky McKinley and Megan Levine. Cubic spline interpolation. *College of the Redwoods*, 45(1):1049–1060, 1998.

- [32] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint arXiv:2010.02502*, 2020.
- [33] 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.
- [34] Guan-Horng Liu, Arash Vahdat, De-An Huang, Evangelos A Theodorou, Weili Nie, and Anima Anandkumar. I²sb: Image-to-image schrödinger bridge. *arXiv preprint arXiv:2302.05872*, 2023.
- [35] Bo Li, Kaitao Xue, Bin Liu, and Yu-Kun Lai. Bbdm: Image-to-image translation with brownian bridge diffusion models. In *Proceedings of the IEEE/CVF conference on computer vision and pattern Recognition*, pages 1952–1961, 2023.
- [36] Liangbo Linus Shen, Jeremy D Keenan, Noor Chahal, Abu Tahir Taha, Jasmeet Saroya, Chu Jian Ma, Mengyuan Sun, Daphne Yang, Catherine Psaras, Jacquelyn Callander, et al. Metformin for the minimization of geographic atrophy progression (metformin): A randomized trial. *Ophthalmology Science*, 4(3):100440, 2024.
- [37] Chen Liu, Matthew Amodio, Liangbo L. Shen, Feng Gao, Arman Avesta, Sanjay Aneja, Jay C. Wang, Lucian V. Del Priore, and Smita Krishnaswamy. CUTS: A Deep Learning and Topological Framework for Multigranular Unsupervised Medical Image Segmentation. In *proceedings of Medical Image Computing and Computer Assisted Intervention – MICCAI 2024*, volume LNCS 15008. Springer Nature Switzerland, October 2024.
- [38] Aaron Carass, Snehashis Roy, Amod Jog, Jennifer L Cuzzocreo, Elizabeth Magrath, Adrian Gherman, Julia Button, James Nguyen, Ferran Prados, Carole H Sudre, et al. Longitudinal multiple sclerosis lesion segmentation: resource and challenge. *NeuroImage*, 148:77–102, 2017.
- [39] Yannick Suter, Urspeter Knecht, Waldo Valenzuela, Michelle Notter, Ekkehard Hewer, Philippe Schucht, Roland Wiest, and Mauricio Reyes. The lumiere dataset: Longitudinal glioblastoma mri with expert rano evaluation. *Scientific data*, 9(1):768, 2022.
- [40] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in neural information processing systems*, 34:8780–8794, 2021.
- [41] Kevin R Moon, David Van Dijk, Zheng Wang, Scott Gigante, Daniel B Burkhardt, William S Chen, Kristina Yim, Antonia van den Elzen, Matthew J Hirn, Ronald R Coifman, et al. Visualizing structure and transitions in high-dimensional biological data. *Nature biotechnology*, 37(12):1482–1492, 2019.
- [42] Tianqi Chen and Carlos Guestrin. Xgboost: A scalable tree boosting system. In *Proceedings of the 22nd acm sigkdd international conference on knowledge discovery and data mining*, pages 785–794, 2016.
- [43] Alex Graves and Alex Graves. Long short-term memory. *Supervised sequence labelling with recurrent neural networks*, pages 37–45, 2012.
- [44] Ling Zhang, Le Lu, Xiaosong Wang, Robert M Zhu, Mohammadhadi Bagheri, Ronald M Summers, and Jianhua Yao. Spatio-temporal convolutional lstms for tumor growth prediction by learning 4d longitudinal patient data. *IEEE transactions on medical imaging*, 39(4):1114–1126, 2019.
- [45] Sharmin Pathan and Yi Hong. Predictive image regression for longitudinal studies with missing data. *arXiv preprint arXiv:1808.07553*, 2018.
- [46] Nal Kalchbrenner, Aäron Oord, Karen Simonyan, Ivo Danihelka, Oriol Vinyals, Alex Graves, and Koray Kavukcuoglu. Video pixel networks. In *International Conference on Machine Learning*, pages 1771–1779. PMLR, 2017.

- [47] Yunbo Wang, Mingsheng Long, Jianmin Wang, Zhifeng Gao, and Philip S Yu. Predrnn: Recurrent neural networks for predictive learning using spatiotemporal lstms. *Advances in neural information processing systems*, 30, 2017.
- [48] Ruben Villegas, Jimei Yang, Seunghoon Hong, Xunyu Lin, and Honglak Lee. Decomposing motion and content for natural video sequence prediction. *arXiv preprint arXiv:1706.08033*, 2017.
- [49] Omer Bar-Tal, Hila Chefer, Omer Tov, Charles Herrmann, Roni Paiss, Shiran Zada, Ariel Ephrat, Junhwa Hur, Yuanzhen Li, Tomer Michaeli, et al. Lumiere: A space-time diffusion model for video generation. *arXiv preprint arXiv:2401.12945*, 2024.
- [50] Tianyu Han, Jakob Nikolas Kather, Federico Pedersoli, Markus Zimmermann, Sebastian Keil, Maximilian Schulze-Hagen, Marc Terwoelbeck, Peter Isfort, Christoph Haarburger, Fabian Kiessling, et al. Image prediction of disease progression for osteoarthritis by style-based manifold extrapolation. *Nature Machine Intelligence*, 4(11):1029–1039, 2022.
- [51] Dmitrii Lachinov, Arunava Chakravarty, Christoph Grechenig, Ursula Schmidt-Erfurth, and Hrvoje Bogunović. Learning spatio-temporal model of disease progression with neuralodes from longitudinal volumetric data. *IEEE Transactions on Medical Imaging*, 2023.
- [52] Julia Mai, Dmitrii Lachinov, Gregor S Reiter, Sophie Riedl, Christoph Grechenig, Hrvoje Bogunovic, and Ursula Schmidt-Erfurth. Deep learning-based prediction of individual geographic atrophy progression from a single baseline oct. *Ophthalmology Science*, 4(4):100466, 2024.
- [53] Neil P Oxtoby, Alexandra L Young, David M Cash, Tammie LS Benzinger, Anne M Fagan, John C Morris, Randall J Bateman, Nick C Fox, Jonathan M Schott, and Daniel C Alexander. Data-driven models of dominantly-inherited alzheimer’s disease progression. *Brain*, 141(5):1529–1544, 2018.
- [54] Matías Nicolás Bossa and Hichem Sahli. A multidimensional ode-based model of alzheimer’s disease progression. *Scientific Reports*, 13(1):3162, 2023.
- [55] Ting Dang, Jing Han, Tong Xia, Erika Bondareva, Chloë Siegele-Brown, Jagmohan Chauhan, Andreas Grammenos, Dimitris Spathis, Pietro Cicuta, and Cecilia Mascolo. Conditional neural ode processes for individual disease progression forecasting: a case study on covid-19. In *Proceedings of the 29th ACM SIGKDD Conference on Knowledge Discovery and Data Mining*, pages 3914–3925, 2023.
- [56] Fabio Carrara, Roberto Caldelli, Fabrizio Falchi, and Giuseppe Amato. On the robustness to adversarial examples of neural ode image classifiers. In *2019 IEEE International Workshop on Information Forensics and Security (WIFS)*, pages 1–6. IEEE, 2019.
- [57] Hans Pinckaers and Geert Litjens. Neural ordinary differential equations for semantic segmentation of individual colon glands. *arXiv preprint arXiv:1910.10470*, 2019.
- [58] Jintao Ru, Beichen Lu, Buran Chen, Jialin Shi, Gaoxiang Chen, Meihao Wang, Zhifang Pan, Yezhi Lin, Zhihong Gao, Jiejie Zhou, et al. Attention guided neural ode network for breast tumor segmentation in medical images. *Computers in Biology and Medicine*, 159:106884, 2023.
- [59] Shubin Wang, Yuanyuan Chen, and Zhang Yi. nmode-unet: A novel network for semantic segmentation of medical images. *Applied Sciences*, 14(1):411, 2024.
- [60] David Snider R. Kent Nagle, Edward B. Saff. *Fundamentals of Differential Equations and Boundary Value Problems*. Pearson Education, 2011.
- [61] Alexander Tong, Jessie Huang, Guy Wolf, David Van Dijk, and Smita Krishnaswamy. Trajectorynet: A dynamic optimal transport network for modeling cellular dynamics. In *International conference on machine learning*, pages 9526–9536. PMLR, 2020.

- [62] Guillaume Huguët, Daniel Sumner Magruder, Alexander Tong, Oluwadamilola Fasina, Manik Kuchroo, Guy Wolf, and Smita Krishnaswamy. Manifold interpolating optimal-transport flows for trajectory inference. *Advances in neural information processing systems*, 35:29705–29718, 2022.
- [63] Alexander Tong, Nikolay Malkin, Kilian Fatras, Lazar Atanackovic, Yanlei Zhang, Guillaume Huguët, Guy Wolf, and Yoshua Bengio. Simulation-free schrödinger bridges via score and flow matching. *arXiv preprint arXiv:2307.03672*, 2023.
- [64] Trang Nguyen, Alexander Tong, Kanika Madan, Yoshua Bengio, and Dianbo Liu. Causal discovery in gene regulatory networks with gflownet: Towards scalability in large systems. In *NeurIPS 2023 Generative AI and Biology (GenBio) Workshop*, 2023.
- [65] María Ramos Zapatero, Alexander Tong, James W Opzooomer, Rhianna O’Sullivan, Ferran Cardoso Rodriguez, Jahangir Sufi, Petra Vlckova, Callum Nattress, Xiao Qin, Jeroen Claus, et al. Trellis tree-based analysis reveals stromal regulation of patient-derived organoid drug responses. *Cell*, 186(25):5606–5619, 2023.
- [66] Lazar Atanackovic, Alexander Tong, Bo Wang, Leo J Lee, Yoshua Bengio, and Jason S Hartford. Dyngfn: Towards bayesian inference of gene regulatory networks with gflownets. *Advances in Neural Information Processing Systems*, 36, 2024.
- [67] Jialin Chen, Jan Eric Lenssen, Aosong Feng, Weihua Hu, Matthias Fey, Leandros Tassioulas, Jure Leskovec, and Rex Ying. From similarity to superiority: Channel clustering for time series forecasting. *arXiv preprint arXiv:2404.01340*, 2024.
- [68] Tingsong Xiao, Zelin Xu, Wenchong He, Jim Su, Yupu Zhang, Raymond Opoku, Ronald Ison, Jason Petho, Jiang Bian, Patrick Tighe, et al. Xtsformer: cross-temporal-scale transformer for irregular time event prediction. *arXiv preprint arXiv:2402.02258*, 2024.
- [69] Omer Bar-Tal, Hila Chefer, Omer Tov, Charles Herrmann, Roni Paiss, Shiran Zada, Ariel Ephrat, Junhwa Hur, Yuanzhen Li, Tomer Michaeli, et al. Lumiere: A space-time diffusion model for video generation. *arXiv preprint arXiv:2401.12945*, 2024.
- [70] Jonathan Ho, William Chan, Chitwan Saharia, Jay Whang, Ruiqi Gao, Alexey Gritsenko, Diederik P Kingma, Ben Poole, Mohammad Norouzi, David J Fleet, et al. Imagen video: High definition video generation with diffusion models. *arXiv preprint arXiv:2210.02303*, 2022.
- [71] Tobias Höppe, Arash Mehrjou, Stefan Bauer, Didrik Nielsen, and Andrea Dittadi. Diffusion models for video prediction and infilling. *arXiv preprint arXiv:2206.07696*, 2022.
- [72] Zhen Xing, Qijun Feng, Haoran Chen, Qi Dai, Han Hu, Hang Xu, Zuxuan Wu, and Yu-Gang Jiang. A survey on video diffusion models. *arXiv preprint arXiv:2310.10647*, 2023.
- [73] Maximilian Pfau, Moritz Lindner, Lukas Goerdt, Sarah Thiele, Jennifer Nadal, Matthias Schmid, Steffen Schmitz-Valckenberg, Srinivas R Sadda, Frank G Holz, Monika Fleckenstein, et al. Prognostic value of shape-descriptive factors for the progression of geographic atrophy secondary to age-related macular degeneration. *Retina*, 39(8):1527–1540, 2019.
- [74] Liangbo L Shen, Mengyuan Sun, Aneesha Ahluwalia, Benjamin K Young, Michael M Park, and Lucian V Del Priore. Geographic atrophy growth is strongly related to lesion perimeter: unifying effects of lesion area, number, and circularity on growth. *Ophthalmology Retina*, 5(9):868–878, 2021.
- [75] Liangbo L Shen, Mengyuan Sun, Aneesha Ahluwalia, Michael M Park, Benjamin K Young, and Lucian V Del Priore. Local progression kinetics of geographic atrophy depends upon the border location. *Investigative Ophthalmology & Visual Science*, 62(13):28–28, 2021.
- [76] Shah Hussain, Iqra Mubeen, Niamat Ullah, Syed Shahab Ud Din Shah, Bakhtawar Abduljalil Khan, Muhammad Zahoor, Riaz Ullah, Farhat Ali Khan, and Mujeeb A Sultan. Modern diagnostic imaging technique applications and risk factors in the medical field: a review. *BioMed research international*, 2022(1):5164970, 2022.

- [77] Chuqin Huang, Yanda Cheng, Wenhan Zheng, Robert W Bing, Huijuan Zhang, Isabel Komornicki, Linda M Harris, Praveen R Arany, Saptarshi Chakraborty, Qifa Zhou, et al. Dual-scan photoacoustic tomography for the imaging of vascular structure on foot. *IEEE Transactions on Ultrasonics, Ferroelectrics, and Frequency Control*, 2023.
- [78] Andre Esteva, Brett Kuprel, Roberto A Novoa, Justin Ko, Susan M Swetter, Helen M Blau, and Sebastian Thrun. Dermatologist-level classification of skin cancer with deep neural networks. *nature*, 542(7639):115–118, 2017.
- [79] Yuanzhou Wei, Dan Zhang, Meiyao Gao, Yuanhao Tian, Ya He, Bolin Huang, and Changyang Zheng. Breast cancer prediction based on machine learning. *Journal of Software Engineering and Applications*, 16(8):348–360, 2023.
- [80] Yuanzhou Wei, Dan Zhang, Meiyao Gao, Aliya Mulati, Changyang Zheng, and Bolin Huang. Skin cancer detection based on machine learning. *Journal of Knowledge Learning and Science Technology ISSN: 2959-6386 (online)*, 3(2):72–86, 2024.
- [81] Xinyu Dong, Rachel Wong, Weimin Lyu, Kayley Abell-Hart, Jianyuan Deng, Yinan Liu, Janos G Hajagos, Richard N Rosenthal, Chao Chen, and Fusheng Wang. An integrated lstm-heterorgnn model for interpretable opioid overdose risk prediction. *Artificial intelligence in medicine*, 135:102439, 2023.
- [82] Davide Placido, Bo Yuan, Jessica X Hjaltelin, Chunlei Zheng, Amalie D Haue, Piotr J Chmura, Chen Yuan, Jihye Kim, Renato Umeton, Gregory Antell, et al. A deep learning algorithm to predict risk of pancreatic cancer from disease trajectories. *Nature medicine*, 29(5):1113–1122, 2023.
- [83] Moloud Abdar, Farhad Pourpanah, Sadiq Hussain, Dana Rezazadegan, Li Liu, Mohammad Ghavamzadeh, Paul Fieguth, Xiaochun Cao, Abbas Khosravi, U Rajendra Acharya, et al. A review of uncertainty quantification in deep learning: Techniques, applications and challenges. *Information fusion*, 76:243–297, 2021.
- [84] Danqi Liao, Chen Liu, Benjamin W Christensen, Alexander Tong, Guillaume Huguet, Guy Wolf, Maximilian Nickel, Ian Adelstein, and Smita Krishnaswamy. Assessing neural network representations during training using noise-resilient diffusion spectral entropy. In *2024 58th Annual Conference on Information Sciences and Systems (CISS)*, pages 1–6. IEEE, 2024.
- [85] Karin T Kirchhoff, Bernard J Hammes, Karen A Kehl, Linda A Briggs, and Roger L Brown. Effect of a disease-specific planning intervention on surrogate understanding of patient goals for future medical treatment. *Journal of the American Geriatrics Society*, 58(7):1233–1240, 2010.
- [86] Jingdi Chen, Tian Lan, and Carlee Joe-Wong. Rgmcomm: Return gap minimization via discrete communications in multi-agent reinforcement learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 38, pages 17327–17336, 2024.
- [87] Caitao Zhan, Himanshu Gupta, and Mark Hillery. Optimizing initial state of detector sensors in quantum sensor networks. *ACM Transactions on Quantum Computing*, 5(2):1–25, 2024.
- [88] Jingdi Chen and Tian Lan. Minimizing return gaps with discrete communications in decentralized pomdp. *arXiv preprint arXiv:2308.03358*, 2023.
- [89] Cheng Jin, Heng Yu, Jia Ke, Peirong Ding, Yongju Yi, Xiaofeng Jiang, Xin Duan, Jinghua Tang, Daniel T Chang, Xiaojian Wu, et al. Predicting treatment response from longitudinal images using multi-task deep learning. *Nature communications*, 12(1):1851, 2021.
- [90] Zongxing Xie, Hanrui Wang, Song Han, Elinor Schoenfeld, and Fan Ye. Deepvs: A deep learning approach for rf-based vital signs sensing. In *Proceedings of the 13th ACM international conference on bioinformatics, computational biology and health informatics*, pages 1–5, 2022.
- [91] Zongxing Xie, Bing Zhou, Xi Cheng, Elinor Schoenfeld, and Fan Ye. Vitalhub: Robust, non-touch multi-user vital signs monitoring using depth camera-aided uwb. In *2021 IEEE 9th International Conference on Healthcare Informatics (ICHI)*, pages 320–329. IEEE, 2021.

- [92] Linwei Fan, Fan Zhang, Hui Fan, and Caiming Zhang. Brief review of image denoising techniques. *Visual Computing for Industry, Biomedicine, and Art*, 2(1):7, 2019.
- [93] Chunwei Tian, Lunke Fei, Wenxian Zheng, Yong Xu, Wangmeng Zuo, and Chia-Wen Lin. Deep learning on image denoising: An overview. *Neural Networks*, 131:251–275, 2020.
- [94] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Image super-resolution using deep convolutional networks. *IEEE transactions on pattern analysis and machine intelligence*, 38(2):295–307, 2015.
- [95] Chao Dong, Chen Change Loy, Kaiming He, and Xiaoou Tang. Learning a deep convolutional network for image super-resolution. In *Computer Vision—ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6–12, 2014, Proceedings, Part IV 13*, pages 184–199. Springer, 2014.
- [96] Boyang Wang, Fengyu Yang, Xihang Yu, Chao Zhang, and Hanbin Zhao. Apisr: Anime production inspired real-world anime super-resolution. *arXiv preprint arXiv:2403.01598*, 2024.
- [97] Boyang Wang, Bowen Liu, Shiyu Liu, and Fengyu Yang. Vcizr: Blind single image super-resolution with video compression synthetic data. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision*, pages 4302–4312, 2024.
- [98] Yunmei Chen, Hongcheng Liu, Xiaojing Ye, and Qingchao Zhang. Learnable descent algorithm for nonsmooth nonconvex image reconstruction. *SIAM Journal on Imaging Sciences*, 14(4):1532–1564, 2021.
- [99] Wanyu Bian, Albert Jang, and Fang Liu. Multi-task magnetic resonance imaging reconstruction using meta-learning. *arXiv preprint arXiv:2403.19966*, 2024.
- [100] Matthew J Muckley, Bruno Riemenschneider, Alireza Radmanesh, Sunwoo Kim, Geunu Jeong, Jingyu Ko, Yohan Jun, Hyungseob Shin, Dosik Hwang, Mahmoud Mostapha, et al. Results of the 2020 fastmri challenge for machine learning mr image reconstruction. *IEEE transactions on medical imaging*, 40(9):2306–2317, 2021.
- [101] Wanyu Bian, Albert Jang, and Fang Liu. Improving quantitative mri using self-supervised deep learning with model reinforcement: Demonstration for rapid t1 mapping. *Magnetic Resonance in Medicine*, 2024.
- [102] Chi Ding, Qingchao Zhang, Ge Wang, Xiaojing Ye, and Yunmei Chen. Learned alternating minimization algorithm for dual-domain sparse-view ct reconstruction. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 173–183. Springer, 2023.
- [103] Wanyu Bian, Qingchao Zhang, Xiaojing Ye, and Yunmei Chen. A learnable variational model for joint multimodal mri reconstruction and synthesis. In *International Conference on Medical Image Computing and Computer-Assisted Intervention*, pages 354–364. Springer, 2022.
- [104] Chen Liu, Nanyan Zhu, Haoran Sun, Junhao Zhang, Xinyang Feng, Sabrina Gjerswold-Selleck, Dipika Sikka, Xuemin Zhu, Xueqing Liu, Tal Nuriel, et al. Deep learning of mri contrast enhancement for mapping cerebral blood volume from single-modal non-contrast scans of aging and alzheimer’s disease brains. *Frontiers in Aging Neuroscience*, 14:923673, 2022.
- [105] Jens Kleesiek, Jan Nikolas Morshuis, Fabian Isensee, Katerina Deike-Hofmann, Daniel Paech, Philipp Kickingereder, Ullrich Köthe, Carsten Rother, Michael Forsting, Wolfgang Wick, et al. Can virtual contrast enhancement in brain mri replace gadolinium?: a feasibility study. *Investigative radiology*, 54(10):653–660, 2019.
- [106] Xueshen Li, Hongshan Liu, Xiaoyu Song, Brigitta C Brott, Silvio H Litovsky, and Yu Gan. Generating virtual histology staining of human coronary oct images using transformer-based neural network. In *Diagnostic and Therapeutic Applications of Light in Cardiology 2024*, volume 12819, page 1281903. SPIE, 2024.

- [107] Nanyan Zhu, Chen Liu, Xinyang Feng, Dipika Sikka, Sabrina Gjerswold-Selleck, Scott A Small, and Jia Guo. Deep learning identifies neuroimaging signatures of alzheimer’s disease using structural and synthesized functional mri data. In *2021 IEEE 18th International Symposium on Biomedical Imaging (ISBI)*, pages 216–220. IEEE, 2021.
- [108] Haoran Sun, Xueqing Liu, Xinyang Feng, Chen Liu, Nanyan Zhu, Sabrina J Gjerswold-Selleck, Hong-Jian Wei, Pavan S Upadhyayula, Angeliki Mela, Cheng-Chia Wu, et al. Substituting gadolinium in brain mri using deepcontrast. In *2020 IEEE 17th International Symposium on Biomedical Imaging (ISBI)*, pages 908–912. IEEE, 2020.
- [109] Wenchao Zhang, Chong Fu, Yu Zheng, Fangyuan Zhang, Yanli Zhao, and Chiu-Wing Sham. Hsnet: A hybrid semantic network for polyp segmentation. *Computers in biology and medicine*, 150:106173, 2022.
- [110] Nanyan Zhu, Chen Liu, Britney Forsyth, Zakary S Singer, Andrew F Laine, Tal Danino, and Jia Guo. Segmentation with residual attention u-net and an edge-enhancement approach preserves cell shape features. In *2022 44th Annual International Conference of the IEEE Engineering in Medicine & Biology Society (EMBC)*, pages 2115–2118. IEEE, 2022.
- [111] Zishun Feng, Dong Nie, Li Wang, and Dinggang Shen. Semi-supervised learning for pelvic mr image segmentation based on multi-task residual fully convolutional networks. In *2018 IEEE 15th International Symposium on Biomedical Imaging (ISBI 2018)*, pages 885–888. IEEE, 2018.
- [112] Haoyu Xie, Chong Fu, Xu Zheng, Yu Zheng, Chiu-Wing Sham, and Xingwei Wang. Adversarial co-training for semantic segmentation over medical images. *Computers in biology and medicine*, 157:106736, 2023.
- [113] Ziyan Li, Jianjiang Feng, Zishun Feng, Yunqiang An, Yang Gao, Bin Lu, and Jie Zhou. Lumen segmentation of aortic dissection with cascaded convolutional network. In *Statistical Atlases and Computational Models of the Heart. Atrial Segmentation and LV Quantification Challenges: 9th International Workshop, STACOM 2018, Held in Conjunction with MICCAI 2018, Granada, Spain, September 16, 2018, Revised Selected Papers 9*, pages 122–130. Springer, 2019.
- [114] Honghui Liu, Jianjiang Feng, Zishun Feng, Jiwen Lu, and Jie Zhou. Left atrium segmentation in ct volumes with fully convolutional networks. In *Deep Learning in Medical Image Analysis and Multimodal Learning for Clinical Decision Support: Third International Workshop, DLMIA 2017, and 7th International Workshop, ML-CDS 2017, Held in Conjunction with MICCAI 2017, Québec City, QC, Canada, September 14, Proceedings 3*, pages 39–46. Springer, 2017.
- [115] Ricky TQ Chen. torchdiffEq. URL <https://github.com/rtqichen/torchdiffEq>, 124, 2018.
- [116] Alexander Buslaev, Vladimir I Iglovikov, Eugene Khvedchenya, Alex Parinov, Mikhail Druzhinin, and Alexandr A Kalinin. Albuumentations: fast and flexible image augmentations. *Information*, 11(2):125, 2020.
- [117] Ilya Loshchilov and Frank Hutter. Decoupled weight decay regularization. *arXiv preprint arXiv:1711.05101*, 2017.
- [118] Stefan Van der Walt, Johannes L Schönberger, Juan Nunez-Iglesias, François Boulogne, Joshua D Warner, Neil Yager, Emmanuelle Gouillart, and Tony Yu. scikit-image: image processing in python. *PeerJ*, 2:e453, 2014.
- [119] Fabian Isensee, Paul F Jaeger, Simon AA Kohl, Jens Petersen, and Klaus H Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature methods*, 18(2):203–211, 2021.

Technical Appendices for ImageFlowNet: Forecasting Multiscale Image-Level Trajectories of Disease Progression with Irregularly-Sampled Longitudinal Medical Images

Table of Contents

A	Propositions and Proofs	20
A.1	Proposition 4.1	20
A.2	Proposition 4.2	22
B	Additional Background on Longitudinal Image Data	23
C	Additional Background on Why Time-Awareness is Important	24
D	Image Registration	25
D.1	Retinal Images	25
D.2	Brain Multiple Sclerosis Images	25
D.3	Brain Glioblastoma Images	25
E	Implementation Details	26
F	Evaluation Metrics	27

A Propositions and Proofs

A.1 Proposition 4.1

Proposition A.1. *Let f_θ be a continuous function that satisfies the Lipschitz continuity and linear growth conditions. Also, let the initial state $y(t_0) = y_0$ satisfy the finite second moment requirement ($\mathbb{E}[|y(t_0)|^2] < \infty$). Suppose $z(t_0)$ is the latent representation learned by ImageFlowNet at the initial state corresponding to t_0 . Then, our neural ODEs (Eqn (3a)) are at least as expressive as the original neural ODEs (Eqn (1a)), and their solutions capture the same dynamics.*

We recall the two dynamic systems for original neural ODEs and our ODEs:

Original neural ODEs:

$$\frac{dy(\tau)}{d\tau} = f_\theta(y(\tau), \tau), \quad f_\theta : \mathbb{R}^n \times [0, T] \rightarrow \mathbb{R}^n$$

Our neural ODEs, with (1) superscript $\cdot^{(b)}$ omitted without loss of generality, (2) z_τ equivalently replaced by $z(\tau)$ for notation consistency, and (3) f_θ replaced by \tilde{f}_θ for distinction:

$$\frac{dz(\tau)}{d\tau} = \tilde{f}_\theta(z(\tau)), \quad \tilde{f}_\theta : \mathbb{R}^m \rightarrow \mathbb{R}^m$$

Proof.

Theorem A.2 (Picard-Lindelöf [60]). *Let $D \subset \mathbb{R}^n$ be an open set, and let $f : D \times [0, T] \rightarrow \mathbb{R}^n$ be a continuous function that satisfies a Lipschitz condition in y uniformly in τ . Then, for any initial condition $y(t_0) = y_0$, there exists a unique solution to the initial value problem:*

$$\frac{dy(\tau)}{d\tau} = f(y(\tau), \tau), \quad y(t_0) = y_0.$$

Lipschitz Condition:

$$|f_\theta(y_1, \tau) - f_\theta(y_2, \tau)| \leq L|y_1 - y_2|$$

Linear Growth Condition:

$$|f_\theta(y, \tau)| \leq K(1 + |y|)$$

Given these conditions, both the original neural ODE and the Latent Space Neural ODE have unique strong solutions.

Since both the original ODE and the Latent Space Neural ODE have unique solutions, we could then construct a bijective and sufficiently smooth mapping $h : \mathbb{R}^n \times [0, T] \rightarrow \mathbb{R}^m$ such that $z(\tau) = h(y(\tau), \tau)$.

We define a function $h(y, \tau)$ that maps the state $y(\tau)$ and time τ to a new latent state $z(\tau)$ as

$$h(y, \tau) := y(\tau) \oplus \tau,$$

where \oplus denotes the concatenation of the state and time.

Then, as h is bijective, the inverse function h^{-1} maps $z(\tau)$ back to $y(\tau)$ and τ . Given $h(y, \tau) = y \oplus \tau$, the inverse is:

$$h^{-1}(z) = (y(z_{\text{time}}), z_{\text{time}})$$

By the chain rule, the derivative of $z(\tau)$ with respect to τ is:

$$\frac{dz(\tau)}{d\tau} = \frac{\partial h}{\partial y} \frac{dy}{d\tau} + \frac{\partial h}{\partial \tau}$$

Substituting the ODE for $y(\tau)$, we get:

$$\frac{dz(\tau)}{d\tau} = \frac{\partial h}{\partial y} f_{\theta}(y(\tau), \tau) + \frac{\partial h}{\partial \tau}$$

We can then simply define the function \tilde{f}_{θ} in the latent space such that it incorporates the dynamics from the original space:

$$\tilde{f}_{\theta}(z(\tau)) := \frac{\partial h}{\partial y} f_{\theta}(y(\tau), \tau) + \frac{\partial h}{\partial \tau}$$

The universal approximation theorem ensures that there exists a neural network parameterized by θ that can approximate any continuous function, including $\tilde{f}_{\theta}(z(\tau))$.

Existence of Equivalent Function Since the neural network can approximate $\tilde{f}_{\theta}(z(\tau))$, there exists a function $\tilde{f}_{\theta}(z(\tau))$ in the latent space that can represent the same system behavior governed by $f_{\theta}(y(\tau), \tau)$ in the original space.

Proving Equivalence: Given $z(\tau) = h(y(\tau), \tau)$ and the corresponding functions f_{θ} and \tilde{f}_{θ} , we have shown that the new ODE formulation:

$$\frac{dz(\tau)}{d\tau} = \tilde{f}_{\theta}(z(\tau))$$

captures the same dynamics as the original ODE:

$$\frac{dy(\tau)}{d\tau} = f_{\theta}(y(\tau), \tau)$$

□

A.2 Proposition 4.2

Proposition A.3. *If we consider an image as a distribution over a 2D grid, ImageFlowNet is equivalently solving a dynamic optimal transport problem, as it meets three essential criteria: (1) matching the density, (2) smoothing the dynamics, and (3) minimizing the transport cost, where the ground distance is the Euclidean distance in the latent joint embedding space.*

Proof. ImageFlowNet can alternatively be viewed in the context of a dynamic optimal transport framework, which aims to determine the optimal plan π to transport mass from an initial distribution μ to a target distribution ν for a fixed state interval $[\tau_i, \tau_j]$. The task meets three requirements of dynamic optimal transport: (1) matching the density, (2) smoothing the dynamics, and (3) minimizing the transport cost. The ground distance in the latent joint embedding space is the Euclidean distance.

Matching the density The image is a 2D grid, and the distribution for the pixel intensities is μ at τ_i and ν at τ_j on this grid. μ and ν are defined on measure space $\mathcal{X} \subset \mathbb{R}^n$ and $\mathcal{Y} \subset \mathbb{R}^n$ respectively. The set of all joint probability measures on $\mathcal{X} \times \mathcal{Y}$ is denoted as $\Pi(\mu, \nu)$ and $c(x, y)$ is the cost of moving a mass unit from the original distribution μ at state τ_i to the target distribution ν at state τ_j . Then, the distance between the two distributions μ and ν is the p-Wasserstein distance:

$$W(\mu, \nu)_p := \left(\inf_{\pi \in \Pi(\mu, \nu)} \int_{\mathcal{X} \times \mathcal{Y}} c(x, y) d\pi(x, y) \right)^{\frac{1}{p}}, \text{ where } p \geq 1$$

Benamou & Brenier [27] present a dynamic view of optimal transport, which links to differential equations. For the state interval $[\tau_i, \tau_j]$, there is a smooth and status-dependent density $P(z, \tau) \geq 0$ with $\int_{\mathbb{R}^n} P(z, \tau) dz = 1, \forall \tau \in [\tau_i, \tau_j]$, and a velocity fields $f(z, \tau)$ that obeys the continuity equation:

$$\partial_\tau P + \nabla \cdot (Pf) = 0, \text{ with } \tau \in [\tau_i, \tau_j] \text{ and } z \in \mathbb{R}^n, \text{ where } P(\cdot, \tau_i) = \mu, P(\cdot, \tau_j) = \nu$$

Smoothing the dynamics The velocity fields $f(z, \tau)$ follows the Lipschitz condition $|f(z_1, \tau) - f(z_2, \tau)| \leq L|z_1 - z_2|$ where $L > 0$, which ensures a smooth and controlled transport process. With the following setup, Benamou & Brenier [27] show that the Wasserstein distance with order 2 (W_2) is:

$$W(\mu, \nu)_2^2 = \inf_{(p, f)} \int_{\mathbb{R}^n} \int_{\tau_i}^{\tau_j} P(z, \tau) \|f(z, \tau)\|^2 d\tau dz$$

Minimizing the transport cost Based on the main theorems in [61, 62], this problem aims to find the trajectory f that minimizes the transport cost on the path space \mathbb{R}^n , we define the ground distance in the latent joint embedding space to be the Euclidean distance:

$$W(\mu, \nu)_2^2 = \inf_f \mathbb{E} \left[\int_{\tau_i}^{\tau_j} \|f(z_\tau, \tau)\|^2 d\tau \right] \text{ s.t. } \frac{\partial z(\tau)}{\partial \tau} = f_\theta(z_\tau, \tau), z_{\tau_i} \sim \mu, z_{\tau_j} \sim \nu$$

Here, f_θ follows the ODE or the SDE.

With the above setups, ImageFlowNet is equivalent to a dynamic optimal transport problem trying to match the density at different states. □

B Additional Background on Longitudinal Image Data

Longitudinal image datasets, including but not limited to retinal images or even medical images, often come with several challenges: ① high dimensionality, ② temporal sparsity, ③ sampling irregularity, and ④ spatial misalignment.

① **High Dimensionality** is intrinsic to image data. For images with height of H pixels, width of W pixels and C image channels, the dimensionality of the data is $\mathbb{R}^{H \times W \times C}$, which can easily go beyond a hundred thousand dimensions: a small image of $256 \times 256 \times 3$ has 196.6 thousand dimensions. Such high dimensionality is rarely encountered by most methods in time series prediction and temporal dynamics modeling [62–68].




Figure S1: Temporal sparsity, sampling irregularity and spatial misalignment in longitudinal images.

② **Temporal Sparsity** is especially common in longitudinal images in healthcare, as images are usually acquired at separate visits of the patient, where the time gap can be several months or years. In contrast, a relatively well-studied adjacent field is video data [69–71], where the frame rate can easily be 60 Hz or higher. This renders our data of interest easily 10^8 times sparser compared to the better studied video data.

③ **Sampling Irregularity** is also ubiquitous in clinical practice, both *within* and *among* longitudinal image series. *Within-series* irregularity means that the visits are not necessarily evenly distributed for the same patient over time. *Among-series* irregularity means that different patients do not follow the same readmission schedule either — *in terms of both time intervals and number of visits*. Times for visits can significantly vary based on doctors’ evaluation of the condition, the availability of doctors and/or imaging facilities, and the patient’s own preferences, among others. This feature defies the assumptions of most methods that require regular sampling or common sampling [62, 72].

④ **Spatial Misalignment** is often seen in longitudinal medical images too. Indeed, it is almost impossible to enforce pixel-perfect alignment of images acquired at different visits. Luckily, this problem can be addressed by image registration without any compounding effect with the temporal sparsity or sampling irregularity issues. See Appendix D for an illustration of image registration.

Temporal sparsity, sampling irregularity, and spatial misalignment are illustrated in Figure S1. These properties and challenges listed above lead to a **fairly unique area of research** that is **largely underexplored but highly interesting** to healthcare professionals.

Consider retinal imaging as an example. Most existing approaches to estimate disease progression in retinal imaging data do not operate in the image space, but rather in a vector space of a few clinical features extracted from the images. Examples of these derived statistics include the area of geographic atrophy lesions [73], the number of lesions [74], the lesion perimeter [73], its prior

observed growth rate [74], the presence and pattern of hyperfluorescence around the border of a lesion [75]. Although these approaches have been effective, they compress the rich context in the images to just a few metrics, and the output is an oversimplified representation of the disease states. This simplification overlooks the nuanced variations and complexities that are discarded during the feature extraction process and limits the interpretability of the output to a few preselected scalar-valued features.

In contrast, our proposed ImageFlowNet capitalizes on the extensive information available in the image to provide a nuanced representation of future conditions and also addresses the limitations of traditional metrics-based methodologies by offering a more dynamic and detailed visualization of disease progression. This method gives healthcare professionals an intuitive understanding of the expected progression of the disease and allows them to provide patients with a visual forecast that goes beyond mere numerical data.

We hope that our method can establish a new standard in the discipline and potentially transform clinical practices in areas including but not limited to ophthalmology or neurology, with the help of the latest imaging and measurement techniques [76, 77] as well as computational tools for disease diagnosis [78–80], risk prediction [81, 82], uncertainty quantification [83, 84], planning [85–88], and patient care [89–91].

C Additional Background on Why Time-Awareness is Important

Solving our problem outlined in Section 2 with deep learning requires designing and optimizing a model $\mathcal{F} : (\mathbb{R}^{H \times W \times C}, \mathbb{R}, \mathbb{R}) \rightarrow \mathbb{R}^{H \times W \times C}$, such that $\hat{x}_j = \mathcal{F}(x_i, t_i, t_j)$ and $\hat{x}_j \approx x_j$.

In most existing image-to-image tasks, the mapping between each pair of input x_i and output x_j obeys the same transformation rules, and hence their models are designed to be time-agnostic. For example, in denoising [92, 93], x_j is the noise-free version of x_i ; in super-resolution [94–97], x_j is higher in resolution than x_i by a fixed factor; in reconstruction [98–103], x_j is the transformed version of x_i through a fixed set of rules guided by physics; in contrast mapping [104–108], x_j represents the effect of staining or contrast agents when applied to x_i ; and in segmentation [109–114], x_j returns a label map describing the anatomical or functional segments in x_i . For these purposes, time-agnostic models, such as UNet or most diffusion models¹ remain competitive.

However, in our scenario, the output image is a function of both the input image and time. Given the same input image x_i , it will not end up at the same output image if the time interval changes. An image showing a disease 2 years after onset may look very different compared to 2 days after onset. In such cases, attempting to solve this problem using a model without time-modeling capabilities would be fundamentally ill-posed. In short, the spatial-temporal problem requires a spatial-temporal solution, which inspired our development of ImageFlowNet.

¹While diffusion models have modules that can encode time, many variants are used in a time-agnostic manner for tasks like denoising or super-resolution, where “time” is no different from “iteration”.

D Image Registration

D.1 Retinal Images

For all images, we extracted descriptive keypoints with SuperRetina [28], a high-quality keypoint detector trained on retinal images. Then we identified the keypoint correspondences for each image pair in each longitudinal series with a k -nearest-neighbor matcher and considered any image pair that has at least 15 keypoint correspondences a successful match. Next, we selected the image that produced the most successful matches as the “anchor image”. Finally, we aligned all images in the longitudinal series towards the “anchor image” using perspective transformation so that the degree of freedom is constrained to the adjustment of camera angle or position. As a post-processing step, for each longitudinal series, we cropped all images with the biggest common foreground square so that no image contained any background pixel outside the retina region.

The image registration process for a pair of images from the same longitudinal series is illustrated in Figure S2. It can be seen that all veins are aligned in the resulting images while atrophy borders are not. This is expected from perspective transformation and is exactly desirable for our task.




Figure S2: Our image registration pipeline. (A) Moving and fixed images come from the same eye at different time points. (B) SuperRetina is used to detect consistent and descriptive keypoints. (C) Keypoints are matched by descriptor similarity and filtered by distance heuristics. (D) The moving image is aligned under the constraint of a perspective transformation.

D.2 Brain Multiple Sclerosis Images

These images were already registered. No additional work was done.

D.3 Brain Glioblastoma Images

We used the scans in the “DeepBraTumIA” folders, which were registered to a common atlas, but the registration did not adequately align the scans in each longitudinal series. We used the Python tool from ANTS [29] to perform *Affine* followed by *Diffeomorphic* registration with [4, 2, 1] iterations to align each scan towards the first scan in series.

E Implementation Details

Architectures The proposed ImageFlowNet combines UNet and Neural ODEs. The UNet model follows the time-conditional UNet implementation in Guided Diffusion [40]. Neural ODEs are implemented with torchdiffeq [115].

Data Augmentation We used the albumentations package [116] to perform flipping, shifting, scaling, rotation, random brightness, and random contrast. We also make the UNet training a denoising process by adding random Gaussian noise to the input.

Hyperparameters and training details All experiments were performed on a SLURM server, where each job was allocated either an NVIDIA A100 GPU, an NVIDIA A5000 GPU, or an NVIDIA RTX 3090 GPU. All jobs can be completed within 2-5 days on a single GPU with 8 CPU cores. T-Diffusion usually takes the longest to train. ImageFlowNet_{SDE} variant may require a 40-GB GPU (sometime that will still hit an OOM error if running too many function evaluations in the SDE) while all other methods can be trained on a 20-GB GPU. Experiments shared the same set of hyperparameters: learning rate = 0.0001, batch size = 64, number of epochs = 120. Adam with decoupled weight decay (AdamW) [117] optimizer was used, along with a cosine annealing learning rate scheduler with linear warmup. We used an exponential moving average (EMA) with decay rate of 0.9 on the ImageFlowNet models.

To accommodate the GPU VRAM limits, we used gradient aggregation to trade efficiency for space while achieving the desired effective batch size — we used an actual batch size of 1, scaled the loss by $\frac{1}{64}$, and updated the weights every 64 batches.

Training of the segmentation networks are described in the next section (Evaluation Metrics).

F Evaluation Metrics

The evaluation metrics cover image similarity, residual magnitude, and atrophy similarity.

Image similarity We measure the image similarity between the real future image x_j and the predicted future image \hat{x}_j using peak signal-to-noise ratio (PSNR) and structural similarity index (SSIM). These two metrics are widely used in image-to-image tasks such as super-resolution, denoising, inpainting, etc.

PSNR is a normalized version of the mean squared error between two images that takes into account the dynamic range of the image data. The formula is given by Eqn (6).

$$\text{PSNR}(x_a, x_b) = 10 \log_{10} \left(\frac{R}{\text{MSE}(x_a, x_b)} \right), \text{ where} \quad (6)$$

R is the common dynamic range of the images

$$\text{MSE}(x_a, x_b) = \frac{1}{H \times W} \sum_{h \in H, w \in W} \|x_a^{(h,w)} - x_b^{(h,w)}\|^2$$

SSIM measures the similarity between two images by describing the perceived change in structural information. The formula is given by Eqn (7). We used the implementation in Scikit-image [118].

$$\text{SSIM}(x_a, x_b) = \frac{(2\mu_{x_a}\mu_{x_b} + c_1)(2\sigma_{x_a x_b} + c_2)}{(\mu_{x_a}^2 + \mu_{x_b}^2 + c_1)(\sigma_{x_a}^2 + \sigma_{x_b}^2 + c_2)}, \text{ where} \quad (7)$$

μ_{x_a} is the pixel sample mean of x_a

μ_{x_b} is the pixel sample mean of x_b

$\sigma_{x_b}^2$ is the variance of x_b

$\sigma_{x_b}^2$ is the variance of x_b

$\sigma_{x_b x_b}$ is the covariance of x_a and x_b

$c_1 = (0.01R)^2, c_2 = (0.03R)^2$

R is the common dynamic range of the images

Residual magnitude We evaluated the magnitude of the residual maps $\hat{x}_j - x_j$ using the mean average error (MAE) and the mean squared error (MSE).

Atrophy similarity We also want to emphasize the precise representation of the atrophy region. To this end, the simplest metric is the dice similarity coefficient (DSC) and Hausdorff distance (HD) of the binarized atrophy regions. DSC and HD between two binary masks X and Y are given by Eqn (8) and Eqn (9), respectively. For HD, we used the implementation in Scikit-image [118].

$$\text{DSC}(X, Y) = \frac{|X \cap Y|}{|X| + |Y|} \quad (8)$$

$$\text{HD}(X, Y) = \max \left\{ \sup_{x \in X} d(x, Y), \sup_{y \in Y} d(X, y) \right\} \quad (9)$$

To perform atrophy segmentation, we separately trained three auxiliary image segmentation network on all images, one for each dataset. All retinal images have their atrophy regions labeled by ophthalmologists. All brain images have associated segmentation maps from the dataset providers. These segmentators that we trained have an nn-UNet [119] architecture and were trained with an AdamW [117] optimizer at an initial learning rate of 0.001 for 120 epochs. With these networks, we can segment the atrophy regions in both the real future image x_j and the predicted future image \hat{x}_j . DSC and HD can be computed on the segmentation masks between each pair of interest.