# PINN-EM: Physics-Guided Disease Progression Model of Geographic Atrophy

**Dmitrii Lachinov**
Institute of Artificial Intelligence
Medical University of Vienna
1090 Vienna, Austria
dmitrii.lachinov@meduniwien.ac.at

**Thomas Pinetz**
Institute of Artificial Intelligence
Medical University of Vienna
1090 Vienna, Austria
thomas.pinetz@meduniwien.ac.at

**Hrvoje Bogunović**
Institute of Artificial Intelligence
Medical University of Vienna
1090 Vienna, Austria
hrvoje.bogunovic@meduniwien.ac.at

## Abstract

Predicting disease progression from longitudinal data is an important open challenge in biomedical research. In this paper, we focus on Geographic Atrophy (GA), a late dry form of Age-Related Macular Degeneration (AMD), and aim to construct a personalizable spatio-temporal disease progression model. Using a series of retinal Optical Coherence Tomography (OCT) scans, we infer the coefficients of a parametrized Partial Differential Equation (PDE) that best describes the true progression dynamics. This recovered PDE acts as a soft constraint, guiding an Implicit Neural Representation (INR) to accurately extrapolate future GA segmentation maps. To enable efficient training, we propose an iterative method, PINN-EM, specifically designed to recover the coefficients of non-linear PDEs from observational data. At each iteration, the method decouples the objective into a PDE coefficient fitting step and a data fitting step, drawing inspiration from the Expectation-Maximization algorithm. We evaluated PINN-EM on the complex task of GA progression modeling, where patients exhibit high variance in growth patterns and biological rates. The proposed spatio-temporal model substantially outperformed existing baselines, demonstrating the strong extrapolation capabilities of PDE-constrained INR models.

## 1 Introduction

Predicting disease progression from medical imaging data remains a challenging and largely unsolved problem Puglisi et al. [2025], Lachinov et al. [2024], Petersen et al. [2021], Ezhov et al. [2019] with a limited number of prior works in *spatio-temporal* biomedical domains (e.g. 138 documents in Scopus). The main challenges in addressing disease progression are the high non-linearity of the progression dynamics, heterogeneity of the normal and pathological alterations, and the stochastic nature of the underlying processes. Depending on the strength of the underlying assumptions, disease progression models can be roughly classified into *analytical*, forcing physical constraints on the model, e.g. FEM models; and *empirical*, inferring hidden dependencies from the data, e.g. Linear Regression, nnUNet Isensee et al. [2024], FNO Li et al. [2020], RNN Nguyen et al. [2020] based approaches.

The Third Austrian Symposium on AI and Vision (AIROV25).

*Analytical* models typically describe the disease progression as a dynamical system, utilizing prior knowledge in the form of PDEs. They excel at providing robust frameworks but often lack flexibility in modeling real high-variance data. Conversely, *empirical* models infer the progression dynamics directly from the observed data, incorporating little to no prior assumption. While these models easily capture hidden correlations, they often struggle with generalization and extrapolation.

In this work, we aim to bring the benefits of both *analytical* and *empirical* models together through the use of Physics Informed Neural Networks (PINNs) Raissi et al. [2019]. We expand the class of *spatio-temporal* disease models with INRs of the progression profile constrained by a surrogate PDE. Specifically, we apply this to construct a spatio-temporal disease progression model of Geographic Atrophy, characterized by the gradual loss of the Photoreceptor (PR) layer and the underlying Retinal Pigment Epithelium (RPE) cells. From a series of prior Optical Coherence Tomography scans, we attempt to predict the disease segmentation at arbitrary future time points.

Our contributions are twofold. First, we introduce a novel optimization framework for PINNs to estimate unknown coefficients in PDEs from observational data. Adapting the Expectation-Maximization algorithm, our approach iteratively refines parameter estimates and neural network weights to enhance convergence. Second, we leverage the proposed PINN-EM framework to train a spatio-temporal disease progression model, encoding disease trajectories as PDE-constrained Implicit Neural Representations, outperforming prior state-of-the-art predictive methods in clinical applications.

## 2 Methods

We consider a series of prior 3D OCT scans of a single patient and extract an en-face maps of GA segmentation and the PR layer thickness. We represent GA segmentation and PR thickness as a parametric function $u_\theta(x,t) : \mathbb{R}^{N+1} \to \mathbb{R}^M$ with parameters $\theta$, that maps spatial coordinates and time to the measured values. We then recover a parametrized differential operator $\mathcal{N}[u_\theta, \lambda] = 0$ that acts as a constraint guiding the INR.

### 2.1 PDE Selection

Following biological insights that GA represents a reaction-diffusion type process affecting the photoreceptor and RPE layers, we hypothesize that GA progression can be modelled using reaction-diffusion equations. For GA growth, we define the set of parameters $\lambda$ and differential operators $\mathcal{N}[u_\theta]$ and $\mathcal{B}$ as following:

$$\begin{aligned}\mathcal{N}[u_\theta] &= [\nabla^2 u_\theta, u_\theta(1-u_\theta), \partial_t u_\theta]^T \\ \lambda &= [\text{diag}(d), A, -1]^T \\ \mathcal{B} &= \nabla_x u_\theta(x,t) \cdot \hat{n}\end{aligned}$$

where $u_\theta$ is the approximate solution produced by a neural network and $\hat{n}$ is the normal directed outward from the boundary $\partial\Omega$, $d$ is a vector in $\mathbb{R}^M$ is a diagonal of the diffusion matrix, $A$ is a matrix in $\mathbb{R}^{M\times M}$ with reaction coefficients. Operator $\mathcal{N}[u_\theta, \lambda]$ is set as a dot product $\mathcal{N}[u_\theta]^T\lambda$.

### 2.2 PINN-EM Optimization Framework

The traditional PINNs approach minimizes a combination of empirical data loss, boundary conditions, and PDE residual loss. We propose interpreting this task as a constrained optimization problem, aiming to minimize the data approximation error subject to the system of differential equations being zero over the domain. To estimate the constraint values over the domain $\Omega$, we rely on sampling.

We decouple the problem at each iteration $k$ into three steps (Fig. 1):

1. **Sampling:** We sample the domain and boundary points for expectation estimation during the following two steps;
2. **Expectation (Parameter Update):** We estimate the value of the PDE parameters $\lambda$ by solving the least squares problem $\lambda^* = argmin_\lambda \mathbb{E}_{x,t\sim\Omega}[\|\mathcal{N}[u_\theta]^T(x,t)\lambda\|_2^2]$ and maintain an exponential running average $\lambda_k = (1-\beta)\lambda^* + \beta\lambda_{k-1}$.

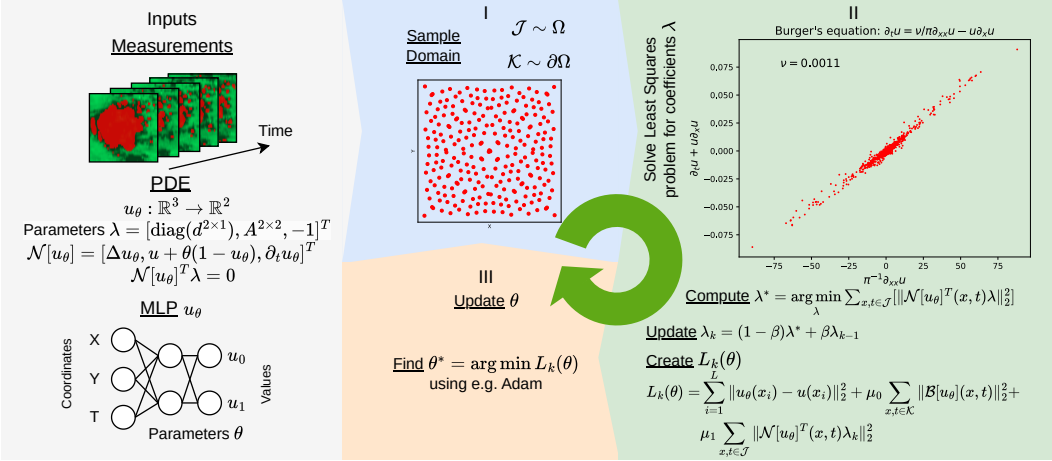


Figure 1: We consider a series of spatio-temporal measurements (thickness of PR layer at multiple visits), parameterized PDE (reaction-diffusion), and parameterized solution $u_\theta$ (neural network). We repeat three steps iteratively: (I) we sample the domain and boundary points for unsupervised loss; (II) we solve the least squares problem for parameters $\lambda$, update the estimate of $\lambda$ and create the optimization objective $L_k(\theta)$, where $\{x_i, t_i, u(x_i, t_i)\}_1^L$ are observed measurements, $\mu_0$ and $\mu_1$ are scalar hyperparameters controlling strength of the constraints; (III) we optimize function $L_k$ and update parameters $\theta$ of the neural network. The example uses Burger's equation with a single parameter $\nu$ to aid visualisation.

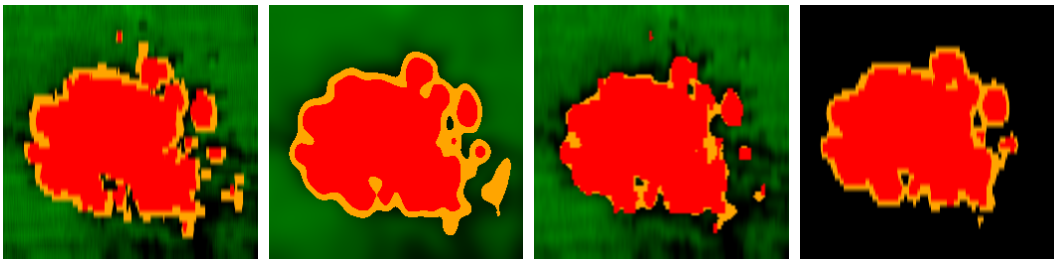

(a) Ground truth (b) Proposed (Hybrid) (c) FEM (Physical) (d) NODE (Data-driven)

Figure 2: Qualitative comparison of the Geographic Atrophy (GA) progression models. GA at baseline visit is depicted in red, GA progression over one year is in orange, and photoreceptor thickness at the baseline is shown in green. NODE utilizes raw imaging data, hence it has no photoreceptor component. The animated version can be found in the supplementary materials.

3. **Maximization (Neural Network Update):** Given $\lambda_k$, we construct an unconstrained objective function $L_k(\theta)$ (Fig. 1) combining empirical data loss, boundary conditions, and the PDE residual constraints. We then optimize the parameters $\theta$ using the Adam optimizer.

## 3 Experiments and Results

We evaluated the proposed method on the complex task of predicting GA progression using real-world clinical imaging data (Figure 2).

*Data and Experimental Setup* We used a longitudinal medical imaging dataset from a clinical observational study of patients with GA. The dataset features the progression of 75 eyes of 51 patients with 553 scans in total, consisting of 3D OCT scans of the retina. In the experiments, we utilize PR thickness and RPE integrity maps extracted automatically from 3D OCTs. We select the baseline visit such that it has at least 4 prior visits and a 1-year follow-up visit. The proposed method utilizes the prior visits to train the PINN-EM model individually per patient, extrapolating the disease map to exactly one year in the future. The model is evaluated on the task of predicting the segmentation map of GA at the primary endpoint of 12 months. The evaluation metrics incorporate the Dice index of

the entire GA lesion, the Dice index of newly converted GA lesions (Dice growth), as well as both Pearson's correlation $r$ and the coefficient of determination $R^2$ regarding the GA growth rate.

*Statistical Analysis* For paired data, e.g. Dice and Dice of the growth, we employed one-tailed paired sample t-test. For $r$ and $R^2$ we performed one-tailed independent sample t-test over 1000 bootstrap estimates. Asterisk (*) denotes significance value $p < 0.05$.

*Clinical Performance* The predictive performance on GA progression over one year is detailed in Table 1. To serve as a ceiling performance indicator, we included models trained utilizing the posterior knowledge from the follow-up visit.

Table 1: Results on estimating spatio-temporal progression of Geographic Atrophy. The + signifies the visits used for inference, e.g., 3 visits at 180, 360 and 540 days before the baseline - day 0; * denotes that the result is significantly better than all of the baselines.

| Model | -540 | -360 | -180 | 0 | 360 | Test: Day 360 | | | |
|---|---|---|---|---|---|---|---|---|---|
| | | | | | | r | $R^2$ | Dice | Dice growth |
| *Data-driven* | | | | | | | | | |
| Linear Regression | | + | | + | | 0.76 | 0.51 | n/a | n/a |
| Linear Growth | | + | | + | | 0.67 | 0.06 | $\underline{0.91 \pm 0.06}$ | $0.39 \pm 0.26$ |
| nnUNet Isensee et al. [2024] | | | | + | | 0.07 | -1.38 | $0.85 \pm 0.14$ | $0.45 \pm 0.15$ |
| RNN Nguyen et al. [2020] | | | | + | | 0.57 | -0.54 | $0.87 \pm 0.09$ | $0.50 \pm 0.12$ |
| Taylor Gigon et al. [2021] | | | | + | | 0.62 | -0.45 | $0.85 \pm 0.13$ | $0.45 \pm 0.14$ |
| FNO Li et al. [2020] | | | | + | | 0.65 | 0.35 | $\underline{0.91 \pm 0.05}$ | $0.55 \pm 0.11$ |
| NODE Lachinov et al. [2024] | | | | + | | 0.68 | 0.23 | $0.90 \pm 0.07$ | $0.57 \pm 0.12$ |
| *Physical* | | | | | | | | | |
| FEM | | + | | + | | 0.65 | 0.14 | $0.89 \pm 0.08$ | $0.50 \pm 0.13$ |
| *Proposed (Hybrid)* | | | | | | | | | |
| PINN (L-BFGS) | | + | | + | | 0.76 | 0.48 | $\underline{0.91 \pm 0.05}$ | $0.52 \pm 0.15$ |
| PINN (Adam) | | + | | + | | 0.79 | $\underline{0.56}$ | $\underline{0.91 \pm 0.06}$ | $0.50 \pm 0.16$ |
| PINN (EM) | | + | | + | | $\underline{0.83}^*$ | $\underline{0.56}$ | $\underline{0.91 \pm 0.10}$ | $\underline{0.61 \pm 0.12}^*$ |
| *With posterior knowledge* | | | | | | | | | |
| Linear Regression | | + | | + | + | 0.91 | 0.83 | n/a | n/a |
| Linear Growth | | + | | + | + | 0.94 | 0.86 | $0.92 \pm 0.04$ | $0.52 \pm 0.20$ |
| FEM | | + | | + | + | 0.97 | 0.92 | $0.91 \pm 0.05$ | $0.59 \pm 0.15$ |

The proposed PINN-EM algorithm, while having access only to the prior visits, was able to surpass the performance of the overfitted FEM and Linear Growth models in the Dice index of the growth region, achieving a score of 0.61. This outcome can be attributed to multiple factors, such as the soft constraints imposed on the implicit representation model, in contrast to FEM with hard constraints, and the suitability of the reaction-diffusion equation for modeling GA progression. The FEM framework is strictly bound by numerical error minimization, resulting in inflexibility when explaining newly formed lesions that were absent in initial conditions. Conversely, PINNs accommodate such occurrences by allowing minor perturbations in the initial conditions to evolve into novel lesion areas. Overall, the proposed optimization scheme substantially outperformed existing baselines in clinical predictive accuracy, especially distinguishing itself when forecasting across extended periods with supplementary prior visits.

## 4 Conclusion

This abstract summarizes spatio-temporal disease progression method, that uses multiple prior visits for learning an Implicit Neural Representation constrained by a parameterized PDE. Applied to GA progression, the method accurately extrapolated spatial changes and outperformed existing purely analytical and empirical approaches.

## References


Ivan Ezhov et al. Neural parameters estimation for brain tumor growth modeling. In *MICCAI 2019*, pages 787–795, 2019.

Anthony Gigon et al. Personalized atrophy risk mapping in age-related macular degeneration. *Translational Vision Science & Technology*, 10, 2021.

Fabian Isensee et al. nnU-net revisited: A call for rigorous validation in 3D medical image segmentation. In *MICCAI 2024*, pages 488–498. Springer Nature Switzerland, 2024.

D. Lachinov et al. Learning spatio-temporal model of disease progression with neuralodes from longitudinal volumetric data. *IEEE Transactions on Medical Imaging*, 43(3):1165–1179, 2024.

Zong-Yi Li et al. Fourier neural operator for parametric partial differential equations. *ArXiv*, abs/2010.08895, 2020.

Minh Nguyen et al. Predicting Alzheimer's disease progression using deep recurrent neural networks. *NeuroImage*, 222:117203, 2020. ISSN 1053-8119.

Jens Petersen et al. Continuous-time deep glioma growth models. In *MICCAI*, 2021.

Lemuel Puglisi, Daniel C. Alexander, and Daniele Ravì. Brain latent progression: Individual-based spatiotemporal disease progression on 3d brain mris via latent diffusion. *Medical image analysis*, 106:103734 – 103734, 2025.

M. Raissi, P. Perdikaris, and G.E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019. ISSN 0021-9991.

Scopus. Number of published articles with spatio-temporal, disease progression, medical keywords. URL https://www.scopus.com/results/results.uri?sort=plf-f&src=s&s=TITLE-ABS-KEY+%28+%28+%22spatio-temporal%22+OR+%22spatiotemporal%22+%29+AND+%28+%22disease+progression%22+OR+%22longitudinal+prediction%22+OR+%22trajectory+modeling%22+%29+AND+%28+%22medical+imaging%22+OR+%22MRI%22+OR+%22CT%22+OR+%22OCT%22+%29+%29&limit=10.