

Reproducible PINN Framework for Patient-Specific Modeling and Biophysical Parameter Inference in Glioblastoma

Juliette Vanderhaeghen¹^a, Cyril Corbet²^b François P. Duhoux³^c Christophe Pierreux⁴^d and Sébastien Jodogne¹^e

¹*ICTEAM, Louvain School of Engineering, UCLouvain, Belgium*

²*Institut de Recherche Expérimentale et Clinique (IREC), UCLouvain, Belgium*

³*Institut Roi Albert II, Department of Medical Oncology, Cliniques universitaires Saint-Luc, Belgium*

⁴*Cell Biology Unit, de Duve Institute, UCLouvain, Belgium*
juliette.vanderhaeghen@uclouvain.be

Keywords: Physics-Informed Neural Networks, Tumor Growth, Biophysical Parameters Inference.

Abstract: Cancer remains a major global health challenge, motivating personalized treatments. PDE-based models can capture tumor dynamics and enable patient-specific predictions, but traditional solvers like FDM or FEM can be computationally costly and require extensive calibration, while purely data-driven neural networks often lack interpretability. Physics-Informed Neural Networks (PINNs) address these limitations by embedding PDE constraints, supporting both forward simulations and inverse parameter estimation. We model glioblastoma growth using a reproducible, open-source PINN framework based on the Fisher–KPP equation. A systematic hyperparameter study evaluates architecture, activation functions, optimizers, learning rates, batching, and sampling strategies. Experiments on synthetic tumors show accurate dynamics and reliable recovery of biophysical parameters. We further provide a standalone Python implementation, transparent datasets, and practical guidelines for reproducible research in personalized oncology.

1 INTRODUCTION

Cancer remains a major global health challenge (Filho et al., 2025). Tumors exhibit strong heterogeneity in size, invasion patterns, and therapy response, complicating treatment design and motivating patient-specific strategies (Zhang et al., 2025b). Mathematical models offer a way to quantify this variability, ranging from ordinary differential equations (ODEs) for overall behavior (Lima et al., 2022), to partial differential equations (PDEs) for spatiotemporal dynamics (Murray, 2012), and agent-based models for cell-level interactions (Poleszczuk et al., 2016). This work focuses on PDE-based tumor models, which capture spatial growth and are naturally compatible with imaging data (Wu et al., 2022; Lorenzo et al., 2024). A key step for clinical applicability is solving the inverse problem of estimating patient-specific pa-

rameters from limited observations. These parameters are typically unknown and exhibit substantial inter-patient variability (Swanson et al., 2008), making accurate estimation essential.

While numerical solvers are accurate, they require complex meshing, high computational cost, and careful calibration (Harish et al., 2021; Lorenzo et al., 2024). In contrast, purely data-driven neural networks can learn complex patterns but lack interpretability and perform poorly with scarce data (Linardatos et al., 2020; Vinuesa and Beril, 2021). Physics-informed neural networks (PINNs) address these limitations by embedding the governing PDEs directly into the loss function, combining clinical observations with physical laws during training (Raissi et al., 2019). This enables efficient simulations, supports inverse parameter estimation, and maintains consistency with underlying tumor biology.

We apply PINNs to both forward and inverse problems for modeling glioblastoma growth in two spatial dimensions using representative brain slices. The novelty of this work is a systematic and reproducible analysis of PINN design choices for tu-

^a <https://orcid.org/0009-0003-6289-7167>

^b <https://orcid.org/0000-0001-8795-9131>

^c <https://orcid.org/0000-0002-5429-7888>

^d <https://orcid.org/0000-0002-9647-2816>

^e <https://orcid.org/0000-0001-6685-7398>

mor growth, combined with a fully standalone, open-source implementation. Our contributions are three-fold: (i) We independently reproduce prior results and provide open-source access to the code, promoting transparency and reproducibility; (ii) We introduce a standalone, Python-based TensorFlow framework without reliance on proprietary software or third-party libraries, lowering technical barriers to future clinical adoption; (iii) We systematically analyze the impact of network architecture and hyperparameters on accuracy. Validation on synthetic tumors with known parameters, simulated on real MRI scans, demonstrates the feasibility and potential of the approach for patient-specific modeling.

The paper reviews existing approaches to tumor growth modeling, describes the Fisher–KPP equation and the proposed open-source PINN framework¹ for both forward and inverse problems, presents results, and discusses model performance and future directions for PDE-based oncology modeling.

2 RELATED WORK

Modeling tumor growth remains a significant challenge, and recent studies have proposed a range of approaches that can be broadly classified into traditional numerical methods, data-driven techniques based on deep learning, and, more recently, on PINNs.

Classical numerical methods such as the finite difference method (FDM) and finite element method (FEM) have long been standard tools for solving PDEs (Peiró and Spencer, 2005). In tumor growth modeling, FDM has been used, for example, to simulate avascular breast tumors (Maganin et al., 2020), as well as in patient-specific settings such as the two-species glioma model of Hormuth et al., calibrated using multiparametric MRI (Hormuth, 2021). FEM has been applied in a similar context, for example by Wu et al. to model breast tumor growth (Wu et al., 2022). Despite their success, these methods face major limitations: high computational cost in 3D, complex mesh generation, and reliance on longitudinal imaging data that are often scarce in clinical settings (Harish et al., 2021; Lorenzo et al., 2024).

Neural networks methods have also been explored for tumor growth prediction, leveraging their ability to approximate nonlinear mappings directly from imaging data. Zhang et al. proposed a spatio-temporal ConvLSTM architecture to model 3D spatial consistency and temporal dynamics from longitudinal imaging (Zhang et al., 2020), while more recent

work combines ConvLSTM with transformers to predict lung tumor evolution (Ma et al., 2023). Although these approaches enable rapid prediction, they often lack interpretability, generalize poorly under limited data, and may produce unrealistic outputs due to the absence of explicit physical constraints (Linardatos et al., 2020; Vinuesa and Beril, 2021).

PINNs incorporate PDE residuals directly into their loss, reducing data requirements and enforcing physical consistency (Raissi et al., 2019). While increasingly adopted in other engineering domains and emerging in biomedical areas such as cardiology (Mitchell et al., 2024), their use in tumor growth modeling remains limited. To the best of our knowledge, only a small number of studies have applied PINNs to the Fisher–KPP equation for modeling tumor growth. Notably, Zhang et al. demonstrated patient-specific parameter estimation from a single 3D MRI, highlighting the potential of PINNs for minimally invasive and personalized neuro-oncology modeling (Zhang et al., 2025a). Beyond reaction–diffusion systems, Rodrigues et al. also explored PINNs for modeling multicellular tumor spheroid growth under the Verhulst and Montroll frameworks (Rodrigues, 2024).

Despite these advances, deploying PINNs in biomedical contexts remains challenging. Existing implementations for glioblastoma growth often involve complex pipelines, limiting accessibility and clinical applicability. Moreover, PINN training is highly sensitive to network architecture, optimizer choice, and hyperparameter selection. Balancing data fidelity, PDE residuals, and boundary conditions leads to ill-conditioned, multi-objective losses that hinder convergence (Roquemen-E and Mosquera-L, 2025).

3 METHODS

To overcome the complexity of existing PINN pipelines and training instabilities, this work proposes a lightweight and open-source framework for tumor growth modeling using PINNs. The framework is implemented exclusively in TensorFlow and distributed with complete code to ensure transparency and reproducibility. This section introduces the mathematical model and the data processing pipeline.

3.1 Tumor Growth Model

This study uses the Fisher–Kolmogorov–Petrovsky–Piskunov (Fisher–KPP) equation, a classical reaction–diffusion model (Fisher, 1937) widely applied in population dynamics and later adapted to tumor growth (Murray, 2012).

¹Code available at https://github.com/juliettevdh/pinn_tumor_growth_framework

The Fisher–KPP equation combines two key processes: diffusion, representing spatial cell migration, and a logistic reaction term modeling proliferation limited by environmental capacity. The paper focuses on glioma growth in 2D using representative brain slices, with tumor dynamics governed by the Fisher–KPP equation:

$$\begin{cases} \frac{\partial u(\mathbf{x}, t)}{\partial t} = D(\mathbf{x}) \nabla^2 u(\mathbf{x}, t) + \rho u(\mathbf{x}, t)(1 - u(\mathbf{x}, t)), \\ u(\mathbf{x}, 0) = u_0(\mathbf{x}), \\ n \cdot \nabla u(\mathbf{x}, t) = 0 \text{ on } \partial\Omega. \end{cases} \quad (1)$$

Here, $u(\mathbf{x}, t)$ denotes the tumor cell density at location $\mathbf{x} = (x, y) \in \mathbb{R}^2$ and time t . The computational domain Ω models the brain, with homogeneous Neumann boundary conditions on $\partial\Omega$ enforcing zero flux across the skull. The diffusion coefficient $D(\mathbf{x})$ [mm²/day] varies spatially, reflecting heterogeneous tumor invasiveness, while the proliferation rate ρ [day⁻¹] is assumed constant, characterizing overall tumor growth. The initial cell density is modeled as a scaled two-dimensional Gaussian distribution:

$$u(\mathbf{x}, 0) = 0.5 \cdot \exp\left(-200 \cdot \left[(x - x_0)^2 + (y - y_0)^2\right]\right), \quad (2)$$

where (x_0, y_0) is the tumor center. Following Zhang et al., the diffusion coefficient is defined as:

$$D(\mathbf{x}) = D_w \cdot P_w(\mathbf{x}) + D_g \cdot P_g(\mathbf{x}), \quad (3)$$

where $P_w(\mathbf{x})$ and $P_g(\mathbf{x})$ denote the fraction of white and gray matter, respectively, and D_w and D_g [mm²/day] are tissue-specific diffusion constants (Zhang et al., 2025a). To reflect lower mobility of tumor cells in gray matter, we assume a ratio $c = D_g/D_w = 0.1$ (Zhang et al., 2025a). For the forward problem, model parameters are taken from Tracqui et al., with the proliferation rate set to $\rho = 0.012$ [day⁻¹] (Tracqui et al., 1995). The diffusion coefficient in white matter is fixed at $D_w = 0.013$ [mm²/day]. Using Fisher’s approximation (Fisher, 1937), this yields a tumor invasion speed corresponding to roughly 9.13 mm per year, consistent with reported low-grade glioma growth rates (Swanson et al., 2008). For the inverse problem, D_w and ρ are treated as unknown, patient-specific parameters. A key advantage of combining the Fisher–KPP framework with PINNs is the ability to infer these parameters directly from sparse tumor observations, enabling personalized growth predictions and treatment planning.

3.2 Physics-Informed Neural Networks

PINNs extend standard neural networks by enforcing both data consistency and compliance with the governing PDE during training. Let $u(\mathbf{x}, t)$ denote the




Figure 1: PINN architecture for forward and inverse settings in brain tumor growth modeling application.

neural network approximation of the solution. The residual associated with a general PDE is defined as:

$$r(\mathbf{x}, t) = \frac{\partial u(\mathbf{x}, t)}{\partial t} + \mathcal{N}[u(\mathbf{x}, t); \lambda], \quad (4)$$

where $\mathcal{N}$ is the nonlinear differential operator, and λ denotes the physical parameters of the system (Raissi et al., 2019). In our application, the residual associated with the Fisher–KPP equation can be derived from Equations 1 and 3 as follows:

$$r(\mathbf{x}, t) = \frac{\partial u(\mathbf{x}, t)}{\partial t} - D_w \nabla \cdot (P(\mathbf{x}) \nabla u(\mathbf{x}, t)) - \rho u(\mathbf{x}, t)(1 - u(\mathbf{x}, t)), \quad (5)$$

where $P(\mathbf{x}) = P_w(\mathbf{x}) + c \cdot P_g(\mathbf{x})$ encodes the patient brain geometry. The learning task is then formulated as the minimization of a composite loss function, which combines the data mismatch with the PDE residual:

$$MSE = MSE_u + MSE_r, \quad (6)$$

where:

$$MSE_u = \frac{1}{N_u} \sum_{i=1}^{N_u} |\hat{u}(\mathbf{x}_u^i, t_u^i) - u^i|^2, \quad (7)$$

$$MSE_r = \frac{1}{N_r} \sum_{i=1}^{N_r} |r(\mathbf{x}_r^i, t_r^i)|^2. \quad (8)$$

Here, $\{\mathbf{x}_u^i, t_u^i, u^i\}_{i=1}^{N_u}$ denotes sampling locations with available data, while $\{\mathbf{x}_r^i, t_r^i\}_{i=1}^{N_r}$ represents collocation points sampled within the spatiotemporal domain. These collocation points are artificially generated and allow the model to enforce the governing equations without direct observations. This formulation naturally accommodates both forward and inverse problems. In forward problems, the parameters λ are known, and the goal is to approximate the solution $u(\mathbf{x}, t)$. In inverse problems, partial observations of the solution are available, and the parameters λ are treated as unknowns to be inferred from the data. A schematic view of the neural network architecture used to approximate the solution is shown in Figure 1.




Figure 2: Distribution of training points for the forward problem. Left: Collocation points centered around tumor region. Right: Initial condition data points at $t = 0$ within the circular region; points outside the brain sampled over time to enforce boundary conditions. Colors indicate $u(\mathbf{x}, t)$ values.

3.3 Data Processing

To provide meaningful inputs to the PINN, brain imaging data² are preprocessed via anatomical masking and spatial normalization. Brain geometry is modeled using the diffuse domain method, which introduces a phase-field function $\phi(\mathbf{x})$ to encode anatomical boundaries on a rectangular domain (Zhang et al., 2025a; Li et al., 2009). The field satisfies $\phi(\mathbf{x}) \approx 1$ inside brain tissue and $\phi(\mathbf{x}) \approx 0$ outside, with a narrow diffuse interface replacing sharp boundaries. This avoids explicit segmentation and complex meshing while restricting tumor growth to biologically plausible regions. The original images (rescaled to 150×180 pixels) are mapped to the computational domain $\Omega = [-1, 1] \times [-1, 1]$ to provide compatibility with the input space expected by the PINN. Then, D_w is rescaled to maintain physical consistency:

$$D_{\text{norm}} = D_w \cdot \frac{1}{2} \left[\left(\frac{2}{L_x} \right)^2 + \left(\frac{2}{L_y} \right)^2 \right], \quad (9)$$

where $L_x = 150$ and $L_y = 180$. The temporal coordinate is similarly normalized to $\tau \in [0, 1]$ via $\tau = t/t_{\text{max}}$, where t_{max} corresponds to the total simulation time in days used in the FDM. Under these transformations, the Fisher-KPP residual enforced by the PINN becomes:

$$r(\mathbf{x}, \tau) = \frac{\partial(\phi u(\mathbf{x}, \tau))}{\partial \tau} - t_{\text{max}} \left(D_{\text{norm}} \nabla \cdot (\phi P(\mathbf{x}) \nabla u(\mathbf{x}, \tau)) + \phi r u(\mathbf{x}, \tau) (1 - u(\mathbf{x}, \tau)) \right), \quad (10)$$

which simultaneously encodes anatomical constraints and normalized spatiotemporal tumor dynamics. With normalized inputs, the PINN is trained using a

²This work uses NIfTI images of Patient 1 slice 171 publicly available at: <https://github.com/Rayzhangzirui/pinnngbm/tree/main>




Figure 3: Alternative distribution of training points explored in the hyperparameter study. Left: Collocation points uniformly sampled. Right: Initial condition data points at $t = 0$ uniformly sampled; points outside the brain sampled over time for boundary enforcement. Colors indicate $u(\mathbf{x}, t)$ values.

targeted sampling strategy. For the forward problem, three point types are used (Figure 2): data points outside the brain, initial condition points, and collocation points, with the latter two concentrated near the tumor origin. An alternative sampling strategy is explored in the hyperparameter study (Figure 3). In the inverse problem, collocation points and data points extracted from the synthetic tumor solutions are used to infer model parameters.

4 RESULTS

We systematically evaluated network architectures, optimizers, and hyperparameters for the forward problem, using glioblastoma as a representative case. The goal is to provide practical guidelines for robust and efficient PINN training in biomedical contexts. Results for forward and inverse problem are then conducted. Training data were generated by simulating tumor growth with an FDM solver applied to the Fisher-KPP equation. This synthetic solution serves three purposes: providing the initial condition for the forward problem, reference data for the inverse problem (inferring PDE parameters), and a benchmark for evaluating prediction accuracy. Prediction accuracy is quantified using absolute and relative errors between the PINN predicted tumor density $\hat{u}_{\text{pred}}(\mathbf{x}, t)$ and the reference FDM solution $u_{\text{fdm}}(\mathbf{x}, t)$ at a selected time t_p . The absolute error is given by:

$$\text{Error}(\mathbf{x}, t_p) = |\hat{u}_{\text{pred}}(\mathbf{x}, t_p) - u_{\text{fdm}}(\mathbf{x}, t_p)| \quad (11)$$

Finally, the relative error at location $\mathbf{x}$ is defined as:

$$\begin{cases} 0, & \text{if } u_{\text{fdm}}(\mathbf{x}, t_p), \hat{u}_{\text{pred}}(\mathbf{x}, t_p) < 0.01, \\ \frac{|\hat{u}_{\text{pred}}(\mathbf{x}, t_p) - u_{\text{fdm}}(\mathbf{x}, t_p)|}{u_{\text{fdm}}(\mathbf{x}, t_p) + \varepsilon}, & \text{otherwise,} \end{cases} \quad (12)$$

where $\varepsilon = 10^{-4}$ ensures numerical stability in regions of very low tumor density.

Table 1: Effect of varying key PINN hyperparameters on model accuracy and computational cost. In this table, 5×32 indicates 5 layers with 32 neurons.

	Modified parameters	Rel. Err. [%]	Max Abs. Err.	Computation Time
<i>Baseline</i>	5×32 , tanh, Adam, 10^{-3} , $\alpha=\beta=\gamma=1$	0.0165	0.0294	15min 5s
<i>Architecture</i>	3×32	0.0639	0.0417	9min 20s
	5×128	0.0392	0.0061	19min 59s
	5×32 (without batching)	0.0120	0.0285	70min 7s
<i>Loss weights</i>	$\alpha=0.1, \beta=1, \gamma=1$	0.2079	0.0569	14min 59s
	$\alpha=1, \beta=10, \gamma=1$	0.0362	0.0282	15min 3s
	$\alpha=1, \beta=1, \gamma=10$	0.1813	0.0411	15min 4s
<i>Activation</i>	swish	0.0175	0.0349	18min 2s
	sigmoid	0.4232	0.0526	15min 19s
<i>Optimizer</i>	SGD	49.51	0.6876	13min 58s
	RMSprop	0.4101	0.049	14min 28s
	AdamW	0.0364	0.0275	15min 16s
<i>Learning rate</i>	10^{-2}	2.0216	0.8632	15min 3s
	10^{-4}	0.2413	0.0603	15min 2s
<i>Sampling</i>	Centered with less points	0.0207	0.0261	15min 2s
	Distribution of points across all domain	0.4453	0.1397	15min 6s

4.1 Hyperparameter Study

To evaluate the impact of architectural and training choices on glioblastoma growth modeling, we conducted a hyperparameter study for the forward problem across seven factors: network architecture, activation function, optimizer, learning rate, loss weighting, sampling strategy, and batch size.

Each experiment varied one factor while keeping others fixed. Architectural tests included 3–5 hidden layers with 32–128 neurons each (default activation tanh), alongside alternative activations (swish, sigmoid), optimizers (Adam, AdamW, RMSprop, SGD), and learning rates (10^{-2} – 10^{-4}). Loss weights α , β , γ balanced contributions from the initial condition, PDE residual, and boundary terms. Sampling strategies varied the number and distribution of collocation, initial, and boundary points. All models except one used batching, which reduced computation time by a factor of five without loss of accuracy. Batched training used random batches of 256 points per set per iteration (see Figure 2). Models were trained for 120,000 epochs with fixed $D_w = 0.013$ [mm²/day] and $\rho = 0.012$ [day⁻¹]. Performance was evaluated using relative and absolute errors. Table 1 reports representative configurations; other results are discussed.

Increasing network depth or width improves accuracy at the cost of longer training times. Optimizer choice is critical: Adam and AdamW consistently perform best, while RMSprop is less accurate and SGD performs poorly. A learning rate of 10^{-3} balances performance. While tanh ensures stable training; swish yields similar accuracy, whereas sigmoid underperforms. Balanced loss weights give the best results, while overemphasizing terms degrades accuracy. If batching is not used, the sampling strategy influence accuracy: distributed sampling can lead to

severe accuracy loss. Reducing the number of collocation points accelerates training without compromising accuracy. With batching, these effects largely disappear, and the distribution of points has less impact. Based on this, the optimal configuration for the forward problem is a 5-layer network with 32 neurons per layer, Adam optimizer (learning rate 10^{-3}), tanh activation, tumor-centered sampling, and batch size 256. For the inverse problem, we used a batch size of 512, swish activation, and distributed Latin hypercube sampling.

4.2 Forward Problem

To assess accuracy of the proposed PINN framework, we first analyze the forward problem. Figure 4 presents a comparison of glioblastoma growth simulations in a brain slice, as computed using PINNs and FDM. Three distinct times are reported: $t = 0$ day (first row), $t = 100$ days (second row), and $t = 200$ days (third row). The first two columns show the PINN and FDM solutions, while the third and fourth columns display the relative and absolute error maps.

At early times ($t = 0$ and $t = 100$ days), the PINN predictions closely match the FDM reference, and even at $t = 200$ days, absolute errors remain low despite small peripheral discrepancies.

4.3 Inverse Problem

This section focuses on the inverse problem. In this case, the PINN is used to estimate unknown biophysical parameters D_w and ρ from the synthetic solution.

Figure 5 monitors the coupled convergence of the biophysical parameters D_w and ρ during PINN training. The proliferation rate ρ is identified robustly




Figure 4: Comparison of PINN predictions with FDM reference solutions for forward simulations of glioblastoma growth. Row 1: Initial condition; Row 2: $t = 100$ days; Row 3: $t = 200$ days. Column 1: PINN predictions; Column 2: FDM reference solutions; Column 3: Relative error maps highlighting discrepancies between the two methods; Column 4: Absolute error.




Figure 5: Coupled convergence of biophysical parameters during PINN training. A: Evolution of the estimated diffusion coefficient D_w during training. B: Evolution of the estimated proliferation rate ρ during training.

Table 2: Relative error of PINN-estimated parameters compared to ground-truth values.

Estimated parameters	Error [%]	Epochs
Unknown ρ and D_w	0.75 and 3.52	500,000

and rapidly with a very small error of 0.75%, whereas D_w is recovered more gradually. After 500,000 epochs, the relative error for D_w is 3.52%. These results demonstrate that the PINN framework can successfully identify key parameters in the Fisher-KPP model. Accurate parameter recovery translates into tumor profiles that closely match the FDM reference solution. In Figure 6, the PINN predictions at $t = 0$, $t = 100$, and $t = 200$ days remain closely aligned with the FDM solution, showing reliable forward reconstructions based on the inferred parameters.

5 DISCUSSION

The first contribution of this study is a comprehensive analysis of PINN hyperparameters and architecture for the forward problem. Batched training reduced computation time fivefold without compromising accuracy as well as overfitting, highlighting its efficiency for biomedical modeling. While common tanh achieved the lowest relative error (0.0165%), swish performed comparably (0.0175%), indicating the potential of alternative activation functions.

Unlike FDM solvers, which efficiently simulate tumor evolution for known parameters (forward problem) but do not support parameter inference (inverse problem), PINNs enforce the PDE residual across the




Figure 6: Comparison of PINN predictions with the FDM reference solution for the inverse problem of glioblastoma growth. This figure uses the same conventions as Figure 4.

domain, enabling simultaneous solution of the equations and estimation of unknown parameters. In our experiments, both the proliferation rate ρ and diffusion coefficient D_w were successfully recovered. The proliferation rate ρ was rapidly identified (with 0.75% error), while D_w (with 3.52% error) converged more slowly due to its small normalized value reducing sensitivity in the loss function. While estimating D_w is more challenging, the resulting PINN predictions remain close to the ground truth, as shown in Figure 6.

These findings support reported evidence of the potential of PINNs for parameter recovery in biomedical applications (Zhang et al., 2025a), although our approach differs by avoiding grid search and adimensionalization. Our goal is to provide complementary insights and a clear framework for reproducible application and further development. Future work could extend these results to multiple slices and 3D volumes. Key challenges in translating this PINN framework to real patient data include imaging noise and sparsity, limited resolution, and discrepancies between the simplified tumor model and patient-specific tumor biology. Several strategies could enhance efficiency and clinical relevance. Bayesian PINNs enable uncertainty quantification (Yang et al., 2021), while domain decomposition and parallel training can accelerate computation (Jagtap et al., 2020). Moreover, expanding PINNs to other cancers is important. Incorporating treatment effects, biological mark-

ers (HER2, Ki67, etc.), and microenvironmental factors represents also a major opportunity. While FEM solvers have explored treatment integration (Wu et al., 2022), extending it to PINNs could support robust, predictive frameworks for treatment optimization.

6 CONCLUSION

This study presents results for both forward and inverse glioblastoma growth problems using PINNs. Under the Fisher-KPP model, PINNs accurately reproduce tumor dynamics and reliably recover diffusion (3.52% error) and proliferation (0.75% error) parameters, confirming that combining data-driven approaches with PDE constraints can disentangle these processes. A key contribution is a fully open-source Python framework for PINN-based tumor modeling, designed for accessibility and reproducibility. Our systematic hyperparameter study also highlights the benefits of batching and offers practical guidance for robust and efficient PINN configuration.

Future work will focus on extending the framework to 3D settings, improving robustness, computational efficiency, uncertainty quantification, and validating performance on real patient data. Incorporating treatment effects and extending the approach to other cancer types are necessary steps to assess the versatility and clinical relevance of PINNs.

ACKNOWLEDGMENTS

This work was supported by the ARC Project BiD-iTwins of the Fédération Wallonie-Bruxelles (Belgium), under Grant 23/28-129.

REFERENCES

- Filho, A. M., Laversanne, M., Ferlay, J., Colombet, M., Piñeros, M., Znaor, A., Parkin, D. M., Soerjomataram, I., and Bray, F. (2025). The GLOBOCAN 2022 cancer estimates: Data sources, methods, and a snapshot of the cancer burden worldwide. *International Journal of Cancer*.
- Fisher, R. A. (1937). The wave of advance of advantageous genes. *Annals of Eugenics*.
- Harish, S., Annanth, V. K., Abinash, M., Kannan, K., Agarwal, S., and Mohana, N. (2021). A study on numerical methodologies in solving fluid flow and heat transfer problems. *IOP Conference Series: Earth and Environmental Science*.
- Hormuth, David A, e. a. (2021). Image-based personalization of computational models for predicting response of high-grade glioma to chemoradiation. *Scientific Reports*.
- Jagtap, A. D., Kharazmi, E., and Karniadakis, G. E. (2020). Conservative physics-informed neural networks on discrete domains for conservation laws: Applications to forward and inverse problems. *Computer Methods in Applied Mechanics and Engineering*.
- Li, X., Lowengrub, J., Rätz, A., and Voigt, A. (2009). Solving pdes in complex geometries: a diffuse domain approach. *Communications in mathematical sciences*.
- Lima, E. A., Wyde, R. A., Sorace, A. G., and Yankeelov, T. E. (2022). Optimizing combination therapy in a murine model of her2+ breast cancer. *Computer methods in applied mechanics and engineering*.
- Linardatos, P., Papastefanopoulos, V., and Kotsiantis, S. (2020). Explainable ai: A review of machine learning interpretability methods. *Entropy*.
- Lorenzo, G., Ahmed, S. R., Hormuth, D. A., Vaughn, B., Kalpathy-Cramer, J., Solorio, L., Yankeelov, T. E., and Gomez, H. (2024). Patient-specific, mechanistic models of tumor growth incorporating artificial intelligence and big data. *Annual Review of Biomedical Engineering*.
- Ma, M., Zhang, X., Li, Y., Wang, X., Zhang, R., Wang, Y., Sun, P., Wang, X., and Sun, X. (2023). Convlstm coordinated longitudinal transformer under spatio-temporal features for tumor growth prediction. *Computers in Biology and Medicine*.
- Maganin, J., Maria-Lopes-Romeiro, N., Rodrigues-Cirilo, E., and Laerte-Natti, P. (2020). Simulation of a mathematical model of tumoral growth using finite differences. *Brazilian Journal of Development*.
- Mitchell, D., Cai, S., Qian, Y., Myzelev, E., Kumbhat, A., Li, H., and Lu, L. (2024). Transfer learning on physics-informed neural networks for tracking the hemodynamics in the evolving false lumen of dissected aorta. *Nexus*.
- Murray, J. D. (2012). Glioblastoma brain tumours: Estimating the time from brain tumour initiation and resolution of a patient survival anomaly after similar treatment protocols. *Journal of Biological Dynamics*.
- Peiró, J. and Spencer, S. (2005). Finite difference, finite element and finite volume methods for partial differential equations. *Handbook of Materials Modeling*.
- Poleszczuk, J., Macklin, P., and Enderling, H. (2016). Agent-based modeling of cancer stem cell driven solid tumor growth. In *Stem cell heterogeneity: Methods and protocols*.
- Raissi, M., Perdikaris, P., and Karniadakis, G. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*.
- Rodrigues, J. A. (2024). Using physics-informed neural networks for tumor cell growth modeling. *Mathematics*.
- Roquemen-E, V. and Mosquera-L, C. (2025). Recent advancements and applications of physics-informed machine learning in biomedical research. *Current Opinion in Biomedical Engineering*.
- Swanson, K. R., Rostomily, R. C., and Alvord, E. (2008). A mathematical modelling tool for predicting survival of individual patients following resection of glioblastoma: a proof of principle. *British journal of cancer*.
- Tracqui, P., Cruywagen, G., Woodward, D., Bartoo, G., Murray, J., and Alvord Jr, E. (1995). A mathematical model of glioma growth: The effect of chemotherapy on spatio-temporal growth. *Cell Proliferation*.
- Vinuesa, R. and Beril, S. (2021). Interpretable deep-learning models to help achieve the sustainable development goals. *Nature Machine Intelligence*.
- Wu, C., Jarrett, A. M., Zhou, Z., Elshafeey, N., and Adrada, B. E. (2022). Mri-based digital models forecast patient-specific treatment responses to neoadjuvant chemotherapy in triple-negative breast cancer. *Cancer Research*.
- Yang, L., Meng, X., and Karniadakis, G. E. (2021). B-pinns: Bayesian physics-informed neural networks for forward and inverse pde problems with noisy data. *Journal of Computational Physics*.
- Zhang, L., Le, L., Wang, X., M. Zhu, R., Bagheri, M., M. Summers, R., and Yao, J. (2020). Spatio-temporal convolutional lstms for tumor growth prediction by learning 4d longitudinal patient data. *IEEE Transactions on Medical Imaging*.
- Zhang, R. Z., Ezhov, I., Balcerak, M., Zhu, A., Wiestler, B., Menze, B., and Lowengrub, J. S. (2025a). Personalized predictions of glioblastoma infiltration: Mathematical models, physics-informed neural networks and multimodal scans. *Medical Image Analysis*.
- Zhang, X., Xie, J., Yang, Z., Yu, C. K. W., Hu, Y., and Qin, J. (2025b). Tumour heterogeneity and personalized treatment screening based on single-cell transcriptomics. *Computational and Structural Biotechnology Journal*.