
Train Forwards, Optimize Backwards: Neural Surrogates for Personalized Medical Simulations

Jonas Weidner

Technical University of Munich
Munich Center for Machine Learning
j.weidner@tum.de

Ivan Ezhov

Technical University of Munich

Michal Balcerak

University of Zurich

Björn Menze

University of Zurich

Daniel Rückert

Technical University of Munich
Imperial College London

Benedikt Wiestler

Technical University of Munich
Munich Center for Machine Learning

Abstract

Personalizing biophysical simulations to individual patients remains a major computational bottleneck, as traditional optimization requires repeated runs of costly numerical solvers. We present a differentiable system that replaces the forward simulation with a neural surrogate, providing a fast and accurate approximation of the underlying biophysical model. The surrogate’s differentiability enables efficient gradient-based inversion of patient-specific parameters, even when the original solver is non-differentiable. Applied to a 3D finite-difference model of brain tumor growth, our method achieves clinically relevant accuracy while reducing optimization time from days to seconds. This demonstrates how differentiable surrogates can serve as core components of broader differentiable systems for scientific machine learning.

1 Introduction

In medical research, computational modeling underpins many applications, from finite element models of brain tumor growth to patient-specific biophysical simulations, which are central to advancing personalized medicine [10, 12]. These models capture complex physiological processes but are computationally expensive, often requiring hours for a single forward run. Consequently, parameter fitting and calibration to patient data, essential for personalization, remain challenging inverse problems [5]. Traditional optimization methods such as Bayesian or evolutionary strategies [21, 14] demand numerous forward evaluations, making clinical integration infeasible. Neural surrogates have emerged to mitigate this limitation by approximating input-output mappings with significant speedups [11, 17]. They are used in medical imaging and physical modeling [19, 13], yet they typically focus on forward predictions rather than solving inverse problems. Physics-Informed Neural Networks (PINNs) [4, 7] partially address this but often suffer from instability and difficulty handling discontinuities [23]. Related approaches that embed physics as soft constraints [2, 9] or rely on differentiable simulators [11] face similar limitations in generalizability and applicability. We propose to use differentiable neural surrogates to solve inverse problems in personalized medical simulations [22, 1]. Specifically, we demonstrate how to transform a non-differentiable biophysical simulator for brain tumor growth into a differentiable system using a learned surrogate, bridging numerical modeling and optimization.¹

¹This manuscript extends a prior short version [22] by incorporating comparative optimization strategies, a broader differentiable-systems framing, and an explicit hybrid initialization. Code: github.com/jonasw247/train-forwards-optimize-backwards.

2 Methods

We aim to personalize biophysical tumor models by fitting simulated tumor concentrations to individual patients’ magnetic resonance images (MRI). Accurate personalization is essential for radiotherapy planning, particularly to infer tumor infiltration in regions not visible in the images. By estimating patient-specific tumor growth coefficients, we seek to optimize treatment targets beyond the radiologically defined tumor boundaries (Figure 1). Typically, the final stage of personalization introduces an imaging function that transforms the continuous tumor-cell concentration into a binary segmentation to compare it to MRI. Earlier studies tested several hand-crafted imaging functions and calibrated some of their parameters [22, 2, 14, 21]. To remove this additional source of uncertainty and focus squarely on the inverse problem of calibrating the biophysical model, we perform parameter estimation directly against the simulated “ground-truth” concentration.




Figure 1: **(a)** The goal is to estimate the invisible 3-D brain tumor concentration crucial for radiation treatment planning based on the visible segmentation of the MRI and biophysical constraints. Therefore, we simulate the tumor growth. The numerical solver inputs the PDE coefficients and the brain tissue to predict a tumor concentration. Optimizing those coefficients with classical methods is slow and thus limits the clinical adaptation of complex brain tumor models. **(b)** We introduce a precise and differentiable neural surrogate enabling fast personalization.

We apply the widely used Fisher-Kolmogorov equation for brain tumor growth

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

where $c(\mathbf{x}, t)$ is the tumor-cell density, $D(\mathbf{x})$ the tissue-dependent diffusion coefficient, and ρ is the proliferation rate. The brain volume is segmented into white matter (WM), grey matter (GM) and cerebrospinal fluid (CSF) resulting in tissue-specific diffusion: $D_{\text{GM}} = 0.1 D_{\text{WM}}$, $D_{\text{CSF}} = 0$ as tumor cells infiltrate GM more slowly and are assumed not to invade CSF [6]. The forward solver inputs the proliferation rate ρ , diffusion coefficient D_{WM} , seed location (x, y, z) , and total growth time T , i.e. the parameter set $\theta_{\text{orig}} = \{x, y, z, \rho, D_{\text{WM}}, T\}$. With only a single clinical scan, θ_{orig} is not identifiable; for example, $(\rho \uparrow, T \downarrow)$ and $(\rho \downarrow, T \uparrow)$ may yield identical tumor extents. Following [14], we unify T into two scale-free parameters $\mu_D = \sqrt{D_{\text{WM}} T}$, $\mu_\rho = \sqrt{\rho T}$, and optimise the reduced set $\theta = \{x, y, z, \mu_D, \mu_\rho\}$. Given θ , the numerical solver returns the 3-D tumor cell-density field $c(\mathbf{x}; \theta)$.

We train a neural surrogate that approximates the forward solver of the reaction–diffusion PDE. During training, the network receives as input the PDE coefficients together with the relevant side constraints, the initial tumor seed, tissue mask, and boundary conditions. The output is the final tumor-cell concentration obtained by the numerical solver. The surrogate learns the mapping: $f_\phi : (\theta, \mathcal{B}) \mapsto c(\mathbf{x})$, where θ denotes the set of biophysical coefficients, ϕ denotes the surrogate weights and $\mathcal{B}$ the boundary data. Because the forward problem is well-posed, this supervised learning task is stable and admits a unique solution, orders of magnitude faster than the classical solver. Empirically, U-Net architectures have proven effective for this task [8, 18]. Recent studies further show that substituting the standard convolutional blocks in a U-Net with ConvNeXt blocks yields superior performance on reaction–diffusion problems [15, 17]. We therefore adopt this ConvNeXt-U-Net variant. The PDE coefficients are first processed by a fully connected layer, and the resulting activations are added channel-wise to the U-Net bottleneck. Injecting the coefficients at this most abstract level allows the latent representation to modulate spatial features in a physics-informed manner.

For inverse optimization, we evaluate four approaches. **CMA-ES classical solver:** As a classical baseline, we employ the CMA-ES (cma Python package) to optimize the tumor-concentration model. To ensure a fair comparison with our neural surrogate, we deliberately use the vanilla algorithm, foregoing common refinements such as multi-resolution search schemes or machine-learning priors [21]. **CMA-ES neural surrogate:** To disentangle the benefits of gradient information from those of the fast forward model, we also run CMA-ES directly on the neural surrogate. In this setting, the expensive numerical solver is replaced by the surrogate, yet optimization remains entirely derivative-free. The gradients of the network are intentionally ignored. This ablation allows us to assess how much of the performance gain stems from the surrogate’s speed alone versus the use of gradient-based inversion. **Direct inverse prediction model:** An obvious solution is the prediction of the coefficients directly by a network (Figure 2a), as done by [6]. For this approach, we used a ConvNeXt [15] architecture similar to the encoder of the forward neural surrogate mode. The network inputs the tumor concentration and the brain tissue to predict the PDE coefficients. **GB neural surrogate:** Once the network has converged, we freeze its weights and exploit its end-to-end differentiability for inverse optimization. Treating f_ϕ as a differentiable function of θ , we compute the exact gradient of the ℓ_2 loss between the predicted and observed tumor maps. This gradient is then supplied to a memory-efficient quasi-Newton optimiser (L-BFGS) that iteratively updates θ to minimise the discrepancy (Figure 2b). **Inverse Prediction and GB Optimization:** We integrate the two complementary inversion strategies, direct inverse prediction and the gradient-based neural surrogate, into a single hybrid pipeline. Concretely, we first execute the direct inverse model to obtain a coarse but physically plausible estimate θ_{DI} of the PDE coefficients. This estimate is then used to initialise the gradient-based optimiser that operates on the neural surrogate. Starting from θ_{DI} provides the quasi-Newton solver with a point already close to the true optimum, reducing the risk of entrapment in poor local minima.

After the GB optimisation terminates, we evaluate both candidate solutions, the original direct-inverse estimate and the refined GB surrogate estimate. The parameter set that achieves the lower loss is retained as the *final* prediction. Optimization was terminated once convergence was achieved, i.e. when the change in MSE between successive iterations dropped below 10^{-7} .




Figure 2: Overview of our neural surrogate optimization. The goal is to estimate the optimal PDE coefficients that explain the brain tumor. (a) The direct prediction model estimates the PDE coefficients that explain the given tumor concentration, thus it solves the inverse prediction within one inference run. (b) A differentiable neural surrogate model is trained on the forward tumor simulation problem. It inputs the PDE coefficients and the brain tissue and outputs a tumor concentration. Finally, the patient-specific coefficients are estimated using gradient-based optimization.

We evaluate surrogate- and solver-generated tumor fields with five complementary metrics. Mean-squared error (MSE) captures the total energy of numerical deviations and harshly penalizes large local mistakes, while mean absolute error (MAE) reports the voxel-wise bias in the same physical units as the data. Because raw errors scale with lesion size, we also report volume-normalized MSE and MAE, dividing each by the ground-truth tumor volume so results remain comparable across samples with different tumor extents. Normalized cross-correlation (NCC) measures how well the spatial pattern of low- and high-concentration regions aligns, independent of any global scale or offset. This metric is particularly interesting for radiotherapy planning, which acts on the relative radiation distribution.

We used a dataset generated by the TUMORGROWTHTOOLKIT [2], designed to match tumor sizes seen in the BRATS dataset [16]. We trained on 28,000 samples, validated on 1,000, and tested on

500. Due to cost constraints, we evaluated the full-resolution CMA-ES only on a subset of 25. We used an NVIDIA Quadro RTX 8000 and an AMD EPYC 7313 16-Core Processor for all experiments.

3 Results

The goal of our approach is to determine the optimal set of coefficients that describes the patient’s tumor, resulting in a fitting tumor concentration (Table 1). The classical pipeline that couples the numerical solver with CMA-ES attains a low simulation error, even though we found that the recovered coefficients deviate substantially from the ground truth. This apparent contradiction arises from the ill-posed nature of the inverse problem: multiple parameter combinations can reproduce the observed tumor distribution similarly well, enabling the optimiser to explain the data without converging to the true coefficients. The results of the surrogate-based methods and direct inverse prediction show worse performance compared to the classical solver. By investigating the per-sample errors, we observed a heavily skewed distribution for the optimization approaches, resulting in the larger standard error (Table 1). A few failed optimization runs dominate the mean values. By applying our hybrid approach, which utilizes the direct inverse prediction as initialization for the subsequent gradient-based optimization, we can dramatically reduce those failure cases. This results in significantly better performance, even outperforming classical optimization on the clinically relevant NCC metric, with a focus on relative differences.

Table 1: Main findings. Performance comparison for the final forward run with a classical solver following strict physical constraints. We compare mean squared error (MSE), mean absolute error (MAE), also normalized by the ground truth tumor volume (MSE / MAE Normalized), the normalized cross correlation (NCC), and the runtime. We report the mean and the standard error. **Bold** indicates best, underlined indicates second best.

Model	MSE ($\downarrow$) ($\times 10^{-3}$)	MAE ($\downarrow$) ($\times 10^{-3}$)	MSE norm. ($\downarrow$) ($\times 10^{-8}$)	MAE norm. ($\downarrow$) ($\times 10^{-8}$)	NCC ($\uparrow$) ($\times 10^{-1}$)	Runtime ($\downarrow$) (min)
CMA-ES classical solver	0.34 ± 0.08	1.55 ± 0.24	0.77 ± 0.22	4.00 ± 0.70	<u>8.95 ± 0.37</u>	2300
CMA-ES Neural Surrogate	5.42 ± 0.59	8.10 ± 0.65	4.59 ± 0.43	8.07 ± 0.47	7.81 ± 0.13	2.5
Direct Inverse Prediction	4.81 ± 0.20	8.87 ± 0.32	5.02 ± 0.23	10.08 ± 0.40	8.47 ± 0.04	0.1
GB Neural Surrogate	5.91 ± 0.65	8.53 ± 0.73	4.88 ± 0.48	7.99 ± 0.53	8.19 ± 0.11	<u>0.7</u>
(Ours) Inverse Prediction and GB Optimization	<u>1.35 ± 0.19</u>	<u>3.29 ± 0.29</u>	<u>1.36 ± 0.15</u>	<u>4.05 ± 0.25</u>	9.15 ± 0.10	0.8

4 Discussion

We find that differentiable neural surrogates can drastically accelerate the personalization of medical simulations. We have demonstrated that learning the well-posed forward path and optimizing it works better than predicting the ill-posed problem directly. By enabling efficient gradient-based optimization for the inverse problem, we achieved a speedup from days to seconds in calibrating brain tumor models, for this "train forwards, optimize backwards" approach. However, occasional optimization runs converged to suboptimal local minima, impacting mean error metrics. The optimal surrogate architecture may vary for different physical systems, and performance inherently depends on the training dataset’s quality and scope. Furthermore, our current approach largely sidesteps the complexities of explicit imaging functions by focusing on synthetic cases.

In ongoing work, we focus on enhancing optimization robustness and extending our framework to further medical simulation tasks, including more complex biophysical models with additional parameters that are currently infeasible to optimize due to their computational cost [20, 3]. Additionally, we are exploring its application to other medical domains, such as modeling Alzheimer’s disease progression and multi-body mechanical spine modeling. Especially for high-dimensional problems, we expect gradient-based optimization to provide a decisive advantage. We believe that differentiable neural surrogates represent a transformative step toward rapid and precise personalized medicine.

References

- [1] Matan Atad, Gabriel Gruber, Marx Ribeiro, Luis Fernando Nicolini, Robert Graf, Hendrik Möller, Kati Nispel, Ivan Ezhov, Daniel Rueckert, and Jan S Kirschke. Neural network surrogate and projected gradient descent for fast and reliable finite element model calibration: A case study on an intervertebral disc. *Computers in Biology and Medicine*, 186:109646, 2025.
- [2] Michal Balcerak, Ivan Ezhov, Petr Karnakov, Sergey Litvinov, Petros Koumoutsakos, Jonas Weidner, Ray Zirui Zhang, John S Lowengrub, Bene Wiestler, and Bjoern Menze. Individualizing glioma radiotherapy planning by optimization of a data and physics informed discrete loss. *arXiv preprint arXiv:2312.05063*, 2023.
- [3] Thomas Bortfeld and Gregory Buti. Modeling the propagation of tumor fronts with shortest path and diffusion models—implications for the definition of the clinical target volume. *Physics in Medicine & Biology*, 67(15):155014, 2022.
- [4] Salvatore Cuomo, Vincenzo Schiano Di Cola, Fabio Giampaolo, Gianluigi Rozza, Maziar Raissi, and Francesco Piccialli. Scientific machine learning through physics-informed neural networks: Where we are and what’s next. *Journal of Scientific Computing*, 92(3):88, 2022.
- [5] Manlio De Domenico, Luca Allegri, Guido Caldarelli, Valeria d’Andrea, Barbara Di Camillo, Luis M Rocha, Jordan Rozum, Riccardo Sbarbati, and Francesco Zambelli. Challenges and opportunities for digital twins in precision medicine from a complex systems perspective. *npj Digital Medicine*, 8(1):37, 2025.
- [6] Ivan Ezhov, Kevin Scibilia, Katharina Franitza, Felix Steinbauer, Suprosanna Shit, Lucas Zimmer, Jana Lipkova, Florian Kofler, Johannes C Paetzold, Luca Canalini, et al. Learn-morphinfer: a new way of solving the inverse problem for brain tumor modeling. *Medical Image Analysis*, 83:102672, 2023.
- [7] Jeremías Garay, Jocelyn Dunstan, Sergio Uribe, and Francisco Sahli Costabal. Physics-informed neural networks for parameter estimation in blood flow models. *Computers in Biology and Medicine*, 178:108706, 2024.
- [8] Zeineb Haouari, Jonas Weidner, Ivan Ezhov, Aswathi Varma, Daniel Rueckert, Bjoern Menze, and Benedikt Wiestler. Efficient deep learning-based forward solvers for brain tumor growth models. In *BVM Workshop*, pages 57–62. Springer, 2025.
- [9] Petr Karnakov, Sergey Litvinov, and Petros Koumoutsakos. Optimizing a discrete loss (odil) to solve forward and inverse problems for partial differential equations using machine learning tools. *arXiv preprint arXiv:2205.04611*, 2022.
- [10] Evangelia Katsoulakis, Qi Wang, Huanmei Wu, Leili Shahriyari, Richard Fletcher, Jinwei Liu, Luke Achenie, Hongfang Liu, Pamela Jackson, Ying Xiao, et al. Digital twins for health: a scoping review. *NPJ digital medicine*, 7(1):77, 2024.
- [11] Felix Koehler, Simon Niedermayr, Nils Thuerey, et al. Apebench: A benchmark for autoregressive neural emulators of pdes. *Advances in Neural Information Processing Systems*, 37:120252–120310, 2024.
- [12] Keying Kuang, Frances Dean, Jack B Jedlicki, David Ouyang, Anthony Philippakis, David Sontag, and Ahmed M Alaa. Med-real2sim: Non-invasive medical digital twins using physics-informed self-supervised learning. *Advances in Neural Information Processing Systems*, 37:5757–5788, 2024.
- [13] Lei Li, Julia Camps, Blanca Rodriguez, and Vicente Grau. Solving the inverse problem of electrocardiography for cardiac digital twins: A survey. *IEEE Reviews in Biomedical Engineering*, 2024.
- [14] Jana Lipkova, Panagiotis Angelikopoulos, Stephen Wu, Esther Alberts, Benedikt Wiestler, Christian Diehl, Christine Preibisch, Thomas Pyka, Stephanie E Combs, Panagiotis Hadjidoukas, et al. Personalized radiotherapy design for glioblastoma: integrating mathematical tumor models, multimodal scans, and bayesian inference. *IEEE transactions on medical imaging*, 38(8):1875–1884, 2019.

- [15] 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.
- [16] Bjoern H Menze, Andras Jakab, Stefan Bauer, Jayashree Kalpathy-Cramer, Keyvan Farahani, Justin Kirby, Yuliya Burren, Nicole Porz, Johannes Slotboom, Roland Wiest, et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014.
- [17] Ruben Ohana, Michael McCabe, Lucas Meyer, Rudy Morel, Fruzsina Agocs, Miguel Beneitez, Marsha Berger, Blakesly Burkhart, Stuart Dalziel, Drummond Fielding, et al. The well: a large-scale collection of diverse physics simulations for machine learning. *Advances in Neural Information Processing Systems*, 37:44989–45037, 2024.
- [18] 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.
- [19] Matteo Salvador, Marina Strocchi, Francesco Regazzoni, Christoph M Augustin, Luca Dede’, Steven A Niederer, and Alfio Quarteroni. Whole-heart electromechanical simulations using latent neural ordinary differential equations. *NPJ Digital Medicine*, 7(1):90, 2024.
- [20] Shashank Subramanian, Ali Ghafouri, Klaudius Matthias Scheufele, Naveen Himthani, Christos Davatzikos, and George Biros. Ensemble inversion for brain tumor growth models with mass effect. *IEEE Transactions on Medical Imaging*, 42(4):982–995, 2022.
- [21] Jonas Weidner, Ivan Ezhov, Michal Balcerak, Marie-Christin Metz, Sergey Litvinov, Sebastian Kaltenbach, Leonhard Feiner, Laurin Lux, Florian Kofler, Jana Lipkova, et al. A learnable prior improves inverse tumor growth modeling. *IEEE Transactions on Medical Imaging*, PP, Nov 2024.
- [22] Jonas Weidner, Ivan Ezhov, Michal Balcerak, Benedikt Wiestler, et al. Rapid personalization of pde-based tumor growth using a differentiable forward model. In *Medical Imaging with Deep Learning*, 2024.
- [23] Ray Zirui Zhang, Ivan Ezhov, Michal Balcerak, Andy Zhu, Benedikt Wiestler, Bjoern Menze, and John S Lowengrub. Personalized predictions of glioblastoma infiltration: Mathematical models, physics-informed neural networks and multimodal scans. *Medical Image Analysis*, 101:103423, 2025.