
Physics-Regularized Multi-Modal Image Assimilation for Brain Tumor Localization

Michal Balcerak^{1*}, Tamaz Amiranashvili^{1,2}, Andreas Wagner²,
Jonas Weidner², Petr Karnakov³, Johannes C. Paetzold⁴, Ivan Ezhov²,
Petros Koumoutsakos³, Benedikt Wiestler^{2†}, Bjoern Menze^{1†}

¹University of Zurich, ²Technical University of Munich
³Harvard University, ⁴Imperial College London

Abstract

Physical models in the form of partial differential equations serve as important priors for many under-constrained problems. One such application is tumor treatment planning, which relies on accurately estimating the spatial distribution of tumor cells within a patient's anatomy. While medical imaging can detect the bulk of a tumor, it cannot capture the full extent of its spread, as low-concentration tumor cells often remain undetectable, particularly in glioblastoma, the most common primary brain tumor. Machine learning approaches struggle to estimate the complete tumor cell distribution due to a lack of appropriate training data. Consequently, most existing methods rely on physics-based simulations to generate anatomically and physiologically plausible estimations. However, these approaches face challenges with complex and unknown initial conditions and are constrained by overly rigid physical models. In this work, we introduce a novel method that integrates data-driven and physics-based cost functions, akin to Physics-Informed Neural Networks (PINNs). However, our approach parametrizes the solution directly on a dynamic discrete mesh, allowing for the effective modeling of complex biomechanical behaviors. Specifically, we propose a unique discretization scheme that quantifies how well the learned spatiotemporal distributions of tumor and brain tissues adhere to their respective growth and elasticity equations. This quantification acts as a regularization term, offering greater flexibility and improved integration of patient data compared to existing models. We demonstrate enhanced coverage of tumor recurrence areas using real-world data from a patient cohort, highlighting the potential of our method to improve model-driven treatment planning for glioblastoma in clinical practice.

1 Introduction

The management of gliomas, particularly glioblastomas, is highly challenging due to their infiltration beyond the tumor margins detectable in medical imaging. This infiltrative behavior complicates the precise personalization of radiotherapy. Current clinical practice applies uniform radiation to a 1.5 cm margin around the active tumor core visible in Magnetic Resonance Imaging (MRI) scans. However, this approach overlooks critical factors such as heterogeneous infiltration patterns, varying brain tissues, and anatomical barriers, failing to account for the highly individualized dynamics of

*Corresponding author: michal.balcerak@uzh.ch

†Contributed equally as senior authors.

tumor spread in each patient. Personalizing radiotherapy to account for this spread remains an unmet clinical need in neuro-oncology [1].

The computational modeling of tumor growth has become a promising solution to this challenge [2, 3, 4, 5], by formalizing the tumor growth process through mathematical models of varying complexity [6]. However, personalizing these models to the limited clinical observational (image) data of a patient remains a formidable challenge. To address it, two main computational approaches have emerged:

Hard-Constrained Physics-Based Models: These models use Partial Differential Equations (PDEs) to constrain the solution space, ensuring that output tumor cell distributions adhere strictly to predefined physical laws [7, 8, 9, 10, 11].

However, since PDEs are approximations of the underlying stochastic processes in complex biological environments, they may overconstrain the system. While these models offer structure and insight, they fail to capture the intricate behaviors inherent to biological processes like tumor growth. Additionally, linking image observations to state variables requires careful consideration. As a result, the full integration of data is hindered by these limitations, reducing the effectiveness of physical models in representing empirical data.

Data-Driven Models: These models [12, 13] offer great flexibility by leveraging access to imaging datasets. They can directly be learned from – and applied to – the patients’ diagnostic scans. However, they often lack rigorous quality control and fail to integrate crucial biological insights, such as tissue-specific infiltration patterns and brain topology. This makes them less reliable for clinical application, especially when there is insufficient data to infer them purely from imaging information.

In Search of a Hybrid Approach: Combining physical insights on spatiotemporal distributions of state variables with increased model flexibility could be highly beneficial. One approach to achieve this is by testing the physics residual in selected locations, as demonstrated in Physics-Informed Neural Networks (PINNs) [14, 15, 16, 17, 18]. Another way involves projecting the learned solution into a latent space that adheres to physics constraints [19, 20]. A different framework, Optimizing a Discrete Loss (ODIL) [21], constructs a residual of PDEs using grid-based discretization. Unlike methods that optimize the weights of Multilayer Perceptrons (MLPs) that implicitly parametrize the unknown fields in the PINNs method, ODIL directly optimizes the discretized unknown fields, which can be faster and more stable due to the local (rather than global, as in MLPs) dependencies of the solution on the learnable weights.

Both ODIL [22] and PINNs [23] have recently been employed to model growth equations conditioned on MRI and metabolic imaging of patients. ODIL advances state-of-the-art radiotherapy planning by relaxing the constraints of the growth equation, thereby enhancing the capture of tumor recurrence. A similar trend is observed in computer vision’s novel view synthesis, where neural radiance field MLPs [24] are being replaced by discrete representations [25, 26, 27].

However, none of these hybrid approaches have modeled brain-tumor interactions dealing with constraints imposed by local structural anatomy, such as the relative composition of brain tissues and their microstructure, or by the impact of gross anatomical deformations visible in structural scans of large tumors. While biomechanical model extensions can address the latter, incorporating tissue elasticity remains challenging due to the unknown initial anatomy and the computational complexity of the tumor growth dynamics. So far, no hybrid approach free from hard PDE constraints has successfully linked biological processes, such as tumor growth, with models of the biological environment, like tissue elasticity. Only a few studies have used PINNs for biomechanical modeling in synthetic scenarios, demonstrating that the displacement field of materials can be learned from observations [28, 29].

The main contributions of this work are as follows:

- We construct discrete physics-based residuals as measures of differences between learned spatiotemporal distributions and those predicted by discretized physical equations governing tumor growth and tissue biomechanics. Using a Lagrangian perspective where particles carry tissue intensities, we project these dynamics onto static Eulerian grids to condition the learnable tumor cell and tissue distributions on available observations. Unlike data-driven models, our hybrid

method remains robust due to physics regularization while being more flexible at adapting data than hard-constrained numerical simulations.

- Our method allows for the estimation of the initial condition of the biomechanical environment. We demonstrate the ability to learn the complex unknown initial condition, which, in our case, is not just the origin of the pathology but also the initial state of the brain tissue anatomy.
- We validate our method* on the downstream task of radiotherapy planning using the largest publicly available dataset known to us, achieving new state-of-the-art performance in capturing tumor recurrence.




Figure 1: Method overview: (a) 3D MRI and PET scans of a glioblastoma patient (b) Preprocessed input includes brain tissue maps, tumor segmentation, and a metabolic map from FET-PET. (c) Tumor cell distribution and brain anatomy inferred using a loss function based on assumptions about physical processes, initial conditions. (d) Outcomes: initial healthy anatomy, spatial tumor cell distribution, and system identification parameters.

Figure 1 provides an overview of our method, showing how we condition the learnable full spatial tumor cell distribution on the patient data and regularize it using soft assumptions about the physics of tumor growth, tissue elasticity, and initial conditions for both the anatomy and the tumor. In addition to multiparametric MRI scans (T1Gd, T1, FLAIR), we incorporate Fluoroethyl-L-tyrosine PET (FET-PET), a specialized form of Positron Emission Tomography (PET), as an additional imaging modality that captures the metabolic activity of tumor cells.

2 Physical Assumptions

Our model integrates physical assumptions to ensure that the learned tumor cell distribution adheres to established biological and mechanical principles. These assumptions are incorporated into the loss function as regularization terms alongside initial condition assumptions and data terms.

2.1 Physical Model

We consider the problem of learning the tumor cell distribution within the domain $\Omega \times (0, 1]$, where $(0, 1]$ represents the assumed time of tumor growth, and the unit cube $\Omega = [0, 1]^3$ describes our spatial domain.

The objective is to determine a tumor cell density function $c(x, t)$, $(x, t) \in \Omega \times (0, 1]$ approximately satisfying the reaction-diffusion-advection equation [30]:

$$\frac{\partial c}{\partial t} = \underbrace{\mathcal{D}c}_{\text{cell migration through diffusion}} + \underbrace{\mathcal{S}c}_{\text{cell proliferation}} - \underbrace{\nabla \cdot (vc)}_{\text{tissue displacements through advection}}, \quad (1)$$

where $\mathcal{D}c = \nabla \cdot (D(\mathbf{m})\nabla c)$ represents the diffusion of the cells with the tissue-dependent diffusion constant $D(\mathbf{m})$, $\mathcal{S}(c) = \rho c(1 - c)$ describes the cell growth with a growth factor ρ , and $\mathbf{v}$ is a

*Code is available at <https://github.com/m1balcerak/PhysRegTumor>

velocity field moving the cells in the interior of the domain due to the tissue's elastic properties and the tumor-induced stresses in the tissue.

The brain tissue vector $\mathbf{m}$, which represents the percentage-wise concentration of different tissue types, is defined as $\mathbf{m}(x, t) = [m_{\text{WM}}(x, t), m_{\text{GM}}(x, t), m_{\text{CSF}}(x, t)]$, where m_{WM} , m_{GM} , and m_{CSF} denote the concentrations of white matter, gray matter, and cerebrospinal fluid, respectively.

The brain tissue vector $\mathbf{m}$ is governed by:

$$\partial_t \mathbf{m} + \text{div}(\mathbf{m} \otimes \mathbf{v}) = 0, \quad (2)$$

where $\otimes$ denotes the outer product. We assume Neumann boundary conditions for both c and $\mathbf{m}$. For c , the no-flux condition applies at the borders of diffusive tissues, specifically the combined white and gray matter. For each tissue component in $\mathbf{m}$, the no-flux boundary is assumed to be the constant brain boundary visible in the MRI scans.

The tumor induced stresses lead to a displacement field $\mathbf{u}$, such that the velocity is given by $\mathbf{v} = \partial_t \mathbf{u}$. We assume quasi-static mechanical equilibrium due to the slow tumor growth rate relative to tissue mechanical responses, leading to:

$$\nabla \cdot \boldsymbol{\sigma}(\mathbf{u}) + \gamma \nabla c = 0, \quad (3)$$

where γ regulates the impact of the tumor on the displacement and is a learnable patient-specific parameter. We apply the Neo-Hookean model [31] to represent hyperelastic behavior in biological tissues, where the stress $\boldsymbol{\sigma}$ is related to the deformation gradient $\mathbf{F}_{ij} = \delta_{ij} + \frac{\partial u_i}{\partial x_j}$ with the Lamé parameters λ and μ (material properties in Appendix G), averaged across different tissue types to accommodate the heterogeneity:

$$\boldsymbol{\sigma} = \frac{\bar{\mu}}{J}(\mathbf{F}\mathbf{F}^T - \mathbf{I}) + \bar{\lambda} \ln(J)\mathbf{I}, \quad (4)$$

with $J = \det(\mathbf{F})$. The averaged Lamé parameters, $\bar{\lambda}$ and $\bar{\mu}$, are computed based on the proportional contributions of constituent tissues, ensuring the stress tensor accurately reflects the composite mechanical properties of the mixed tissues. This methodology facilitates detailed simulations adhering to the elasticity laws relevant to heterogeneous biological environments.

The set of learnable parameters θ_{Dynamics} describing the tumor dynamics, is thus given by $\theta_{\text{Dynamics}} = \{D_{\text{GM}}, R, \rho, \gamma\}$, where D_{GM} and $D_{\text{WM}} = RD_{\text{GM}}$ are the diffusivities in pure gray and white matter regions, which are averaged to the effective diffusivity with a weighted average based on the tissue proportions.

2.2 Discrete Physics Residuals

We use two different approaches to discretize the elasticity equations and the reaction-diffusion-advection equations. For the former, we use a time-dependent grid in space, described by the N_x points $(\mathbf{p}_j^n)_{1 \leq j \leq N_x}$ at time $n\Delta t$, $0 \leq n \leq N_t$, which initially form a uniform grid on Ω . For the latter, we further partition the domain Ω at the n th point in time into N_x cells Ω_i^n , $1 \leq i \leq N_x$, such that $\Omega = \bigcup_{i=1}^{N_x} \Omega_i^n$. Note that the nodes defining the cell boundaries are carried by the particles $\mathbf{p}_j^n$, whose movement, governed by tissue elasticity, gradually deforms the cells Ω_i^n from their initial rectangular shape.

The displacement field $\mathbf{u}^n$ is computed based on the initial ($t = 0$) and current ($t = n\Delta t$) positions of particles $\mathbf{p}^n$ within the tissue by $\mathbf{u}^n = \mathbf{p}^n - \mathbf{p}^0$. Using central finite differences on the initial uniform reference grid, we can consecutively approximate $\nabla \mathbf{u}_j^n$, $\mathbf{F}_j^n$, ∇c_j^n , $\nabla \cdot \boldsymbol{\sigma}_j^n$, and define the tissue residual:

$$L_{\text{Tissue Elasticity PDE}} = \sum_{n=0}^{N_t} \sum_{j=0}^{N_x} (\nabla \cdot \boldsymbol{\sigma}_j^n + \gamma \nabla c_j^n)^2. \quad (5)$$

As the particles, representing grid points, move with the tissue's deformation, they inherently accommodate the advection. Thus, the observed displacement of tumor cells attributed to advection is absorbed into the changes in particle positions. This leads to a simplification of the tumor equation when solved in the moving reference frame of the particles.

We use an implicit Euler scheme combined with a cell-centered finite-volume discretization to discretize the tumor equation. We denote the boundaries between two cells $1 \leq i, j \leq N_x$ by $\Gamma_{ij}^n = \Gamma_{ji}^n = \partial\Omega_i^n \cap \partial\Omega_j^n$. Further, the index set of all neighboring cells of Ω_i^n is given by $N_i = \{j : \partial\Omega_j \cap \partial\Omega_i \neq \emptyset \wedge j \neq i\}$ and assumed to stay independent of n , due to the extend of our deformations (see limitations in Section 7). For a function c_i^n at time $n\Delta t$ at the i th cell, we discretize the diffusion $\mathcal{D}$ and reaction $\mathcal{S}$ operators by:

$$\mathcal{D}[c_i^n, D_i^n] = \sum_{j \in N_i} |\Gamma_{ij}| D_{ij}^n \frac{c_j^n - c_i^n}{\|\mathbf{x}_j^n - \mathbf{x}_i^n\|}, \text{ and } \mathcal{S}[c_i^n] = |\Omega_i^n| \rho c_i^n (1 - c_i^n), \quad (6)$$

where D_{ij}^n is the harmonic mean of D_i^n at the cell boundary, $|\Gamma_{ij}^n|$ denotes the interface area, and $|\Omega_i^n|$ denotes the cell volume. We thus obtain the tumor loss by:

$$L_{\text{Tumor Growth PDE}} = \sum_{n=1}^{N_t} \sum_{i=0}^{N_x} \left(\frac{|\Omega_i^n|}{\Delta t} (c_i^n - c_i^{n-1}) - \mathcal{D}[c_i^n, D_i^n] - \mathcal{S}[c_i^n] \right)^2. \quad (7)$$




Figure 2: Learning process overview: optimization of spatiotemporal distributions of tumor cells and tissues. (a) Initial condition penalties enforce symmetric healthy anatomy, with the initial tumor distribution at $t = 0$ as a small Gaussian blob. (b) Physics penalties regularize dynamics between the initial and the final time. The first row shows gray matter contours with particle positions; the second row shows white matter contours with learned tumor concentrations. (c) Agreement of tumor distribution with anatomical tissues, visible tumor segmentations, and metabolic map after transforming the final tumor distribution through the imaging function.

2.3 Particle-Grid Projections

While a Lagrangian frame representation from the point of view of particles is beneficial for accurately modeling advection, an Eulerian frame is essential for rendering the advected particle states into the image space. Denoting the Eulerian and Lagrangian frames as G and P respectively, an Eulerian field $\mathcal{F}^G$ and a Lagrangian field $\mathcal{F}^P$ are related through:

$$\mathcal{F}_j^P \approx \sum_i w_{ij} \mathcal{F}_i^G, \quad \mathcal{F}_i^G \approx \frac{\sum_j w_{ij} \mathcal{F}_j^P}{\sum_j w_{ij}}, \quad (8)$$

where i indexes grid nodes, j indexes particles, and w_{ij} represents the weight of the trilinear shape function defined on node i and evaluated at the location of particle j .

3 Initial Assumptions

We model the initial conditions of the tumor and brain tissues by incorporating assumptions into our loss functions, ensuring the model aligns with the expected biological and anatomical state of the brain before and at the onset of tumor growth.

3.1 Initial Tumor Distribution

We assume that the initial tumor distribution can be represented by a small Gaussian blob. To enforce this assumption, we penalize the difference between $\mathcal{F}^G(u)^0$ and a Gaussian function centered at $x_0 \in \Omega$. The center of this Gaussian blob, x_0 , is part of the learnable parameters $\theta_{\text{Initial}} = \{x_0, \mathbf{m}^0\}$. The initial condition loss for the tumor cells is defined as:

$$L_{\text{Tumor starts from a Gaussian}} = \sum_{i=0}^{N_x} \left(\mathcal{F}^G(c)_i^0 - D_1 \exp \left(-\frac{(x_i - x_0)^2}{D_2} \right) \right)^2, \quad (9)$$

where $D_1 = 0.5$, $D_2 = 0.02$ are constants and x_i is the position of the center of Ω_i^0 .

3.2 Healthy Anatomy

The deformations are clearly visible in the MRI images and are an integral part of understanding growth dynamics. To obtain a plausible estimate of the unperturbed brain, we need to incorporate an anatomical prior.

A healthy brain is roughly symmetric between hemispheres, as justified by anatomical and functional observations. Imaging techniques such as MRI reveal symmetrical tissue structures, and electroencephalography (EEG) recordings show symmetrical patterns of electrical activity in the brain [32, 33]. This assumption applies to key tissue types, including white matter, gray matter, and cerebrospinal fluid. We construct a loss function, denoted as $L_{\text{Healthy Anatomy}}$, to quantify the asymmetry of the learned $\mathcal{F}^G(\mathbf{m}^0)$. The implementation details are provided in Appendix E. We do not use the unperturbed anatomy explicitly for inference of the tumor cells, however, we need additional constraints to bound γ from Eq. 3 that couples tumor cells with tissue dynamics. While symmetry of the healthy brain is a coarse assumption, we soften it by quantifying the asymmetry in lower resolutions.




Figure 3: Inference overview: (a) Patient's MRI 3D scans. (b) Estimated tissues through non-rigid registration of the average brain, showing white matter and visible tumor segmentations. (c) Learned tumor cell distribution, regularized by the physics residual and aligned with patient data. (d) Learned initial condition of the tissues representing healthy anatomy. (e) Average brain template for reference, rigidly registered to the MRI scan. Notable anatomical differences include the lack of matter passage between the hemispheres next to the tumor, which could affect tumor cell inference results.

4 Imaging Data Reconstruction

Imaging data refers to information directly observed from the patient's MRI and FET-PET 3D scans [34], including visible tumor parts and surrounding brain tissues. As shown in Figure 1, this data includes estimated brain tissue probability maps, visible tumor segmentation, and a metabolic map derived from FET-PET scans, focusing on the enhancing core and edema. MRI scans provide segmentation maps [35], revealing the visible tumor. Using non-rigid atlas image registration with

tumor masking [34], we estimate tissues both outside and hidden by the tumor, including white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF). These observations offer partial information about both the anatomy and the tumor cell distribution.

The imaging function (see Figure 2) [22, 36] bridges simulated tumor cell densities with visible tumor segmentation and brain tissues, as well as the metabolic map. It depends on parameters for thresholding the tumor cell distribution to obtain segmentation, denoted as θ_{Imaging} , which must be learned.

We construct a loss function that minimizes the difference between model predictions and patient observables. Specifically, L_{Core} matches the visible tumor core, L_{Edema} matches the visible edema, and $L_{\text{Metabolic Map}}$ matches the metabolic map profile. Tissue intensities carried by particles are set to reproduce the patient's estimated tissues at $t = 1$. While these intensities are predetermined, particle positions are weakly constrained, so matching the tissues (WM, GM, CSF) to the observed structures outside the tumor is enforced through L_{Tissues} . The constructed penalty terms align our spatiotemporal distributions of tumor cells and tissues with patient data characteristics. Implementation details of the data losses and the imaging function are provided in Appendix D.

5 Combined Loss Function

The loss function $\mathcal{L}_{\text{Total}}$ is constructed to balance the contributions of physics-based regularization, assumptions about the initial conditions, and the imaging empirical data. This facilitates a comprehensive approach to infer the full spatial tumor cell distribution at the time of the patient's data acquisition (see Figure 3 for an overview of the inference process and Figure 2 for the learning process visualization):

$$\begin{aligned} \mathcal{L}_{\text{Total}} = & \underbrace{\alpha_1 L_{\text{Tumor Growth PDE}} + \alpha_2 L_{\text{Tissue Elasticity PDE}}}_{\text{Physical Assumptions}} + \\ & \underbrace{\alpha_3 L_{\text{Tumor Starts From a Gaussian}} + \alpha_4 L_{\text{Healthy Anatomy}}}_{\text{Initial Assumptions}} + \\ & \underbrace{\alpha_5 L_{\text{Core}} + \alpha_6 L_{\text{Edema}} + \alpha_7 L_{\text{Metabolic Map}} + \alpha_8 L_{\text{Tissues}}}_{\text{Imaging Data Reconstruction}}. \end{aligned} \quad (10)$$

All α_* parameters were determined based on experiments with synthetic data involving both single focal and local multi-focal tumors. Extensive ablation studies are provided in Appendix C, and the details of the synthetic experiment setup can be found in Appendix F.

6 Validation and Ablation Study of Radiotherapy Planning

There is no objective ground truth for pre-operative tumor cell distribution, so methods are evaluated through downstream tasks like defining radiation target volumes. An "accurate" method targets areas that later show progression in follow-up MRI, acknowledging that radiation delays but does not eliminate tumor regrowth. An "inaccurate" method spares these areas or targets regions that remain tumor-free. The current "Standard Plan" defines the target volume using a 1.5 cm isosurface around the MRI-segmented tumor core, while our approach generates isosurfaces from learned tumor cell distributions. Figure 4 visualizes the Standard Plan and a plan based on our learned tumor cell distribution. Each model's target volume is simply a redistribution of the target volume of the Standard Plan.

We assess accuracy by measuring the overlap between the target volumes and regions of later tumor progression, calculating the *Recurrence Coverage* [%]. This metric represents the percentage of the tumor progression region that falls within the radiotherapy target volume of a given model, after rigid alignment of both scans.

Additional experiments studying properties of the tumor growth and imaging model on synthetic data, i.e., without relying on indirect validation via radiotherapy planning, are given in Appendix B.




Figure 4: Radiotherapy planning: a) Pre-operative segmentations with white matter concentration as background. b) Distance map from the tumor core segmentation with a 1.5 cm contour and within diffusive tissue, constituting the Standard Plan. c) Our learned tumor cell distribution with the isosurface contour where the enclosed volume equals the total volume of the Standard Plan.

6.1 Implementation details

Our computational model uses a multi-resolution method [37] with additional, coarser grids to accelerate convergence. We employ four levels of grid refinement, resulting in 152,880,048 unknowns. At the finest grid level of $72 \times 72 \times 72 \times 96$, each grid point contains four unknowns: three for particle positions and one for tumor cell density.

The learning process includes system identification of unknowns: θ_{Dynamics} , θ_{Initial} , and θ_{Imaging} , which parameterize tumor growth dynamics, tissue elasticity, initial conditions, and imaging characteristics. Using the Adam optimizer, convergence takes around 3 hours on an NVIDIA RTX A6000 GPU.

6.2 Dataset

We use a publicly available dataset of 58 patients with preoperative MRI scans and preoperative FET-PET imaging, along with follow-up MRI scans at the time of the first visible tumor recurrence [22].

We would like to emphasize that this is the largest publicly available dataset of this kind, i.e., with multi-modal imaging data and follow-up MRI. Its size is common for medical studies and sufficient to motivate clinical trials (see, for example, [38, 39, 40] on clinical investigations into the use of PET for radiotherapy planning). The problem we address is a critical challenge in cancer treatment research that calls for ML solutions. However, data acquisition is costly and time-consuming, and clinical datasets, such as ours with only a few dozen complete observations, are insufficient to support purely data-driven ML approaches (as visible from results reported in Table 1). At the same time, it is this inherent limitation that is encouraging us to incorporate physics-based assumptions and physics-informed regularizations into an ML model to effectively address the driving clinical problem.

The follow-up images with visible tumor recurrences are not used by our method during the optimization or hyper-parameter search at any point; therefore, the dataset can be considered 100% unseen. The decision not to use any recurrence data for hyper-parameter tuning is to avoid diluting the already small dataset.

6.3 Models and Features

Our analysis includes a range of models, each categorized by their distinct capabilities:

- **Numerical Physics Simulations:** Employ methods such as Finite Element Method (FEM) and Finite Difference Method (FDM) to model tumor growth dynamics and account for dynamic tissue behaviors. The dynamic parameters are determined through numerous simulations with varying parameters. The brain tissues' initial conditions are based on average brains.

- **Data-Driven Neural Networks (Unconstrained):** Use Convolutional Neural Networks (CNNs) to directly predict likely recurrence locations, representing a population-based data-driven approach.
- **Data-Driven Neural Networks (Physics-Constrained):** Similar to the above unconstrained approach but limited to finding parameters for physics simulations, incorporating hard physics constraints.
- **Static Grid Discretization:** A previous state-of-the-art method that uses Optimizing a Discrete Loss (ODIL) and static grids to penalize the learned spatiotemporal tumor cell distribution by quantifying the discrepancy with the tumor growth equation, keeping brain anatomy static throughout the entire process.
- **Standard Plan:** Apply uniform safety margins around the tumor core, serving as the baseline for volume determination and reflecting current clinical practice post-tumor resection surgery radiotherapy planning.

Specific implementations within these categories are explained in Appendix A.

6.4 Results

Table 1, specifically the "Recurrence Coverage [%] (Any)" column, and Figure 5a present results where recurrence is defined by any segmentation visible in the follow-up MRI scan. Our method achieved 74.7% recurrence coverage, surpassing the previous state-of-the-art at 72.9% and the Standard Plan at 70.0%. Using a predefined average brain reduced the performance to 73.4%. The physics-constrained data-driven approach slightly outperforms the Standard Plan at 67.1% and significantly outperforms its unconstrained variant at 59.0%. This highlights the advantage of adding physics constraints. Figure 5a shows the distance between 'Greater' and 'Less' categories increased from 26% to 36%.

Table 1, using the "Recurrence Coverage [%] (Enhancing Core)" column, and Figure 5b show results where recurrence is defined by enhancing core segmentation, per RANO guidelines [41]. Our method leads with 89.9% average recurrence coverage, compared to 89.0% for static grid discretization and 87.3% for the Standard Plan. Numerical physics simulations are more robust at 86.2% compared to 84.3% for the physics-constrained data-driven method. The unconstrained data-driven method remains the lowest at 66.8%. Results are close because the enhancing core often occurs near pre-operative MRI segmentations, leading to many methods achieving 100% coverage (seen in 'Equal' in Figure 5b). However, Figure 5b shows an increase in the preference from 19% to 28%.

In both recurrence definitions, mean estimate uncertainty is substantial, as Recurrence Coverage for individual patients can vary significantly. These uncertainties are the standard errors of the mean. Nonetheless, Figure 5 shows that even slight improvements in average Recurrence Coverage consistently enhance radiotherapy outcomes compared to the Standard Plan baseline.

Table 1: Comparison of recurrence segmentation coverage given equal radiation volume.

Model	Recurrence Coverage[%] (Any)	Recurrence Coverage[%] (Enhancing Core)	Dynamical Tissues	Inferable Healthy Anatomy	Population-Based Data-Driven	Physics-Constrained
NN (Unconstrained)	59.0 ± 4.3	66.8 ± 4.9	×	×	✓	×
NN (Physics-Constrained)	70.4 ± 3.7	84.3 ± 3.3	×	×	✓	Hard
Numerical Physics Simulations	67.1 ± 3.8	86.2 ± 3.6	✓	×	×	Hard
Standard Plan	70.0 ± 3.8	87.3 ± 3.6	×	×	×	×
Static Grid Discretization	72.9 ± 3.5	89.0 ± 3.3	×	×	×	Soft
Ours (w/o inferable anatomy)*	73.4 ± 3.2	89.3 ± 2.9	✓	×	×	Soft
Ours	74.7 ± 3.1	89.9 ± 2.7	✓	✓	×	Soft

*We set the initial condition for the brain tissues as a hard constraint, representing the average brain.




Figure 5: Direct patient-by-patient comparisons to the Standard Plan of radiotherapy plans with equal total volumes: "Greater," "Equal," and "Less" refer to the direct numerical comparison of Recurrence Coverage. a) Recurrence is defined as the union of edema, enhancing core, and necrotic core on the follow-up MRI segmentation. b) Recurrence is defined as the enhancing core on the follow-up MRI segmentation.

As shown in Figure 5, the gap between the "Greater" and "Less" categories represents the number of patients who benefit from our method compared to the standard plan. We observe a consistent improvement with our method, demonstrating superior performance beyond just higher mean recurrence coverage, which could be skewed by outliers. The results are statistically significant, as shown in Table 3 in Appendix B.

7 Conclusion

In this work, we present the first successful approach to a soft-constrained physics system-identification problem that combines the complex biological process of tumor growth with the biomechanical environment such as elastic brain tissues. Our method integrates physics-based constraints with multi-modal imaging data to enhance tumor treatment planning for glioblastomas. The method balances adherence to observed patient data and physics-based penalties through a unique discretization scheme, serving as a flexible spatiotemporal regularization term.

Additionally, our method provides estimates of the initial condition, i.e., tumor-free, pre-deformation anatomy. Such healthy brain anatomy can be utilized in various studies, e.g., for aligning with post-surgical scans to assess the extent of tumor removal and detect complications [42, 43, 44].

Overall, our results show that this approach outperforms previous state-of-the-art techniques in covering tumor recurrence areas, with improved performance on real-world patient data. Finally, we want to point out that relaxing physical model constraints has broad applicability beyond glioblastoma localization. This approach can extend to other real-world problems where rigid physics-based models are limited but still relevant.

Limitations: The convergence speed of the method depends on a multiresolution grid, assuming that problems can be represented at lower resolutions. However, this approach may not be directly applicable to PDEs with narrow-phase interfaces. Additionally, while large deformations are allowed, we assume they are not large enough to cause material rupture, which may not be optimally handled in the Lagrangian perspective without modifications to the presented scheme.

Acknowledgments and Disclosure of Funding

We extend our gratitude to John Lowengrub and Ray Zirui Zhang from the University of California, Irvine, for their invaluable insights and discussions. This research was supported by the Helmut Horten Foundation and the DCoMEX project.

References

- [1] Guido Frosina. Radiotherapy of high-grade gliomas: dealing with a stalemate. *Critical Reviews in Oncology/Hematology*, 190:104110, Oct 2023.
- [2] Ender Konukoglu, Olivier Clatz, Bjoern H. Menze, Bram Stieltjes, Marc-André Weber, Emmanuel Mandonnet, Hervé Delingette, and Nicholas Ayache. Image guided personalization of reaction-diffusion type tumor growth models using modified anisotropic eikonal equations. *IEEE Transactions on Medical Imaging*, 29(1):77–95, Jan 2010.

- [3] Jan Unkelbach, Bjoern H Menze, Ender Konukoglu, Florian Dittmann, Matthieu Le, Nicholas Ayache, and Helen A Shih. Radiotherapy planning for glioblastoma based on a tumor growth model: improving target volume delineation. *Physics in Medicine & Biology*, 59(3):747, 2014.
- [4] Russell Rockne, E.C. Alvord, J.K. Rockhill, and K.R. Swanson. A mathematical model for brain tumor response to radiation therapy. *Journal of Mathematical Biology*, 58(4):561–578, 2009.
- [5] Cristian Hoguea, Christos Davatzikos, and George Biros. An image-driven parameter estimation problem for a reaction–diffusion glioma growth model with mass effects. *Journal of Mathematical Biology*, 56:793–825, 2008.
- [6] Vittorio Cristini and John Lowengrub. *Multiscale Modeling of Cancer: An Integrated Experimental and Mathematical Modeling Approach*. Cambridge University Press, 2010.
- [7] S. Subramanian, A. Gholami, and G. Biros. Ensemble inversion for brain tumor growth models with mass effect. *IEEE Transactions on Medical Imaging*, 42(4):982–995, 2023.
- [8] J. Lipkova, P. Angelikopoulos, S. Wu, E. Alberts, B. Wiestler, C. Diehl, C. Preibisch, T. Pyka, S.E. Combs, P. Hadjidakas, K. Van Leemput, P. Koumoutsakos, J. Lowengrub, and B. Menze. Personalized radiotherapy design for glioblastoma: Integrating mathematical tumor models, multimodal scans, and bayesian inference. *IEEE Transactions on Medical Imaging*, 38(8):1875–1884, 2019.
- [9] Shashank Subramanian, Klaudius Scheufele, Miriam Mehl, and George Biros. Where did the tumor start? an inverse solver with sparse localization for tumor growth models. *Inverse problems*, 36(4):045006, 2020.
- [10] Jonas Weidner, Ivan Ezhov, Michal Balcerak, Marie-Christin Metz, Sergey Litvinov, Sebastian Kaltenbach, Leonhard Feiner, Laurin Lux, Florian Kofler, Jana Lipkova, Jonas Latz, Daniel Rueckert, Bjoern Menze, and Benedikt Wiestler. A learnable prior improves inverse tumor growth modeling. *arXiv preprint arXiv:2403.04500*, 2024.
- [11] Olivier Clatz, Maxime Sermesant, Pierre-Yves Bondiau, Hervé Delingette, Simon K. Warfield, Grégoire Malandain, and Nicholas Ayache. Realistic simulation of the 3-d growth of brain tumors in mr images coupling diffusion with biomechanical deformation. *IEEE Transactions on Medical Imaging*, 24(10):1334–1346, Oct 2005.
- [12] Jens Petersen, Fabian Isensee, Gregor Köhler, Paul F Jäger, David Zimmerer, Ulf Neuberger, Wolfgang Wick, Jürgen Debus, Sabine Heiland, Martin Bendszus, et al. Continuous-time deep glioma growth models. In *MICCAI 2021: 24th International Conference, Strasbourg, France*, pages 83–92. Springer, 2021.
- [13] Ivan Ezhov, Kevin Scibilia, Katharina Frantiza, Felix Steinbauer, Suprosanna Shit, Lucas Zimmer, Jana Lipkova, Florian Kofler, Johannes Paetzold, Luca Canalini, Diana Waldmannstetter, Martin J. Menten, Marie Metz, Benedikt Wiestler, and Bjoern Menze. Learn-morph-infer: a new way of solving the inverse problem for brain tumor modeling. *Medical Image Analysis*, page 102672, 2022.
- [14] Alfredo De Goyeneche Macaya, Shreya Ramachandran, Ke Wang, Ekin Karasan, Joseph Y. Cheng, Stella X. Yu, and Michael Lustig. Resonet: Noise-trained physics-informed mri off-resonance correction. *Advances in Neural Information Processing Systems*, 36, 2024.
- [15] Ramansh Sharma and Varun Shankar. Accelerated training of physics-informed neural networks (pinns) using meshless discretizations. *Advances in Neural Information Processing Systems*, 35, 2022.
- [16] Woojin Cho, Kookjin Lee, Donsub Rim, and Noseong Park. Hypernetwork-based meta-learning for low-rank physics-informed neural networks. *Advances in Neural Information Processing Systems*, 36, 2023.
- [17] Junwoo Cho, Seungtae Nam, Hyunmo Yang, Seok-Bae Yun, Youngjoon Hong, and Eunbyung Park. Separable physics-informed neural networks. *Advances in Neural Information Processing Systems*, 36, 2023.

- [18] Aditi Krishnapriyan, Amir Gholami, Shandian Zhe, Robert Kirby, and Michael W. Mahoney. Characterizing possible failure modes in physics-informed neural networks. *Advances in Neural Information Processing Systems*, 34, 2021.
- [19] Tsung-Yen Yang, Justinian Rosca, Karthik Narasimhan, and Peter J Ramadge. Learning physics constrained dynamics using autoencoders. *Advances in Neural Information Processing Systems*, 35, 2022.
- [20] Karolis Martinkus, Jan Ludwiczak, Wei-Ching Liang, Julien Lafrance-Vanasse, Isidro Hotzel, Arvind Rajpal, Yan Wu, Kyunghyun Cho, Richard Bonneau, Vladimir Gligorijevic, et al. Abdifuser: full-atom generation of in-vitro functioning antibodies. *Advances in Neural Information Processing Systems*, 36, 2024.
- [21] Petr Karnakov, Sergey Litvinov, and Petros Koumoutsakos. Solving inverse problems in physics by optimizing a discrete loss: Fast and accurate learning without neural networks. *PNAS Nexus*, 3(1):pgae005, Jan 2024.
- [22] 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.
- [23] Ray Zirui Zhang, Ivan Ezhov, Michal Balcerak, Andy Zhu, Benedikt Wiestler, Bjoern Menze, and John Lowengrub. Personalized predictions of glioblastoma infiltration: Mathematical models, physics-informed neural networks and multimodal scans. *arXiv preprint arXiv:2311.16536*, 2024.
- [24] Ben Mildenhall, Pratul P. Srinivasan, Matthew Tancik, Jonathan T. Barron, Ravi Ramamoorthi, and Ren Ng. Nerf: Representing scenes as neural radiance fields for view synthesis. In *European Conference on Computer Vision*, pages 405–421. Springer, 2020.
- [25] Xuan Li, Yi-Ling Qiao, Peter Yichen Chen, Krishna Murthy Jatavallabhula, Ming Lin, Chenfanfu Jiang, and Chuang Gan. PAC-nerf: Physics augmented continuum neural radiance fields for geometry-agnostic system identification. In *The Eleventh International Conference on Learning Representations*, 2023.
- [26] Bernhard Kerbl, Georgios Kopanas, Thomas Leimkühler, and George Drettakis. 3d gaussian splatting for real-time radiance field rendering. *ACM Transactions on Graphics*, 42(4):1–14, 2023.
- [27] Cheng Sun, Min Sun, and Hwann-Tzong Chen. Direct voxel grid optimization: Super-fast convergence for radiance fields reconstruction. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 5459–5469, 2022.
- [28] Federica Caforio, Francesco Regazzoni, Stefano Pagani, Elias Karabelas, Christoph Augustin, Gundolf Haase, Gernot Plank, and Alfio Quarteroni. Physics-informed neural network estimation of material properties in soft tissue nonlinear biomechanical models. *arXiv preprint arXiv:2312.09787*, 2023.
- [29] Arunabha M. Roy and Rikhi Bose. Physics-aware deep learning framework for linear elasticity. *arXiv preprint arXiv:2302.09668*, 2023.
- [30] S. Subramanian, A. Gholami, and G. Biros. Simulation of glioblastoma growth using a 3d multispecies tumor model with mass effect. *Journal of Mathematical Biology*, 79(3):941–967, 2019.
- [31] Raymond W. Ogden. *Non-linear Elastic Deformations*. Dover Publications, Mineola, New York, 1997.
- [32] Arthur W. Toga and Paul M. Thompson. Mapping brain asymmetry. *Nature Reviews Neuroscience*, 4(1):37–48, 2003.
- [33] Meng-Chuan Lai et al. A shift to randomness of brain oscillations in people with autism. *Biological Psychiatry*, 71(5):410–417, 2012.

- [34] Jana Lipková, Panagiotis Angelikopoulos, Stephen Wu, Esther Alberts, Benedikt Wiestler, Christian Diehl, Christine Preibisch, Thomas Pyka, Stephanie E. Combs, Panagiotis Hadji-doukas, 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.
- [35] Florian Kofler, Christoph Berger, Diana Waldmannstetter, Jana Lipkova, Ivan Ezhov, Giles Tetteh, Jan Kirschke, Claus Zimmer, Benedikt Wiestler, and Bjoern H. Menze. Brats toolkit: translating brats brain tumor segmentation algorithms into clinical and scientific practice. *Frontiers in Neuroscience*, page 125, 2020.
- [36] Bjoern H. Menze, Koen Van Leemput, Antti Honkela, Ender Konukoglu, Marc-André Weber, Nicholas Ayache, and Polina Golland. A generative approach for image-based modeling of tumor growth. *Information Processing in Medical Imaging*, 22:735–747, 2011.
- [37] Petr Karnakov, Sergey Litvinov, and Petros Koumoutsakos. Flow reconstruction by multiresolution optimization of a discrete loss with automatic differentiation. *The European Physical Journal E*, 46(7):59, Jul 2023.
- [38] Stefan Rieken, Daniel Habermehl, Frederik L. Giesel, and et al. Radiation-induced brain tissue damage after radiotherapy for glioma: a problem of growing concern. *Radiotherapy and Oncology*, 109(3):362–367, 2013.
- [39] Per M. Munck Af Rosenschold, Junia Costa, Svend Aage Engelholm, and et al. Impact of [18f]-fluoro-ethyl-tyrosine pet imaging on target definition for radiation therapy of high-grade glioma. *Neuro-Oncology*, 17(5):757–763, 2015.
- [40] Anca-Ligia Grosu, Michael Souvatzoglou, Beate Roper, and et al. Hypoxia imaging with faza-pet: A new approach for target definition in glioblastoma multiforme. *International Journal of Radiation Oncology, Biology, Physics*, 63(2):517–526, 2005.
- [41] Benjamin M. Ellingson, Patrick Y. Wen, and Timothy F. Cloughesy. Modified criteria for radiographic response assessment in glioblastoma clinical trials. *Neurotherapeutics*, 14(2):307–320, Apr 2017.
- [42] Adam T. Kessler and Ami A. Bhatt. Brain tumour post-treatment imaging and treatment-related complications. *Insights into Imaging*, 9(6):1057–1075, Dec 2018.
- [43] Barbara Kiesel, Romana Prihoda, Martin Borkovec, Petra A. Mercea, Ariane Steindl, Anna S. Berghoff, Julia Furtner, Johannes Leitner, Thomas Roetzer, Matthias Preusser, Karl Roessler, and Georg Widhalm. Postoperative magnetic resonance imaging after surgery of brain metastases: Analysis of extent of resection and potential risk factors for incomplete resection. *World Neurosurgery*, 143:e365–e373, 2020.
- [44] Michael W. Gerber et al. Assessing the extent of resection in glioma surgery using intraoperative mri and functional mapping. *Journal of Neurosurgery*, 128(1):87–98, 2018.
- [45] Shashank Subramanian, Amir Gholami, and George Biros. Simulation of glioblastoma growth using a 3d multispecies tumor model with mass effect. *Journal of Mathematical Biology*, 79:941–967, 2019.
- [46] Sergio M. Martin, Daniel Wälchli, Georgios Arampatzis, Athena E. Economides, Petr Karnakov, and Petros Koumoutsakos. Koral: Efficient and scalable software framework for bayesian uncertainty quantification and stochastic optimization. *Computer Methods in Applied Mechanics and Engineering*, 389:114264, 2022.
- [47] Fabian Isensee, Paul F. Jaeger, Simon A.A. Kohl, Jens Petersen, and Klaus H. Maier-Hein. nnu-net: a self-configuring method for deep learning-based biomedical image segmentation. *Nature Methods*, 18:203–211, 2021.
- [48] Michael Weller. Eano guidelines on the diagnosis and treatment of diffuse gliomas of adulthood. *Nature Reviews Clinical Oncology*, 18:170–186, 2021.
- [49] Maximilian Niyazi, Nicolaus Andratschke, Martin Bendszus, Anthony J Chalmers, Sara C Erridge, et al. Estro-eano guideline on target delineation and radiotherapy details for glioblastoma. *Radiotherapy and Oncology*, 184:109663, 2023.

A Baselines Implementation

Here we provide details regarding the specific baseline implementations:

Numerical Physics Simulations: We employ the numerical forward simulation scheme from [45] which was implemented on GPU [13]. The simulations follows the equation 1. We initialize the brain tissues with a rigidly registered average brain. For θ_{Dynamics} , $x_0 \in \theta_{\text{Initial}}$, and θ_{Imaging} , we use the Covariance Matrix Adaptation Evolution Strategy (CMA-ES) method [10, 46].

Data-Driven Neural Networks (Unconstrained): We use nnU-Net [47] which takes brain tissue distribution, FET-PET map, and pre-operative tumor segmentations to directly predict core tumor recurrences. After predicting the recurrences, we construct a distance distribution from the predicted recurrences. This distribution is then thresholded, similar to other methods, to create a binary radiotherapy treatment map of the same volume as the Standard Plan. The network is trained using an 80%/20% train-test split.

Data-Driven Neural Networks (Physics-Constrained): We employ a recently published method [13] that uses CNNs to map pre-operative tumor segmentation to θ_{Dynamics} and runs a numerical simulation following equation 1. This method uses an average brain as the initial anatomy. After inference, the tumor cell distribution is non-rigidly registered to the patient's brain anatomy.

Static Grid Discretization: This method, as implemented in [22], uses estimated brain tissues (4) as the initial condition for the anatomy and later keeps the tissues static, effectively setting $\gamma \in \theta_{\text{Dynamics}}$ in Equation 3 to 0.

Standard Plan: This approach uses a 1.5 cm uniform safety margin around the resection cavity and/or remaining tumor on MRI imaging, aligning with both North American and European guidelines [48, 49].

B Synthetic Results

To show how the method performs on tasks where there is ground truth, our method and GliODIL [16] (previous SOTA, Static Grid) are additionally evaluated on synthetic cases generated by numerical physics solvers under various conditions. See Table 2 for the results and Table 3 for statistical tests.

Input Type	RMSE (1 Focal Origin) [%]		RMSE (3 Focal Origins) [%]	
	Ours	Static Grid [16]	Ours	Static Grid [16]
Tumor Core	10.2 $\pm$ 0.4	12.2 $\pm$ 0.6	13.3 $\pm$ 0.5	13.0 $\pm$ 0.7
Tumor Core, Edema, and Metabolic Map	2.5 $\pm$ 0.3	5.5 $\pm$ 0.8	4.0 $\pm$ 0.5	6.9 $\pm$ 0.5
Tumor Core, Edema, and Metabolic Map, $\theta_{\text{Initial}}, \theta_{\text{Dynamics}}$	0.3 $\pm$ 0.1	3.3 $\pm$ 0.2	0.4 $\pm$ 0.1	4.2 $\pm$ 0.3

Table 2: Synthetic results comparing RMSE for 30 patients with different input types for one and three tumor origins, using our method and Static Grid Discretization (GliODIL [16]). The test dataset was not used for hyper-parameter tuning. Parameter ranges in Appendix D and $\gamma \in [0, 1.5]$ from Eq. 3 in the manuscript. Ours wins because GliODIL cannot deform brain tissues through its static grid limitations.

Comparison	Recurrence (Any)	Recurrence (Enhancing Core)
	p-value	p-value
Ours vs Standard Plan (1.5cm)	0.014	0.00082
Ours vs Static Grid Discretization (GliODIL[16])	0.043	0.039

Table 3: Statistical significance tests comparing recurrence coverage (any and enhancing core) among different models using the Wilcoxon signed-rank test. The data is not normally distributed which we tested using Shapiro-Wilk test resulting in p-values below 1E-12.

C Extended Ablation Study

Additional information regarding the loss coefficients:

α_2 is essential for grid stability. Without it, the grid would be unstable, causing folding and rendering the results meaningless. If set to zero, α_2 can be replaced by a Laplacian-based regularization, as visualized in Figure 6b and Figure 6c, though this clearly results in less biologically plausible deformations.

α_4 regularizes the initial anatomy to bound γ from Eq 3. It should be kept low due to the natural asymmetry in brains. More details are provided in Appendix E.

α_5 , α_6 , α_7 , and α_8 are crucial for grounding PDE solutions in reality (most important, all other losses can be seen simply as regularisation of these components). Imaging data reconstruction loss cannot be calibrated using synthetic studies as all the other alphas. Instead of using recurrence data to calibrate them, we set them such that they provide a balanced contribution to the total loss

α_3 should be kept low since the initial tumor shape is unknown. However, decreasing it too much causes a local minimum with no tumor.

α_1 serves as the main regularization term. It should not be $\rightarrow \infty$ to avoid solutions not flexible enough to properly accommodate patient data (see Numerical Physics Simulations results in Table 1; this setup performs poorly), nor should it be $\rightarrow 0$ to prevent overfitting to imaging data alone, effectively collapsing to standard plan level performance.

Removing elements of the imaging data reconstruction loss introduces ill-posedness and degradation of synthetic results (see Table 4). Removing tumor growth physics regularizations effectively reduces our solution to performance close to the standard plan. Removing particle dynamics (static grid) reduces our model to [16] Static Grid Discretization (“GliODIL”).




Figure 6: Deformation through particle movement: a) Total loss of our model for reference. b) Particle position with tissue elasticity. c) Particle position without tissue elasticity. These particle positions are learned for the pre-operative time. For grid stability in c), additional regularization is required. For $\alpha_2 > 0$ we observe one large localized deformation caused by the tumor, which is more biologically plausible than multiple smaller deformations visible for $\alpha_2 = 0$.

Model Variation	Recurrence (Any)		Recurrence (Enhancing Core)	
	Coverage [%]	IoU*	Coverage [%]	IoU*
Ours, $\alpha_5 = 0, \alpha_6 = 0, \alpha_7 = 0, \alpha_8 = 0$	60.5 ± 4.9	0.15 ± 0.02	83.5 ± 3.7	0.05 ± 0.01
Ours, $\alpha_6 = 0, \alpha_7 = 0$	63.5 ± 4.0	0.15 ± 0.03	86.5 ± 3.6	0.06 ± 0.02
Ours, $\alpha_3 = 0$	70.0 ± 3.8	0.16 ± 0.02	87.4 ± 3.4	0.06 ± 0.02
Standard Plan (1.5cm)	70.0 ± 3.8	0.16 ± 0.02	87.3 ± 3.6	0.04 ± 0.01
Ours, $\alpha_1 = 0$	70.9 ± 3.8	0.16 ± 0.02	87.6 ± 3.4	0.05 ± 0.02
Ours, $\alpha_2 = 0$	71.9 ± 3.8	0.15 ± 0.02	88.1 ± 3.4	0.05 ± 0.02
Ours, $\alpha_4 = 0$	72.8 ± 3.9	0.16 ± 0.02	89.0 ± 3.4	0.06 ± 0.02
Static Grid Discretization (GliODIL, [16])	72.9 ± 3.5	0.16 ± 0.02	89.0 ± 3.3	0.06 ± 0.02
Ours	74.7 ± 3.1	0.19 ± 0.02	89.9 ± 2.7	0.08 ± 0.02

Table 4: Ablation study showing the impact of different model variations on recurrence coverage and IoU. Large deviations account for various sizes of tumors and recurrences, yet the results remain consistent and statistically significant (Table 3). *IoU heavily penalizes the fact that the radiotherapy volumes are equal between baselines and much larger than typical recurrences. Equal volume is to make comparisons with the Standard Plan without volume/recall trade-offs.

D Imaging Model

The core of the method’s alignment with the data is encapsulated by the loss function components, which associate the simulated outputs with key imaging traits such as the tumor core, surrounding edema, metabolic activity detected through FET-PET imaging, and visible brain tissues. This loss function depends on two learnable parameters $\theta_{\text{Imaging}} = \{\theta_{\text{down}}, \theta_{\text{up}}\}$ and is expressed as:

$$L_{\text{Imaging Data Reconstruction}} = \alpha_5 L_{\text{Core}}(c, \theta_{\text{up}}) + \alpha_6 L_{\text{Edema}}(c, \theta_{\text{down}}, \theta_{\text{up}}) + \alpha_7 L_{\text{Metabolic Map}}(c) + \alpha_8 L_{\text{Tissues}}(\mathbf{m}). \quad (11)$$

The loss function is composed of individual terms corresponding to distinct anatomical features:

- L_{Core} relates tumor cell concentrations above the threshold θ_{up} to the tumor core region.
- L_{Edema} delineates the edema area surrounding the tumor, regulated by the lower and upper thresholds θ_{down} and θ_{up} .
- $L_{\text{Metabolic Map}}$ assesses the metabolic activity as indicated by FET-PET signals within the edema and core regions, using a simple correlation metric, L_{Tissues} connects visible brain tissues with those inferred by the method.

These adaptive parameters $\{\theta_{\text{down}}, \theta_{\text{up}}\}$ enable the model to accommodate variations in MRI/FET-PET imaging contrasts and noise levels.

We adopt sigmoid functions to portray the gradational transitions observed at tumor region margins. The sigmoid, $\sigma(x)$, is specified as:

$$\sigma(x) = \frac{1}{1 + e^{-\beta x}}. \quad (12)$$

Here, β modulates the steepness of the transition and is set to $\beta = 50$. For the tumor core:

$$L_{\text{Core}}(c, \theta_{\text{up}}) = \sum_{i=0}^{N_x} \sigma(\theta_{\text{up}} - c_i^{N_t} - \alpha), \quad (13)$$

for the edema:

$$L_{\text{Edema}}(c, \theta_{\text{down}}, \theta_{\text{up}}) = \sum_{i=0}^{N_x} \sigma(\theta_{\text{down}} - c_i^{N_t} - \alpha) + (1 - \sigma(\theta_{\text{up}} - c_i^{N_t} + \alpha)), \quad (14)$$

where α offsets the thresholds and is set to $\alpha = 0.05$.

The metabolic activity within the tumor is evaluated by the loss term L_{PET} , which measures the correlation between the simulated metabolic signal and actual FET-PET scan observations:

$$L_{\text{PET}}(c) = 1 - \text{corr}(c^{N_t}, p^{\text{PET}}). \quad (15)$$

Here, $\text{corr}(c^{N_t}, p^{\text{PET}})$ represents the Pearson correlation coefficient between the predicted tumor cell densities c^{N_t} and the restricted FET-PET signal p^{PET} , across all voxels within the edema and enhancing core regions, providing a direct measure of how well the model predictions align with the observed metabolic profiles.

E Assumption of Brain Symmetry at Initial Time Without the Tumor

Our model employs a multi-resolution analysis to assess brain symmetry, particularly useful since the brain's symmetry is more evident at coarser resolutions. This approach is applied to key tissue types: white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF), each analyzed at full, 2x, and 4x downsampled levels to accommodate both the brain's structural nuances and computational efficiency.

The process for each tissue type involves converting tissue data into particles, then calculating symmetry loss at each resolution. This calculation entails:

1. Downsample the tissue representation, using average pooling
2. Divide the tissue's downsampled representation along its height, reflecting the brain's hemispherical division.
3. Mirror one hemisphere and compare it to the other, quantifying symmetry by calculating the mean absolute difference between them.

The total symmetry loss accumulates across all tissues and resolutions to measure deviation from an ideal symmetric state:

$$L_{\text{Healthy Brain}}(\mathbf{m}) = \sum_{k=1}^3 \sum_{\kappa=0}^2 \text{calculate_symmetry_loss}(m_k, \text{scale_factor} = 2^{\kappa}). \quad (16)$$

The multi-resolution approach not only addresses the limitations of assuming brain symmetry at detailed levels but also streamlines computational efforts. The smoother landscape at downsampled resolutions is especially favorable for particle movement through gradient descent.

F Synthetic Experiments

To calibrate the loss function weights, we generated a synthetic dataset for single focal and localized multi-focal tumors by solving PDEs using a finite difference method numerical solver [13]. An average brain model was used to represent the spatial distribution of brain tissues. Tumor growth model parameters, elasticity, imaging model parameters, and 1 or 3 focal locations were varied using uniform random distributions.

The metric used was the RMSE between the learned and ground truth tumor cell distributions. In total, 100 synthetic patients were generated. Table 5 provides the details of the parameter ranges used for generating the synthetic single focal and multi-focal tumor datasets.

Table 5: Parameter ranges for generating synthetic single focal and multi-focal tumor datasets

Shared parameters		
Parameter	Min	Max
D_w	0.035	0.2
γ	0	1.5
ρ	0.035	0.2
R	10	30
θ_{necro}	0.70	0.85
θ_{up}	0.45	0.60
θ_{down}	0.15	0.35
T_{sim}		100
Single focal tumor center (mm)		
(x_0, y_0, z_0)	57.6	96
Multi-focal tumor centers (mm)		
Tumor 1 center (x_0^1, y_0^1, z_0^1)	57.6	96
Tumor 2 center (x_0^2, y_0^2, z_0^2)	$(x_0^1, y_0^1, z_0^1) \pm 9.6$	
Tumor 3 center (x_0^3, y_0^3, z_0^3)	$(x_0^2, y_0^2, z_0^2) \pm 9.6$	

G Material Parameters for Biological Tissues

The table below summarizes the Young's modulus and Poisson's ratio for different tissue types used in the hyperelastic modeling of biological tissues.

Table 6: Young's modulus and Poisson's ratio for biological tissues.

Tissue Type	Young's Modulus (Pa)	Poisson's Ratio
Gray Matter (GM)	2100	0.4
White Matter (WM)	2100	0.4
Cerebrospinal Fluid (CSF)	100	0.1
Tumor	8000	0.45

NeurIPS Paper Checklist

1. Claims

Question: Do the main claims made in the abstract and introduction accurately reflect the paper's contributions and scope?

Answer: [\[Yes\]](#)

Justification: The claims are justified. The physics loss is clearly formulated and the results are visualized.

Guidelines:

- The answer NA means that the abstract and introduction do not include the claims made in the paper.
- The abstract and/or introduction should clearly state the claims made, including the contributions made in the paper and important assumptions and limitations. A No or NA answer to this question will not be perceived well by the reviewers.
- The claims made should match theoretical and experimental results, and reflect how much the results can be expected to generalize to other settings.
- It is fine to include aspirational goals as motivation as long as it is clear that these goals are not attained by the paper.

2. Limitations

Question: Does the paper discuss the limitations of the work performed by the authors?

Answer: [\[Yes\]](#)

Justification: Yes, we have addressed the multi-resolution grid, which operates under the assumption that the problem at hand can be reliably approximated at lower-than-target resolutions.

Guidelines:

- The answer NA means that the paper has no limitation while the answer No means that the paper has limitations, but those are not discussed in the paper.
- The authors are encouraged to create a separate "Limitations" section in their paper.
- The paper should point out any strong assumptions and how robust the results are to violations of these assumptions (e.g., independence assumptions, noiseless settings, model well-specification, asymptotic approximations only holding locally). The authors should reflect on how these assumptions might be violated in practice and what the implications would be.
- The authors should reflect on the scope of the claims made, e.g., if the approach was only tested on a few datasets or with a few runs. In general, empirical results often depend on implicit assumptions, which should be articulated.
- The authors should reflect on the factors that influence the performance of the approach. For example, a facial recognition algorithm may perform poorly when image resolution is low or images are taken in low lighting. Or a speech-to-text system might not be used reliably to provide closed captions for online lectures because it fails to handle technical jargon.
- The authors should discuss the computational efficiency of the proposed algorithms and how they scale with dataset size.
- If applicable, the authors should discuss possible limitations of their approach to address problems of privacy and fairness.
- While the authors might fear that complete honesty about limitations might be used by reviewers as grounds for rejection, a worse outcome might be that reviewers discover limitations that aren't acknowledged in the paper. The authors should use their best judgment and recognize that individual actions in favor of transparency play an important role in developing norms that preserve the integrity of the community. Reviewers will be specifically instructed to not penalize honesty concerning limitations.

3. Theory Assumptions and Proofs

Question: For each theoretical result, does the paper provide the full set of assumptions and a complete (and correct) proof?

Answer: [NA]

Justification: There is no theoretical results in a strict sense. It is easy to see that if $\Delta t \rightarrow 0, \Delta x \rightarrow 0$ the RHS of the discrete residual converges to the original PDEs.

Guidelines:

- The answer NA means that the paper does not include theoretical results.
- All the theorems, formulas, and proofs in the paper should be numbered and cross-referenced.
- All assumptions should be clearly stated or referenced in the statement of any theorems.
- The proofs can either appear in the main paper or the supplemental material, but if they appear in the supplemental material, the authors are encouraged to provide a short proof sketch to provide intuition.
- Inversely, any informal proof provided in the core of the paper should be complemented by formal proofs provided in appendix or supplemental material.
- Theorems and Lemmas that the proof relies upon should be properly referenced.

4. Experimental Result Reproducibility

Question: Does the paper fully disclose all the information needed to reproduce the main experimental results of the paper to the extent that it affects the main claims and/or conclusions of the paper (regardless of whether the code and data are provided or not)?

Answer: [Yes]

Justification: All components of the loss function are clearly formulated. While access to the specific code and hyper-parameter calibration scheme, which utilizes synthetic data generation, might yield slightly different results, these variations should not affect the overall claims.

Guidelines:

- The answer NA means that the paper does not include experiments.
- If the paper includes experiments, a No answer to this question will not be perceived well by the reviewers: Making the paper reproducible is important, regardless of whether the code and data are provided or not.
- If the contribution is a dataset and/or model, the authors should describe the steps taken to make their results reproducible or verifiable.
- Depending on the contribution, reproducibility can be accomplished in various ways. For example, if the contribution is a novel architecture, describing the architecture fully might suffice, or if the contribution is a specific model and empirical evaluation, it may be necessary to either make it possible for others to replicate the model with the same dataset, or provide access to the model. In general, releasing code and data is often one good way to accomplish this, but reproducibility can also be provided via detailed instructions for how to replicate the results, access to a hosted model (e.g., in the case of a large language model), releasing of a model checkpoint, or other means that are appropriate to the research performed.
- While NeurIPS does not require releasing code, the conference does require all submissions to provide some reasonable avenue for reproducibility, which may depend on the nature of the contribution. For example
 - (a) If the contribution is primarily a new algorithm, the paper should make it clear how to reproduce that algorithm.
 - (b) If the contribution is primarily a new model architecture, the paper should describe the architecture clearly and fully.
 - (c) If the contribution is a new model (e.g., a large language model), then there should either be a way to access this model for reproducing the results or a way to reproduce the model (e.g., with an open-source dataset or instructions for how to construct the dataset).
 - (d) We recognize that reproducibility may be tricky in some cases, in which case authors are welcome to describe the particular way they provide for reproducibility. In the case of closed-source models, it may be that access to the model is limited in some way (e.g., to registered users), but it should be possible for other researchers to have some path to reproducing or verifying the results.

5. Open access to data and code

Question: Does the paper provide open access to the data and code, with sufficient instructions to faithfully reproduce the main experimental results, as described in supplemental material?

Answer: [\[Yes\]](#)

Justification: Both data and code are available to reproduce the main results.

Guidelines:

- The answer NA means that paper does not include experiments requiring code.
- Please see the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- While we encourage the release of code and data, we understand that this might not be possible, so “No” is an acceptable answer. Papers cannot be rejected simply for not including code, unless this is central to the contribution (e.g., for a new open-source benchmark).
- The instructions should contain the exact command and environment needed to run to reproduce the results. See the NeurIPS code and data submission guidelines (<https://nips.cc/public/guides/CodeSubmissionPolicy>) for more details.
- The authors should provide instructions on data access and preparation, including how to access the raw data, preprocessed data, intermediate data, and generated data, etc.
- The authors should provide scripts to reproduce all experimental results for the new proposed method and baselines. If only a subset of experiments are reproducible, they should state which ones are omitted from the script and why.
- At submission time, to preserve anonymity, the authors should release anonymized versions (if applicable).
- Providing as much information as possible in supplemental material (appended to the paper) is recommended, but including URLs to data and code is permitted.

6. Experimental Setting/Details

Question: Does the paper specify all the training and test details (e.g., data splits, hyperparameters, how they were chosen, type of optimizer, etc.) necessary to understand the results?

Answer: [\[Yes\]](#)

Justification: Yes, all major aspects are provided.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The experimental setting should be presented in the core of the paper to a level of detail that is necessary to appreciate the results and make sense of them.
- The full details can be provided either with the code, in appendix, or as supplemental material.

7. Experiment Statistical Significance

Question: Does the paper report error bars suitably and correctly defined or other appropriate information about the statistical significance of the experiments?

Answer: [\[Yes\]](#)

Justification: Yes - this information is provided.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The authors should answer "Yes" if the results are accompanied by error bars, confidence intervals, or statistical significance tests, at least for the experiments that support the main claims of the paper.
- The factors of variability that the error bars are capturing should be clearly stated (for example, train/test split, initialization, random drawing of some parameter, or overall run with given experimental conditions).

- The method for calculating the error bars should be explained (closed form formula, call to a library function, bootstrap, etc.)
- The assumptions made should be given (e.g., Normally distributed errors).
- It should be clear whether the error bar is the standard deviation or the standard error of the mean.
- It is OK to report 1-sigma error bars, but one should state it. The authors should preferably report a 2-sigma error bar than state that they have a 96% CI, if the hypothesis of Normality of errors is not verified.
- For asymmetric distributions, the authors should be careful not to show in tables or figures symmetric error bars that would yield results that are out of range (e.g. negative error rates).
- If error bars are reported in tables or plots, The authors should explain in the text how they were calculated and reference the corresponding figures or tables in the text.

8. Experiments Compute Resources

Question: For each experiment, does the paper provide sufficient information on the computer resources (type of compute workers, memory, time of execution) needed to reproduce the experiments?

Answer: [Yes]

Justification: Yes - this information is provided.

Guidelines:

- The answer NA means that the paper does not include experiments.
- The paper should indicate the type of compute workers CPU or GPU, internal cluster, or cloud provider, including relevant memory and storage.
- The paper should provide the amount of compute required for each of the individual experimental runs as well as estimate the total compute.
- The paper should disclose whether the full research project required more compute than the experiments reported in the paper (e.g., preliminary or failed experiments that didn't make it into the paper).

9. Code Of Ethics

Question: Does the research conducted in the paper conform, in every respect, with the NeurIPS Code of Ethics <https://neurips.cc/public/EthicsGuidelines>?

Answer: [Yes]

Justification: We are here to help, no cause harm.

Guidelines:

- The answer NA means that the authors have not reviewed the NeurIPS Code of Ethics.
- If the authors answer No, they should explain the special circumstances that require a deviation from the Code of Ethics.
- The authors should make sure to preserve anonymity (e.g., if there is a special consideration due to laws or regulations in their jurisdiction).

10. Broader Impacts

Question: Does the paper discuss both potential positive societal impacts and negative societal impacts of the work performed?

Answer: [NA]

Justification: There are no clear societal impacts, except for the potential for improved healthcare, which is difficult to assess without clinical trials. The work also has the potential to translate to other areas.

Guidelines:

- The answer NA means that there is no societal impact of the work performed.
- If the authors answer NA or No, they should explain why their work has no societal impact or why the paper does not address societal impact.

- Examples of negative societal impacts include potential malicious or unintended uses (e.g., disinformation, generating fake profiles, surveillance), fairness considerations (e.g., deployment of technologies that could make decisions that unfairly impact specific groups), privacy considerations, and security considerations.
- The conference expects that many papers will be foundational research and not tied to particular applications, let alone deployments. However, if there is a direct path to any negative applications, the authors should point it out. For example, it is legitimate to point out that an improvement in the quality of generative models could be used to generate deepfakes for disinformation. On the other hand, it is not needed to point out that a generic algorithm for optimizing neural networks could enable people to train models that generate Deepfakes faster.
- The authors should consider possible harms that could arise when the technology is being used as intended and functioning correctly, harms that could arise when the technology is being used as intended but gives incorrect results, and harms following from (intentional or unintentional) misuse of the technology.
- If there are negative societal impacts, the authors could also discuss possible mitigation strategies (e.g., gated release of models, providing defenses in addition to attacks, mechanisms for monitoring misuse, mechanisms to monitor how a system learns from feedback over time, improving the efficiency and accessibility of ML).

11. Safeguards

Question: Does the paper describe safeguards that have been put in place for responsible release of data or models that have a high risk for misuse (e.g., pretrained language models, image generators, or scraped datasets)?

Answer: [NA]

Justification: The method is evaluated on a publicly available dataset and does not inherently require safeguards.

Guidelines:

- The answer NA means that the paper poses no such risks.
- Released models that have a high risk for misuse or dual-use should be released with necessary safeguards to allow for controlled use of the model, for example by requiring that users adhere to usage guidelines or restrictions to access the model or implementing safety filters.
- Datasets that have been scraped from the Internet could pose safety risks. The authors should describe how they avoided releasing unsafe images.
- We recognize that providing effective safeguards is challenging, and many papers do not require this, but we encourage authors to take this into account and make a best faith effort.

12. Licenses for existing assets

Question: Are the creators or original owners of assets (e.g., code, data, models), used in the paper, properly credited and are the license and terms of use explicitly mentioned and properly respected?

Answer: [Yes]

Justification: Yes - the dataset authors are credited with a licence.

Guidelines:

- The answer NA means that the paper does not use existing assets.
- The authors should cite the original paper that produced the code package or dataset.
- The authors should state which version of the asset is used and, if possible, include a URL.
- The name of the license (e.g., CC-BY 4.0) should be included for each asset.
- For scraped data from a particular source (e.g., website), the copyright and terms of service of that source should be provided.

- If assets are released, the license, copyright information, and terms of use in the package should be provided. For popular datasets, paperswithcode.com/datasets has curated licenses for some datasets. Their licensing guide can help determine the license of a dataset.
- For existing datasets that are re-packaged, both the original license and the license of the derived asset (if it has changed) should be provided.
- If this information is not available online, the authors are encouraged to reach out to the asset's creators.

13. New Assets

Question: Are new assets introduced in the paper well documented and is the documentation provided alongside the assets?

Answer: [NA]

Justification: No new assets. (other than later provided code)

Guidelines:

- The answer NA means that the paper does not release new assets.
- Researchers should communicate the details of the dataset/code/model as part of their submissions via structured templates. This includes details about training, license, limitations, etc.
- The paper should discuss whether and how consent was obtained from people whose asset is used.
- At submission time, remember to anonymize your assets (if applicable). You can either create an anonymized URL or include an anonymized zip file.

14. Crowdsourcing and Research with Human Subjects

Question: For crowdsourcing experiments and research with human subjects, does the paper include the full text of instructions given to participants and screenshots, if applicable, as well as details about compensation (if any)?

Answer: [NA]

Justification: The authors were not involved in data gathering. The patients protocols are explained at the cited source.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.
- Including this information in the supplemental material is fine, but if the main contribution of the paper involves human subjects, then as much detail as possible should be included in the main paper.
- According to the NeurIPS Code of Ethics, workers involved in data collection, curation, or other labor should be paid at least the minimum wage in the country of the data collector.

15. Institutional Review Board (IRB) Approvals or Equivalent for Research with Human Subjects

Question: Does the paper describe potential risks incurred by study participants, whether such risks were disclosed to the subjects, and whether Institutional Review Board (IRB) approvals (or an equivalent approval/review based on the requirements of your country or institution) were obtained?

Answer: [Yes]

Justification: The authors were not involved in data gathering. The patients protocols are explained at the cited source.

Guidelines:

- The answer NA means that the paper does not involve crowdsourcing nor research with human subjects.

- Depending on the country in which research is conducted, IRB approval (or equivalent) may be required for any human subjects research. If you obtained IRB approval, you should clearly state this in the paper.
- We recognize that the procedures for this may vary significantly between institutions and locations, and we expect authors to adhere to the NeurIPS Code of Ethics and the guidelines for their institution.
- For initial submissions, do not include any information that would break anonymity (if applicable), such as the institution conducting the review.