

Predictive Digital Twins with Quantified Uncertainty for Patient-Specific Decision Making in Oncology

Graham Pash^{a,*}, Umberto Villa^{a,c}, David A. Hormuth II^{a,b},
Thomas E. Yankeelov^{a,b,c,d,e}, Karen Willcox^a

^a*Oden Institute for Computational Engineering and Sciences, The University of Texas at
Austin, Austin, TX, 78712, USA*

^b*UT Austin Cancer Research Center, The University of Texas at Austin, Austin, TX, 78712, USA*

^c*Department of Biomedical Engineering, The University of Texas at Austin, Austin, TX, 78712, USA*

^d*Department of Diagnostic Medicine, The University of Texas at Austin Cancer Institutes, Austin, TX, 78712, USA*

^e*Department of Imaging Physics, The University of Texas M.D. Anderson Cancer
Center, Houston, TX, 77230, USA*

Abstract

Quantifying the uncertainty in predictive models is critical for establishing trust and enabling risk-informed decision making for personalized medicine. In contrast to one-size-fits-all approaches that seek to mitigate risk at the population level, digital twins enable personalized modeling thereby potentially improving individual patient outcomes. Realizing digital twins in biomedicine requires scalable and efficient methods to integrate patient data with mechanistic models of disease progression. This study develops an end-to-end data-to-decisions methodology that combines longitudinal non-invasive imaging data with mechanistic models to estimate and predict spatiotemporal tumor progression accounting for patient-specific anatomy. Through the solution of a statistical inverse problem, imaging data inform the spatially varying parameters of a reaction-diffusion model of tumor progression. An efficient parallel implementation of the forward model coupled with a scalable approximation of the Bayesian posterior distribution enables rigorous, but tractable, quantification of uncertainty due to the sparse, noisy measurements. The methodology is verified on a virtual patient with synthetic data to control for model inadequacy, noise level, and the frequency of data collection. The application to decision-making is illustrated by evaluating the importance of imaging frequency and formulating an optimal experimental design question. The clinical relevance is demonstrated through a model validation study on a cohort of patients with publicly available longitudinal imaging data.

Keywords: Digital Twins, Computational Oncology, Uncertainty Quantification, Bayesian Inverse Problems, PDE-Constrained Optimization

1. Introduction

Predictive digital twins are poised to make an impact in the burgeoning field of precision medicine by coupling mathematical models with patient-specific data. While it is not feasible to perform multiple *in vivo* trials on an individual patient, mathematical and computational modeling augment the traditional clinical trial approach by enabling *in silico* trialing and assessment of potential

*Corresponding author.

Email address: gtpash@utexas.edu (Graham Pash)

interventions in a personalized manner [1]. The heterogeneity in cancer physiology [2] and patient response to a particular therapy [3, 4] means that there are many interventions that may work on average for a population, but are ineffective for an individual [5]. Moreover, assessment of patient-specific response in the clinical setting is dependent on monitoring radiological and clinical changes over long timelines after the conclusion of therapy [6, 7]. An effectively realized digital twin will account for inter-patient variability in disease presentation and progression, while accelerating the feedback loop from data acquisition to clinical action [8, 9, 10, 11]. Figure 1 illustrates the bidirectional flow of information between a cancer patient and their digital twin. Observational data are collected and combined with mechanistic models of the spatio-temporal tumor growth to update forecasts of prognosis, which are in turn used to guide clinical decisions between visits. In this work, we develop a scalable mathematical formulation and efficient computational framework of this digital twin data-to-decisions feedback loop, demonstrated by application to high-grade gliomas.

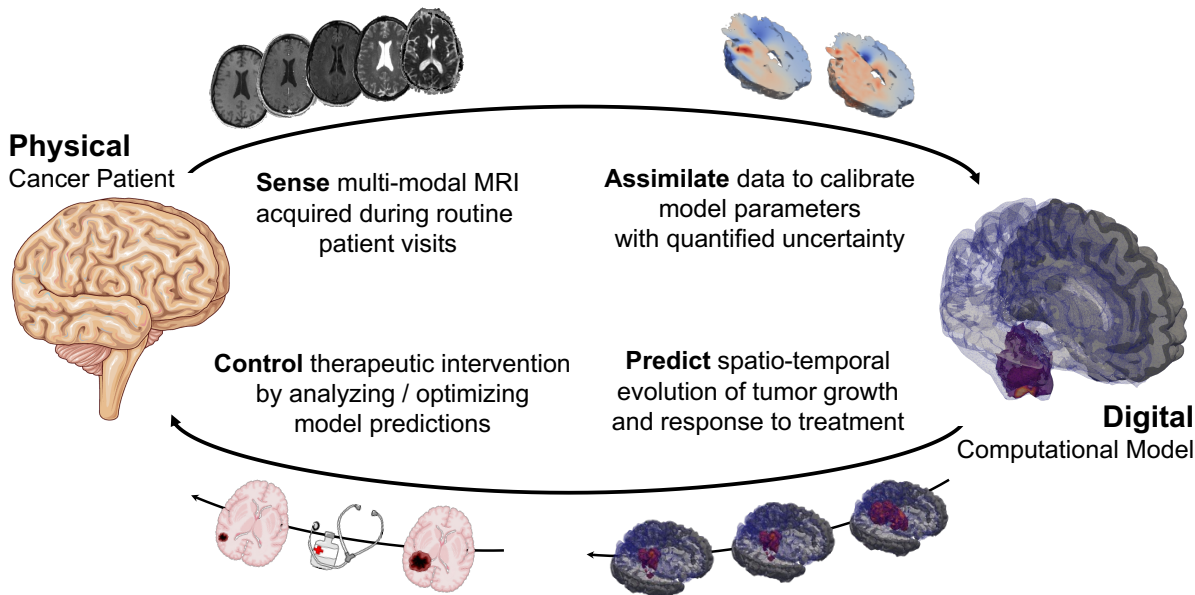


Figure 1: Illustration of our digital twin workflow for a cancer patient. Observational data are integrated with mechanistic models of tumor growth to update the computational representation of the dynamics with quantified uncertainty. The calibrated model is used to make probabilistic forecasts of tumor progression accounting for patient response to therapy. In turn, these forecasts guide clinical decision making. The model may be used to assess what-if scenarios for alternative interventions or to optimize therapy directly by (for example) adjusting dose level or schedule.

High-grade gliomas comprise the majority of malignant brain tumor cases and are characterized by their aggressive, invasive, and heterogeneous nature [12]. Despite aggressive standard-of-care treatment [13] consisting of multi-modal treatments (surgical resection, radiotherapy, and chemotherapy), overall prognosis remains poor and survival rates remain low [14]. For a digital twin to be effective in this setting, the underlying models incorporating biological mechanisms must be able to reliably predict the complex spatio-temporal tumor dynamics. In particular, the variability of intra-tumoral properties presents a significant modeling challenge, requiring a spatial, and often high-dimensional, characterization of the model parameters driving growth to account for the variety of observed dynamics. One successful line of work employs image-based modeling to resolve the spatio-temporal tumor dynamics and leverages non-invasive medical imaging to calibrate models of tumor growth [15, 16, 17, 18, 19, 20, 21]. These data can be collected where patients receive care, both within and outside of academic research oriented facilities [22], indicating the potential for broad impact. However, measurement noise, scarcity of observational data, and inadequacy of

models to capture the biological complexity of tumor growth all contribute to significant modeling uncertainty [23].

Uncertainty quantification plays a critical role in the development of a digital twin, helping establish trust in models and enabling risk estimation for robust decision making [24]. The high-consequence nature of decisions in personalized medicine underscores the need for mathematical tools to rigorously account for uncertainty. There is a growing literature on the Bayesian calibration of image-based mechanistic models of tumor growth [25, 26, 27, 28, 29]. However, the computational complexity of characterizing the posterior with expensive, nonlinear forward models remains a key challenge [30]. Low-dimensional representations of model parameters enable tractable exploration of the posterior and have shown promise in assessing alternative treatment plans in human patients [27]. An alternative approach approximates the posterior directly, maintaining spatial resolution in the model parameters and demonstrating impressive predictive power in a murine model of brain cancer [29]. It remains to be shown that model parameterizations retaining sufficient spatial resolution to represent intra-tumoral heterogeneity can be scaled to the complex anatomies and long time horizons in human brain cancers. We address these modeling and computational challenges, with a particular focus on the advancement of scalable algorithms and efficient model implementations to enable the solution of high-dimensional Bayesian inverse problems in computational oncology.

Our central contribution is the development of a mathematical formulation for an end-to-end Bayesian data-to-decisions pipeline to serve as the basis for the development of digital twins in oncology. Clinical magnetic resonance imaging (MRI) data are used to generate patient-specific computational geometries and extract information about the tumor biology. A computational framework is established for tractable Bayesian calibration of high-fidelity three-dimensional spatio-temporal models of tumor growth to the clinical data. Efficient, parallel, and flexible finite element representations of the governing equations are paired with state-of-the-art optimization algorithms. We numerically verify the ability to calibrate high-dimensional model parameters to MRI data by reconstructing known spatial heterogeneity in a virtual patient. Furthermore, we investigate the impact of alternative imaging schedules on predictive performance and pose the optimal experimental design question of *when* to observe. We demonstrate clinical relevance through a validation study on clinical data, identifying key aspects of model inadequacy.

The remainder of this paper is structured as follows. Section 2 introduces the mathematical model of tumor growth and discusses the image processing pipeline. The scalable Bayesian methodology is presented in Section 3. Section 4 verifies the inverse problem methodology through an *in silico* case study, illustrates the connection to clinical decision making, and formulates an optimal experimental design question. Section 5 establishes the real world utility of the framework applied to a cohort of patients with publicly available clinical data. The model’s predictive ability is quantified along with potential limitations and extensions. Conclusions are drawn in Section 6.

2. Image-based mechanistic modeling of tumors

This section provides necessary background on the mathematical modeling of tumor growth as well as the medical imaging data used for calibration. The governing partial differential equation (PDE) model is derived in Section 2.1 detailing tumor invasion, proliferation, and response to chemoradiation. Section 2.2 describes a computational pipeline incorporating longitudinal image registration, tumor state estimation, and generation of a patient-specific computational geometry in the context of modeling glioma growth in the human brain.

2.1. Mechanistic modeling of high-grade gliomas

Mathematical models of tumor growth primarily focus on two key characteristics: invasion of the tumor into the surrounding healthy tissue [31] and the proliferation of the existing tumor [32]. Tumor invasion is typically modeled as a diffusion process and proliferation is modeled as a logistic growth process subject to a biological or physical carrying capacity. These phenomenological assumptions give rise to a semi-linear parabolic reaction-diffusion PDE [33, 34], which has proven successful in modeling the growth of solid tumors in a variety of organs [35, 36, 37, 38]. Furthermore, it is common to assume a known carrying capacity of the tissue and model the tumor cellularity as a volume fraction. Henceforth, tumor cellularity and tumor volume fraction are used interchangeably unless explicitly noted. The model is given by:

$$\begin{aligned} \frac{\partial u}{\partial t} - \nabla \cdot (D\nabla u) - \kappa u(1 - u) &= f \quad \text{in } \Omega \times (t_0, t_f) \\ u(x, t_0) &= u_0 \quad \text{in } \Omega \\ \nabla u \cdot \eta &= 0 \quad \text{on } \partial\Omega \times (t_0, t_f) \end{aligned} \tag{1}$$

Here, $u(x, t)$ is tumor volume fraction which varies in time t and over spatial coordinates x , $D(x)$ is the diffusion coefficient field, $\kappa(x)$ is the proliferation rate coefficient field, (t_0, t_f) is the simulation window, $\Omega \subset \mathbb{R}^n$ is the bounded spatial domain with dimension $n \in \{2, 3\}$, and η is the outward unit normal to the boundary $\partial\Omega$. The homogeneous Neumann boundary condition stems from the assumption that the tumor does not grow beyond the boundary of the domain, for example, the skull. The initial cellularity $u_0(x)$ is determined from MRI data following the procedure in Section 2.2. The source term f is used to model the effect of treatment.

The standard-of-care therapy for high grade gliomas incorporates both highly conformal radiotherapy and systemic chemotherapy [13]. Let $\mathcal{T}_{\text{rt}} = \{\tau_{k,\text{rt}}\}_k$ denote the collection of times at which radiotherapy is applied. Similarly, let $\mathcal{T}_{\text{ct}} = \{\tau_{k,\text{ct}}\}_k$ be the collection of times at which chemotherapy is administered. Together these define the treatment regimen. We make the common modeling assumption that the radiotherapy effect is instantaneous, killing some cells at the moment of treatment and with no lasting or time-delayed effects. This assumption results in a model of the form

$$f_{\text{rt}}(u, z_{\text{rt}}, t) = \begin{cases} 0 & \text{for } t \notin \mathcal{T}_{\text{rt}}, \\ -\gamma(1 - S_{\text{rt}}(z_{\text{rt}}))u & \text{for } t \in \mathcal{T}_{\text{rt}}, \end{cases} \tag{2}$$

where z_{rt} is the applied radiation dose, γ is a positive parameter to ensure appropriate dimensionality [39], and S_{rt} is the surviving fraction of the tumor after application of the therapy. We compute the surviving fraction with the well-established linear-quadratic (LQ) model relating applied dosage to radiotherapy induced cell death [40, 41]. The LQ model calculates the surviving fraction as

$$S_{\text{rt}}(z_{\text{rt}}) = \exp(-\alpha_{\text{rt}}z_{\text{rt}} - \beta_{\text{rt}}z_{\text{rt}}^2), \tag{3}$$

where $\alpha_{\text{rt}} > 0$ and $\beta_{\text{rt}} > 0$ are parameters describing the radiosensitivity of the tissue [42, 43]. To capture the effect of chemotherapy, we assume a known drug efficacy $\alpha_{\text{ct}} > 0$ and clearance rate $\beta_{\text{ct}} > 0$ and employ a decaying exponential model of the form [44],

$$f_{\text{ct}}(u, z_{\text{ct}}, t) = \begin{cases} 0 & \text{for } t < \tau_{0,\text{ct}}, \\ -\alpha_{\text{ct}}(z_{\text{ct}}) \sum_k \exp(-\beta_{\text{ct}}(t - \tau_{k,\text{ct}}))u & \text{for } t \geq \tau_{0,\text{ct}}. \end{cases} \tag{4}$$

The drug efficacy α_{ct} may be modulated by the spatial distribution of the drug concentration z_{ct} when data are available as in [44], though we choose to model the term as a constant surviving

fraction acting homogeneously in the domain as in [20]. Finally, we combine the two effects as the source term in Eq. 1,

$$f(u, z_{rt}, z_{ct}, t) = f_{rt}(u, z_{rt}, t) + f_{ct}(u, z_{ct}, t). \quad (5)$$

We employ the finite element method (FEM) [45, 46] for our spatial discretization. Among other advantages, FEM admits flexibility in the spatial discretization and is able to conform to the complex geometries arising in biomedical applications. The variational formulation for the model Eq. (1) is derived in Appendix A and the implicit Euler method is used for the time discretization. We discuss a computational pipeline for meshing human brain anatomy in the following section.

2.2. Computational pipeline: From medical imaging to computational models

Biomedical imaging, especially MRI [47, 48], plays a critical role in the diagnosis, treatment planning, and management of a variety of tumors, such as those in the brain [49], the breast [50], and the prostate [51]. In our digital twin framework, we consider this to be the primary source of observational data coming from the clinic. To integrate imaging data with the biophysical models described in Section 2.1, we must estimate the tumor volume fraction, i.e. the state variable of the PDE, and define an appropriate computational domain. Figure 2 provides an overview of such a computational pipeline for the human brain.

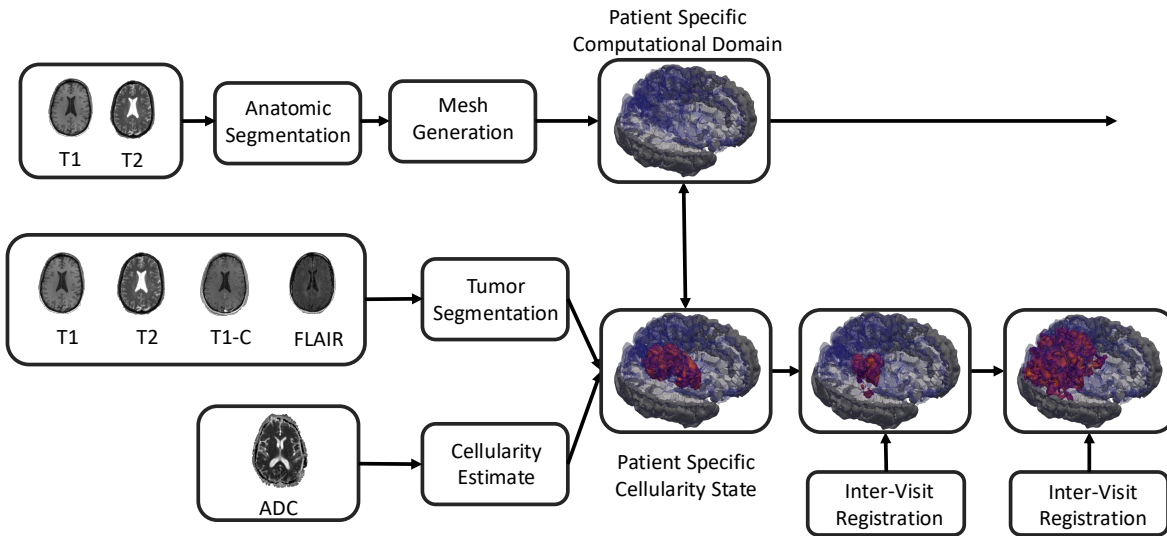


Figure 2: The computational pipeline: anatomic segmentation and mesh generation, cellularity estimation, and longitudinal registration. Together T_1 - and T_2 -weighted scans are used to generate a computational domain tailored to the patient’s anatomy. Tumor cellular density is estimated by combining ADC imaging with tumor segmentations. When radiologist defined segmentations are not available, automated tools utilizing T_1 pre- and post-contrast imaging along with T_2 -weighted and FLAIR modalities are used to develop appropriate regions of interest.

Expert segmentation of tumor regions of interest (ROIs) are used to determine tumor extent, when available. However, semi-automated approaches leveraging routinely collected MRI data are improving [52] and may be used when expert segmentation is unavailable. The apparent diffusion coefficient (ADC) quantifies the rate of diffusion of water molecules, which is inversely proportional

to cellular density. The ADC can be used to estimate the tumor cellularity [21, 53] via

$$d(\bar{x}, t) = \frac{\text{ADC}_w - \text{ADC}(\bar{x}, t)}{\text{ADC}_w - \text{ADC}_{\min}}, \quad (6)$$

where ADC_w is the apparent diffusion coefficient of free water [20], $\text{ADC}(\bar{x}, t)$ is the measured apparent diffusion coefficient at time t in the voxel with spatial coordinate \bar{x} , and ADC_{\min} is the minimum value recorded within the tumor. This approach generates observations of the tumor cellularity, $d(\bar{x}, t)$, every time the patient is imaged. We note that while it is not possible to directly measure cellular diffusion *in vivo*, serial diffusion imaging has also been used in attempts to directly inform diffusion and proliferation rates [54, 55]. To ensure a consistent frame of reference, longitudinally collected data is rigidly registered to a baseline MRI using a tool such as `elastix` [56, 57]. Similarly, multiple modalities within a visit may be registered using a mutual information criterion to account for varying image resolution and contrast [56, 58]. For an overview of medical image registration, readers are referred to [59].

To showcase the approach in application to brain cancer, we consider two publicly available datasets: UPENN-GBM [60, 61] and IvyGAP [62, 63]. Both datasets contain pre- and post-contrast T_1 -weighted, T_2 -weighted, and FLAIR images. The UPENN-GBM dataset also provides expert segmentation of the tumor ROI, while a semi-automated approach [64] is used to generate tumor ROIs for the IvyGAP dataset. Additionally, the IvyGAP dataset contains diffusion-weighted imaging (DWI) and ADC data as well as longitudinal imaging data, enabling retrospective studies with clinical data. Brief descriptions of collected modalities and their connection to biophysical modeling are summarized in Table 1, adapted from [65].

Table 1: Summary of acquired MRI and their utility.

Imaging technique	Major utility in brain tumor imaging
Pre- and post-contrast T_1	Anatomy, necrosis, enhancing tumor
T_2 / FLAIR	Anatomy, edema, non-enhancing tumor
DWI / ADC	Cellularity estimation

To construct a reference computational geometry for the biological domain upon which to simulate the growth of the tumor, we build upon the work of [66]. We use T_1 - and T_2 -weighted images for segmentation of the gray matter, white matter and cerebral aqueduct using the `FreeSurfer` neuroimaging software [67]. To ensure reliable segmentation when large lesions are present, the T_1 and T_2 images are first repaired with virtual brain grafting [68]. Tetrahedral meshing of the resultant surfaces is performed with `CGAL` [69]. Similar approaches have provided a basis for high-fidelity simulation of a variety of physical phenomena in the brain from the circulation of cerebrospinal fluid [70, 71, 72] to neurodegenerative disease progression [73, 74].

3. Scalable Bayesian model calibration

A fundamental capability of a digital twin is the ability to integrate and assimilate observational data with mathematical and computational models. This section details a scalable methodology to estimate model parameters in the presence of uncertainty. We formulate this process as a Bayesian inverse problem, seek an approximation to the posterior distribution, and detail how to propagate uncertainty through the forward model to generate probabilistic forecasts of tumor growth and clinically relevant quantities of interest.

3.1. Formulation

We pose the calibration of the mechanistic model Eq. (1) to the imaging data as an inverse problem: from observational data $\mathbf{d} = [d(\bar{x}, t_1), d(\bar{x}, t_2), \dots, d(\bar{x}, t_{n_t})] \in \mathbb{R}^{n_d \times n_t}$, we infer the values of the unknown parameter(s) $m \in \mathcal{M}$ for some suitable function space \mathcal{M} . Here, n_d is the dimensionality of a single data point, i.e., the number of voxels in an MRI image, and n_t is the number of visits at which tumor volume fraction estimates are acquired. We choose the parameterization $m = [m_D, m_\kappa] = [\log(D), \log(\kappa)]$ to preserve positivity of the model coefficients $D = \exp(m_D), \kappa = \exp(m_\kappa)$. The objective is to estimate the posterior distribution of the model parameters, given by Bayes' theorem [75] as:

$$\frac{d\nu_{\text{post}}}{d\nu_{\text{pr}}} \propto \pi_{\text{like}}(\mathbf{d}|m), \quad (7)$$

where $d\nu_{\text{post}}/d\nu_{\text{pr}}$ denotes the Radon-Nikodym derivative [76] of the posterior measure ν_{post} with respect to the prior measure ν_{pr} . The likelihood π_{like} is specified by the choice of noise model.

3.2. Likelihood

To account for measurement uncertainty in the pipeline outlined in Section 2.2, we assume an additive Gaussian noise model,

$$d(\bar{x}, t_i) = \mathcal{B}(u(x, t_i)) + \varepsilon_i, \quad (8)$$

where $\mathcal{B} : \mathcal{U} \rightarrow \mathbb{R}^{n_d}$ is an observation operator that extracts the observable from the state $u \in \mathcal{U}$, for some suitable Hilbert space \mathcal{U} . For example, an observable may be the voxel-resolved tumor cellularity. Additionally, the observation operator appropriately interpolates between the spatial coordinates of the FEM representation x and the spatial coordinates of the voxel data \bar{x} . The additive Gaussian noise is assumed to have zero mean and covariance $\mathbf{\Gamma}_{\text{noise}} \in \mathbb{R}^{n_d \times n_d}$, such that $\varepsilon_i \sim \mathcal{N}(0, \mathbf{\Gamma}_{\text{noise}})$. Further, we assume uniformity of the noise, $\mathbf{\Gamma}_{\text{noise}} = \sigma_{\text{noise}}^2 \mathbf{I}$, for some unknown noise variance σ_{noise}^2 . Note that we consider statistical uncertainty to be the main source of error and do not account explicitly for discretization error, which is well controlled relative to the noise in the data. We define the parameter-to-observable (PtO) map:

$$\mathcal{F} : \mathcal{M} \rightarrow \mathbb{R}^{n_d \times n_t} = [\mathcal{B}(u(x, t_1)), \mathcal{B}(u(x, t_2)), \dots, \mathcal{B}(u(x, t_{n_t}))], \quad \text{s.t. } r(u, m) = 0, \quad (9)$$

where $r(u, m) = 0$ is the residual form of the governing PDE Eq. (1), explicitly noting the dependence on the state and parameter. Evaluating the PtO map \mathcal{F} involves the solution of the forward model Eq. (1) for the state $u \in \mathcal{U}$ and the application of the observation operator \mathcal{B} to extract the observable at each time that data is acquired. The probability density function of the likelihood may be expressed as

$$\pi_{\text{like}}(\mathbf{d}|m) \propto \exp\{-\Phi(m; \mathbf{d})\}, \quad (10)$$

where $\Phi(m; \mathbf{d})$ is the negative log-likelihood. For the additive Gaussian case with multiple measurements, we sum the negative log-likelihoods for each observation,

$$\Phi(m, \mathbf{d}) := \sum_{i=1}^{n_t} \frac{1}{2} \|\mathcal{F}_i(m) - \mathbf{d}_i\|_{\mathbf{\Gamma}_{\text{noise}}^{-1}}^2, \quad (11)$$

where \mathcal{F}_i denotes the i -th component of the PtO map, that is, $\mathcal{F}_i := \mathcal{B}(u(x, t_i))$.

3.3. Prior distribution

The prior distribution encodes structural knowledge of the parameters and is taken to be a Gaussian random field, with mean m_{pr} and covariance operator \mathcal{C}_{pr} , such that $m \sim \mathcal{N}(m_{\text{pr}}, \mathcal{C}_{\text{pr}})$. This modeling choice has desirable properties for analysis and computation, and has been used to model various physical systems including tumor growth in murine models [29]. In particular, this selection ensures a well-posed posterior [75] and implies a prior measure of the form

$$d\nu_{\text{pr}}(m) \propto \exp \left\{ -\frac{1}{2} \|m - m_{\text{pr}}\|_{\mathcal{C}_{\text{pr}}^{-1}}^2 \right\}. \quad (12)$$

Values for the mean m_{pr} may be taken from the literature or determined by deterministic calibration to a cohort.

The seminal work [77] established a link between Matérn covariance operators and precision operators defined by PDEs. Following [75, 78] the covariance operator is taken to be the inverse of an elliptic operator of the form

$$\mathcal{C}_{\text{pr}} := \mathcal{A}^{-2} = (-\gamma\Delta + \delta I)^{-2} \quad (13)$$

with hyperparameters $\gamma > 0$ and $\delta > 0$ that control the correlation length ρ and the pointwise variance σ^2 of the operator [77]. Specifically, the hyperparameters are related by

$$\delta = \frac{\sqrt{2}}{\sigma\rho\sqrt{\pi}} \quad , \quad \gamma = \frac{\rho}{4\sigma\sqrt{2\pi}}. \quad (14)$$

A Robin boundary condition is applied to the operator \mathcal{C}_{pr} to reduce boundary artifacts [79].

We assume the log coefficient field parameters m to be independent and model each as a Gaussian random field as described above. That is, $m_D \sim \mathcal{N}(m_{\text{pr},D}, \mathcal{C}_{\text{pr},D})$ and $m_\kappa \sim \mathcal{N}(m_{\text{pr},\kappa}, \mathcal{C}_{\text{pr},\kappa})$. Thus, $m_{\text{pr}} := [m_{\text{pr},D}, m_{\text{pr},\kappa}]$ and \mathcal{C}_{pr} has a block diagonal structure with blocks $\mathcal{C}_{\text{pr},D}$ and $\mathcal{C}_{\text{pr},\kappa}$, respectively.

3.4. Low-rank Laplace approximation to the posterior distribution

We reiterate that we model the parameters as spatial fields which are high-dimensional upon discretization. For example, when using first-order Lagrange elements, each discretized parameter has size equal to the number of vertices in the mesh. The computational cost to thoroughly explore the posterior distribution with standard methods, such as Markov chain Monte Carlo (MCMC), is prohibitive in this setting. We instead leverage the Laplace approximation to the posterior [80, 81] for a tractable solution,

$$\nu_{\text{post}}^{\text{LA}} \propto \mathcal{N}(m_{\text{MAP}}, \mathcal{C}_{\text{post}}). \quad (15)$$

The Laplace approximation is a Gaussian approximation to the posterior centered at the maximum *a posteriori* (MAP) point m_{MAP} with covariance $\mathcal{C}_{\text{post}}$ equal to the inverse of the negative log-posterior Hessian evaluated at m_{MAP} . This approximation is exact in the case of linear inverse problems and will serve as a cost-effective surrogate for the nonlinear problem considered in this work. Building the Laplace approximation requires one to compute the MAP point and calculate the covariance.

We compute the MAP estimate by minimizing the negative log-posterior

$$m_{\text{MAP}} := \arg \min_{m \in \mathcal{M}} (-\log \nu_{\text{post}}(m|\mathbf{d})). \quad (16)$$

Recalling the expression for the likelihood Eq. (10) and the prior Eq. (12), we may write the posterior as

$$\nu_{\text{post}}(m|\mathbf{d}) \propto \left\{ \underbrace{-\sum_i \frac{1}{2} \|\mathcal{F}_i(m) - \mathbf{d}_i\|_{\mathbf{\Gamma}_{\text{noise}}^{-1}}^2}_{\text{data misfit}} - \underbrace{\frac{1}{2} \|m - m_{\text{pr}}\|_{\mathcal{C}_{\text{pr}}^{-1}}^2}_{\text{prior}} \right\}. \quad (17)$$

Thus, computing the MAP point amounts to solving a deterministic inverse problem with the log-likelihood Eq. (11) playing the role of the data misfit term and the prior Eq. (12) acting as regularization.

Inexact Newton-Krylov algorithms have been shown to efficiently solve the optimization problem Eq. (16) [82]. Early termination of conjugate gradient (CG) iterations is used to inexactly solve the Newton system using the Eisenstat-Walker criterion to prevent over-solving [83] and the Steihaug criterion to avoid negative curvature [84]. Globalization is performed with an Armijo backtracking line search to guarantee global convergence [85]. For a wide class of nonlinear inverse problems, the number of outer Newton iterations and inner CG iterations is independent of the mesh size and hence parameter dimension [86]. This is a consequence of the Newton solver, the compactness of the Hessian of the data misfit term in Eq. (17), and preconditioning by the inverse of the regularization operator. The gradient and Hessian action are obtained using the adjoint method, thereby limiting the number of expensive PDE solves that are required. These expressions are derived in Appendix A with the Lagrangian formalism [87, 88].

The covariance of the Laplace approximation to the posterior is given by the inverse of the Hessian of the negative log-posterior,

$$\mathcal{C}_{\text{post}} := \mathcal{H}(m_{\text{MAP}})^{-1} = (\mathcal{H}_{\text{misfit}}(m_{\text{MAP}}) + \mathcal{C}_{\text{pr}}^{-1})^{-1}, \quad (18)$$

where $\mathcal{H}_{\text{misfit}}$ denotes the Hessian of the negative log-likelihood. Upon discretization, we have

$$\mathbf{\Gamma}_{\text{post}} = (\mathbf{H}_{\text{misfit}}(m_{\text{MAP}}) + \mathbf{\Gamma}_{\text{pr}}^{-1})^{-1} \quad (19)$$

where $\mathbf{\Gamma}_{\text{post}}$ is the discretization of $\mathcal{C}_{\text{post}}$, $\mathbf{\Gamma}_{\text{pr}}$ is the discretization of \mathcal{C}_{pr} , and $\mathbf{H}_{\text{misfit}}(m_{\text{MAP}})$ is the discretized Hessian of the negative log-likelihood at the MAP estimate. The construction of $\mathbf{H}_{\text{misfit}}$ for large-scale applications is prohibitive, let alone the inversion required to form $\mathbf{\Gamma}_{\text{post}}$. Instead, we leverage the fact that in many cases, the spectral decay of $\mathbf{H}_{\text{misfit}}$ is rapid as the data contain limited information about the (infinite-dimensional) parameters. Thus we adopt a low-rank correction to the prior covariance to approximate the posterior as in [81]. Randomized algorithms [89] are used to estimate the leading eigenpairs $\{(\lambda_j, v_j)\}_j^k$ of the generalized eigenvalue problem

$$\mathbf{H}_{\text{misfit}} v_j = \lambda_j \mathbf{\Gamma}_{\text{pr}}^{-1} v_j. \quad (20)$$

The eigenvectors associated with the dominant eigenvalues represent the directions in parameter space most informed by the data. The covariance matrix in Eq. (19) is approximated using the Sherman-Morrison-Woodbury identity as in [81],

$$\mathbf{\Gamma}_{\text{post}} \approx \mathbf{\Gamma}_{\text{pr}} - \sum_{j=1}^k \frac{\lambda_j}{1 + \lambda_j} v_j v_j^T. \quad (21)$$

Since the data inform the model in an finite dimensional subspace, k is much smaller and independent of the discretized parameter dimension [30]. Thus, the low-rank construction of the Laplace approximation to the posterior distribution is an algorithmically scalable approach for high-dimensional Bayesian inverse problems. We employ `hippylib` [78, 90, 91] for implementation of the Newton-CG solver as well as the Matérn prior.

3.5. Propagation of uncertainty

Ultimately, we wish to use the model (1) to *predict* the future evolution of the tumor and compute meaningful quantities of interest (QoIs). Furthermore, we are interested in the propagating uncertainty in the parameters through the forward model to generate predictive distributions of the QoIs. To this end, we define the prediction map, $\mathcal{F}_{\text{pred}}$, where the tumor state is computed at some final time t_{pred} ,

$$\mathcal{F}_{\text{pred}} : \mathcal{M} \rightarrow \mathcal{U} := u(x, t_{\text{pred}}), \quad \text{s.t. } r(u, m; u_0, (t_0, t_{\text{pred}})) = 0. \quad (22)$$

Once more r denotes the PDE model in residual form, this time explicitly noting the dependence on parameter m , simulation window (t_0, t_{pred}) and initial condition u_0 . For example, in the prediction setting one may wish to estimate the initial condition u_0 from the last acquired MRI observation and predict over a new simulation window representing, say, one month into the future. Further, consider a generic quantity of interest $q : \mathcal{U} \rightarrow \mathcal{Q}$ for some suitable space \mathcal{Q} . We are interested in computing the pushforward of a measure through the composition of the QoI map with the prediction map, $(q \circ \mathcal{F}_{\text{pred}})_{\#}\nu$. Here $\#$ denotes the pushforward of a measurable function [92]. Mechanically, we sample $m \sim \nu$ and compute $q(\mathcal{F}_{\text{pred}}(m))$. This is called the prior predictive distribution when $\nu = \nu_{\text{pr}}$ and the posterior predictive distribution when $\nu = \nu_{\text{post}}$.

We consider four QoIs: total tumor cellularity, tumor volume, the concordance correlation coefficient (CCC), and the Dice similarity coefficient. We first define a threshold for the measurable tumor region, $\hat{u}(x, t) := u(x, t) > \bar{u}$. The threshold \bar{u} is the volume fraction at which the lesion is measurable and is taken to be 0.1 for this study. The total tumor cellularity is computed as the integration of the measurable tumor volume fraction over the domain, $q_{\text{TTC}} := \int_{\Omega} \hat{u} dx$, and may optionally be scaled by the carrying capacity to recover the number of tumor cells. The tumor volume is computed by integrating an indicator function for regions with measurable tumor over the domain, $q_{\text{TV}} := \int_{\Omega} \mathbb{1}_{\hat{u}} dx$. The CCC [93] measures correlation between two datasets, with a penalty for deviating from the line of unity. This is used to quantify voxel-wise agreement, for example, between the observed data and the model prediction. The Dice similarity coefficient [94, 95] measures the degree of spatial overlap and is used to compare predicted measurable tumor with the observed tumor. The Dice coefficient is defined as $2|X \cap Y|/(|X| + |Y|)$ for two sets X, Y with $|\cdot|$ denoting the cardinality, in this case the sets are the indicator functions for the predicted measurable tumor \hat{u} and the true, observed tumor u^{\dagger} .

4. Data-to-decisions: Demonstration on a virtual patient

As an initial demonstration of the digital twin bidirectional feedback loop illustrated in Figure 1, we deploy the methodology on a virtual patient. This setting allows for control of the experimental setup, including the underlying tumor growth mechanism, noise level, and frequency of observation. We first detail the data generation process and verify the inverse problem methodology in the controlled setting. Further, we assess the potential benefit of alternative imaging regimens on forecasts of tumor growth. To support the envisioned clinical deployment of the approach, we also establish the computational tractability.

4.1. Synthetic data generation

We utilize the publicly available UPENN-GBM dataset [60, 61]. In particular, we generate a computational geometry of subject 101 following the procedure outlined in Section 2.2. Observational data is synthesized by solving Eq. (1) forward from the initial condition for two weeks of untreated growth followed by the Stupp protocol [13] (six weeks of concurrent radiotherapy and

chemotherapy), followed by another month of untreated growth, that is, without adjuvant temozolomide. Three different scenarios are considered in which observational data are collected daily, weekly, or fortnightly. The prediction window is one month after the last observation at the conclusion of therapy. Snapshots of the tumor progression are shown in Figure 3.

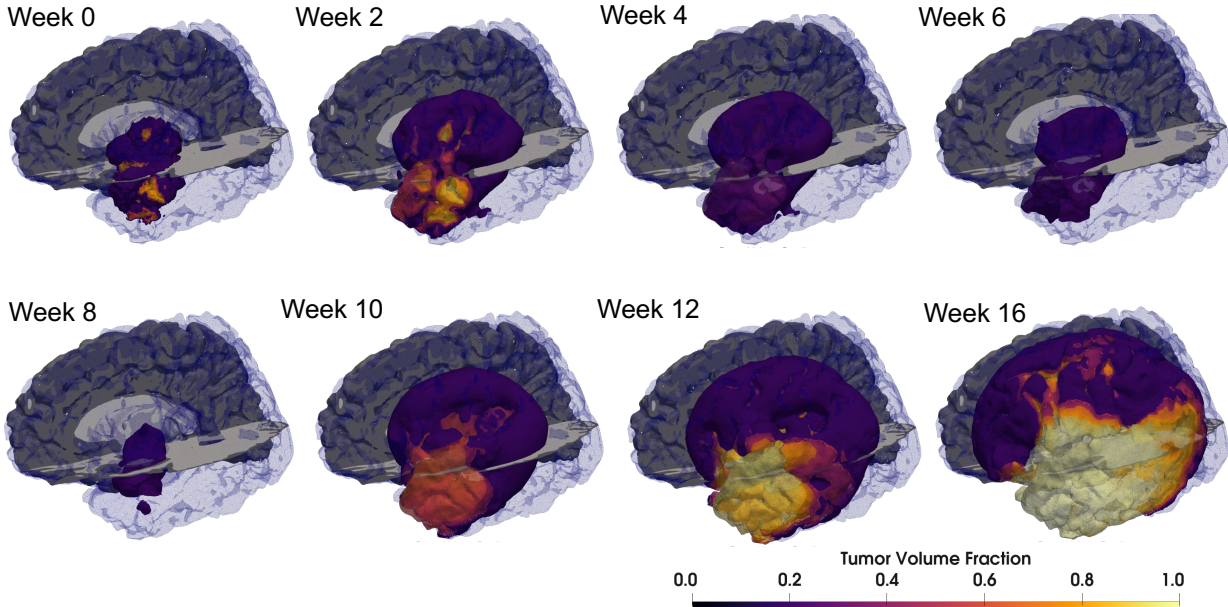


Figure 3: Snapshots of synthetic tumor progression for UPENN-GBM subject 101. Note the heterogeneous initial state and the retreat while under therapy (weeks 2-8). When therapy ends, tumor recurrence is swift and the tumor extent is larger at the prediction time (week 16) than during the imaging window (up to week 12).

The coarsest computational mesh is comprised of 108,569 vertices and is uniformly refined resulting in a finer mesh of 827,490 vertices which is then used for data generation. First order Lagrange elements are used for the state, resulting in an equal number of degrees of freedom (DOFs). Since the UPENN-GBM dataset lacks ADC estimates, we follow [96] and take the tumor volume fraction is taken to be 0.8 and 0.16 in the enhancing and non-enhancing regions of the tumor, respectively. The non-enhancing region is characterized by invasive and diffuse disease and the lower cellularity is chosen to account for the well-documented difficulties in separating proliferative tumor from other processes, such as vasogenic edema [62]. The underlying ground truth parameters are a diffusion coefficient of $0.03 \text{ mm}^2/\text{day}$ in gray matter and $0.3 \text{ mm}^2/\text{day}$ in white matter to account for preferential growth along fiber bundles in white matter [97], with a proliferation rate of 0.15 day^{-1} in the whole domain. These values are selected to be in biologically realistic ranges based on previous model calibration studies [20]. The radiosensitivity parameter ratio is set to $\alpha/\beta = 10 \text{ Gy}$ following values reported in the literature [42]. Additionally, we fix $\alpha_{rt} = 0.025 \text{ Gy}^{-1}$ and $\alpha_{ct} = 0.1 \text{ day}^{-1}$. We assume $\beta_{ct} = 24 \ln(2)/1.8 \text{ days}$ for the clearance rate of the temozolomide chemotherapy agent [98]. The model Eq. (1) is discretized and solved as detailed in Section 2.1. To generate observations, the finite-element representation of the state is interpolated onto the native T_1 voxel image space and the resultant voxel measurements are polluted with 2% Gaussian noise. A mesh convergence study with two successive refinements of the coarsest mesh was performed to estimate the discretization error, which was found to be within noise level applied to the data, such that the statistical uncertainty dominates.

4.2. Model verification

To avoid an inverse crime [99], the inverse problem is solved using the coarser mesh. First order Lagrange elements are used for the state, parameter, and adjoint variables. Additionally, unlike the true parameter which depends on the underlying gray / white matter tissue, the modeled log diffusivity field does not explicitly incorporate this information. The inversion parameters m_D, m_κ are modeled as Gaussian random fields as outlined in Section 3.3. The hyper-parameters defining the prior are reported in Table 2.

Table 2: Hyper-parameters for the Bayesian calibration of UPENN-GBM Subject 101.

Prior mean and variance of parameters				
m_D		m_κ		
log(mm^3/day)		log($1/day$)		
Mean	Variance	Mean	Variance	
-1.30	0.05	-1.00	0.02	
Spatial correlation lengths				
ρ_D (mm)		ρ_κ (mm)		
180		180		

The MAP reconstructions of the diffusion and proliferation coefficient fields for the daily imaging case are presented in Figure 4. In the reconstructed diffusion field, we observe that spatial heterogeneity arising from the underlying tissue type is captured. In the reconstructed proliferation rate, we do not see this effect, which is consistent with the true underlying parameter that is also homogeneous in the domain. Note in both cases that the MAP reconstruction is only well-informed where the tumor is active and the prior regularization dominates elsewhere. Furthermore, visualizing the spatial structure of the variance reduction stemming from the data-informed correction to the prior, that is, the diagonal of the low-rank term in (21), sheds light into the role of the observation operator and data collection strategy. In particular, the posterior features corrections where there is tumor is present.

While the underlying structure in the true parameters is reconstructed, inference of the exact values appears challenging. The diffusion and reaction terms drive tumor dynamics in a complementary fashion, exacerbating the ill-posedness of the inverse problem. The mutual information between the inferred parameters was computed under the Laplace approximation and corresponds to an estimated correlation coefficient [100, 101] of $\rho \approx 0.51$, indicating some level of correlation between the parameters. Theoretical analysis of the reaction-diffusion model has long focused on traveling wave solutions [102, 103, 104] and such techniques have been applied to suggest alternative parameterizations used to feasibly identify diffusivity in different tissues [105]. Data-driven methods offer a complementary approach for reparameterization and have been applied to similar models [106, 107] to alleviate structural identifiability [108]. Additionally, identifiability studies have been carried out for similar forms of the widely adopted reaction-diffusion model [16, 19, 109] indicating practical identifiability under the noise level in clinical data [110, 111]. Indeed, the reconstructed parameters lead to excellent reconstruction of the QoIs, as we will show in the next section.

4.3. Assessing the value of additional imaging

MRI data are highly informative, but expensive to collect and the resulting lack of information to assess response to therapy and progression may contribute to sub-optimal outcomes. One cannot accurately calibrate a model without data, nor can a clinician assess a patient’s condition. Thus,

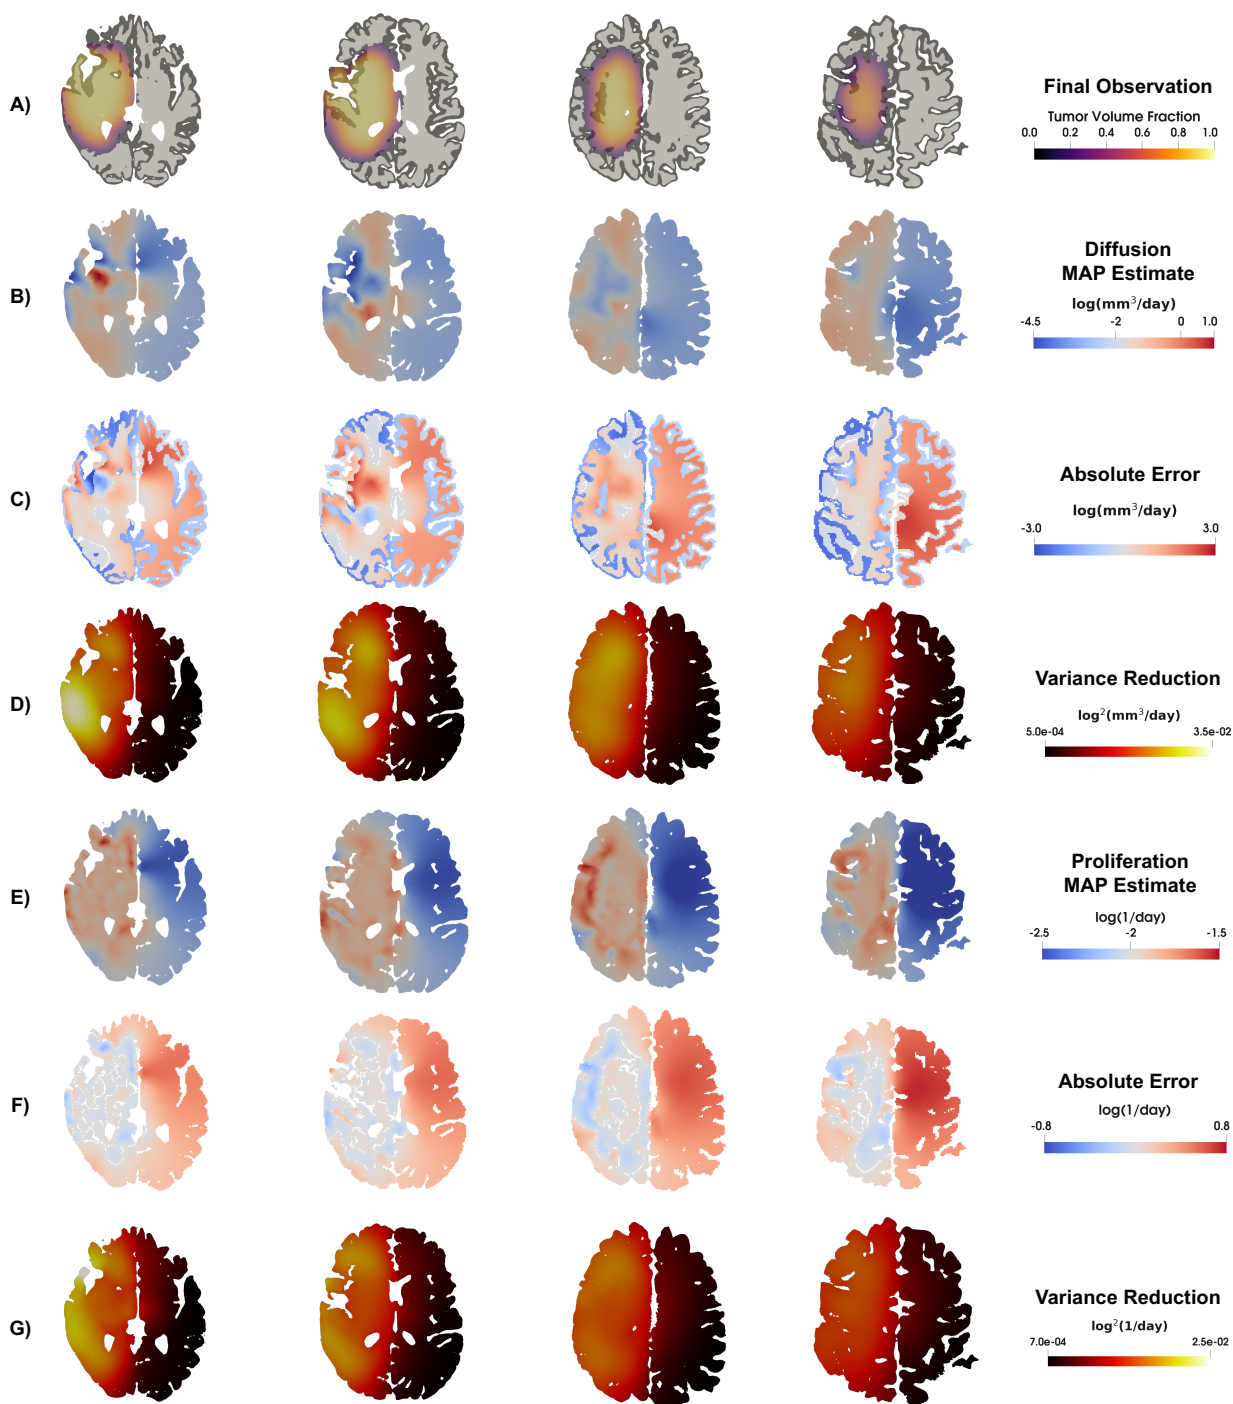


Figure 4: A: axial slices showing white and gray matter segmentation with tumor state at the final observation (week 12). B: MAP reconstruction of the log-diffusion field. C: absolute error in the reconstructed log-diffusion field. D: pointwise reduction in variance of the log-diffusion posterior distribution. E: MAP reconstruction of the log-reaction field. F: absolute error in the reconstructed log-reaction field. G: pointwise reduction in variance of the log-reaction posterior distribution. Each log parameter has discretized dimension 108,569. Note the inferred heterogeneity in the reconstructed log-diffusion field and spatial structure of the error and variance reduction indicating the parameters are well informed where the tumor is active.

there is a clear tension between the costs associated with additional imaging and the benefit of the additional information, for example, improving accuracy and reducing uncertainty in computational models used to predict patient prognosis. More accurate and precise models would offer better insights when tailoring interventions. We seek to better understand the effect of data availability on both parameter estimation and the prediction of clinically relevant QoIs.

Recall that Eq. (21) expresses the posterior covariance as a low-rank update to the prior covariance through the prior-orthogonal eigenvectors of the data-misfit Hessian given by Eq. (20). The associated eigenvalues quantify how much information is extracted from the data, while the spatial structure of the eigenvectors sheds insight into the directions in parameter space which are informed by the data. The spectrum of the prior-orthogonal Hessian, defined by Eq. (20), is computed for the first 50 eigenvalues and plotted for the three imaging frequency scenarios in Figure 5. As expected, the magnitude of eigenvalues is larger when the patient is observed more frequently, indicating larger corrections to the prior covariance when calculating the low-rank approximation to the posterior covariance (21). In other words, the directions in parameter space represented by the associated eigenvectors are better informed with more data. We remark that when the inverse problem is solved on the finer mesh, the spectrum for the two successively finer meshes lie on top of each other, indicating mesh independence of the spectrum.




Figure 5: Spectral decay of the prior-orthogonal Hessian for various imaging frequencies. The larger eigenvalues associated with more frequent imaging indicate more information gain in the corresponding eigenvectors.

Figure 6 displays several of the associated eigenvectors, when imaged daily. The localization of meaningful information observed in the MAP is again evident in the eigenvectors corresponding to the largest eigenvalues; that is, the modes most informed by the data. Additionally, we can see that eigenvectors corresponding to smaller eigenvalues capture higher frequency content from the data, though they do not impart as large of a correction to the prior covariance.

We assess the impact on the model’s predictive performance by computing two quantities of interest: the relative error in total tumor volume and concordance correlation coefficient, calculated as outlined in Section 3.5. The pushforward is computed by drawing 500 Monte Carlo samples from the prior, $\mathbf{m} \sim \nu_{\text{pr}}$, as well as the Laplace approximation to the posterior for the three cases, $\mathbf{m} \sim \nu_{\text{post}}^{\text{LA}}$. The sampled parameters are then used to simulate tumor growth from the end of the imaging to the one month prediction; i.e., we compute the pushforwards $(q_{\text{TV}} \circ \mathcal{F}_{\text{pred}})_{\#} \nu$ and $(q_{\text{CCC}} \circ \mathcal{F}_{\text{pred}})_{\#} \nu$, with $\nu = \nu_{\text{pr}}$ or $\nu_{\text{post}}^{\text{LA}}$. The pushforward distributions are reported in Figure 7 for




Figure 6: Prior-orthogonal eigenvectors of the data-misfit Hessian, $\mathbf{H}_{\text{misfit}}$, for UPENN-GBM subject 101 when imaged daily corresponding to the log-diffusion (top row) and log-reaction (bottom row) model parameters. From left to right the eigenvectors are shown in decreasing order of the corresponding eigenvalue, so that directions in parameter space most informed by the data are shown on the left. The different eigenvectors localize different regions of the brain, with those corresponding to smaller eigenvalues showing higher spatial frequency while also being less informed by the data.

the two QoIs.




Figure 7: Predictive distributions of (a) relative error in total tumor volume and (b) concordance correlation coefficient for various imaging schedules. The prior performs significantly ($p < 0.001$) worse in both QoIs than the calibrated models and is omitted from (b) to focus on the relative performance. Overall, calibration to patient data improves model predictive quality and significantly reduces variance ($p < 0.001$) in predictive distributions when compared to sampling from the prior distribution. All calibrated models capture the general tumor growth as measured by the volume, while the concordance correlation coefficient shows model accuracy increasing with additional observational data.

We observe that the predictive performance of the model improves in both QoIs as the patient is imaged more frequently and all models calibrated with patient-specific data significantly ($p < 0.001$) outperform the prior as determined by a Mann-Whitney U-test [112]. As anticipated, more information leads to better prediction of tumor status and significantly reduced variance in both QoIs ($p < 0.001$) in the predictive distributions compared to the prior as determined with Levene’s test [112]. For this particular study, the prior distribution underpredicts the true tumor volume

while all calibrations capture the true value. The CCC indicates an ordering between the calibrated models, with additional imaging improving performance, though there is a diminishing return as was observed in Figure 5.

These results illustrate that there is a tradeoff between information gain and imaging frequency. An exciting direction for future work would be to formulate an optimal experimental design problem to optimally choose the imaging frequency balancing improved model predictive quality with the cost of imaging. Related approaches have shown promise for one-dimensional radiotherapy planning [113, 114], however, in the PDE setting reduced order modeling [115] and multifidelity methods [116] will be crucial to ensure tractability of the resultant optimization problem. Moreover, coupling the experimental design question to therapy optimization [117, 118, 119] will be crucial for assessing potential clinical impact. In the presence of uncertainty, this will require characterization of risk and optimization of statistical properties of the QoIs [24].

4.4. Enabling computational tractability in clinically relevant timeframes

It is crucial to deliver insights in a clinically relevant timeframe. While the methods developed in Section 3.4 scale with the intrinsic difficulty of the inverse problem rather than the discretization, the computational burden of the nonlinear forward PDE solve and gradient and Hessian actions is still of concern. We need an efficient, parallelized implementation of the forward model to complement the algorithmic scalability. To this end, we interface with PETSc [120] to extend `hIPPYlib` [78] and `FEniCS` [121] to enable fast, distributed solution on high-performance computing resources. In particular, by interfacing directly with PETSc, one gains access to high-quality solvers and a wide variety of preconditioners useful for fast solution of the forward model; for example, we utilize `BoomerAMG` [122]. Furthermore, implementation of the forward model in `FEniCS` reduces the overhead associated with exploring models other than the reaction-diffusion Eq. (1) to the ability to write the variational form.

Scalability studies are performed on the Frontera supercomputer at the Texas Advanced Computing Center (TACC) [123] and are reported in Figure 8. For the scalability study, a tumor seed defined by a Gaussian function is grown within a 100 mm^3 domain with 1 million DOFs (for the intra-node case) and 4 million DOFs (for the inter-node case). First-order Lagrange elements are used for the state, parameter, and adjoint variable. An implicit Euler discretization is used in time with a step size of one day. We achieve near perfect strong scaling within node as well as good scaling across nodes up to 32 nodes. Similar performance is achieved for the linear adjoint solves and is reported in Appendix B. Strong scalability is of the greatest importance in this context as we wish to accelerate computation on a pre-defined, fixed computational domain. However, a weak-scaling study was also performed and demonstrated adequate scalability, with results reported in Appendix B. The code for the forward and inverse solvers is publicly available at: <https://github.com/gtpash/dt4co>.

5. Model validation and uncertainty quantification in the clinical setting

Model validation studies play a key role in establishing the clinical relevance of predictive science and certifying the computational models underpinning a digital twin. We further demonstrate the high-dimensional Bayesian model calibration process on a subset of patients with historical longitudinal MRI data from the publicly available IvyGAP dataset [62, 63]. We first detail the data and experimental setup. Performance across the cohort is then assessed by computing pushforward distributions for relevant QoIs in a variety of prediction settings. Finally, we address the role of model inadequacy and give an outlook.




Figure 8: Strong scaling of the forward solve obtaining a one-month prediction in (a) 42.2 seconds on one node for a benchmark problem with 1 million DOFs and (b) 12.2 seconds on 32 nodes for a benchmark problem with 4 million DOFs.

5.1. Experimental setup and Bayesian model calibration

The IvyGAP dataset contains MRI data (including the types discussed in Section 2.2) acquired during the course of the patient’s treatment. Additionally, the dataset contains information about the radiotherapy and chemotherapy treatment schedules that the patients received. As the IvyGAP dataset does not contain tumor segmentations, the ONCOhabitats [124] tool is used to generate an initial segmentation of the enhancing and non-enhancing tumor regions. These regions are then manually corrected. Tumor cellularity and volume fraction are computed as in Section 2.2. In all results, each patient’s tumor state is estimated from data at the first imaging time, $d(\bar{x}, t_0)$. If there was a resection event, we use the first post-resection image in the dataset as the baseline. Additionally, we restrict the computational domain to only the hemisphere containing the lesion since all of the tumors considered in this study are unifocal and do not invade the contralateral hemisphere. This results in meshes with approximately 250,000 vertices for all cases in this study. The imaging and treatment timelines for the cohort are shown in Figure 9. We observe a wide variation in both imaging schedules and treatment regimens.

We infer the log diffusivity field, m_D , and the log proliferation rate, m_κ , of the reaction-diffusion model Eq. (1). Additionally, the log diffusivity field is assumed to be uncorrelated across the gray and white matter interface $m_D = m_{D,gm}\chi_{gm} + m_{D,wm}(1 - \chi_{gm})$ where the gray matter indicator function χ_{gm} is obtained from the tissue segmentation derived in Section 2.2. Each log parameter is modeled using the Gaussian random field prior detailed in Section 3.3. To determine the prior mean and variance for each of the parameters, an initial calibration of the cohort is performed where the parameters are modeled as scalar quantities. The correlation length is set to approximately the largest dimension of the domain, with larger correlation lengths allowed in white matter. The variance of the noise is taken to be equivalent to 6.25% noise, following the reported value in [29] for a similar data-acquisition and tumor volume fraction estimation pipeline. The hyper-parameters defining the priors are summarized in Table 3.

For model validation, the last image in the dataset is withheld from the calibration and is set aside to serve as a prediction target. Numerical experiments were performed on four nodes of TACC’s Frontera supercomputer. An implicit Euler discretization was used in time, with a time step of one day. The chemoradiation model Eq. (5) is fixed for all patients and implemented as in Section 4, with the exception that $\alpha_{ct} = 0.18 \text{ day}^{-1}$ to match the calibrated therapy effects reported in the literature [20]. A maximum of 50 Newton iterations are used to compute the MAP point by solving the nonlinear optimization problem Eq. (16) using the relative decrease in the norm of the gradient as a convergence criterion. The required time to compute the MAP point ranged

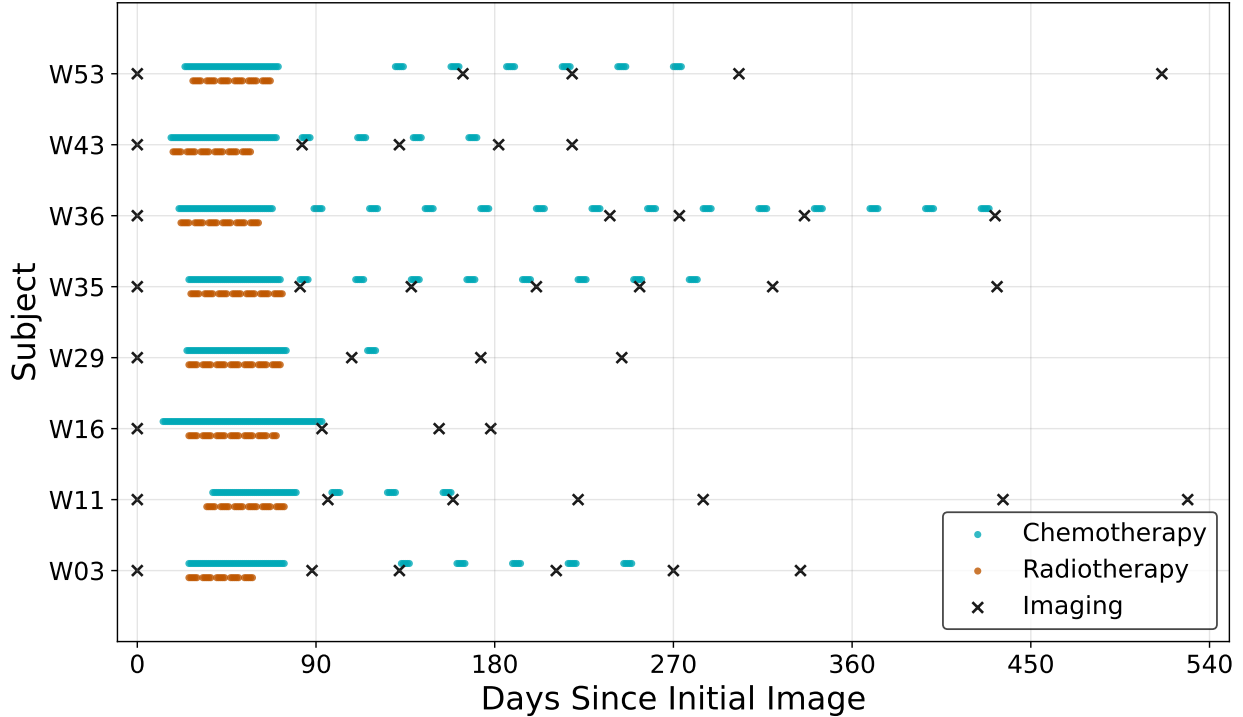


Figure 9: Imaging and treatment timelines for the IvyGAP cohort used for the model validation study. There is considerable variation amongst the patients in the number and frequency of data collection as well as the treatment schedules received.

Table 3: Estimated hyper-parameters for the Bayesian calibration.

Prior mean and variance of parameters							
m_κ		m_D		$m_{D,wm}$		$m_{D,gm}$	
$\log(1/day)$		$\log(mm^3/day)$		$\log(mm^3/day)$		$\log(mm^3/day)$	
Mean	Variance	Mean	Variance	Mean	Variance	Mean	Variance
-1.230	0.040	-1.167	0.115	-0.991	0.115	-1.467	0.115
Spatial correlation lengths and noise variance							
$\rho_\kappa (mm)$		$\rho_{D,gm} (mm)$		$\rho_{D,wm} (mm)$		σ_{noise}^2	
180		180		360		3.9e-3	

extensively due to the varied lengths of the patient imaging timelines and the number of required inner CG iterations, with a median wall clock time of 14.2 hours (IQR 10.4-18.3 hours). The low-rank approximation of the posterior covariance $\mathbf{\Gamma}_{post}$ given by Eq. (18) was computed with $r = 50$ and an oversampling factor of 10 for the randomized eigensolver and took approximately three hours in all cases.

5.2. Predictive performance on cohort

Model performance is assessed on the cohort by generating posterior predictive distributions of QoIs. In particular, the Dice similarity coefficient and relative error in total tumor cellularity are used and computed as described in Section 3.5. To mimic the envisioned clinical deployment, the initial condition is taken to be the penultimate scan (the last data seen during model calibration) and the inferred parameters are used to make a prediction for comparison at the time last scan. That

is, u_0 is determined from $d(\bar{x}, t_{n_t-1})$ and the simulation window is (t_{n_t-1}, t_f) with t_{n_t-1} denoting the time the second to last measurement. Figure 10 reports the predictive distributions, with 500 parameter samples from both the prior distribution and the low-rank based Laplace approximation to the posterior. Summary statistics are tabulated and reported in Appendix C. The reduced variance of the low-rank Laplace approximations is evident across both the cohort for both QoIs. While not all patients see a clear separation from the prior, such as W11 or W35, the posterior predictive distributions based on the Laplace approximation typically exhibit better spatial agreement with the observed tumor extent, as measured by the Dice coefficient. Additionally, the posterior predictive distributions appear to capture the total tumor cellularity well. Overall, the subject-specific calibrated model demonstrate a much improved ability to capture tumor progression.

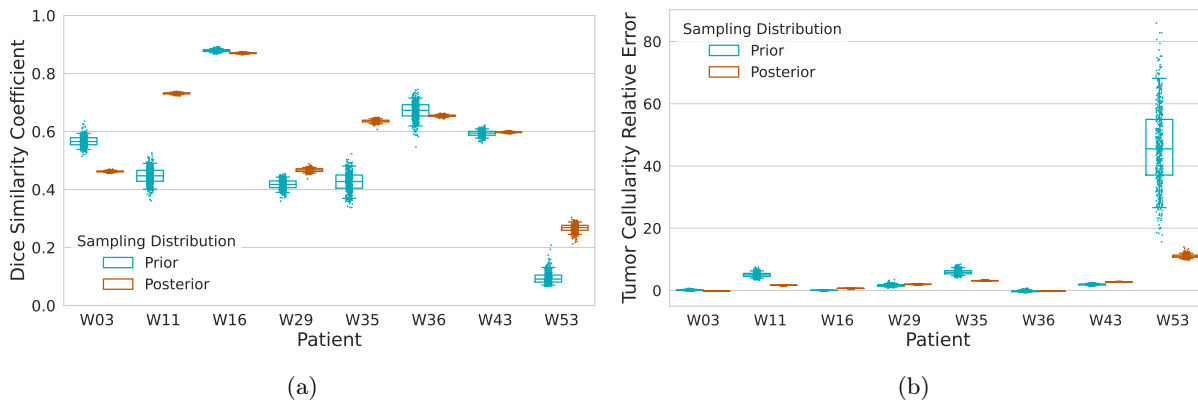


Figure 10: Prior and posterior predictive distributions of (a) Dice similarity coefficient and (b) total tumor cellularity for the IvyGAP patient cohort. The simulation window is from the last scan used for model calibration to the last acquired image.

To assess the stability of the inferred parameters, we consider the case where prediction is made from the first visit through the calibration window to the final observation. The prior and posterior pushforward distributions are presented in Figure 11. As before, summary statistics are reported in Appendix C. While tumor cores are unlikely to change over short horizons, the longer simulation window exacerbates accumulated errors as the model diverges from the observed data, and manifests in increased variance in the pushforwards. This also highlights structural deficiencies in the models calibrated to intermediate observations. We observe that in many cases, there is only a slight drop in the predictive performance of the posteriors over this longer horizon and conclude that the calibration does in fact provide a robust prediction of both the tumor shape (as measured by the Dice coefficient) and the intra-tumoral heterogeneity (as measured by the total tumor cellularity). This is in stark contrast to the prior distribution, where model inadequacy compounds over the long prediction horizon manifesting in large variability and drop in predictive performance.

5.3. Model inadequacy and opportunities

The predictive performance across the cohort was shown to be both stable over long time horizons and meaningfully informed by the data, especially given the sparse data collection. However, the model is far from a perfect match with reality. To better understand the limitations of the model and calibration, we focus on subject W43 as a representative patient. The patient’s disease progression is visualized in Figure 12. The drastic change in the tumor extent during the course of therapy, and especially between the last two snapshots clearly poses a significant modeling challenge. In particular, there are two phenomena that are difficult for the formulated model to capture: (1) the

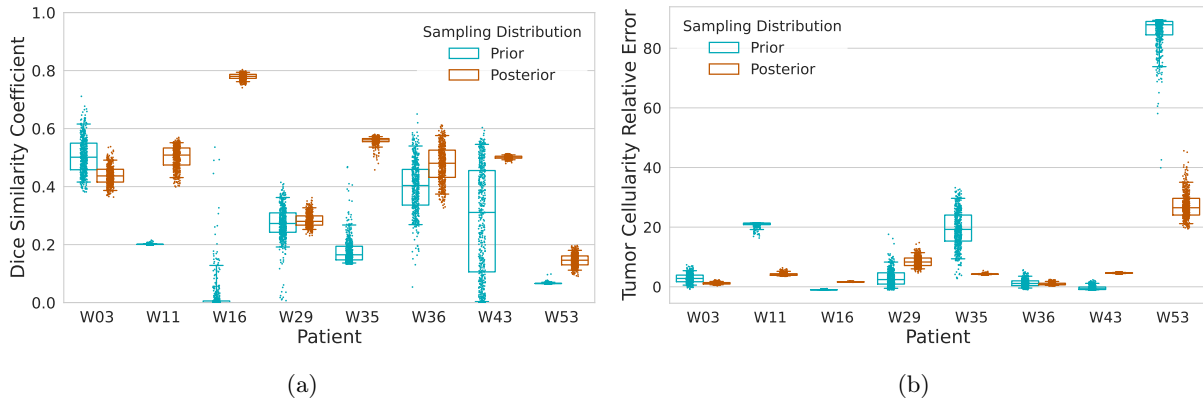


Figure 11: Posterior predictive distributions of (a) Dice similarity coefficient and (b) total tumor cellularity for the IvyGAP patient cohort. The simulation window is from the first acquired scan until the last acquired image, through the calibration window.

apparent retreat of the tumor in other regions while not under active treatment and (2) the rapid progression of disease into new regions.

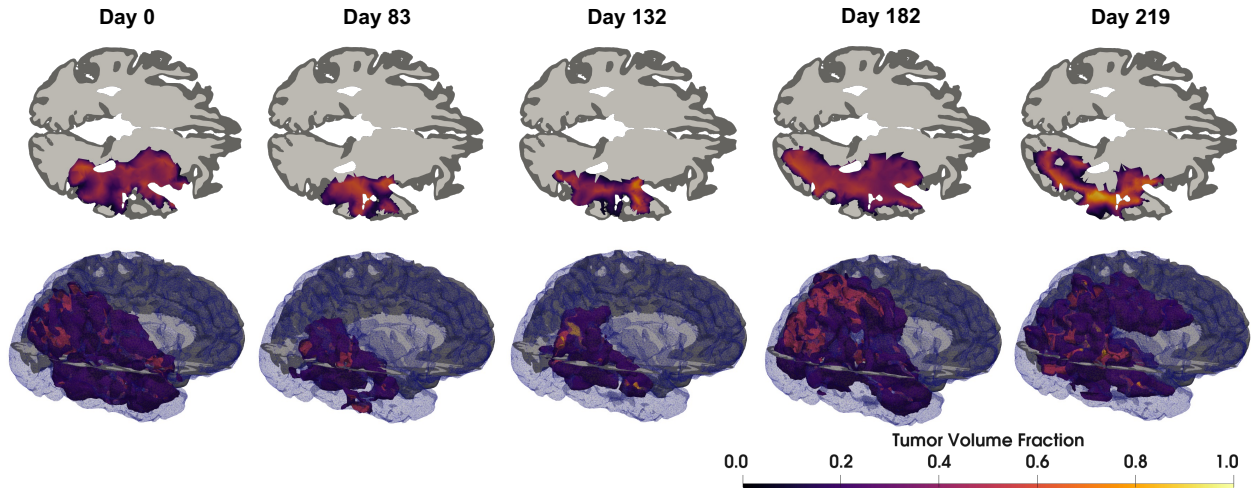


Figure 12: Observed disease progression of IvyGAP patient W43 throughout the course of treatment. The top row is an axial cross-sectional view near the middle of the tumor. The second row is a volumetric rendering of the computational domain, tumor, and axial slice. Observe the strong effect of chemoradiation treatment in the first three snapshots. In the final two snapshots, note the apparent tumor retreat in some areas despite the lack of therapy, while there is aggressive invasion in other regions.

Change in the non-enhancing tumor region is one biological explanation for the apparent retreat of the tumor between the final two snapshots. The non-enhancing tumor region is comprised of both edema and infiltration and it is notoriously difficult to accurately estimate the tumor cellularity [62]. For instance, changes in steroid dose and schedule can impact the presentation of the non-enhancing region of the tumor. Since the enhancing and non-enhancing regions are combined to estimate the tumor ROI, large variations may appear as aggressive invasion or rapid retreat in the data. The reaction-diffusion model employed may fail to capture these fluctuations in the non-enhancing tumor component since it cannot adequately resolve the complex dynamics related to treatment and biology in this region. Alternative imaging modalities that better resolve intra-tumoral heterogeneity should

be considered to improve data fidelity [125].

Furthermore, the fixed treatment model is not appropriate for every patient and contributes to model inadequacy. Consider the initial-to-final prediction case. The MAP point prediction of the tumor progression is presented in Figure 13. For this specific patient, we observe that the specified chemoradiation therapy model is too strong, leading to an unrealistic near total predicted remission at the second visit. Accounting for model errors as an additional source of uncertainty may widen the posterior credible interval when the model is inadequate [99]. Furthermore, patient-specific calibration of treatment models may improve predictive quality, but must be balanced with identifiability of the model parameters given the dearth of observational data.

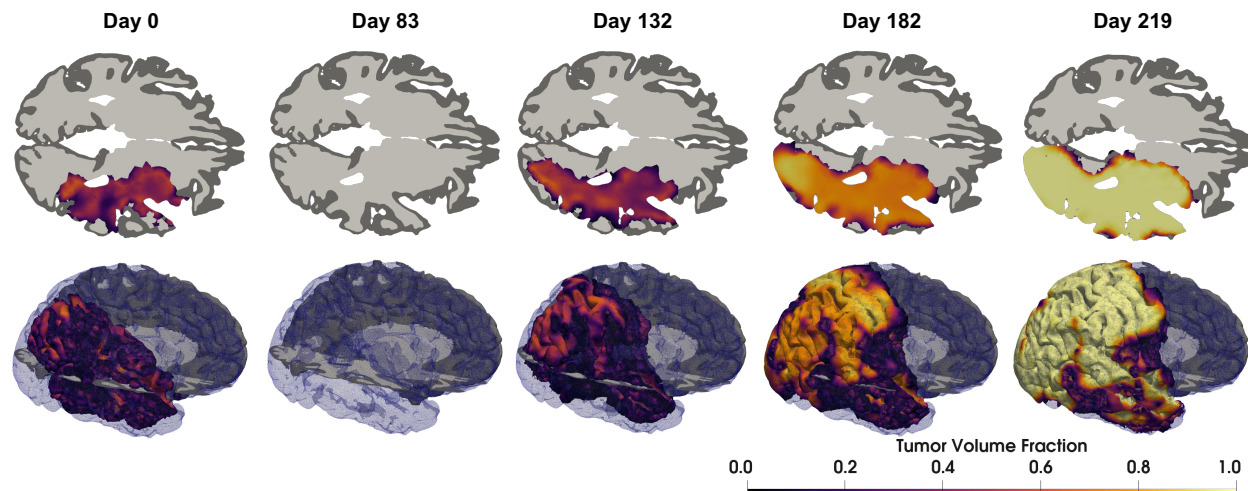


Figure 13: Predicted tumor progression from the first visit to the last visit for patient W43 using the MAP point. The top row is an axial cross-sectional view near the middle of the tumor. The second row is a volumetric rendering of the computational domain, tumor, and axial slice. The model overpredicts response to therapy compared to the observed data in this case (c.f. Figure 12). Additionally, the model overpredicts the recurrence and struggles to capture internal heterogeneity of the tumor at the prediction time.

While the total tumor cellularity results in Figure 10 indicate potentially better performance of the calibrated model to predict the internal growth dynamics of the tumor core, both the MAP and prior mean generate qualitatively similar predictions of tumor shape and extent for the last-to-final case, as shown in Figure 14. This highlights the challenging nature of extrapolation and motivates further development of predictive models to better capture these effects. One approach may be to generate patient-specific priors based on genetic subtype as recurrent gliomas typically present with similar molecular structure [126] and epigenetic classification has demonstrated prognostic value [127]. The reaction-diffusion model of tumor growth combines multiple tumor development mechanisms into its two biophysical parameters and is certainly inadequate to capture the full heterogeneity present. More complex models are necessary to improve accuracy, incorporating mass effect [128, 129], multi-species [20, 130], metabolic processes [131, 132], or vascular structure [133, 134], amongst others. Many of these models can readily be brought into the presented framework. Furthermore, development and calibration of more sophisticated models for radio- and chemo-therapy will be necessary to increase fidelity for optimization of therapeutic regimens.

6. Conclusion

This work developed an end-to-end Bayesian framework for the integration of MRI data with mathematical models of tumor growth to enable digital twins in precision oncology. High-fidelity

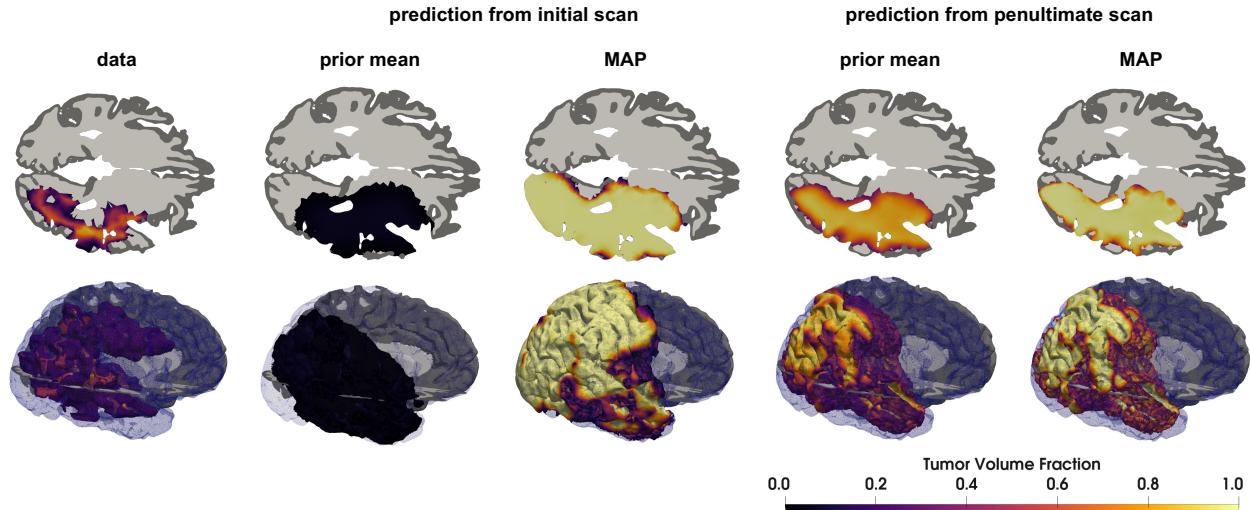


Figure 14: Predicted final tumor state for IvyGAP patient W43 using the MAP point and prior mean both from the initial and penultimate scans as initial condition. The top row is an axial cross-sectional view near the middle of the tumor. The second row is a volumetric rendering of the computational domain, tumor, and axial slice. The prior mean generates a poor prediction from the initial scan indicating that it is not appropriate for this patient. The short prediction horizon from the penultimate scan leads to qualitatively similar structure between the prior and calibrated models. The calibrated models generate qualitatively similar predictions both over the long and short prediction horizons.

computational representations of the patient anatomy are generated from MRI data. Patient-specific biophysical model parameters are inferred from longitudinally collected MRI measurements and account for uncertainty in the data acquisition and processing. The spatially varying model parameters are high-dimensional upon discretization and pose a significant computational challenge. To ensure tractability of the inverse problem, an efficient parallel implementation of the forward model was coupled with a scalable adjoint-based Newton-Krylov optimization algorithm. To overcome the prohibitive computational cost of performing MCMC to explore the posterior distribution, we employ the Laplace approximation to the posterior and make use of a low-rank approximation for rapid sampling. The Bayesian approach rigorously quantifies uncertainties in the inferred parameters, which are then propagated to predict future tumor growth and key clinical quantities of interest. The broad utility of the proposed approach was demonstrated by application to patients from two publicly available datasets. Moreover, the flexible implementation of the framework could readily be translated to other other organs, such as the breast [21] and prostate [135]. To the best of our knowledge, this is the first development and demonstration of such an end-to-end Bayesian pipeline on clinical data accounting for the complex anatomy of the human brain.

The methodology was verified to accurately capture intra-tumoral spatial heterogeneity in an *in silico* case study where the underlying growth mechanism was controlled. Patient-specific posteriors were shown to improve model predictive quality of key clinical metrics, while reducing variance in the model parameters. Alternate imaging schedules were implemented and used to quantitatively assess the information gain and value of additional imaging. Given the significant cost of MRI, these results motivate an interesting optimal experimental design question regarding how frequently MRI data should be collected as well as *when* patients should be imaged. To assess the reliability of the approach in the clinic, the model was calibrated to a cohort of patients using publicly available clinical data. The posterior predictive distributions demonstrated robust agreement both of tumor shape, as measured by the Dice coefficient, and structure, as measured by the relative error in

total tumor cellularity. This validation study also revealed future directions for model improvement and development to address key questions of model inadequacy to improve realism and predictive power. Together, these two studies establish the capacity of predictive science in oncology to provide a strong foundation for the development of digital twins.

As predictive science is adopted into clinical workflows, the demands placed upon high-fidelity simulations may outstrip the computational capabilities of even the most efficient, scalable algorithms and powerful hardware. Clinical timelines can rapidly evolve and temporal variations in clinical measurements may necessitate swift feedback and decision support from digital twins in complex scenarios on a compressed time frame. Moreover, truly realizing the potential of medical digital twins extends beyond patient-specific forecasting to personalized therapy optimization and data collection strategies which account for uncertainty [11]. These tasks may be formulated as optimization under uncertainty [136] or optimal experimental design [137] problems, which in turn require many queries of the forward model, and potentially, multiple solutions of the inverse problem. The development of surrogate models is a crucial and complementary research thrust. Such surrogates reduce the marginal cost of queries to the model, and would allow clinicians and practitioners to more quickly and thoroughly explore the design space and assess future scenarios of disease progression. In particular, scientific machine learning has seen an explosion in recent years [138, 139, 140, 141, 142, 143], and the development of certified data-driven surrogates for dynamical systems governed by PDEs which can be trained with limited data and provide uncertainty quantification is an exciting direction.

7. Acknowledgments

The authors thank Michael Kapteyn for insightful comments on an earlier draft of this work.

GP acknowledges support from the Office of Science, Advanced Scientific Computing Research, U.S. Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0021110. GP and KW acknowledge support from Department of Energy grant DE-SC002317, the National Science Foundation (NSF) grant 2436499, and NSF FDT-Biotech award 2436499. DAH acknowledges funding from CPRIT RP220225, NSF DMS 2436499, and the American Cancer Society IRG-21-135-01-IRG. TEY acknowledges funding from the National Cancer Institute R01CA235800, U24CA226110, U01CA174706. TEY is a CPRIT Scholar in Cancer Research. The authors acknowledge the Texas Advanced Computing Center (TACC) at The University of Texas at Austin for providing computational resources (Frontera) that have contributed to the research results reported within this paper. URL: <http://www.tacc.utexas.edu>

This report was prepared as an account of work sponsored by an agency of the United States Government. Neither the United States Government nor any agency thereof, nor any of their employees, makes any warranty, express or implied, or assumes any legal liability or responsibility for the accuracy, completeness, or usefulness of any information, apparatus, product, or process disclosed, or represents that its use would not infringe privately owned rights. Reference herein to any specific commercial product, process, or service by trade name, trademark, manufacturer, or otherwise does not necessarily constitute or imply its endorsement, recommendation, or favoring by the United States Government or any agency thereof.

8. CRediT authorship contribution statement

GP: Conceptualization, Methodology, Investigation, Software, Writing - original draft. UV: Conceptualization, Methodology, Writing - review & editing. DAH: Conceptualization, Data curation,

Writing - review & editing. TEY: Conceptualization, Writing - review & editing. KW: Conceptualization, Supervision, Writing - review & editing.

References

- [1] National Academies of Sciences, Engineering, and Medicine, Foundational research gaps and future directions for digital twins, National Academies Press, 2024.
- [2] A. Omuro, L. M. DeAngelis, Glioblastoma and other malignant gliomas: a clinical review, *JAMA* 310 (17) (2013) 1842–1850.
- [3] D. J. Aum, D. H. Kim, T. L. Beaumont, E. C. Leuthardt, G. P. Dunn, A. H. Kim, Molecular and cellular heterogeneity: The hallmark of glioblastoma, *Neurosurgical Focus* 37 (6) (2014) E11.
- [4] R. P. Hill, R. G. Bristow, A. Fyles, M. Koritzinsky, M. Milosevic, B. G. Wouters, Hypoxia and predicting radiation response, in: *Seminars in Radiation Oncology*, Vol. 25, Elsevier, 2015, pp. 260–272.
- [5] T. E. Yankeelov, D. A. Hormuth, E. A. Lima, G. Lorenzo, C. Wu, L. C. Okereke, G. M. Rauch, A. M. Venkatesan, C. Chung, Designing clinical trials for patients who are not average, *Iscience* 27 (1) (2024).
- [6] P. Y. Wen, D. R. Macdonald, D. A. Reardon, T. F. Cloughesy, A. G. Sorensen, E. Galanis, J. DeGroot, W. Wick, M. R. Gilbert, A. B. Lassman, et al., Updated response assessment criteria for high-grade gliomas: response assessment in neuro-oncology working group, *Journal of Clinical Oncology* 28 (11) (2010) 1963–1972.
- [7] U. N. Chukwueke, P. Y. Wen, Use of the response assessment in neuro-oncology (RANO) criteria in clinical trials and clinical practice, *CNS Oncology* 8 (1) (2019) CNS28.
- [8] T. Hernandez-Boussard, P. Macklin, E. J. Greenspan, A. L. Gryshuk, E. Stahlberg, T. Syeda-Mahmood, I. Shmulevich, Digital twins for predictive oncology will be a paradigm shift for precision cancer care, *Nature Medicine* 27 (12) (2021) 2065–2066.
- [9] C. Wu, G. Lorenzo, D. A. Hormuth, E. A. Lima, K. P. Slavkova, J. C. DiCarlo, J. Virostko, C. M. Phillips, D. Patt, C. Chung, T. E. Yankeelov, Integrating mechanism-based modeling with biomedical imaging to build practical digital twins for clinical oncology, *Biophysics Reviews* 3 (2) (2022).
- [10] C. A. Taylor, K. Petersen, N. Xiao, M. Sinclair, Y. Bai, S. R. Lynch, A. UpdePac, M. Schaap, Patient-specific modeling of blood flow in the coronary arteries, *Computer Methods in Applied Mechanics and Engineering* 417 (2023) 116414.
- [11] R. Laubenbacher, B. Mehrad, I. Shmulevich, N. Trayanova, Digital twins in medicine, *Nature Computational Science* 4 (3) (2024) 184–191.
- [12] Y. Wang, T. Jiang, Understanding high grade glioma: molecular mechanism, therapy and comprehensive management, *Cancer Letters* 331 (2) (2013) 139–146.
- [13] R. Stupp, W. P. Mason, M. J. Van Den Bent, M. Weller, B. Fisher, M. J. Taphoorn, K. Belanger, A. A. Brandes, C. Marosi, U. Bogdahn, et al., Radiotherapy plus Concomitant and Adjuvant Temozolomide for Glioblastoma, *New England Journal of Medicine* 352 (10) (2005) 987–996.

- [14] A. C. Tan, D. M. Ashley, G. Y. López, M. Malinzak, H. S. Friedman, M. Khasraw, Management of glioblastoma: State of the art and future directions, *CA: a Cancer Journal for Clinicians* 70 (4) (2020) 299–312.
- [15] C. Hoguea, C. Davatzikos, G. Biros, An image-driven parameter estimation problem for a reaction–diffusion glioma growth model with mass effects, *Journal of Mathematical Biology* 56 (2008) 793–825.
- [16] P. R. Jackson, J. Juliano, A. Hawkins-Daarud, R. C. Rockne, K. R. Swanson, Patient-specific mathematical neuro-oncology: using a simple proliferation and invasion tumor model to inform clinical practice, *Bulletin of Mathematical Biology* 77 (5) (2015) 846–856.
- [17] J. T. Oden, E. A. Lima, R. C. Almeida, Y. Feng, M. N. Rylander, D. Fuentes, D. Faghihi, M. M. Rahman, M. DeWitt, M. Gadde, et al., Toward predictive multiscale modeling of vascular tumor growth: computational and experimental oncology for tumor prediction, *Archives of Computational Methods in Engineering* 23 (2016) 735–779.
- [18] K. Scheufele, S. Subramanian, A. Mang, G. Biros, M. Mehl, Image-driven biophysical tumor growth model calibration, *SIAM Journal on Scientific Computing* 42 (3) (2020) B549.
- [19] A. Mang, S. Bakas, S. Subramanian, C. Davatzikos, G. Biros, Integrated biophysical modeling and image analysis: application to neuro-oncology, *Annual Review of Biomedical Engineering* 22 (1) (2020) 309–341.
- [20] D. A. Hormuth II, K. A. Al Feghali, A. M. Elliott, T. E. Yankeelov, C. Chung, Image-based personalization of computational models for predicting response of high-grade glioma to chemoradiation, *Scientific Reports* 11 (1) (2021) 8520.
- [21] A. M. Jarrett, A. S. Kazerouni, C. Wu, J. Virostko, A. G. Sorace, J. C. DiCarlo, D. A. Hormuth, D. A. Ekrut, D. Patt, B. Goodgame, et al., Quantitative magnetic resonance imaging and tumor forecasting of breast cancer patients in the community setting, *Nature Protocols* 16 (11) (2021) 5309–5338.
- [22] M. S. Copur, R. Ramaekers, M. Gönen, M. Gulzow, R. Hadenfeldt, C. Fuller, J. Scott, S. Einspahr, H. Benzel, M. Mickey, et al., ReCAP: Impact of the National Cancer Institute Community cancer centers program on clinical trial and related activities at a community cancer Center in Rural Nebraska, *Journal of Oncology Practice* 12 (1) (2016) 67–68.
- [23] A. Hawkins-Daarud, S. Prudhomme, K. G. van der Zee, J. T. Oden, Bayesian calibration, validation, and uncertainty quantification of diffuse interface models of tumor growth, *Journal of Mathematical Biology* 67 (6) (2013) 1457–1485.
- [24] D. P. Kouri, T. M. Surowiec, Risk-averse PDE-constrained optimization using the conditional value-at-risk, *SIAM Journal on Optimization* 26 (1) (2016) 365–396.
- [25] M. Lê, H. Delingette, J. Kalpathy-Cramer, E. R. Gerstner, T. Batchelor, J. Unkelbach, N. Ayache, MRI based Bayesian personalization of a tumor growth model, *IEEE Transactions on Medical Imaging* 35 (10) (2016) 2329–2339.
- [26] E. Lima, J. Oden, B. Wohlmuth, A. Shahmoradi, D. Hormuth II, T. Yankeelov, L. Scarabosio, T. Horger, Selection and validation of predictive models of radiation effects on tumor growth based on noninvasive imaging data, *Computer Methods in Applied Mechanics and Engineering* 327 (2017) 277–305.

- [27] J. Lipková, P. Angelikopoulos, S. Wu, E. Alberts, B. Wiestler, C. Diehl, C. Preibisch, T. Pyka, S. E. Combs, P. Hadjidakas, et al., Personalized radiotherapy design for glioblastoma: integrating mathematical tumor models, multimodal scans, and Bayesian inference, *IEEE Transactions on Medical Imaging* 38 (8) (2019) 1875–1884.
- [28] A. Chaudhuri, G. Pash, D. A. Hormuth, G. Lorenzo, M. Kapteyn, C. Wu, E. A. Lima, T. E. Yankeelov, K. Willcox, Predictive digital twin for optimizing patient-specific radiotherapy regimens under uncertainty in high-grade gliomas, *Frontiers in Artificial Intelligence* 6 (2023) 1222612.
- [29] B. Liang, J. Tan, L. Lozenski, D. A. Hormuth, T. E. Yankeelov, U. Villa, D. Faghihi, Bayesian inference of tissue heterogeneity for individualized prediction of glioma growth, *IEEE Transactions on Medical Imaging* 42 (10) (2023) 2865–2875.
- [30] O. Ghattas, K. Willcox, Learning physics-based models from data: perspectives from inverse problems and model reduction, *Acta Numerica* 30 (2021) 445–554.
- [31] J. Alfonso, K. Talkenberger, M. Seifert, B. Klink, A. Hawkins-Daarud, K. Swanson, H. Hatzikirou, A. Deutsch, The biology and mathematical modelling of glioma invasion: a review, *Journal of the Royal Society Interface* 14 (136) (2017) 20170490.
- [32] M. Chaplain, Avascular growth, angiogenesis and vascular growth in solid tumours: The mathematical modelling of the stages of tumour development, *Mathematical and Computer Modelling* 23 (6) (1996) 47–87.
- [33] J. D. Murray, *Mathematical Biology I: An introduction*, Vol. 17, Springer Science & Business Media, 2002.
- [34] J. D. Murray, *Mathematical Biology II: Spatial Models and Biomedical Applications*, Vol. 18, Springer Science & Business Media, 2003.
- [35] O. Clatz, M. Sermesant, P.-Y. Bondiau, H. Delingette, S. K. Warfield, G. Malandain, N. 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) (2005) 1334–1346.
- [36] H. L. Harpold, E. C. Alvord Jr, K. R. Swanson, The evolution of mathematical modeling of glioma proliferation and invasion, *Journal of Neuropathology & Experimental Neurology* 66 (1) (2007) 1–9.
- [37] I. Garg, M. I. Miga, Preliminary investigation of the inhibitory effects of mechanical stress in tumor growth, in: *Medical Imaging 2008: Visualization, Image-Guided Procedures, and Modeling*, Vol. 6918, SPIE, 2008, pp. 854–864.
- [38] D. A. Hormuth II, J. A. Weis, S. L. Barnes, M. I. Miga, E. C. Rericha, V. Quaranta, T. E. Yankeelov, Predicting in vivo glioma growth with the reaction diffusion equation constrained by quantitative magnetic resonance imaging data, *Physical Biology* 12 (4) (2015) 046006.
- [39] G. Borasi, A. Nahum, Modelling the radiotherapy effect in the reaction-diffusion equation, *Physica Medica* 32 (9) (2016) 1175–1179.
- [40] B. Douglas, J. Fowler, The effect of multiple small doses of X rays on skin reactions in the mouse and a basic interpretation, *Radiation Research* 66 (2) (1976) 401–426.

- [41] S. J. McMahon, The linear quadratic model: usage, interpretation and challenges, *Physics in Medicine & Biology* 64 (1) (2018) 01TR01.
- [42] R. Rockne, E. Alvord, J. Rockhill, K. Swanson, A mathematical model for brain tumor response to radiation therapy, *Journal of Mathematical Biology* 58 (2009) 561–578.
- [43] D. A. Hormuth, A. M. Jarrett, T. E. Yankeelov, Forecasting tumor and vasculature response dynamics to radiation therapy via image based mathematical modeling, *Radiation Oncology* 15 (2020) 1–14.
- [44] A. M. Jarrett, D. A. Hormuth II, C. Wu, A. S. Kazerouni, D. A. Ekrut, J. Virostko, A. G. Sorace, J. C. DiCarlo, J. Kowalski, D. Patt, B. Goodgame, S. Avery, T. E. Yankeelov, Evaluating patient-specific neoadjuvant regimens for breast cancer via a mathematical model constrained by quantitative magnetic resonance imaging data, *Neoplasia* 22 (12) (2020) 820–830.
- [45] T. J. Hughes, *The finite element method: linear static and dynamic finite element analysis*, Courier Corporation, 2003.
- [46] J. T. Oden, *Finite elements of nonlinear continua*, Courier Corporation, 2006.
- [47] T. E. Yankeelov, J. C. Gore, Dynamic contrast enhanced magnetic resonance imaging in oncology: theory, data acquisition, analysis, and examples, *Current Medical Imaging* 3 (2) (2007) 91–107.
- [48] A. R. Padhani, G. Liu, D. Mu-Koh, T. L. Chenevert, H. C. Thoeny, T. Takahara, A. Dzik-Jurasz, B. D. Ross, M. Van Cauteren, D. Collins, D. A. Hammoud, G. J. S. Rustin, B. Taouli, P. L. Choyke, Diffusion-weighted magnetic resonance imaging as a cancer biomarker: consensus and recommendations, *Neoplasia* 11 (2) (2009) 102–125.
- [49] J. D. Bernstock, S. E. Gary, N. Klinger, P. A. Valdes, W. Ibn Essayed, H. E. Olsen, G. Chagoya, G. Elsayed, D. Yamashita, P. Schuss, et al., Standard clinical approaches and emerging modalities for glioblastoma imaging, *Neuro-Oncology Advances* 4 (1) (2022) vda080.
- [50] A. Karellas, S. Vedantham, Breast cancer imaging: a perspective for the next decade, *Medical Physics* 35 (11) (2008) 4878–4897.
- [51] H. Hricak, P. L. Choyke, S. C. Eberhardt, S. A. Leibel, P. T. Scardino, Imaging prostate cancer: a multidisciplinary perspective, *Radiology* 243 (1) (2007) 28–53.
- [52] F. Kofler, C. Berger, D. Waldmannstetter, J. Lipkova, I. Ezhov, G. Tetteh, J. Kirschke, C. Zimmer, B. Wiestler, B. H. Menze, Brats toolkit: translating brats brain tumor segmentation algorithms into clinical and scientific practice, *Frontiers in Neuroscience* 14 (2020) 125.
- [53] B. M. Ellingson, P. S. LaViolette, S. D. Rand, M. G. Malkin, J. M. Connelly, W. M. Mueller, R. W. Prost, K. M. Schmainda, Spatially quantifying microscopic tumor invasion and proliferation using a voxel-wise solution to a glioma growth model and serial diffusion MRI, *Magnetic Resonance in Medicine* 65 (4) (2011) 1131–1143.
- [54] S. Angeli, K. E. Emblem, P. Due-Tonnessen, T. Stylianopoulos, Towards patient-specific modeling of brain tumor growth and formation of secondary nodes guided by DTI-MRI, *NeuroImage: Clinical* 20 (2018) 664–673.

- [55] B. M. Ellingson, T. F. Cloughesy, A. Lai, P. L. Nghiemphu, W. B. Pope, Cell invasion, motility, and proliferation level estimate (CIMPLE) maps derived from serial diffusion MR images in recurrent glioblastoma treated with bevacizumab, *Journal of Neuro-Oncology* 105 (1) (2011) 91–101.
- [56] S. Klein, M. Staring, K. Murphy, M. A. Viergever, J. P. Pluim, Elastix: a toolbox for intensity-based medical image registration, *IEEE Transactions on Medical Imaging* 29 (1) (2009) 196–205.
- [57] D. P. Shamonin, E. E. Bron, B. P. Lelieveldt, M. Smits, S. Klein, M. Staring, A. D. N. Initiative, Fast parallel image registration on CPU and GPU for diagnostic classification of Alzheimer’s disease, *Frontiers in Neuroinformatics* 7 (2014) 50.
- [58] P. Thévenaz, M. Unser, Optimization of mutual information for multiresolution image registration, *IEEE Transactions on Image Processing* 9 (12) (2000) 2083–2099.
- [59] M. A. Viergever, J. A. Maintz, S. Klein, K. Murphy, M. Staring, J. P. Pluim, A survey of medical image registration—under review (2016).
- [60] S. Bakas, C. Sako, H. Akbari, M. Bilello, A. Sotiras, G. Shukla, J. D. Rudie, N. F. Santamaría, A. F. Kazerooni, S. Pati, et al., The University of Pennsylvania glioblastoma (UPENN-GBM) cohort: advanced MRI, clinical, genomics, & radiomics, *Scientific Data* 9 (1) (2022) 453.
- [61] S. Bakas, C. Sako, H. Akbari, M. Bilello, A. Sotiras, G. Shukla, J. D. Rudie, N. Flores Santamaria, A. Fathi Kazerooni, S. Pati, S. Rathore, A. Mamourian, S. M. Ha, W. Parker, J. Doshi, U. Baid, M. Bergman, Z. A. Binder, R. Verma, C. Davatzikos, Multiparametric magnetic resonance imaging (mpMRI) scans for de novo glioblastoma (GBM) patients from the University of Pennsylvania Health System (UPENN-GBM), Dataset (2021). doi:10.7937/TCIA.709X-DN49. URL <https://doi.org/10.7937/TCIA.709X-DN49>
- [62] R. B. Puchalski, N. Shah, J. Miller, R. Dalley, S. R. Nomura, J.-G. Yoon, K. A. Smith, M. Lankerovich, D. Bertagnolli, K. Bickley, et al., An anatomic transcriptional atlas of human glioblastoma, *Science* 360 (6389) (2018) 660–663.
- [63] N. Shah, X. Feng, M. Lankerovich, R. B. Puchalski, B. Keogh, Data from Ivy Glioblastoma Atlas Project (IvyGAP), Dataset (2016). doi:10.7937/K9/TCIA.2016.XLwaN6nL. URL <https://doi.org/10.7937/K9/TCIA.2016.XLwaN6nL>
- [64] J. Juan-Albarracín, E. Fuster-Garcia, J. V. Manjon, M. Robles, F. Aparici, L. Martí-Bonmatí, J. M. Garcia-Gomez, Automated glioblastoma segmentation based on a multiparametric structured unsupervised classification, *PloS One* 10 (5) (2015) e0125143.
- [65] M. C. Mabray, R. F. Barajas, S. Cha, Modern brain tumor imaging, *Brain Tumor Research and Treatment* 3 (1) (2015) 8–23.
- [66] K.-A. Mardal, M. E. Rognes, T. B. Thompson, L. M. Valnes, *Mathematical modeling of the human brain: from magnetic resonance images to finite element simulation*, Vol. 10, Springer Nature, 2022.
- [67] B. Fischl, FreeSurfer, *Neuroimage* 62 (2) (2012) 774–781.

- [68] A. M. Radwan, L. Emsell, J. Blommaert, A. Zhylyka, S. Kovacs, T. Theys, N. Sollmann, P. Dupont, S. Sunaert, Virtual brain grafting: Enabling whole brain parcellation in the presence of large lesions, *NeuroImage* 229 (2021) 117731.
- [69] A. Fabri, S. Pion, CGAL: The Computational Geometry Algorithms Library, in: *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*, 2009, pp. 538–539.
- [70] M. Croci, V. Vinje, M. E. Rognes, Fast uncertainty quantification of tracer distribution in the brain interstitial fluid with multilevel and quasi Monte Carlo, *International Journal for Numerical Methods in Biomedical Engineering* 37 (1) (2021) e3412.
- [71] M. Corti, P. F. Antonietti, L. Dede', A. M. Quarteroni, Numerical modeling of the brain poromechanics by high-order discontinuous Galerkin methods, *Mathematical Models and Methods in Applied Sciences* 33 (08) (2023) 1577–1609.
- [72] L. W. Dreyer, A. Eklund, M. E. Rognes, J. Malm, S. Qvarlander, K.-H. Støverud, K.-A. Mardal, V. Vinje, Modeling CSF circulation and the glymphatic system during infusion using subject specific intracranial pressures and brain geometries, *Fluids and Barriers of the CNS* 21 (1) (2024) 82.
- [73] M. Corti, F. Bonizzoni, A. M. Quarteroni, P. F. Antonietti, et al., Discontinuous Galerkin methods for Fisher–Kolmogorov equation with application to α -synuclein spreading in Parkinson’s disease, *Computer Methods in Applied Mechanics and Engineering* 417 (2023) 116450.
- [74] T. B. Thompson, B. Z. Vigil, R. S. Young, Alzheimer’s disease and the mathematical mind, *Brain Multiphysics* 6 (2024) 100094.
- [75] A. M. Stuart, Inverse problems: A Bayesian perspective, *Acta Numerica* 19 (2010) 451–559.
- [76] D. Williams, *Probability with Martingales*, Cambridge University Press, 1991.
- [77] F. Lindgren, H. Rue, J. Lindström, An explicit link between Gaussian fields and Gaussian Markov random fields: the stochastic partial differential equation approach, *Journal of the Royal Statistical Society Series B: Statistical Methodology* 73 (4) (2011) 423–498.
- [78] U. Villa, N. Petra, O. Ghattas, HIPPLYlib: an extensible software framework for large-scale inverse problems governed by PDEs: Part I: deterministic inversion and linearized Bayesian inference, *ACM Transactions on Mathematical Software* 47 (2) (2021) 1–34.
- [79] Y. Daon, G. Stadler, Mitigating the influence of the boundary on PDE-based covariance operators, *arXiv preprint arXiv:1610.05280* (2016).
- [80] T. Bui-Thanh, O. Ghattas, J. Martin, G. Stadler, A computational framework for infinite-dimensional Bayesian inverse problems Part I: The linearized case, with application to global seismic inversion, *SIAM Journal on Scientific Computing* 35 (6) (2013) A2494–A2523.
- [81] T. Isaac, N. Petra, G. Stadler, O. Ghattas, Scalable and efficient algorithms for the propagation of uncertainty from data through inference to prediction for large-scale problems, with application to flow of the Antarctic ice sheet, *Journal of Computational Physics* 296 (2015) 348–368.

- [82] N. Petra, H. Zhu, G. Stadler, T. J. Hughes, O. Ghattas, An inexact Gauss-Newton method for inversion of basal sliding and rheology parameters in a nonlinear Stokes ice sheet model, *Journal of Glaciology* 58 (211) (2012) 889–903.
- [83] S. C. Eisenstat, H. F. Walker, Choosing the forcing terms in an inexact Newton method, *SIAM Journal on Scientific Computing* 17 (1) (1996) 16–32.
- [84] T. Steihaug, Local and superlinear convergence for truncated iterated projections methods, *Mathematical Programming* 27 (1983) 176–190.
- [85] J. Nocedal, S. J. Wright, *Numerical Optimization*, Springer, 1999.
- [86] M. Heinkenschloss, Mesh independence for nonlinear least squares problems with norm constraints, *SIAM Journal on Optimization* 3 (1) (1993) 81–117.
- [87] F. Tröltzsch, *Optimal control of partial differential equations: theory, methods, and applications*, Vol. 112, American Mathematical Society, 2010.
- [88] A. Manzoni, A. Quarteroni, S. Salsa, *Optimal control of partial differential equations*, Springer, 2021.
- [89] N. Halko, P.-G. Martinsson, J. A. Tropp, Finding structure with randomness: Probabilistic algorithms for constructing approximate matrix decompositions, *SIAM Review* 53 (2) (2011) 217–288.
- [90] U. Villa, N. Petra, O. Ghattas, hIPPYlib: An extensible software framework for large-scale inverse problems, *Journal of Open Source Software* 3 (30) (2018).
- [91] K.-T. Kim, U. Villa, M. Parno, Y. Marzouk, O. Ghattas, N. Petra, hIPPYlib-MUQ: A bayesian inference software framework for integration of data with complex predictive models under uncertainty, *ACM Transactions on Mathematical Software* 49 (2) (2023) 1–31.
- [92] V. I. Bogachev, M. A. S. Ruas, *Measure theory*, Vol. 1, Springer, 2007.
- [93] I. Lawrence, K. Lin, A concordance correlation coefficient to evaluate reproducibility, *Biometrics* (1989) 255–268.
- [94] L. R. Dice, Measures of the amount of ecologic association between species, *Ecology* 26 (3) (1945) 297–302.
- [95] T. Sorensen, A method of establishing groups of equal amplitude in plant sociology based on similarity of species content and its application to analyses of the vegetation on Danish commons, *Biologiske Skrifter* 5 (1948) 1–34.
- [96] K. R. Swanson, R. C. Rostomily, E. Alvord, A mathematical modelling tool for predicting survival of individual patients following resection of glioblastoma: a proof of principle, *British Journal of Cancer* 98 (1) (2008) 113–119.
- [97] P. G. Gritsenko, O. Ilina, P. Friedl, Interstitial guidance of cancer invasion, *The Journal of Pathology* 226 (2) (2012) 185–199.
- [98] J. Wesolowski, P. Rajdev, S. Mukherji, Temozolomide (temodar), *American Journal of Neuroradiology* 31 (8) (2010) 1383–1384.

- [99] J. Kaipio, E. Somersalo, Statistical inverse problems: discretization, model reduction and inverse crimes, *Journal of Computational and Applied Mathematics* 198 (2) (2007) 493–504.
- [100] E. H. Linfoot, An Informational Measure of Correlation, *Information and Control* 1 (1) (1957) 85–89.
- [101] J. T. Kent, Information gain and a general measure of correlation, *Biometrika* 70 (1) (1983) 163–173.
- [102] J. Canosa, On a Nonlinear Diffusion Equation Describing Population Growth, *IBM Journal of Research and Development* 17 (4) (1973) 307–313.
- [103] R. A. Fisher, The wave of advance of advantageous genes, *Annals of Eugenics* 7 (4) (1937) 355–369.
- [104] A. Kolmogoroff, I. Petrovsky, N. Piscounoff, Study of the diffusion equation with growth of the quantity of matter and its application to a biology problem, in: *Dynamics of Curved Fronts*, Elsevier, 1988, pp. 105–130.
- [105] E. Konukoglu, O. Clatz, B. H. Menze, B. Stieltjes, M.-A. Weber, E. Mandonnet, H. Delingette, N. Ayache, Image Guided Personalization of Reaction-Diffusion Type Tumor Growth Models using Modified Anisotropic Eikonal Equations, *IEEE Transactions on Medical Imaging* 29 (1) (2009) 77–95.
- [106] M. J. Simpson, Data-informed model reduction for inference and prediction from non-identifiable models, *Journal of Theoretical Biology* 611 (2025) 112155.
- [107] M. J. Simpson, R. E. Baker, Parameter Identifiability, Parameter Estimation, and Model Prediction for Differential Equation Models, *SIAM Review* 68 (1) (2026) 153–171.
- [108] F.-G. Wieland, A. L. Hauber, M. Rosenblatt, C. Tönsing, J. Timmer, On structural and practical identifiability, *Current Opinion in Systems Biology* 25 (2021) 60–69.
- [109] A. Gholami, A. Mang, G. Biros, An inverse problem formulation for parameter estimation of a reaction–diffusion model of low grade gliomas, *Journal of mathematical biology* 72 (1) (2016) 409–433.
- [110] K. C. Hiremath, K. Atakishi, E. A. Lima, M. Farhat, B. Panthi, H. Langshaw, M. D. Shanker, W. Talpur, S. Thrower, J. Goldman, et al., Identifiability and model selection frameworks for models of high-grade glioma response to chemoradiation, *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 383 (2293) (2025).
- [111] C. M. Phillips, E. A. Lima, C. Wu, A. M. Jarrett, Z. Zhou, N. Elshafeey, J. Ma, G. Rauch, T. E. Yankeelov, Assessing the identifiability of model selection frameworks for the prediction of patient outcomes in the clinical breast cancer setting, *Journal of Computational Science* 69 (2023) 102006.
- [112] L. Sachs, *Applied Statistics: A Handbook of Techniques*, Springer Science & Business Media, 2012.
- [113] H. Cho, A. L. Lewis, K. M. Storey, Bayesian information-theoretic calibration of radiotherapy sensitivity parameters for informing effective scanning protocols in cancer, *Journal of Clinical Medicine* 9 (10) (2020) 3208.

- [114] H. Cho, A. L. Lewis, K. M. Storey, A. C. Zittle, An adaptive information-theoretic experimental design procedure for high-to-low fidelity calibration of prostate cancer models, *Oncology* 8 (2023) 10.
- [115] P. Benner, S. Gugercin, K. Willcox, A survey of projection-based model reduction methods for parametric dynamical systems, *SIAM Review* 57 (4) (2015) 483–531.
- [116] B. Peherstorfer, K. Willcox, M. Gunzburger, Survey of multifidelity methods in uncertainty propagation, inference, and optimization, *SIAM Review* 60 (3) (2018) 550–591.
- [117] D. Fuentes, Y. Feng, A. Elliott, A. Shetty, R. J. McNichols, J. T. Oden, R. J. Stafford, Adaptive real-time bioheat transfer models for computer-driven MR-guided laser induced thermal therapy, *IEEE Transactions on Biomedical Engineering* 57 (5) (2010) 1024–1030.
- [118] S. J. Fahrenholtz, T. Y. Moon, M. Franco, D. Medina, S. Danish, A. Gowda, A. Shetty, F. Maier, J. D. Hazle, R. J. Stafford, T. Warburton, D. Fuentes, A model evaluation study for treatment planning of laser-induced thermal therapy, *International Journal of Hyperthermia* 31 (7) (2015) 705–714.
- [119] P. Colli, H. Gomez, G. Lorenzo, G. Marinoschi, A. Reali, E. Rocca, Optimal control of cytotoxic and antiangiogenic therapies on prostate cancer growth, *Mathematical Models and Methods in Applied Sciences* 31 (07) (2021) 1419–1468.
- [120] L. D. Dalcin, R. R. Paz, P. A. Kler, A. Cosimo, Parallel distributed computing using Python, *Advances in Water Resources* 34 (9) (2011) 1124–1139.
- [121] M. Alnæs, J. Blechta, J. Hake, A. Johansson, B. Kehlet, A. Logg, C. Richardson, J. Ring, M. E. Rognes, G. N. Wells, The FEniCS project version 1.5, *Archive of Numerical Software* 3 (100) (2015).
- [122] V. E. Henson, U. M. Yang, BoomerAMG: A parallel algebraic multigrid solver and preconditioner, *Applied Numerical Mathematics* 41 (1) (2002) 155–177.
- [123] D. Stanzione, J. West, R. T. Evans, T. Minyard, O. Ghattas, D. K. Panda, Frontera: The Evolution of Leadership Computing at the National Science Foundation, in: *Practice and Experience in Advanced Research Computing 2020: Catch the Wave*, PEARC '20, Association for Computing Machinery, New York, NY, USA, 2020, p. 106–111. doi:10.1145/3311790.3396656.
URL <https://doi.org/10.1145/3311790.3396656>
- [124] J. Juan-Albarracín, E. Fuster-Garcia, G. A. García-Ferrando, J. M. García-Gómez, ONCOhabitats: A system for glioblastoma heterogeneity assessment through MRI, *International Journal of Medical Informatics* 128 (2019) 53–61.
- [125] L. S. Hu, A. Hawkins-Daarud, L. Wang, J. Li, K. R. Swanson, Imaging of intratumoral heterogeneity in high-grade glioma, *Cancer Letters* 477 (2020) 97–106.
- [126] F. P. Barthel, K. C. Johnson, F. S. Varn, A. D. Moskalik, G. Tanner, E. Kocakavuk, K. J. Anderson, O. Abiola, K. Aldape, K. D. Alfaro, et al., Longitudinal molecular trajectories of diffuse glioma in adults, *Nature* 576 (7785) (2019) 112–120.

- [127] M. Ceccarelli, F. P. Barthel, T. M. Malta, T. S. Sabedot, S. R. Salama, B. A. Murray, O. Morozova, Y. Newton, A. Radenbaugh, S. M. Pagnotta, et al., Molecular profiling reveals biologically discrete subsets and pathways of progression in diffuse glioma, *Cell* 164 (3) (2016) 550–563.
- [128] C. Hoguea, C. Davatzikos, G. Biros, Modeling glioma growth and mass effect in 3D MR images of the brain, in: *International Conference on Medical Image Computing and Computer-Assisted Intervention*, Springer, 2007, pp. 642–650.
- [129] D. A. Hormuth, S. L. Eldridge, J. A. Weis, M. I. Miga, T. E. Yankeelov, Mechanically Coupled Reaction-Diffusion Model to Predict Glioma Growth: Methodological Details, *Cancer Systems Biology: Methods and Protocols* (2018) 225–241.
- [130] M. Fritz, P. K. Jha, T. Köppl, J. T. Oden, B. Wohlmuth, Analysis of a new multispecies tumor growth model coupling 3d phase-fields with a 1d vascular network, *Nonlinear Analysis: Real World Applications* 61 (2021) 103331.
- [131] R. Laubenbacher, V. Hower, A. Jarrah, S. V. Torti, V. Shulaev, P. Mendes, F. M. Torti, S. Akman, A systems biology view of cancer, *Biochimica et Biophysica Acta (BBA)-Reviews on Cancer* 1796 (2) (2009) 129–139.
- [132] M. A. Clarke, J. Fisher, Executable cancer models: successes and challenges, *Nature Reviews Cancer* 20 (6) (2020) 343–354.
- [133] D. A. Hormuth, A. M. Jarrett, X. Feng, T. E. Yankeelov, Calibrating a predictive model of tumor growth and angiogenesis with quantitative MRI, *Annals of Biomedical Engineering* 47 (2019) 1539–1551.
- [134] M. Fritz, P. K. Jha, T. Köppl, J. T. Oden, A. Wagner, B. Wohlmuth, Modeling and simulation of vascular tumors embedded in evolving capillary networks, *Computer Methods in Applied Mechanics and Engineering* 384 (2021) 113975.
- [135] G. Lorenzo, T. J. Hughes, P. Dominguez-Frojan, A. Reali, H. Gomez, Computer simulations suggest that prostate enlargement due to benign prostatic hyperplasia mechanically impedes prostate cancer growth, *Proceedings of the National Academy of Sciences* 116 (4) (2019) 1152–1161.
- [136] M. Heinkenschloss, D. P. Kouri, Optimization problems governed by systems of PDEs with uncertainties, *Acta Numerica* 34 (2025) 491–577.
- [137] X. Huan, J. Jagalur, Y. Marzouk, Optimal experimental design: Formulations and computations, *Acta Numerica* 33 (2024) 715–840.
- [138] K. Azizzadenesheli, N. Kovachki, Z. Li, M. Liu-Schiaffini, J. Kossaifi, A. Anandkumar, Neural operators for accelerating scientific simulations and design, *Nature Reviews Physics* 6 (5) (2024) 320–328.
- [139] Y. Chen, D. Xiu, Learning stochastic dynamical system via flow map operator, *Journal of Computational Physics* 508 (2024) 112984.
- [140] I. Ezhov, T. Mot, S. Shit, J. Lipkova, J. C. Paetzold, F. Kofler, C. Pellegrini, M. Kollovich, F. Navarro, H. Li, et al., Geometry-Aware Neural Solver for Fast Bayesian Calibration of Brain Tumor Models, *IEEE Transactions on Medical Imaging* 41 (5) (2021) 1269–1278.

- [141] B. Kramer, B. Peherstorfer, K. E. Willcox, Learning Nonlinear Reduced Models from Data with Operator Inference, *Annual Review of Fluid Mechanics* 56 (1) (2024) 521–548.
- [142] N. Kovachki, Z. Li, B. Liu, K. Azizzadenesheli, K. Bhattacharya, A. Stuart, A. Anandkumar, Neural Operator: Learning Maps Between Function Spaces With Applications to PDEs, *Journal of Machine Learning Research* 24 (89) (2023) 1–97.
- [143] F. Regazzoni, S. Pagani, M. Salvador, L. Dede', A. Quarteroni, Learning the intrinsic dynamics of spatio-temporal processes through Latent Dynamics Networks, *Nature Communications* 15 (1) (2024) 1834.

Appendix A. Model derivatives

We develop the variational formulation for the reaction-diffusion model with therapy presented in Section 2.1. We will leverage the Lagrangian formalism to compute the adjoint and gradient expression for the model equations [87, 88]. We will also compute the Hessian action for the model equations. Consider the governing equation (1), restated here for convenience,

$$\begin{aligned} \frac{\partial u}{\partial t} - \nabla \cdot (e^{m_D} \nabla u) - e^{m_\kappa} u(1-u) &= f(u) & \text{in } \Omega \times (t_0, t_f) \\ u(x, t_0) &= u_0 & \text{in } \Omega \\ \nabla u \cdot \eta &= 0 & \text{on } \partial\Omega \times (t_0, t_f). \end{aligned}$$

We will first fix notation. Here,

(t_0, t_f)	The observation time window.
Ω	The spatial domain in \mathbb{R}^2 or \mathbb{R}^3 with boundary $\partial\Omega$.
$u(x, t) \in \mathcal{X}$	The state, tumor volume fraction.
$m_D(x) \in H^1(\Omega)$	The diffusion parameter.
$m_\kappa(x) \in H^1(\Omega)$	The proliferation rate field.
$f(x, u, t) \in L^2(H^{-1}(\Omega); (t_0, t_f))$	The source term.
$v \in \mathcal{V} := L^2(H^1(\Omega); (t_0, t_f))$	A test function.

The space for the state variable $\mathcal{X} := \{u \in L^2(H^1(\Omega); (t_0, t_f)) \mid u_t \in L^2(H^1(\Omega); (t_0, t_f))\}$ is chosen to accommodate the required regularity for the spatial and temporal derivatives as well as the nonlinear term. Here the temporal derivative is denoted $u_t := \partial u / \partial t$. For control of the quadratic nonlinearity, we require $u \in L^2(L^4(\Omega) \cap H^1(\Omega); (0, t_f))$ to ensure that $u^2 \in L^2(L^2(\Omega); (t_0, t_f))$. However, appealing to a Sobolev Embedding Theorem, the embedding $H^1(\Omega) \hookrightarrow L^4(\Omega)$ is continuous so that $u \in L^2(H^1(\Omega); (t_0, t_f))$ is sufficient (see Sec. 4.2.1. and Theorem 7.1 in [87]).

Appendix A.1. Variational Formulation

The weak form of the reaction-diffusion equation is obtained by multiplying the governing equation (1) by a test function v and integrating over the space-time domain $\Omega \times (0, t_f)$.

$$\int_{t_0}^{t_f} \int_{\Omega} u_t v \, dx \, dt - \int_{t_0}^{t_f} \int_{\Omega} \nabla \cdot (e^{m_D} \nabla u) v \, dx \, dt - \int_{t_0}^{t_f} \int_{\Omega} e^{m_\kappa} u(1-u) \, dx \, dt - \int_{t_0}^{t_f} \int_{\Omega} f v \, dx \, dt = 0$$

Straightforward application of the Divergence theorem along with the homogeneous Neumann boundary conditions yields the weak form,

$$\begin{aligned} \int_{t_0}^{t_f} \int_{\Omega} u_t v \, dx \, dt + \int_{t_0}^{t_f} \int_{\Omega} e^{m_D} \nabla u \cdot \nabla v \, dx \, dt - \int_{t_0}^{t_f} \int_{\partial\Omega} \eta \cdot (e^{m_D} \nabla u) v \, ds \, dt \\ - \int_{t_0}^{t_f} \int_{\Omega} e^{m_\kappa} u(1-u) v - f v \, dx \, dt = 0 \end{aligned}$$

$$\int_{t_0}^{t_f} \int_{\Omega} u_t v \, dx \, dt + \int_{t_0}^{t_f} \int_{\Omega} e^{m_D} \nabla u \cdot \nabla v \, dx \, dt - \int_{t_0}^{t_f} \int_{\Omega} e^{m_\kappa} u(1-u) v \, dx \, dt - \int_{t_0}^{t_f} \int_{\Omega} f v \, dx \, dt = 0.$$

The development proceeded with a generic source term to account for therapy, but with the treatment model (5) given by (2) and (4), one would have the last term given by

$$\int_{t_0}^{t_f} \int_{\Omega} f v \, dx \, dt = \int_{t_0}^{t_f} \int_{\Omega} \gamma(1 - S_{\text{rt}}(z)) u v \, dx \, dt + \int_{t_0}^{t_f} \int_{\Omega} \alpha_{\text{ct}} \sum_k \exp(-\beta_{\text{ct}}(t - \tau_{k,\text{ct}})) u v \, dx \, dt$$

with appropriate application in time according to the therapy schedules \mathcal{T}_{rt} and \mathcal{T}_{ct} .

Appendix A.2. Lagrangian

It is assumed that the state variable comes from a space with sufficient regularity, i.e., \mathcal{X} or $W_2^{1,1}(\Omega)$, and the parameters m similarly come from a space with sufficient regularity \mathcal{M} , i.e., $H^1(\Omega) \times H^1(\Omega)$. The forward model is given by the PDE and is a mapping $\mathcal{F} : \mathcal{X} \times \mathcal{M} \rightarrow \mathcal{Y}$. The adjoint variable $p \in \mathcal{P} := \mathcal{Y}^*$, the dual of the action of the forward model. A suitable space is the previously defined \mathcal{V} . For an extended discussion, the reader is referred to Chapter 3 of [87]. The inverse problem for the computation of the MAP point (16) is stated as:

$$\min_{m \in \mathcal{M}} \phi(m_D, m_\kappa) = \min_{m \in \mathcal{M}} -\log \nu_{\text{post}}(m|\mathbf{d}) = \min_{m \in \mathcal{M}} \underbrace{\frac{1}{2\sigma_{\text{noise}}^2} \int_{t_0}^{t_f} \int_{\Omega} (\mathcal{B}u - \mathbf{d})^2 dx dt}_{\text{data misfit}} + \underbrace{\mathcal{R}(m)}_{\text{regularization}}.$$

While the data misfit functional is specified, the exact form of the regularization functional is not, as we prefer to derive the action in generality. Moreover, the observation operator \mathcal{B} may incorporate a collection of Dirac deltas in time, to compare the state with the finitely many data collected over the simulation window. For the Gaussian random field priors developed in Section 3 analysis is available in the literature, in particular [80, 78].

Appendix A.3. Gradient expression

Here we derive expressions for the gradient of ϕ with respect to the model parameters m_D and m_κ using the formal Lagrange method. We first form the Lagrangian functional, \mathcal{L}^g (where the superscript g denotes the role in deriving the gradient), that combines the regularized data misfit $\phi(m)$ with the weak form of the model equations. The Lagrangian functional is,

$$\begin{aligned} \mathcal{L}^g(u, p, m_D, m_\kappa) := & \frac{1}{2} \int_{t_0}^{t_f} \int_{\Omega} (\mathcal{B}u - d)^2 dx dt + \mathcal{R}(m_D, m_\kappa) \\ & + \int_{t_0}^{t_f} \int_{\Omega} \left[u_t p + e^{m_D} \nabla u \cdot \nabla p - e^{m_\kappa} u(1-u)p - f(u)p \right] dx dt \end{aligned}$$

for functions $(u, p, (m_D, m_\kappa)) \in \mathcal{X} \times \mathcal{P} \times \mathcal{M}$.

Appendix A.3.1. The variational formulation

Taking variations of \mathcal{L}^g with respect to $p \in \mathcal{P}$ and requiring them to vanish for all admissible variations \hat{p} simply recovers the weak form of the forward reaction-diffusion model previously derived. That is, $\delta_p \mathcal{L}^g = 0$ for all $\hat{p} \in \mathcal{P}$ yields the weak form of the PDE.

$$\int_{t_0}^{t_f} \int_{\Omega} \frac{\partial u}{\partial t} \hat{p} - e^{m_D} \nabla u \cdot \nabla \hat{p} - e^{m_\kappa} u(1-u)\hat{p} - f(u)\hat{p} dx dt = 0$$

Appendix A.3.2. The adjoint equation

Next, we require that variations of \mathcal{L}^g with respect to the state u vanish for all admissible variations $\hat{u} \in \mathcal{X}$, that is $\delta_u \mathcal{L}^g = 0$ for all $\hat{u} \in \mathcal{X}$. This will yield the weak form of the adjoint PDE. We have,

$$\int_{t_0}^{t_f} \int_{\Omega} \frac{1}{\sigma_{\text{noise}}^2} (\mathcal{B}u - d)(\mathcal{B}\hat{u}) dx dt + \int_{t_0}^{t_f} \int_{\Omega} \left[p \frac{\partial \hat{u}}{\partial t} + e^{m_D} \nabla \hat{u} \cdot \nabla p - e^{m_\kappa} p \hat{u} + 2e^{m_\kappa} p u \hat{u} \right] dx dt - \int_{t_0}^{t_f} \int_{\Omega} \delta_u f(u) \hat{u} p dx dt$$

for all $\hat{u} \in \mathcal{X}$. For the specific treatment models that we consider, we have

$$\int_{t_0}^{t_f} \int_{\Omega} \delta_u f(u) \hat{u} p dx dt = \int_{t_0}^{t_f} \int_{\Omega} \gamma(1 - S_{\text{rt}}(z)) \hat{u} p dx dt + \int_{t_0}^{t_f} \int_{\Omega} \alpha_{\text{ct}} \sum_k \exp(-\beta_{\text{ct}}(t - \tau_{k,\text{ct}})) \hat{u} p dx dt$$

By appropriate integration by parts in time and space to remove derivatives of \hat{u} and arguing the arbitrariness of \hat{u} , we can arrive at the strong form of the adjoint equation, a terminal boundary value problem given by,

$$\begin{aligned} -\frac{\partial p}{\partial t} - \nabla \cdot (e^{m_D} \nabla p) - (e^{m_\kappa} + 2e^{m_\kappa} u) p \\ = -\frac{1}{\sigma_{\text{noise}}^2} \mathcal{B}^*(\mathcal{B}u - d) - \gamma(1 - S_{\text{rt}}(z))p - \alpha_{\text{ct}} \sum_k \exp(-\beta_{\text{ct}}(t - \tau_{k,\text{ct}}))p \quad \text{in } \Omega \times (t_0, t_f) \\ p(x, t_f) = 0 \quad \text{in } \Omega \\ e^{m_D} \nabla p \cdot \eta = 0 \quad \text{on } \partial\Omega \times (t_0, t_f) \end{aligned}$$

Appendix A.3.3. The gradient expression

Finally, we derive expressions for the gradient, the Fréchet derivative of ϕ with respect to m_D and m_κ , denoted $\mathfrak{D}_{m_D} \phi$ and $\mathfrak{D}_{m_\kappa} \phi$, respectively. We consider the variations of the Lagrangian with respect to the parameters. The Fréchet derivative of $\phi(m_D, m_\kappa)$ with respect to m_D in an arbitrary direction $\hat{m}_D \in H^1(\Omega)$ evaluated at (m_D, m_κ) is given by $\delta_{m_D} \mathcal{L}^g$, that is

$$\mathfrak{D}_{m_D} \phi(m_D, m_\kappa, \hat{m}_D) := \delta_{m_D} \mathcal{R}(m_D, m_\kappa) + \int_{t_0}^{t_f} \int_{\Omega} \hat{m}_D e^{m_D} \nabla u \cdot \nabla p \, dx \, dt.$$

Similarly, the Fréchet derivative of $\phi(m_D, m_\kappa)$ with respect to m_κ in an arbitrary direction $\hat{m}_\kappa \in H^1(\Omega)$ evaluated at (m_D, m_κ) is given by $\delta_{m_\kappa} \mathcal{L}^g$,

$$\mathfrak{D}_{m_\kappa} \phi(m_D, m_\kappa, \hat{m}_\kappa) := \delta_{m_\kappa} \mathcal{R}(m_D, m_\kappa) + \int_{t_0}^{t_f} \int_{\Omega} \hat{m}_\kappa e^{m_\kappa} u (1 - u) p \, dx \, dt.$$

The gradient \mathfrak{G} with respect to a parameter m is defined as the Riesz representer of the Fréchet derivative of ϕ with respect to a chosen inner product,

$$(\mathfrak{G}_m(m), \hat{m}) := \mathfrak{D}_m \phi(m, \hat{m}) = \delta_m \mathcal{L}^g.$$

In summary, to compute the gradient at (m_D, m_κ) , we

1. Solve the forward model equation for u , given (m_D, m_κ) .
2. Solve the adjoint model equation for p , given (m_D, m_κ) and u .
3. Evaluate the Fréchet derivatives $\delta_{m_D} \mathcal{L}^g$ and $\delta_{m_\kappa} \mathcal{L}^g$, given (m_D, m_κ) , u , and p .

Appendix A.4. Hessian action

We now form the Lagrangian for the Hessian, \mathcal{L}^H (where the superscript H refers to the role in deriving the Hessian action). We retain the notation that $\tilde{\cdot}$ denotes a trial function and $\hat{\cdot}$ denotes a test function. We consider the Hessian action in an arbitrary direction $(\tilde{m}_D, \tilde{m}_\kappa)$, replacing $(\hat{m}_D, \hat{m}_\kappa)$ in the previous section. This is done to preserve the notation $\hat{\cdot}$ for the *current* variations. Similarly, we replace \hat{p} with \tilde{p} and \hat{u} with \tilde{u} . We refer to \tilde{u} and \tilde{p} as the *incremental state* and

incremental adjoint variables, respectively. The Lagrangian functional for the Hessian is,

$$\begin{aligned}
\mathcal{L}^H(u, p, m_D, m_\kappa, \tilde{u}, \tilde{p}, \tilde{m}_D, \tilde{m}_\kappa) := & \\
& \underbrace{\delta_{m_D} \mathcal{R}(m_D, m_\kappa) + \int_{t_0}^{t_f} \int_{\Omega} \tilde{m}_D e^{m_D} \nabla u \cdot \nabla p \, dx \, dt}_{\text{Fréchet derivative with respect to } m_D \text{ in direction } \tilde{m}_D} \\
& + \underbrace{\delta_{m_\kappa} \mathcal{R}(m_D, m_\kappa) + \int_{t_0}^{t_f} \int_{\Omega} \tilde{m}_\kappa e^{m_\kappa} u(1-u)p \, dx \, dt}_{\text{Fréchet derivative with respect to } m_\kappa \text{ in direction } \tilde{m}_\kappa} \\
& + \underbrace{\int_{t_0}^{t_f} \int_{\Omega} \left[\frac{\partial u}{\partial t} \tilde{p} + e^{m_D} \nabla u \cdot \nabla \tilde{p} - e^{m_\kappa} u(1-u)\tilde{p} - f(u)\tilde{p} \right] dx \, dt}_{\text{Weak form of forward model equation}} \\
& + \underbrace{\int_{t_0}^{t_f} \int_{\Omega} \left[\frac{1}{\sigma_{\text{noise}}^2} (\mathcal{B}u - d)\mathcal{B}\tilde{u} + p \frac{\partial \tilde{u}}{\partial t} + e^{m_D} \nabla \tilde{u} \cdot \nabla p - e^{m_\kappa} \tilde{u}p + 2m_\kappa u \tilde{u}p \right] dx \, dt + \delta_u f(u)}_{\text{Weak form of adjoint model equation}},
\end{aligned}$$

where $(u, p, (m_D, m_\kappa), \tilde{u}, \tilde{p}, (\tilde{m}_D, \tilde{m}_\kappa)) \in \mathcal{X} \times \mathcal{P} \times \mathcal{M} \times \mathcal{X} \times \mathcal{P} \times \mathcal{M}$. To derive the expression for the action of the Hessian of ϕ with respect to (m_D, m_κ) in a direction $(\tilde{m}_D, \tilde{m}_\kappa)$ we take variations of \mathcal{L}^H with respect to its arguments.

Appendix A.4.1. The incremental forward equation

Requiring variations of \mathcal{L}^H with respect to the adjoint p to vanish for all admissible variations $\hat{p} \in \mathcal{P}$ yields the *incremental forward* equation: Given (m_D, m_κ) , $(\tilde{m}_D, \tilde{m}_\kappa)$, and u , find the incremental state \tilde{u} such that for all $\hat{p} \in \mathcal{P}$,

$$\int_{t_0}^{t_f} \int_{\Omega} \left[\tilde{m}_D e^{m_D} \nabla u \cdot \nabla \hat{p} + \tilde{m}_\kappa e^{m_\kappa} u(1-u)\hat{p} - e^{m_\kappa} \tilde{u} \hat{p} + 2e^{m_\kappa} u \tilde{u} \hat{p} + \hat{p} \frac{\partial \tilde{u}}{\partial t} + e^{m_D} \nabla \tilde{u} \cdot \nabla \hat{p} \right] dx \, dt = 0.$$

Integration by parts to remove derivatives of \hat{p} and arguing the arbitrariness of \hat{p} yields the strong form of the incremental forward equation,

$$\begin{aligned}
\frac{\partial \tilde{u}}{\partial t} - \nabla \cdot (e^{m_D} \nabla \tilde{u}) - e^{m_\kappa} \tilde{u} + 2e^{m_\kappa} u \tilde{u} &= \tilde{m}_D \nabla \cdot (e^{m_D} \nabla u) - \tilde{m}_\kappa e^{m_\kappa} u(1-u) & \text{in } \Omega \times (t_0, t_f) \\
\tilde{u}(x, t_0) &= 0 & \text{in } \Omega \\
e^{m_D} \nabla \tilde{u} \cdot \eta &= 0 & \text{on } \partial\Omega \times (t_0, t_f)
\end{aligned}$$

Note that the initial condition is identically zero for the incremental state variable in order to satisfy first order optimality conditions (see Theorem 2.22 [87]).

Appendix A.4.2. The incremental adjoint equation

Taking variations of \mathcal{L}^H with respect to the state u and requiring them to vanish for all admissible variations yields the incremental adjoint problem, that is $\delta_u \mathcal{L}^H = 0$ for all $\hat{u} \in \mathcal{X}$,

$$\begin{aligned} & \int_{t_0}^{t_f} \int_{\Omega} \frac{1}{\sigma_{\text{noise}}^2} \mathcal{B}\hat{u} \mathcal{B}\tilde{u} + 2e^{m_\kappa} \hat{u} \tilde{u} \, dx \, dt \\ & + \int_{t_0}^{t_f} \int_{\Omega} \left[\frac{\partial \hat{u}}{\partial t} \tilde{p} + e^{m_D} \nabla \hat{u} \cdot \nabla \tilde{p} - e^{m_\kappa} \hat{u} (1-u) \tilde{p} + e^{m_\kappa} u \hat{u} \tilde{p} \right] dx \, dt \\ & + \int_{t_0}^{t_f} \int_{\Omega} \left[\tilde{m}_\kappa e^{m_\kappa} \hat{u} (1-u) p - \tilde{m}_\kappa e^{m_\kappa} u \hat{u} p + \tilde{m}_D e^{m_D} \nabla \hat{u} \cdot \nabla p \right] dx \, dt = 0 \end{aligned}$$

Note that the source terms considered in this work are linear with respect to the state variable and so $\delta_{uu} f(u) = 0$. However, for treatment terms with nonlinear state dependence, one should take care to account for the required term. Appropriate integration by parts in time and space to remove derivatives of \hat{u} and arguing the arbitrariness of \hat{u} yields the strong form of the incremental adjoint equation,

$$\begin{aligned} & -\frac{\partial \tilde{p}}{\partial t} - \tilde{m}_D \nabla \cdot (e^{m_D} \nabla \tilde{p}) + e^{m_\kappa} (2u-1) \tilde{p} \\ & = -\frac{1}{\sigma_{\text{noise}}^2} \mathcal{B}^* (\mathcal{B}\tilde{u}) - 2e^{m_\kappa} \tilde{u} + \tilde{m}_D \nabla \cdot (e^{m_D} \nabla p) - \tilde{m}_\kappa e^{m_\kappa} p (1-2u) \quad \text{in } \Omega \times (t_0, t_f) \\ & \tilde{p}(x, t_f) = 0 \quad \text{in } \Omega \\ & e^{m_D} \nabla \tilde{p} \cdot \eta = 0 \quad \text{on } \partial\Omega \times (t_0, t_f) \end{aligned}$$

Appendix A.4.3. The Hessian action

Finally, we derive expressions for the action of the Hessian of ϕ with respect to (m_D, m_κ) in a direction $(\tilde{m}_D, \tilde{m}_\kappa)$, that is, the second Fréchet derivative of ϕ , $\mathfrak{D}^2 \phi$. Since we have multiple parameter fields, the Hessian is a block operator and we must consider the mixed derivatives,

$$\mathcal{H}(m_D, m_\kappa) := \begin{bmatrix} \mathcal{H}_{DD} & \mathcal{H}_{D\kappa} \\ \mathcal{H}_{\kappa D} & \mathcal{H}_{\kappa\kappa} \end{bmatrix}.$$

Given (m_D, m_κ) , the action of the first block row of \mathcal{H} in the direction $(\tilde{m}_D, \tilde{m}_\kappa)$ is given by $\delta_{m_D} \mathcal{L}^H$,

$$\begin{aligned} \delta_{m_D} \mathcal{L}^H & := (\hat{m}_D, \mathcal{H}_{DD} \tilde{m}_D) + (\hat{m}_\kappa, \mathcal{H}_{d\kappa} \tilde{m}_\kappa) \\ & = \delta_{m_D}^2 \mathcal{R}(m_D, m_\kappa) + \delta_{m_D} (\delta_{m_\kappa} \mathcal{R}(m_D, m_\kappa)) \\ & + \int_{t_0}^{t_f} \int_{\Omega} \left[\tilde{m}_D (\hat{m}_D e^{m_D} \nabla u \cdot \nabla p) + \hat{m}_D e^{m_D} \nabla \tilde{u} \cdot \nabla p + \hat{m}_D e^{m_D} \nabla u \cdot \nabla \tilde{p} \right] dx \, dt \end{aligned}$$

Similarly, given (m_D, m_κ) , the action of the second block row of \mathcal{H} in the direction $(\tilde{m}_D, \tilde{m}_\kappa)$ is given by $\delta_{m_\kappa} \mathcal{L}^H$,

$$\begin{aligned} \delta_{m_\kappa} \mathcal{L}^H & := (\hat{m}_\kappa, \mathcal{H}_{\kappa d} \tilde{m}_D) + (\hat{m}_\kappa, \mathcal{H}_{\kappa\kappa} \tilde{m}_\kappa) \\ & = \delta_\kappa (\delta_{m_D} \mathcal{R}(m_D, m_\kappa)) + \delta_{m_\kappa}^2 (\mathcal{R}(m_D, m_\kappa)) \\ & + \int_{t_0}^{t_f} \int_{\Omega} \left[\hat{m}_\kappa e^{m_\kappa} (2\tilde{u} u - 1) - \hat{m}_\kappa e^{m_\kappa} u (1-u) \tilde{p} \right] dx \, dt \end{aligned}$$

In summary, to compute the Hessian action at a point in parameter space $(m_D, m_\kappa) \in \mathcal{M}$ in a direction $(\tilde{m}_D, \tilde{m}_\kappa)$, we

1. Solve the forward model for u , given (m_D, m_κ) .
2. Solve the adjoint equation for p , given (m_D, m_κ) and u .
3. Solve the incremental forward equation for \tilde{u} , given (m_D, m_κ) , $(\tilde{m}_D, \tilde{m}_\kappa)$, and u .
4. Solve the incremental adjoint equation for \tilde{p} , given (m_D, m_κ) , $(\tilde{m}_D, \tilde{m}_\kappa)$, u , p , and \tilde{u} .
5. Evaluate the Hessian actions, given (m_D, m_κ) , $(\tilde{m}_D, \tilde{m}_\kappa)$, u , p , \tilde{u} , and \tilde{p} .

Appendix B. Additional scaling studies

Since the computational domain developed in Sec. 2.2 is assumed to be fixed at simulation time, we are primarily concerned with the strong scalability of the forward solve. That is, how quickly can a given amount of work be parallelized and solved. While it is not currently useful to refine the mesh past the resolution of the data, higher quality data may be available, thus requiring meshes with additional DOFs to accurately represent the data. To account for our implementation’s ability to handle these cases we also perform a weak scaling study.

Appendix B.1. Strong scaling of the adjoint problem

In Fig. B.15, we report the strong scaling of the adjoint solve for the benchmark problem outlined in Sec. 4.4. Note that the inter-node study requires at least two nodes because of the required memory footprint for the adjoint solve (storage of the state variable). Once more we observe excellent scaling within node and adequate scaling out of node.




Figure B.15: Strong scaling of the adjoint solve in (a) 14.2 seconds on one node for a benchmark problem with 1 million DOFs and (b) 4.1 seconds on 32 nodes for a benchmark problem with 4 million DOFs.

Appendix B.2. Weak scaling

For the weak scaling study, we fix the amount of work per processor to be 300,000 DOFs and test the efficiency by adding additional work and cores in equal proportion until resources are exhausted. Results on a single Frontera CLX node are reported in Fig. B.16, where we observe a reasonably efficient scaling up to 2.4 million DOFs.

Appendix C. Model validation summary statistics

We collect and tabulate summary statistics for the clinical model validation study presented in Section 5.2. In particular, we report the mean, standard deviation, and 90% credible interval (5th to 95th percentile) for both prior distribution and posterior distribution pushforwards.

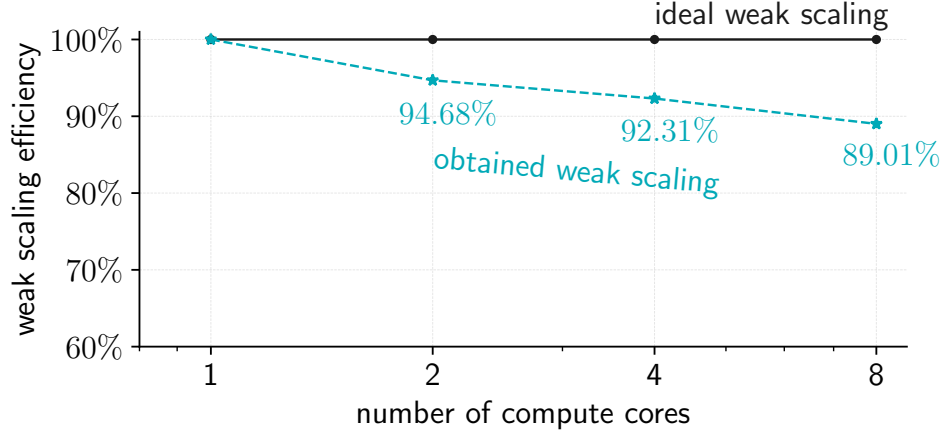


Figure B.16: Weak scaling of the forward solve with 300,000 DOFs per processor on a Frontera CLX node.

Table C.4: Summary of mean, standard deviation, and 90% credible interval by patient for the Dice similarity coefficient in the last-to-final prediction case.

Patient	Prior Distribution			Laplace Approximation		
	Mean	Std. Dev.	Credible Interval	Mean	Std. Dev.	Credible Interval
W03	0.566	0.018	(0.538, 0.593)	0.462	0.0022	(0.459, 0.466)
W11	0.447	0.028	(0.401, 0.490)	0.731	0.0027	(0.727, 0.736)
W16	0.879	0.0046	(0.872, 0.886)	0.870	0.0018	(0.867, 0.873)
W29	0.417	0.016	(0.390, 0.443)	0.467	0.0062	(0.457, 0.476)
W35	0.427	0.032	(0.369, 0.481)	0.636	0.0052	(0.627, 0.644)
W36	0.671	0.030	(0.619, 0.715)	0.654	0.0031	(0.649, 0.659)
W43	0.594	0.010	(0.577, 0.609)	0.597	0.0018	(0.595, 0.600)
W53	0.095	0.020	(0.070, 0.131)	0.267	0.014	(0.245, 0.288)

Table C.5: Summary of mean, standard deviation, and 90% credible interval by patient for the Dice similarity coefficient in the initial-to-final prediction case.

Patient	Prior Distribution			Laplace Approximation		
	Mean	Std. Dev.	Credible Interval	Mean	Std. Dev.	Credible Interval
W03	0.506	0.061	(0.416, 0.616)	0.438	0.031	(0.387, 0.491)
W11	0.201	0.0015	(0.201, 0.203)	0.502	0.038	(0.431, 0.551)
W16	0.019	0.060	(0.000, 0.128)	0.779	0.010	(0.762, 0.794)
W29	0.272	0.060	(0.192, 0.362)	0.285	0.023	(0.252, 0.327)
W35	0.178	0.046	(0.137, 0.268)	0.559	0.013	(0.536, 0.572)
W36	0.398	0.088	(0.269, 0.540)	0.477	0.063	(0.374, 0.576)
W43	0.285	0.186	(0.004, 0.546)	0.501	0.0053	(0.493, 0.509)
W53	0.067	0.0024	(0.066, 0.070)	0.146	0.021	(0.112, 0.180)

Table C.6: Summary of mean, standard deviation, and 90% credible interval by patient for the relative error in total tumor cellularity in the last-to-final prediction case.

Patient	Prior Distribution			Laplace Approximation		
	Mean	Std. Dev.	Credible Interval	Mean	Std. Dev.	Credible Interval
W03	0.122	0.151	(-0.085, 0.391)	-0.137	0.0064	(-0.147, -0.125)
W11	5.041	0.715	(3.949, 6.296)	1.683	0.012	(1.664, 1.704)
W16	0.046	0.053	(-0.039, 0.133)	0.683	0.0077	(0.670, 0.696)
W29	1.640	0.383	(1.088, 2.294)	1.971	0.075	(1.855, 2.102)
W35	5.891	0.781	(4.749, 7.344)	3.146	0.042	(3.079, 3.215)
W36	-0.222	0.239	(-0.550, 0.222)	-0.162	0.0077	(-0.175, -0.150)
W43	1.935	0.213	(1.617, 2.301)	2.766	0.021	(2.731, 2.800)
W53	46.081	12.865	(26.609, 68.141)	11.059	0.567	(10.226, 11.961)

Table C.7: Summary of mean, standard deviation, and 90% credible interval by patient for the relative error in total tumor cellularity in the initial-to-final prediction case.

Patient	Prior Distribution			Laplace Approximation		
	Mean	Std. Dev.	Credible Interval	Mean	Std. Dev.	Credible Interval
W03	2.863	1.496	(0.555, 5.394)	1.175	0.341	(0.685, 1.774)
W11	20.925	0.764	(19.197, 21.370)	4.177	0.481	(3.665, 5.182)
W16	-0.994	0.029	(-1.000, -0.973)	1.637	0.049	(1.562, 1.721)
W29	2.966	2.773	(-0.525, 8.284)	8.449	1.703	(6.016, 11.436)
W35	19.379	6.115	(9.380, 29.660)	4.244	0.136	(4.073, 4.425)
W36	1.270	1.198	(-0.430, 3.541)	0.966	0.420	(0.464, 1.785)
W43	-0.376	0.705	(-0.996, 1.149)	4.618	0.107	(4.444, 4.797)
W53	85.526	5.741	(73.815, 89.354)	27.216	4.374	(21.167, 35.063)