

Predictive Digital Twin for Optimizing Patient-Specific Radiotherapy Regimens under Uncertainty in High-Grade Gliomas

Anirban Chaudhuri¹, Graham Pash¹, David A. Hormuth II^{1,7}, Guillermo Lorenzo^{1,3}, Michael Kapteyn¹, Chengyue Wu¹, Ernesto A. B. F. Lima^{1,4}, Thomas E. Yankeelov^{1,2,5,6,7,8}, and Karen Willcox¹

¹*Oden Institute for Computational Engineering and Sciences, The University of Texas at Austin, TX, USA*

²*Department of Biomedical Engineering, The University of Texas at Austin, Austin, TX, USA*

³*Department of Civil Engineering and Architecture, University of Pavia, Italy*

⁴*Texas Advanced Computing Center, The University of Texas at Austin, Austin, TX, USA*

⁵*Department of Diagnostic Medicine, The University of Texas at Austin, Austin, TX, USA*

⁶*Department of Oncology, The University of Texas at Austin, Austin, TX, USA*

⁷*Livestrong Cancer Institutes, The University of Texas at Austin, Austin, TX, USA*

⁸*Department of Imaging Physics, MD Anderson Cancer Center, Houston, TX, USA*

Abstract

We develop a methodology to create data-driven predictive digital twins for optimal risk-aware clinical decision-making. We illustrate the methodology as an enabler for an anticipatory personalized treatment that accounts for uncertainties in the underlying tumor biology in high-grade gliomas, where heterogeneity in the response to standard-of-care (SOC) radiotherapy contributes to sub-optimal patient outcomes. The digital twin is initialized through prior distributions derived from population-level clinical data in the literature for a mechanistic model’s parameters. Then the digital twin is personalized using Bayesian model calibration for assimilating patient-specific magnetic resonance imaging data. The calibrated digital twin is used to propose optimal radiotherapy treatment regimens by solving a multi-objective risk-based optimization under uncertainty problem. The solution leads to a suite of patient-specific optimal radiotherapy treatment regimens exhibiting varying levels of trade-off between the two competing clinical objectives: (i) maximizing tumor control (characterized by minimizing the risk of tumor volume growth) and (ii) minimizing the toxicity from radiotherapy. The proposed digital twin framework is illustrated by generating an *in silico* cohort of 100 patients with high-grade glioma growth and response properties typically observed in the literature. For the same total radiation dose as the SOC, the personalized treatment regimens lead to median increase in tumor time to progression of around six days. Alternatively, for the same level of tumor control as the SOC, the digital twin provides optimal treatment options that lead to a median reduction in radiation dose by 16.7% (10 Gy) compared to SOC total dose of 60 Gy. The range of optimal solutions also provide options with increased doses for patients with aggressive cancer, where SOC does not lead to sufficient tumor control.

Keywords: digital twin, risk-aware clinical decision-making, personalized tumor forecasts, uncertainty quantification, adaptive radiotherapy, mathematical oncology, brain cancer

1 Introduction

A digital twin can be defined as a mathematical model (or a collection of models) that provides a virtual representation of a specific physical object (e.g., a tumor), updates its status by assimilating object-specific

data (e.g., imaging and clinical measurements of tumor growth and radiotherapy response), predicts the behavior of the object under external actions (e.g., treatments), and enables decision-making to optimize the future behavior of the object (e.g., design an optimal radiotherapy plan maximizing tumor control and minimizing toxicities) [33, 45, 78, 57, 36]. Recent work explores the use of digital twins in healthcare and medicine to perform simulations of cardiovascular diseases [10, 51], enable virtual reality for surgery [1], and enable improved decision-making in clinical oncology [42, 18, 78]. However, most of the work on digital twins in medicine relies on deterministic implementations. Predictive digital twins account for uncertainty through a Bayesian framework [33] and provide a computational environment to support personalized risk-based management of solid tumors supported by computer forecasts of biologically-inspired mechanistic models representing these diseases and their treatments. We propose a patient-specific predictive digital twin that can address the three critical needs of: (i) accounting for uncertainty in the mechanistic model parameters by continuous integration of incoming patient data, (ii) forecasting the parameter uncertainty to estimate risk associated with the therapeutic outcomes, and (iii) supporting risk-aware clinical decision-making under uncertainty. To illustrate the methodology, we present a predictive digital twin to enable the personalized monitoring and forecast of high-grade glioma (HGG) response to radiotherapy (RT), as well as the design of optimal adaptive RT regimens for individual HGG patients (see Figure 1).

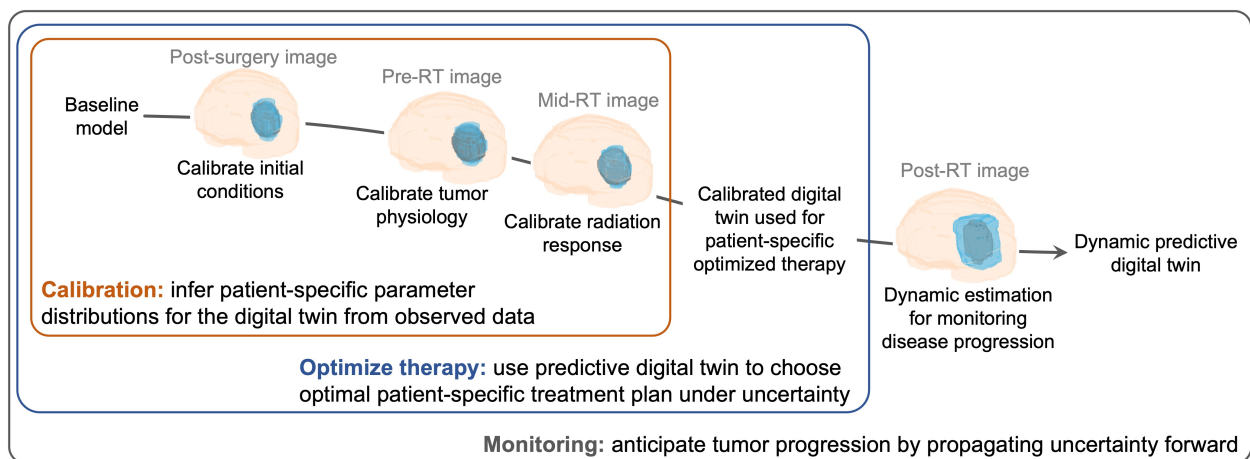


Figure 1: **Creating and evolving an HGG patient-specific predictive digital twin.** Our digital twin methodology is illustrated for the case of HGG growth and response to RT from the post-surgery imaging visit to post-RT monitoring for disease progression. The digital twin is personalized through calibration using MRI data and then used to design an optimal risk-aware treatment regimen under uncertainty. The digital twin also allows for monitoring the disease progression throughout the patient’s treatment and recovery.

HGG is a class of brain tumors that typically exhibit an aggressive, infiltrative behavior as well as high heterogeneity in both tumor physiology and cell composition [49]. Patients with HGG are usually treated surgical resection to reduce the tumor burden and intracranial pressure followed by adjuvant radiotherapy (RT) and chemotherapy to target residual, unresected tumor cells. The standard-of-care (SOC) RT plans account for patient-specific heterogeneity in tumor shape and location through pre-treatment anatomical or structural imaging approaches, such as T_1 - and T_2 -weighted magnetic resonance imaging (MRI), that can identify the residual tumor after surgery and define a surrounding 2–3 cm margin. However, the SOC RT dose and schedule generally conform to the Stupp protocol [68], which is informed from clinical studies

involving large populations, consisting of 60 Gy delivered in 30 fractions of 2 Gy. A fundamental challenge in the treatment of HGG is that response to RT is highly variable from patient to patient due to the inherent heterogeneity in both cellular architecture and tumor micro-environment [3, 19], which may ultimately lead to poor treatment outcomes. This heterogeneity in tumor physiology, growth, and radiation response characteristics translates into uncertainty in the therapeutic outcomes of HGG patients receiving a standard RT protocol. This work addresses the challenge of quantifying the uncertainty in tumor characteristics from limited patient-specific data early in the course of treatment, propagating this uncertainty via mechanistic models to determine the uncertainty in treatment outcomes, and optimizing adaptive RT plans to improve overall survival through individualized predictive digital twins.

Adaptive RT has been regarded as a promising strategy to account for heterogeneity in tumor response. The varying treatment regimens aim at adjusting the dose, number, and timing of the radiation fractions in the prescribed RT plan on a patient-specific basis to account for the individual response of their HGG to radiation. Adaptive RT plans can be guided by quantitative imaging measurements to improve target delineation to deliver radiation, enable the voxelization of the RT plan (i.e., dose-painting), and facilitate the adaptation of the treatment regimen in response to observed tumor dynamics [30, 72, 34, 55]. Another approach for adaptive RT planning is through a predictive framework that can be constructed by leveraging patient-specific computational forecasts of HGG growth and response to RT, which are obtained via computer simulation of biologically-inspired mechanistic models informed by routine clinical and imaging data collected before and during the course of treatment [62, 73, 37, 52, 79, 69, 41, 23, 13, 77]. However, the tumor forecasts obtained with these models are estimated using a deterministic approach. Systematically accounting for data and model uncertainties has shown promise in Bayesian calibration of tumor models [47, 39, 40, 41]. The clinical translation of current computational technologies to forecast HGG response to RT needs to accommodate a probabilistic risk assessment of pathological and therapeutic outcomes to support clinical decision-making along with progressive update of model according to incoming data. Predictive digital twins are a step towards personalized medicine that can address patient-specific clinical decision-making while accounting for the underlying uncertainty.

We construct the digital twin on an MRI-informed biology-inspired mechanistic model, which describes HGG growth and response to RT, and a probabilistic graphical model [33], which accounts for uncertainty quantification and enables the risk assessment of therapeutic outcomes during tumor forecasting and RT optimization. We use prior distributions of the mechanistic model parameters informed from clinical data to initialize the digital twin. The personalization of the digital twin starts with the Bayesian model calibration for assimilating patient-specific MRI data followed by risk-aware optimized RT regimens for each patient. We solve a multi-objective risk-based optimization under uncertainty (OUU) problem to provide a suite of optimal RT regimens that trade-off maximizing tumor control and minimizing the toxicity from RT. We use the α -superquantile risk measure [61] that accounts for the magnitude of exceeding a set threshold in a particular tumor characteristic. The risk-based optimization formulation allows one to account for the patient’s and treating physician’s risk preferences. We illustrate the predictive digital twin using an *in silico* population of HGG patients that is constructed by pooling clinical data of MRI measurements of HGG cellularity and mechanistic parameter values describing the dynamics of HGG growth and RT response from the literature [76, 54, 25]. We investigate varying levels of total dose to analyze the ensuing suite of therapeutic planning options. Our results show that the optimal RT plans can lengthen time to progression (TTP) with respect to the SOC regimen for the same total RT dose. Furthermore, we show that for certain patients, the proposed digital twin can provide optimal RT plans achieving similar tumor control as the SOC regimen, while significantly lower toxicities by lowering the radiation dose. We show that the range

of optimal solutions also provide options for patients with aggressive cancer with increased radiation doses. We demonstrate that the optimal RT plans are non-inferior to the SOC regimen in terms of tumor control and always maintain the total radiation dose within a clinically-admissible range.

The rest of this work is organized as follows. Section 2 describes the proposed digital twin methodology to create a patient-specific predictive digital twin for HGG growth and RT response under uncertainty. Section 3 illustrates the digital twin methodology for a cohort of *in silico* patients. Section 4 discusses the proposed predictive digital twin framework as well as the results of our computational study, the limitations of this work, and future lines of investigation.

2 Patient-specific digital twin methodology

In this section, we describe the methodologies used to create, update, and utilize a patient-specific predictive digital twin. We begin with an overview of the components comprising the predictive digital twin and ground them in the oncology setting in Section 2.1. We then describe the mechanistic model in Section 2.2 followed by the treatment control parameters in Section 2.3 and the generation of observational data for an *in silico* patient in Section 2.4. The Bayesian model calibration, the propagation of uncertainty for patient-specific prognosis, and the multi-objective risk-based optimization under uncertainty problem are described in Sections 2.5, 2.6, and 2.7, respectively. We detail the survival analysis method used to assess the performance of the optimal treatment plans in Section 2.8.

2.1 Predictive digital twin formulation

We adopt the mathematical abstraction for a predictive digital twin proposed in [33]. We formally define a predictive digital twin in terms of the six quantities shown in Figure 2 representing the physical and digital twins: physical state, observational data, control inputs, digital state, quantities of interest, and reward, where each of these six quantities will be considered to vary with time. The *physical twin* refers to the specific patient and the *physical state* represents the physiology and anatomy of the patient, which is only partially and indirectly observable. The *digital twin* is characterized by a computational model or a set of coupled computational models that can represent the physical twin to the desired level. The *digital state* is considered to be the parameterization of the computational models comprising the digital twin. To understand the health of the patient and inform our digital state, we rely upon *observational data* obtained from the physical twin, such as data obtained from MRI. There is a strong relationship between the digital state and the observational data since the accuracy of the digital state depends on the type and quantity of the observational data. The *control inputs* in a clinical setting are the therapeutic decisions that influence the physical state of the patient, such as the dose and scheduling of medical interventions. The control inputs can also comprise scheduling decisions for observational data, such as when a patient comes in for imaging. We will use the predictive digital twin to inform our choice of control inputs. The *quantities of interest* (QOIs) are computational estimates of possibly unobservable patient characteristics, which are evaluated using the computational models underlying the digital twin, such as tumor cell count, time to progression and toxicity. The *reward* is used to quantify the performance of the patient-twin system and can encode the success or failure of the system after applying a specific control, such as the risk of under-treating a specific disease giving a therapeutic plan. These quantities describe an abstract coupled patient-twin system

representing a patient physical twin and their associated computational digital twin.

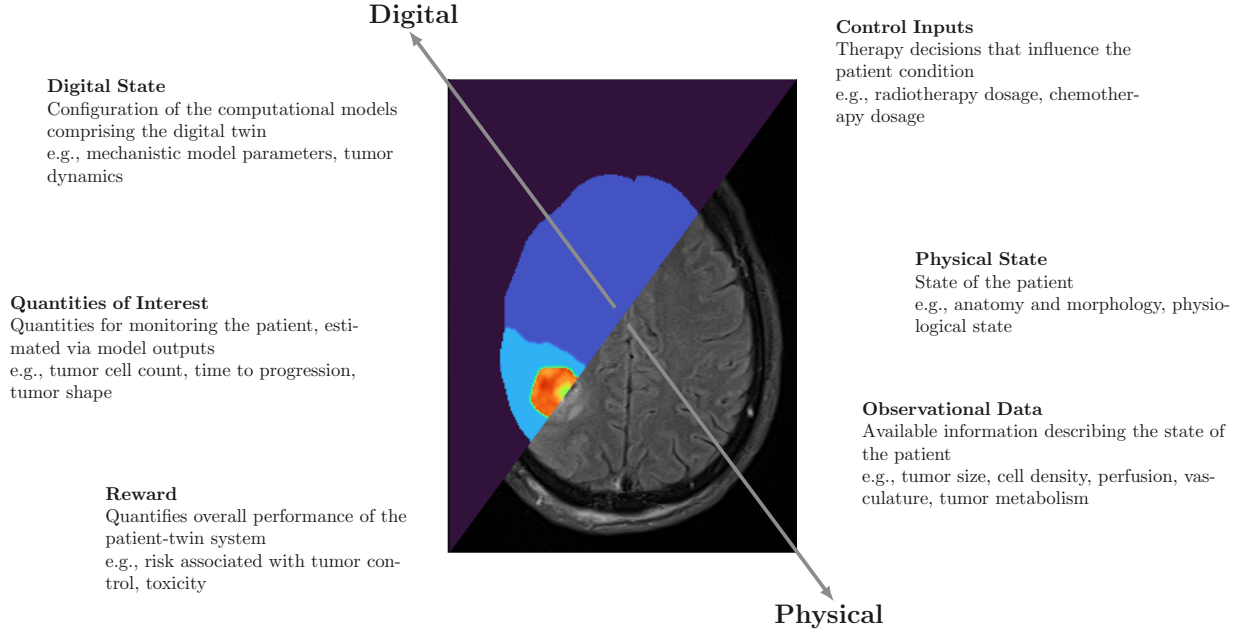


Figure 2: Overview of the components comprising a predictive digital twin in the oncology setting.

2.2 Mechanistic tumor growth model comprising the HGG digital twin

The computational model underlying the predictive digital twin is the logistic growth model of tumor growth governed by the ordinary differential equation (ODE) [4, 31, 22]

$$\frac{dN}{dt} = \rho N \left(1 - \frac{N}{K}\right); \quad (1)$$

$$N(t=0) = N_{\text{initial}},$$

where $N(t)$ is the number of tumor cells at time t , N_{initial} is the initial tumor burden, ρ is the net proliferation rate of the tumor cells, and K is the carrying capacity of the tissue (i.e., the maximum number of tumor cells that can be sustained physically and biologically).

We model the effects of RT and concomitant chemotherapy as discrete treatment events that result in an instantaneous reduction of tumor cell count at treatment time, $N_{\text{post-treatment}}$, to a fraction of the pre-treatment cell count $N_{\text{pre-treatment}}$, given by

$$N_{\text{post-treatment}}(t; u_t) = N_{\text{pre-treatment}}(t)S(u_t), \quad (2)$$

where S is the surviving fraction of tumor cells resulting from a single dose of RT and chemotherapy at time t , which is given by u_t . The surviving fraction is defined by a linear-quadratic model of cell survival [44] with a multiplicative term to account for the concurrent chemotherapy effect as

$$S(u_t) = S_C S_{\text{RT}}(u_t) = S_C \exp(-\alpha_{\text{RT}} u_t - \beta_{\text{RT}} u_t^2), \quad (3)$$

where S_C is the surviving fraction resulting from a single dose of chemotherapy, S_{RT} is the surviving fraction resulting from a single dose of RT, and α_{RT} and β_{RT} are the radiosensitivity parameters.

The entire model formulation given by Eqs. (1)-(3) can be written as

$$\frac{dN}{dt} = f(N; \theta, \mathbf{u}), \quad (4)$$

where $N(t; \theta, \mathbf{u})$ is the number of tumor cells at time t given that $\theta := [\rho, K, N_{\text{initial}}, \alpha_{RT}]^\top \in \Omega \subseteq \mathbb{R}_+^4$ are the patient-specific probabilistic model parameters in the digital state and \mathbf{u} is a vector comprising all the RT treatment doses u_t given at time t . The treatment control considered here adapts the RT dose \mathbf{u} to specific patients as elaborated in Section 2.3. All the parameters necessary to define the ODE model are provided in Table 1. We fix the radiosensitivity parameter ratio to $\alpha_{RT}/\beta_{RT} = 10$ [62] and the surviving fraction resulting from chemotherapy to $S_C = 0.82$ [25]. The system of ODEs in Eq. (4) is solved via a forward Euler scheme with sufficiently small time step size of 0.2 days to ensure numerical stability. The discrete treatment events are applied at the beginning of each day.

Table 1: Mechanistic tumor growth model parameters

Parameter type	Name	Symbol	Units
Probabilistic parameters	Proliferation rate	ρ	day ⁻¹
	Carrying capacity	K	cells
	Initial tumor burden	N_{initial}	cells
	Radiosensitivity parameter	α_{RT}	Gy ⁻¹
Fixed parameters	Surviving fraction due to chemotherapy	S_C	-
	Radiosensitivity parameter ratio	α_{RT}/β_{RT}	Gy

2.3 Treatment control: patient-specific RT treatment regimen

The treatment goal considered for this digital twin illustration is to adapt the RT regimen to specific patients. The Stupp protocol [68] is the current SOC for treatment of HGG. The SOC administers six weeks of treatment featuring five consecutive treatment days each week with a fractionated RT dose of 2 Gy/day. We consider a similar setup as the SOC dose fractionation to define the treatment control $\mathbf{u} \in \mathcal{U} \subseteq \mathbb{R}_+^{n_u}$ as a vector of weekly-fractionated RT dose with $n_u = 6$ being the number of weeks of RT. Under this setup, the SOC is represented by $\mathbf{u}_{\text{SOC}} = [2, 2, 2, 2, 2, 2]$ Gy/day leading to a total dose of $5\|\mathbf{u}_{\text{SOC}}\|_1 = 60$ Gy, where $\|\cdot\|_1$ denotes the ℓ_1 norm (i.e., the summation of the absolute value of the vector entries). The SOC is used as the baseline to compare with the optimized RT treatment regimens. We administer chemotherapy throughout the six weeks even when no RT is prescribed. Note that extensions of this work can consider a finer time scale, such as allowing for a varying dose on a daily basis within each week of the treatment as well as adapting the chemotherapy treatment.

To calibrate the RT-related parameters in the predictive digital twin, we need at least one observation after starting the RT treatment (see Section 2.5 for details on calibration). Thus, we consider that the patients receive the SOC RT plan during this first week and that the intra-treatment MRI is performed at the end of the first week of RT (here, simulated patients as described in Section 2.4). We fix the fractionated

dose of the first week to the SOC value of 2 Gy/day and use the predictive digital twin to optimally control the RT doses for the remaining five weeks (i.e., $\mathbf{u} = [u_1 = 2, u_2, \dots, u_6]$, where $u_i \in [0, 10]$ Gy/day, $i = 2, \dots, 6$). Importantly, this strategy to optimally adjust the RT plan with our predictive digital twin follows the same approach as image-driven adaptive RT methods that are currently being explored in the clinical-research setting [46, 66].

2.4 Simulated patients as proxy for observational clinical data

The physical twin state of each patient is partially and indirectly observed through specific observational data, which are leveraged to inform the digital twin state. In the HGG setting, these observational data can be acquired non-invasively using MRI. Recently, MRI data have been used to calibrate computational models and obtain tumor forecasts [28, 27, 26, 25] after appropriate post-processing. In particular, the MRI data needs to be post-processed to extract the total tumor burden as a cell count, such that it relates to the state variable N of the ODE model in Eq. (4). Since we are collapsing the spatial information of the tumor provided by the MRI data, we incur a volume-wise error in the measurement of the tumor burden. This source of error has been studied by both [50] and [43], although it is difficult to quantify its effect on the model parameterization and ensuing forecasts.

We consider an *in silico* patient cohort where the physical state of the HGG tumors is simulated by generating noisy observations of the solution to the ODE model in Eq. (4) using an underlying “true” parameter set θ_{true} , which is varied for each patient in the cohort. Note that θ_{true} is never seen by the predictive digital twin and the model parameters instead adopt a probabilistic formulation based on the noisy observations generated for each *in silico* patient. We simulate a collection of n_{obs} observations $\mathbf{o} = [o_{t_1}, \dots, o_{t_{n_{\text{obs}}}}]$ at times $\{t_i\}_{i=1}^{n_{\text{obs}}}$ with an additive noise model as

$$o_{t_i} = N(t_i; \theta_{\text{true}}, \mathbf{u}) + \varepsilon_i, \quad (5)$$

where the noise ε_i follows a truncated normal distribution $\mathcal{TN}(0, \sigma^2, -N(t_i; \theta_{\text{true}}, \mathbf{u}), +\infty)$ with a lower truncation bound of $-N(t_i; \theta_{\text{true}}, \mathbf{u})$ to avoid non-physical negative observations. The general truncated normal distribution definition $\mathcal{TN}(\mu, \sigma^2, a, b)$ can be read as μ and σ^2 specifying the mean and variance, respectively, of the general normal distribution with the truncation range as $[a, b]$ [9]. For this work, we assume a constant standard deviation of $\sigma = 2 \times 10^9$ cells, corresponding to a value of 10% of the mean value from population data used to model the initial tumor burden.

2.5 Patient-specific tumor modeling via Bayesian model calibration

We use a Bayesian formulation that combines prior knowledge with observed data to update the probabilistic distribution placed on the model parameters. The observational data from the patient are assimilated by solving an inverse problem to calibrate the parameters in the computational model [5, 67, 48] to the specific patient. The Bayesian framework provides confidence levels for the computational model output. Prior knowledge plays a critical role in the oncology setting as it is an especially data-poor regime owing to the difficulty and expense of collecting quantitative patient information such as MRI data. Priors allow us to inject knowledge of the disease at the population level to inform the modeling process of a specific patient. Priors $\mathcal{P}(\theta)$ for the probabilistic parameters θ are constructed from reported values of clinical data in the

literature. The study by [76] computed the global proliferation rate for a cohort of 31 patients. [54] conducted a study to compute a plausible set of radio-sensitivity parameters for the linear-quadratic model. We define an independent truncated normal distribution based on the values from the literature as our prior distribution as shown in Table 2. A truncated normal distribution is used to enforce positivity as well as to account for physically reasonable upper and lower bounds on the model parameters.

Table 2: Prior distribution for the probabilistic model parameters $\theta := [\rho, K, N_{\text{initial}}, \alpha_{\text{RT}}]^\top$

Parameter	Mean	Standard deviation	Lower bound	Upper bound
ρ	0.09	0.15	0.007	0.25
K	1×10^{11}	2×10^{10}	9×10^{10}	1.8×10^{11}
N_{initial}	1.9×10^{10}	1.2×10^{10}	4.7×10^9	4.7×10^{10}
α_{RT}	0.05	0.025	0.001	0.1

Observational data \mathbf{o} from the patient are assimilated to estimate the updated probability distribution of θ , otherwise known as the posterior distribution, through the Bayesian update formula given by

$$\mathcal{P}(\theta|\mathbf{o}) \propto \mathcal{P}(\mathbf{o}|\theta)\mathcal{P}(\theta), \quad (6)$$

where $\mathcal{P}(\theta|\mathbf{o})$ is the posterior distribution of the digital state conditioned upon the observed data, $\mathcal{P}(\mathbf{o}|\theta)$ is the likelihood of the observed data given a digital state, and $\mathcal{P}(\theta)$ is the prior distribution that encodes our knowledge about the patient’s state before any observations are made.

We consider a scenario with three available MRI observations on days 0 (post-surgery image), 20 (pre-RT image), and 27 (mid-RT image). As noted in Eq. (5), we assume an additive noise model for our observational data generation. The observation on day 0 only provides knowledge about the initial tumor burden N_{initial} for the patient. The observation on day 20 provides information on all the probabilistic model parameters except the radiosensitivity parameter α_{RT} . The observation on day 27 is generated using the SOC fractionated dose and provides knowledge about all the probabilistic model parameters $\{\rho, K, N_{\text{initial}}, \alpha_{\text{RT}}\}$. This leads us to solve the inverse problem for model calibration in two steps:

Step 1: After obtaining the post-surgery image for the patient, the first step updates the posterior of the initial tumor burden to use the truncated normal distribution associated with the observation o_{t_1} at $t_1 = 0$ days as the mean with σ^2 accounting for the measurement error as $N_{\text{initial}} \sim \mathcal{TN}(o_{t_1}, \sigma^2, \max\{0, o_{t_1} - 2\sigma\}, o_{t_1} + 2\sigma)$. The upper and lower truncation bounds are defined by $\pm 2\sigma$ from the observational data. The lower truncation bound is further modified to avoid non-physical negative values of initial tumor burden by using $\max\{0, o_{t_1} - 2\sigma\}$. The updated distribution of N_{initial} is used as the prior distribution in Step 2.

Step 2: The Bayesian inference is performed via Markov Chain Monte Carlo (MCMC) for assimilating the remaining two observations $[o_{t_2}, o_{t_3}]$ from $t_2 = 20$ days and $t_3 = 27$ days to complete the digital twin model calibration for tumor physiology and radiation response. We use the updated distribution of N_{initial} from step 1 and the priors for $\rho, K, \alpha_{\text{RT}}$ given in Table 2. Then the patient-specific posterior distribution for all the probabilistic model parameters are obtained by using MCMC for the remaining two observations. We use the MCMC implementation in the PyMC3 Python package [64] with four chains of 100,000 samples each. The posterior distribution is characterized by 100,000 retained

samples. For a more in-depth discussion of MCMC techniques, the reader is referred to [65] and the references within.

2.6 Patient-specific prognosis via uncertainty propagation

Consider an uncertain digital state defined by the posterior distributions of the model parameters generated by Bayesian model calibration as described in Section 2.5. At any point in time, we can propagate uncertainty forward in time using the model in Eq. (4) to estimate our quantities of interest, which characterize the tumor control resulting from an RT regimen and the treatment toxicities. Hence, a QoI maps realizations of random model parameters $\theta \in \Omega$ and a treatment control $\mathbf{u} \in \mathcal{U}$ to a real value as $M : \mathcal{U} \times \Omega \mapsto \mathbb{R}$. There are a variety of clinically relevant QoIs that could be considered to characterize the tumor control resulting from a treatment regimen. These include the predicted tumor burden (e.g., defined through the tumor cell count) after a set period of time post-treatment, time to progression (TTP) of the tumor cell count beyond a given threshold, or the amount of time that the tumor cell count is kept below a specified threshold. In this work, the QoI for tumor control is based on the TTP [63]. We consider toxicity to be proportional to the total dose of radiation delivered to the patient during RT, i.e. $5\|\mathbf{u}\|_1$. Note that the QoI for toxicity does not depend on θ and, thus, it is a deterministic quantity since we have no uncertainty in the dose administered to a patient.

We define the TTP as the amount of time that the tumor takes to proliferate to a clinically-relevant size after the conclusion of six weeks of RT, which we set as the tumor cell count right before the onset of the RT regimen (i.e., at $t = 20$ days). Hence, to calculate the TTP, we further define a threshold tumor cell count $N_{\text{th}}(\theta) := N(t = 20; \theta, \mathbf{u})$, which does not depend on \mathbf{u} since it is a pre-treatment quantity. Additionally, we denote the day of the conclusion of RT as $t_{\text{post-RT}}$ (here, $t_{\text{post-RT}} = 20 + 6 \times 7 = 62$ days). Since a second round of chemotherapy is generally prescribed to prevent or combat tumor progression three months after the conclusion of RT, we consider a finite time-horizon of $t_{\text{finite}} = 152$ days after surgery for the end of our simulations and the calculation of the TTP. We can now express the TTP as

$$T_{\text{TTP}}(\mathbf{u}, \theta) := \min_{t_{\text{post-RT}} < t \leq t_{\text{finite}}} \{t : N(t; \theta, \mathbf{u}) > N_{\text{th}}(\theta)\} - 20. \quad (7)$$

We want to maximize the TTP, or equivalently, minimize the negative of TTP. Thus, we define our QoI for tumor control as $M(\mathbf{u}, \theta) = -T_{\text{TTP}}(\mathbf{u}, \theta)$ to be compatible with the definition of a general minimization problem (see Section 2.7).

In general, the nonlinear effects of the underlying computational model and the non-parametric distributions of θ prohibit one from obtaining an analytic expression for the distribution of the QoI, hence we rely on the Monte Carlo sampling approach. We sample a realization of the parameter set θ from the posterior distribution describing a patient’s digital state, solve the forward model in Eq. (4) to obtain the TTP, and compute the QoI $M(\mathbf{u}, \theta)$. The process is continued till a desired number of samples n_{MC} of θ are processed to obtain the realizations of TTP as shown in Figure 3a. The distribution of TTP and the connection with $M(\mathbf{u}, \theta)$ is illustrated in Figures 3b and 3c.

To make decisions concerning the patient-specific design of optimal treatment regimens, there are different statistics associated with the distribution of $M(\mathbf{u}, \theta)$ that can be used as the reward function. We use a risk measure as the statistic to characterize the risk in tumor control associated with a given treatment regimen \mathbf{u} . The risk measure is denoted by $\mathcal{R} : \mathcal{U} \times \Omega \mapsto \mathbb{R}$. Different types of risk measures can be used to

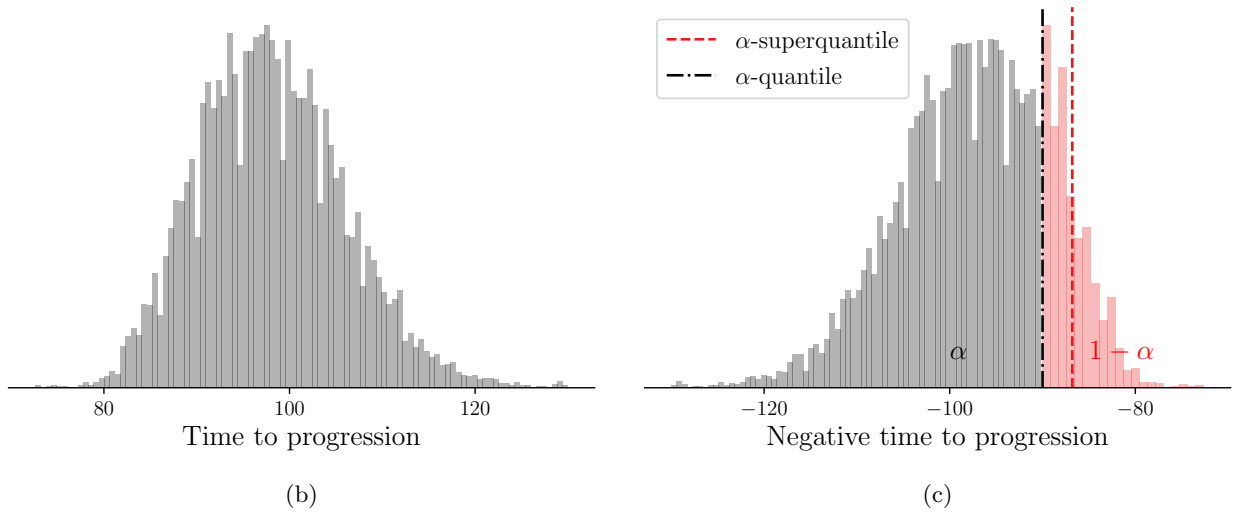
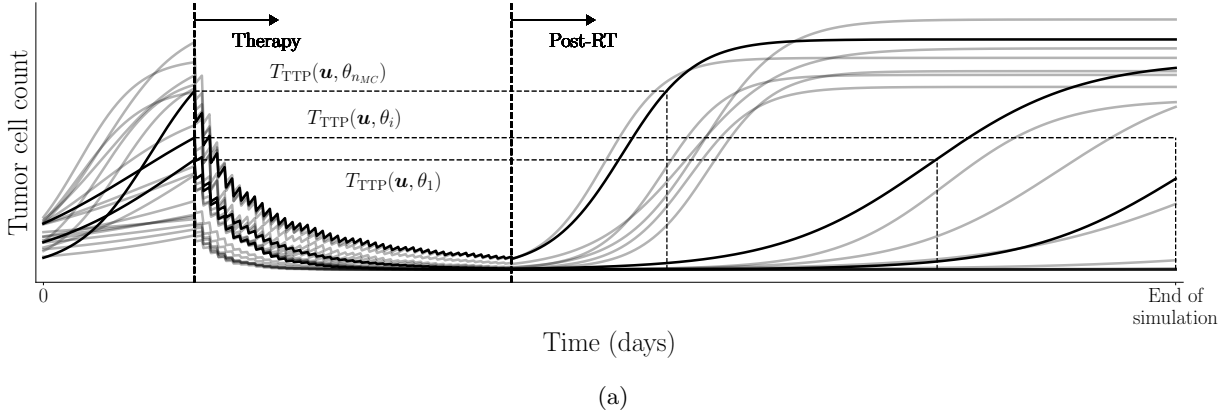


Figure 3: **Illustration of TTP and use of negative TTP as QoI.** (a) Visual representation of TTP for various trajectories obtained from n_{MC} Monte Carlo samples of θ , (b) histogram of samples of $T_{TTP}(\mathbf{u}, \theta)$ and (c) histogram of samples of QoI $-T_{TTP}(\mathbf{u}, \theta)$ along with the estimated risk with $\alpha = 0.95$.

quantify the risk associated with tumor control, such as the probability of exceeding a threshold TTP value and the α -quantile for a set risk level α . We use the α -superquantile risk measure [61, 59, 60, 35], which has certain desirable mathematical properties. The α -superquantile satisfies two notions of certifiability in risk-based OUU [8]: (i) accounting for near-failure and catastrophic failure events, and (ii) preserving convexity of the underlying QoI to aid in convergence guarantees for risk-based optimization formulations. The α -superquantile risk measure accounts for the magnitude of the largest $100(1 - \alpha)\%$ realizations that captures the specified portion of the worst-case scenarios through risk level α . In oncology, it is important to account for the extent to which a tumor progresses beyond a specified threshold or remains in remission but close to the threshold rather than just the frequency of exceeding the threshold. Clinical ramifications of such extreme cases could result in under-treatment of aggressive disease resulting in poor tumor control and early disease progression. This is especially important in high-grade glioma where the prognosis is already overwhelming poor. The α -superquantile risk measure can take into account the magnitude of the TTP exceeding a set threshold to build in a desired level of conservativeness in a data-driven way.

To define the α -superquantile risk measure, we first define the related quantity of α -quantile Q_α for risk

level $\alpha \in (0, 1)$ as

$$Q_\alpha [M(\mathbf{u}, \theta)] := F_{M(\mathbf{u}, \theta)}^{-1}(\alpha), \quad (8)$$

where $F_{M(\mathbf{u}, \theta)}^{-1}$ is the inverse cumulative distribution function of $M(\mathbf{u}, \theta)$. Then, we can define the α -superquantile $\bar{Q}_\alpha[M(\mathbf{u}, \theta)]$ as

$$\bar{Q}_\alpha[M(\mathbf{u}, \theta)] := Q_\alpha[M(\mathbf{u}, \theta)] + \frac{1}{1-\alpha} \mathbb{E} \left[[M(\mathbf{u}, \theta) - Q_\alpha[M(\mathbf{u}, \theta)]]^+ \right], \quad (9)$$

where $[\cdot]^+ = \max(0, \cdot)$. When the cumulative distribution of $M(\mathbf{u}, \theta)$ is continuous, we can view $\bar{Q}_\alpha[M(\mathbf{u}, \theta)]$ as the conditional expectation of the QoI conditioned on the QoI exceeding the α -quantile as $\bar{Q}_\alpha[M(\mathbf{u}, \theta)] = \mathbb{E} [M(\mathbf{u}, \theta) \mid M(\mathbf{u}, \theta) \geq Q_\alpha[M(\mathbf{u}, \theta)]]$. We use Algorithm 1 in Ref. [8] to estimate the α -superquantile through Monte Carlo simulations. Figure 3c shows an illustration of the data-driven conservativeness inferred from the Monte Carlo samples when using α -superquantile compared to α -quantile ($\bar{Q}_\alpha[M(\mathbf{u}, \theta)] > Q_\alpha[M(\mathbf{u}, \theta)]$).

The α -superquantile risk measure is used to formulate a risk-based optimization problem under uncertainty to select the optimal treatment regimen as described in Section 2.7. The α -superquantile risk measure helps in making risk-aware decisions, where the risk preference for each patient can be adjusted by changing the value of α . Hence, larger values of α lead to more risk-averse decisions. We demonstrate our approach with a value of $\alpha = 0.95$.

2.7 Optimal patient-specific RT treatment regimens under uncertainty

In this work, we consider two clinical objectives: (i) minimizing the risk associated with the QoI characterizing the tumor control, and (ii) minimizing toxicity through a proportional quantity to the total RT dose. To define personalized optimal RT regimens achieving these goals, we formulate a multi-objective risk-based optimization problem under uncertainty. This method results in a suite of eligible optimal RT regimens featuring different levels of trade-off between (i) decreasing toxicity by minimizing the total dose, or (ii) minimizing the risk associated with tumor control. The multi-objective formulation is given by

$$\mathbf{u}^* = \arg \min_{\mathbf{u} \in \mathcal{U}} (\mathcal{R}(M(\mathbf{u}, \theta)), 5\|\mathbf{u}\|_1), \quad (10)$$

where optimized therapy comes into effect after the first week of RT and we optimize the RT doses for the remaining five weeks as denoted by $\mathbf{u} = [u_1 = 2, u_2, \dots, u_6]$ (see Section 2.3). Note that we can also explore the trade-off with changing risk preferences by varying the risk level α for the α -superquantile risk measure, as discussed in Section 2.6. An illustration of the predictive digital twin for HGG that shows the specific timeline is shown in Figure 4.

To solve the multi-objective problem in Eq. (10), we reformulate it into a constrained single-objective problem using the ϵ -constraint method [17, 29] as

$$\begin{aligned} \mathbf{u}^* &= \arg \min_{\mathbf{u} \in \mathcal{U}} \mathcal{R}(M(\mathbf{u}, \theta)) \\ &\text{s.t. } 5\|\mathbf{u}\|_1 \leq D_{\max}, \end{aligned} \quad (11)$$

where D_{\max} is the total dose to be delivered. The value of D_{\max} is varied over a range of total doses to obtain the Pareto optimal solutions [29]. Each of the Pareto optimal solutions provide different balance of toxicity and tumor control. Hence, the ϵ -constraint method is suitable for our application since we have a known primary objective of minimizing the risk associated with tumor control, while reducing the dose is




Figure 4: **Predictive digital twin timeline for HGG patients.** The predictive digital twin features a personalized risk-based strategy to optimize the adaptive RT regimen under uncertainty. This strategy is constructed assuming a standard collection of MRI data during the course of the clinical management of HGG after surgery, whereby MRI scans are prescribed after the surgical intervention (day 0), before the onset of RT (day 20), and during the second week of the RT regimen (day 27). On day 27, the digital twin is calibrated using the MRI data and then the risk-aware treatment plan is solved using the calibrated model and deployed for the remaining five weeks of RT.

a secondary objective for controlling toxicity. We also have a structured way to define the range of values for D_{\max} based on the preferred total dose, which solves the challenge of selecting meaningful levels in the ϵ -constraint method. Thus, the specific advantages of the ϵ -constraint method in the context of the multi-objective problem defined in Eq. (10) are: (i) preserving properties of the underlying problem, such as convexity of QOIs and risk measures, and (ii) ease of specification of the number of Pareto optimal treatment plans by selecting the number of D_{\max} values and solving a single-objective constrained optimization problem each time.

The solution of Eq. (11) will always lead to an optimal treatment regimen with active constraint ($5\|\mathbf{u}^*\|_1 = D_{\max}$) since increasing dose directly leads to decreasing risk. However, after an initial analysis of the optimization problem in Eq. (11), we found that it leads to sub-optimal solutions for patients whose HGG exhibits a low tumor cell proliferation rate and/or low initial tumor burden. We observed that these patients do not require the maximum total dose enforced by an active constraint at the optimal solution to achieve the same amount of tumor control after the conclusion of the RT regimen. Typically for such patients, the amount of tumor control corresponds to a maximum possible value of 132 days for the TTP realizations based on the finite time-horizon of 152 days for our simulations. In other words, the same amount of risk associated with tumor control can be obtained while reducing the total dose below D_{\max} for these patients.

To account for such cases, we add an additional penalty term using parameter λ (here, $\lambda = 0.001$) on the total delivered dose as

$$\begin{aligned} \mathbf{u}^* &= \arg \min_{\mathbf{u} \in \mathcal{U}} \mathcal{R}(M(\mathbf{u}, \theta)) + \lambda \|\mathbf{u}\|_1 \\ &\text{s.t. } 5\|\mathbf{u}\|_1 \leq D_{\max}. \end{aligned} \tag{12}$$

The single-objective constrained optimization problem in Eq. (12) is solved for various selections of the total dose threshold, $D_{\max} \in \{40, 50, 60, 70, 80, 100\}$ Gy, to obtain the Pareto front of optimal solutions. The SciPy Python package [74] is used for solving the optimization problem. We use the basin-hopping algorithm [75] for multi-start local optimization with the gradient-free COBYLA (constrained optimization by linear approximation) optimizer [53] from the SciPy package ¹. We use 20 restarts of COBYLA with a maximum of 200 function evaluations in each restart. The risk estimation in each function evaluation during the optimization uses $n_{\text{MC}} = 5000$ forward simulations.

The risk-based multi-objective problem formulation with the specific choice of objectives used here has the following desirable properties:

- accounts for uncertainty during clinical decision-making process to provide multiple patient-specific optimal treatment regimens
- leads to a suite of patient-specific optimal treatment regimens with different levels of trade-offs between tumor control and toxicity from RT to allow more flexibility to consider the patient’s and the treating physician’s preferences
- provides an optimal treatment regimen option for a patient that can achieve similar tumor control as the SOC, but with reduced RT dose to mitigate toxicity effects whenever possible
- provides an optimal treatment regimen option with increased RT dose for patients with aggressive cancer, where the SOC does not provide sufficient tumor control either for patient preference or to allow for further surgical intervention

2.8 Survival analysis

To assess the tumor control achieved with the different optimized treatment regimens for varying D_{\max} and compare them against the SOC, we analyze the time to tumor progression across the *in silico* patient cohort using the Kaplan-Meier estimator [32, 11]. In this case, a right-censor criterion is enforced to account for the finite time-horizon of HGG response simulation at $t_{\text{finite}} = 152$ days, which leads to a maximum TTP value of 132 days. For a cohort of n_p patients, we define the probability of survival to tumor progression at time t as the probability of the α -superquantile value of the TTP being longer than time t . Thus, for each treatment obtained with a different D_{\max} , the α -superquantile value of the TTP of each patient is used to calculate the survival probability as

$$P_S(t) := \prod_{i:t_i \leq t_{\text{finite}}} \left(1 - \frac{d(t_i)}{m(t_i)}\right), \tag{13}$$

where $d(t_i)$ denotes the number of patients with TTP equal to time t_i and $m(t_i)$ are the number of patients with TTP larger than time t_{i-1} for all $i = 0, \dots, n_p$. A further advantage of the probabilistic modeling

¹SciPy Basin-hopping; SciPy COBYLA

approach used in the predictive digital twin is the ability to estimate the variance in survival probability numerically from the samples of the TTP instead of relying on the approximation provided by the Greenwood’s formula [15].

To assess the statistical significance of different treatment plans, the logrank test [6] is used to compare the survival distributions between digital-twin-based treatment plans and the SOC. The logrank test is implemented through the LifeLines Python package [12]. Note that the null hypothesis for this test is that there is no difference between the populations in the probability of survival, so significant p-values indicate that the considered treatment is significantly different from the SOC.

3 Illustrative cohort of *in silico* patients

This section presents an illustration of the digital twin methodology for a cohort of *in silico* patients. We describe the details of the computational study in Section 3.1 followed by the results for Bayesian calibration (Section 3.2), risk-based OUU (Section 3.3), and survival analysis (Section 3.4). The code used for generating the results reported for the predictive digital twin is available at <https://github.com/Willcox-Research-Group/predictive-dtwin-glioma-frontiers>.

3.1 Computational study format

We initialize an *in silico* cohort of 100 patients whose true physical states θ_{true} are determined by sampling from the marginal distributions of the population level priors (see Section 2.4). Noisy measurements are then generated from this cohort using Eq. (5) to mimic the data that would be collected in the clinic. The specific observational and control timeline used for our predictive digital twins is outlined in Figure 4. Note that we are considering a finite time-horizon of three months post-RT ($t_{\text{finite}} = 152$ days) for our simulations and optimization. Thus, the maximum possible value of TTP is 132 days (see Eq. (7)).

We divide the cohort of 100 patients into three groups based on the SOC TTP α -superquantile exhibited by the patient:

- (i) Early progressors: 16 patients with SOC TTP α -superquantile ≤ 1 month after end of RT
- (ii) Intermediate progressors: 62 patients with SOC TTP α -superquantile between 1-3 months after end of RT
- (iii) Late progressors: 22 patients with SOC TTP α -superquantile ≥ 3 months after end of RT

We first analyze the results for Bayesian calibration and risk-aware treatment plans produced by the patient-specific predictive digital twins for three patients with varied growth and response behaviors. The underlying physical state of the three patients is provided in Table 3. Note that these patient parameters are not known to the digital twin and are only used to generate observational data for the case-study. Patients 1 and 2 are intermediate progressors and Patient 3 is an early progressor according to the patient groups defined above. We then analyze the effectiveness of our predictive digital twins over the 100 patient cohort and the three patient groups.

Table 3: Physical state model parameters θ_{true} for simulating *in silico* case-study patients.

Model parameter	Patient 1	Patient 2	Patient 3
Proliferation rate ρ (day^{-1})	1.14e-01	1.09e-01	2.25e-01
Carrying capacity K (cells)	1.17e+11	1.09e+11	1.40e+11
Initial tumor burden N_{initial} (cells)	1.54e+10	2.60e+10	2.62e+10
Radiosensitivity parameter α (Gy^{-1})	1.05e-03	4.58e-02	3.90e-02
Chemotherapy surviving fraction S_C		0.82	
Radiosensitivity parameter ratio α/β (Gy)		10	

3.2 Patient-specific Bayesian calibration of the digital twin

The patient-specific calibrated digital twin is obtained by assimilating the three observational data available for each patient as described in Section 2.5, while the priors of all patients are obtained from population clinical data as described in Table 2. The histograms in Figure 5 show the prior and the posterior distributions for the probabilistic model parameters θ in the digital state for the three case-study patients. We observe that the posterior distributions concentrate around the (unseen) true physical state θ_{true} given in Table 3 (shown by dashed lines in Figure 5). We can also see that the overall uncertainty in the probabilistic parameters reduces compared to the prior distribution. In the clinical context, this could be interpreted as updating our belief about the specific patient’s tumor dynamics as more data is collected during the clinical management of the disease.

To forecast model uncertainty, the uncertainty in the probabilistic model parameters is propagated forward by sampling 10,000 parameter sets from the joint posterior obtained by MCMC and solving the mechanistic model in Eq. (4) to generate the posterior trajectories. Figure 6 shows the posterior trajectories for the three case-study patients with the maximum and minimum bands. The posterior trajectories simulated from the calibrated digital twins exhibit tighter bounds (as shown by the orange posterior vs blue prior shaded regions) compared to the prior trajectories, which indicate a decrease in overall uncertainty. This is a consequence of our updated belief about the patient’s tumor dynamics resulting from the integration of longitudinal data into the digital twin for each individual patient. We also see that the observations for each patient are captured within the shaded region of the respective posterior trajectories, which indicates the range of posterior trajectories. The quantification of uncertainty in model outputs is crucial for estimating the risk of tumor propagation used for risk-aware clinical decision-making.

3.3 Patient-specific treatment optimization under uncertainty

We now present the risk-aware optimal treatment regimens obtained using the calibrated digital twins. We use the α -superquantile risk measure for quantifying risk of tumor growth as described in Section 2.6 and then solve a risk-based multi-objective problem to obtain a suite of patient-specific optimal treatment regimens under uncertainty as described in Section 2.7. Figure 7 shows the results of our patient-specific treatment optimization for the three case-study patients. We plot the TTP α -superquantile against the dose threshold to show the Pareto front indicating the trade-off between the two quantities. Figure 7 shows that the SOC regimen is not the optimal dosing schedule in multiple ways. First, for the applied radiation dose

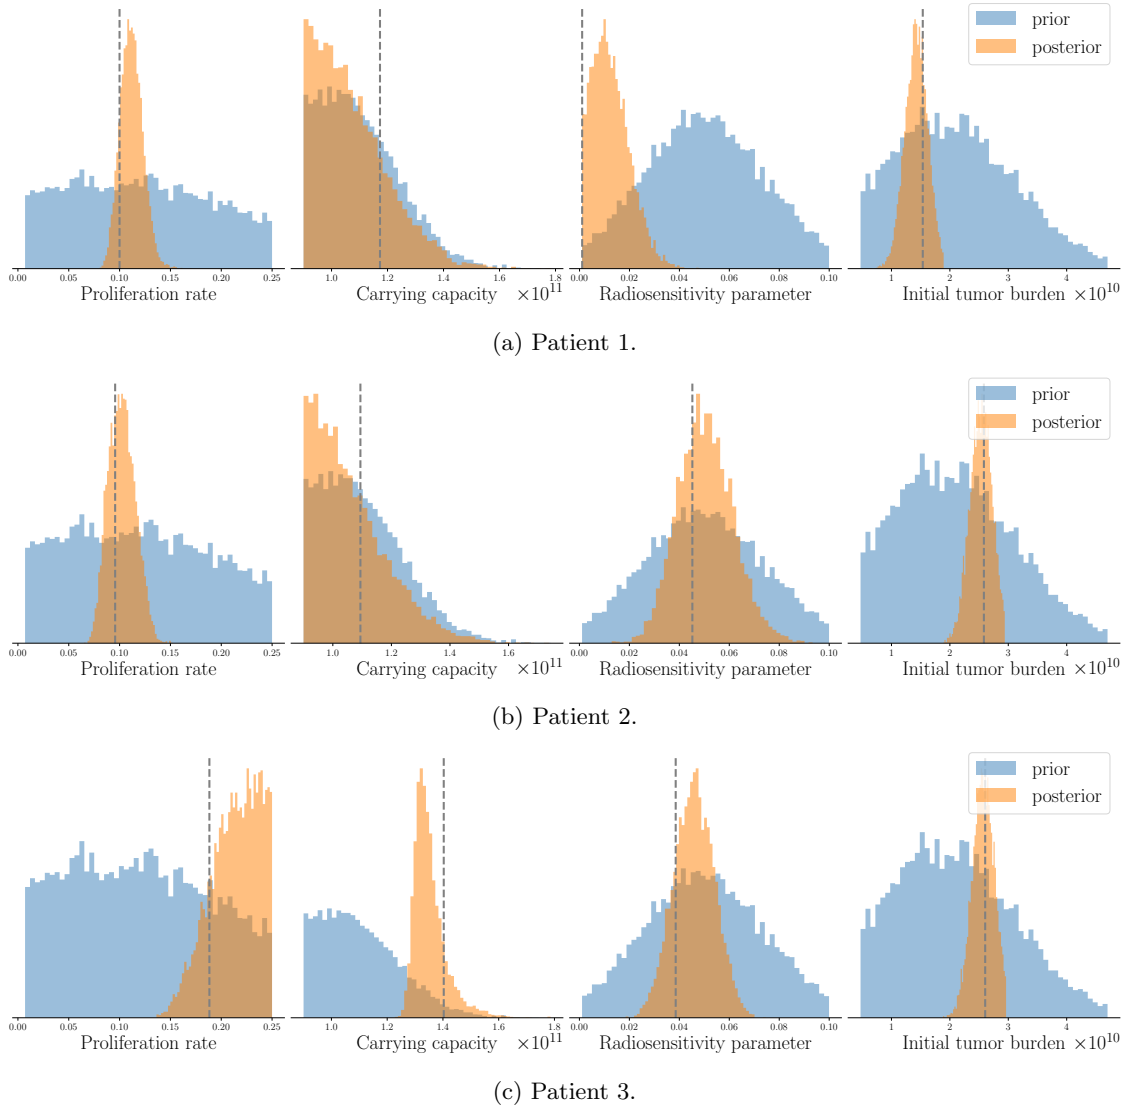




Figure 5: Posterior parameter distributions of the calibrated digital twins after assimilating observed data at the three imaging visits are shown for three patients: (a) Patient 1, (b) Patient 2, and (c) Patient 3. The dashed gray line indicates true parameter values for reference and not seen by the digital twin. Note that the same prior distribution is used for initializing digital twins of all the patients. Posterior distributions concentrate around the unseen true parameters and show reduction in uncertainty compared to prior distributions.

of 60 Gy, the optimal radiotherapy plan can demonstrate superior tumor control in terms of longer TTP. Second, similar or better tumor control can be exerted with a reduced amount of dose compared to SOC, as shown by the TTP values at the 40 Gy and 50 Gy levels. Third, greater tumor control than the SOC regimen can be exerted for total doses ≥ 60 Gy as seen in Figure 7.


Figure 8 shows the entire distribution of the TTP using the optimized treatment regimens compared to the SOC treatment for the three case-study patients. For Patient 1 in Figure 8a, we observe that with total



(a) Patient 1.



(b) Patient 2.



(c) Patient 3.

Figure 6: Calibrated digital twin posterior trajectory after assimilating observed data at the three imaging visits compared to prior trajectory for the three patients: (a) Patient 1, (b) Patient 2, and (c) Patient 3 (median shown by solid line; minimum and maximum shown by dashed lines). Posterior trajectories capture the observed data and reduce uncertainty compared to prior trajectories.

dose constraint $D_{\max} \geq 60$ Gy we see a progressive separation between the TTP distributions obtained from the SOC and the optimized treatment plans as we increase the total dose. The TTP distribution for optimized plans move towards higher TTP values. Thus, the optimal treatment plans lead to TTP distributions that are therapeutically superior than the ones obtained from the SOC. Additionally, considering total doses lower than the 60 Gy administered by the SOC, we observe that there is not much separation in the distributions of TTP. This result highlights the possibility of reducing the total dose administered in the SOC protocol

by 10-20 Gy while obtaining a similar tumor control for Patient 1. A similar trend in separation of TTP distributions is seen for Patient 3, as shown in Figure 8c. In fact, the separation is obvious even at a total dose constraint of 50 Gy, for which the TTP distribution obtained from the risk-aware optimization already indicates superior therapeutic outcomes than the one obtained from the SOC. At higher doses, the TTP distributions from the optimized results shift towards higher TTP values for patient 3 and shows separation from the TTP distribution obtained from the SOC. For Patient 2 in Figure 8b, the most likely outcome is that the disease is controlled within the finite time horizon considered in this study (i.e., $t_{\text{finite}} = 152$ days). However, the optimized treatment plans still increase the probability of treatment success. This can be seen by the higher peaks at the “end of simulation” in Figure 8b.




Figure 7: Pareto front showing the suite of patient-specific optimal therapy solutions obtained from the multi-objective risk-based OOU compared to SOC treatment for the three patients: (a) Patient 1, (b) Patient 2, and (c) Patient 3. The treatment plans from OOU solutions show increased TTP compared to the SOC.

The SOC and optimized treatment regimens for the three case-study patients are visualized in Figure 9, which shows the dose schedule over the six weeks of RT. At lower total dose levels, treatment plans tend to feature one large dose roughly halfway through the six week treatment course. At intermediate total doses, there is a secondary dose similar to the SOC in the second week of treatment and a large dose applied towards the end of the treatment timeframe. Finally, strategies featuring a higher total dose tend to feature one large dose in the second week of treatment and then another dose towards the end of the therapy. This second dose can be larger or lower than the one delivered the second week of treatment, but it tends to be larger than the corresponding SOC dose. Overall, Figure 9 shows that optimized plans tend to leave weeks without RT with only the chemotherapy being administered during that time. This can be viewed as an additional advantage with not requiring patients to come in every week.

We now analyze the performance of the predictive digital twins at the cohort level relative to the SOC and highlight the performance for the three patient groups defined in Section 3.1. To this end, we deploy the predictive digital twins for each of the 100 *in silico* patients to obtain patient-specific optimal RT treatment regimens. Figure 10 shows the boxplot of TTP α -superquantile values obtained for the patient-specific optimal treatment regimens and the SOC treatment. We denote the solution of the optimization problem in Eq. (12) with $D_{\text{max}} = X$ Gy as “OUU: X Gy” in Figure 10. We can see that the optimal treatment regimens lead to better TTP compared to the SOC over the cohort of 100 patients with $D_{\text{max}} \geq 60$ Gy. We further analyze the change in TTP α -superquantile given by $((\text{TTP } \alpha\text{-superquantile from “OUU: X Gy”}) - (\text{TTP } \alpha\text{-superquantile from SOC}))$ for each patient in the different patient groups. Figure 11 shows the boxplots for the change in TTP α -superquantile values for early and intermediate progressors. Figure 11a




Figure 8: Comparing TTP distribution using optimized treatment under uncertainty vs SOC treatment for the three case-study patients: (a) Patient 1, (b) Patient 2, and (c) Patient 3. Note that “end of simulation” refers to the maximum TTP value of 132 days stipulated by $t_{\text{finite}} = 152$ days. We can see the optimized treatment plans lead to separation in TTP distributions compared to the SOC and move towards higher TTP values indicating superior tumor control.

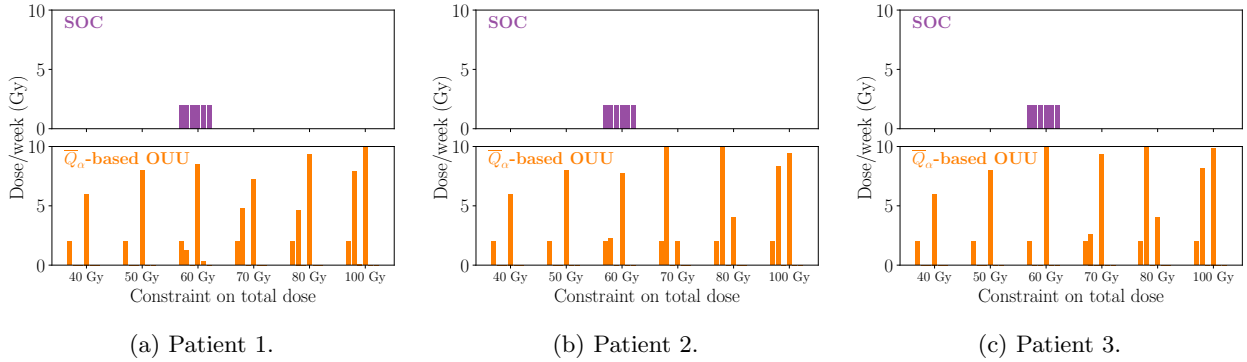


Figure 9: Optimized dose schedules reveal different strategies on a patient-specific level as well as for the solutions along the Pareto front for the three patients: (a) Patient 1, (b) Patient 2, and (c) Patient 3. Generally, optimized dose schedules result in larger doses towards the end of the treatment time frame.

captures the response to optimal treatment for early progressors. For optimal solutions “OUU: 60 Gy”, which use the same total dose as the SOC, we obtain a median change in TTP α -superquantile of +2.4 days with all 16 early progressors having better TTP α -superquantile values compared to the SOC. The

higher dose optimal treatment regimens provide other therapeutic options to the patient and the treating physician with up to +8.5 days median change in TTP α -superquantile. Figure 11b captures the response to optimal treatment for intermediate progressors. In this case, the optimal solution “OUU: 60 Gy” leads to a median change in TTP α -superquantile of +7 days with all 62 intermediate progressors having better TTP α -superquantile values compared to the SOC. The higher dose optimal treatment regimens provide additional therapeutic strategies with up to +21.3 days median change in TTP α -superquantile.




Figure 10: Comparison between optimal treatment regimens under uncertainty and SOC for TTP α -superquantile over the cohort of 100 patients. The solution of the optimization problem in Eq. (12) with $D_{\max} = X$ Gy is indicated as “OUU: X Gy”. Optimal treatment regimens lead to better TTP with $D_{\max} \geq 60$ Gy and similar TTP with $D_{\max} < 60$ Gy when compared to the SOC

The late progressors are not expected to have any reduction in TTP since almost all treatment options reach the TTP α -superquantile maximum value of 132 days. One can see that the definition of late progressors coincides with the finite time horizon considered. Instead, the late progressors can have similar tumor control with reduction in total required dose. Recall that to achieve this therapeutic objective of mitigating the toxicity effects with our digital twin formulation, we use the penalty parameter $\lambda = 0.001$ in Eq. (12) to drive the optimizer to lowest possible dose while obtaining the same level of tumor control, as described in Section 2.7. Figure 12 shows that this approach achieves a reduction in total radiation dose while maintaining a similar treatment efficacy as SOC for the entire cohort of 100 patients as well as the three patient groups. We can see that risk-aware optimal treatment plans can lead to a median of 10 Gy (16.7%) reduction in total dose compared to the SOC over all 100 patients, while keeping the TTP α -superquantile within ± 1 day of that obtained for the SOC. Analyzing the effect on the different patient groups provides additional insight on the possible amount of dose reduction depending on the underlying dynamics of tumor response. For the early and intermediate progressors, we can see a median reduction in dose of 10 Gy and 20 Gy, respectively. However, for the late progressors we can see that the median reduction in dose is 46.7 Gy. This comparatively higher decrease in dose is a direct consequence of slower tumor growth and lower number of tumor cells in the tumors of late progressors within the post-treatment timeframe considered in this study.




Figure 11: Change in TTP α -superquantile when using optimal treatment regimens as compared to SOC for (a) early progressors: 16 patients with TTP α -superquantile less than 1 month after end of RT, and (b) intermediate progressors: 62 patients with TTP α -superquantile between 1-3 months after end of RT. Median values are indicated in text above each box plot. The solution of the optimization problem in Eq. (12) with $D_{\max} = X$ Gy is indicated as “OUU: X Gy”. The median change indicates longer TTP compared to the SOC when optimal treatments are used.

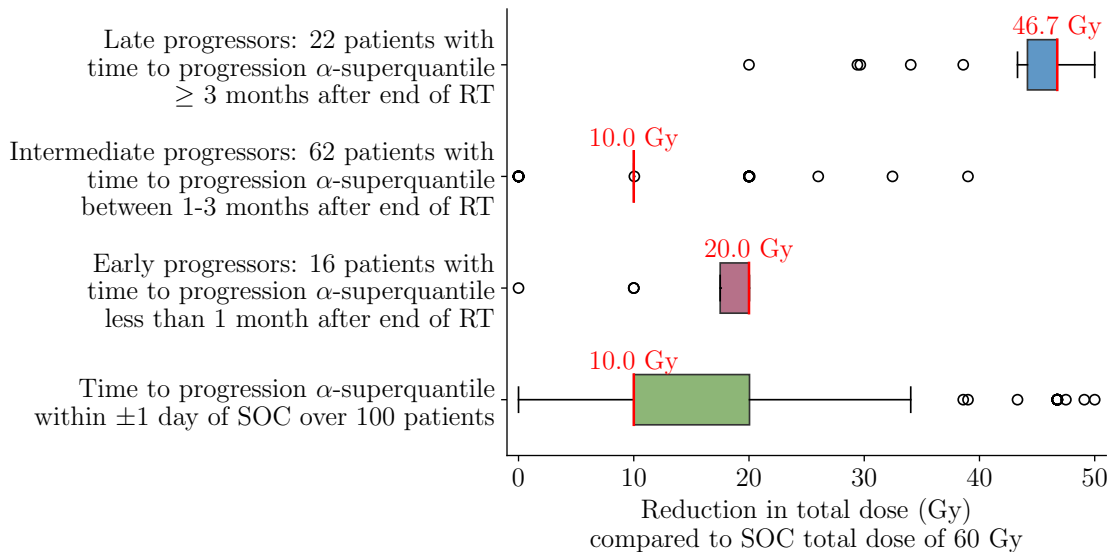


Figure 12: Reduction in total dose compared to SOC total dose of 60 Gy to achieve TTP α -superquantile within ± 1 day of SOC over the cohort of 100 patients and for the different patient groups of early, intermediate, and late progressors. Median values are indicated in red text above each box plot. Optimal treatment plans lead to 16.7% median dose reduction over the 100 patient cohort and 77.8% median dose reduction for the late progressors compared to the SOC.

3.4 Survival analysis of optimal treatments

We now show the results for the Kaplan-Meier survival analysis described in Section 2.8. The results on survival analysis indicate that making optimal decisions at the individual level using patient-specific predictive digital twins provide either the same survivability as SOC plans or improve it across the cohort of patients. The survival curves in Figure 13 show that all optimal therapeutic regimens generated with a total dose threshold greater than or equal to that of the SOC plan of 60 Gy have survival curves that are superior to the corresponding curve obtained with SOC plan. There is a clear gap between the 80 Gy and 100 Gy plans and the SOC with logrank p-values of 0.005 and 0.0002, respectively. For the optimized therapy with maximum allowable dose of 60 Gy, the survival curve is statistically superior to the SOC treatment option with p-value of 0.05. Figure 13 further shows that the 40 Gy plans show almost no difference in tumor control when compared to the SOC. With a p-value of $0.35 \gg 0.05$, we conclude that the optimized treatment plans with maximum allowable dose of 40 Gy is not significantly different than the SOC, which uses 60 Gy.

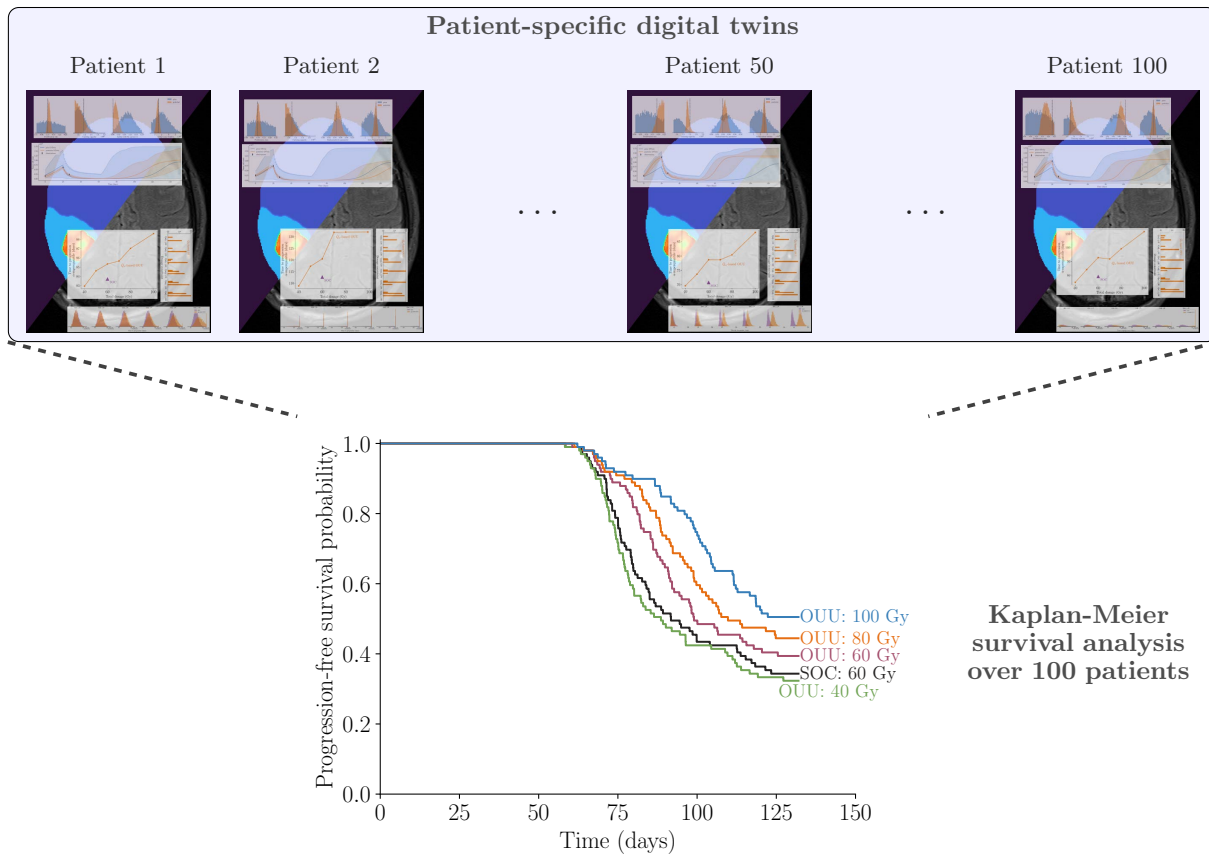


Figure 13: Illustrating patient-specific digital twins for cohort of 100 patients and analyzing the performance of proposed optimal treatment plans vs SOC through Kaplan-Meier survival analysis using the TTP α -superquantile values.

4 Discussion

Our digital twin methodology provides the computational and mathematical foundation to dynamically integrate patient data with any given model of tumor growth and response. The predictive digital twin can provide risk-aware personalized treatment plans and we demonstrate it through optimized RT treatment regimens for HGG. There is a rich literature of modeling glioma growth and response across different scales [2, 22]; however, there are limited approaches that integrate these models with patient-specific data to optimize or adapt therapy in a dynamic fashion while accounting for underlying uncertainty. For development and demonstration of this framework, we employed an ODE model describing the total tumor cell proliferation. Response to radiotherapy and chemotherapy were reduced to an instantaneous effect at the time of treatment. While the model’s biological detail is sufficient to demonstrate from end-to-end the key components of the predictive digital twin, a more biologically complex model may be needed in the clinical setting to account for inter-tumor heterogeneity in treatment delivery, proliferation rates, and resistance to therapy. Brain tumor growth and response to treatment has been described via a more complex reaction-diffusion model [2, 25, 70] that describes the spatial-temporal evolution of tumors cells throughout the brain due to invasion (i.e., the diffusion term) and proliferation (i.e., the reaction term). Our predictive digital twin can be readily adapted to other more complex models and cancer types where the prerequisite spatial-temporal or temporal data are available. We plan to extend the predictive digital twin to a more descriptive partial differential equation model that can extract the spatial information, but comes with a much higher computational cost challenge. These partial differential equation models could incorporate other components of the tumor microenvironment including the vasculature (and the blood brain barrier) which fundamentally influences the delivery and efficacy of chemotherapy and RT [24]. We have previously explored these components at the pre-clinical [27] and clinical levels [25] and it would be natural to apply our digital twin framework to these models. While this digital twin focused on the delivery of RT concurrently with chemotherapy, our predictive digital twin could be adapted to other treatment modalities (e.g., immunotherapy, tumor treatment fields, convection-enhanced delivery) provided the pre-requisite data and models are available. We will also explore the critical issue of when to get the patient in to gather more imaging data.

We stress that the overall motivation of this work was the development of a digital twin framework which could be applied to any biology-based model of tumor growth. However, clinical treatment should ultimately be based upon established clinical guidelines [21]. Indeed, one limitation of this study is the use of an *in silico* patient cohort whose growth and response parameters are sampled from literature values and governed by our model. We would anticipate that when applied to actual patient data, the precise gains between different RT regimen would differ to those presented in this work due to variability in the patient population and limitations of our current model. Therefore, a clinical trial would be necessary to establish the benefit of digital twin adapted RT regimens.

The optimized treatment regimens identified here tended towards a more hypofractionated (i.e., larger dose per day over a shorter period of time) or intermittent paradigm for treatment. While not widely used in the standard-of-care setting for initial treatment, hypofractionated radiotherapy is safe and tolerable [16] and employed in the recurrent setting and for patients with poor prognoses [71]. Compared to standard treatment schedules, hypofractionated radiotherapy has an increased cell kill over a shorter period of time, potential immunogenic effects, but at a potentially increased neuro-toxicity [58, 20]. In the context of first-line treatment following surgery for HGG, the benefits of hypofractionated radiotherapy are less clear but some studies have reported similar tumor control and a palliative benefit due to its condensed treatment schedule [14]. Other modeling approaches from [56, 38, 7] have also identified optimized therapies that do not

conform to the standard treatment paradigm. Recently, in the pre-clinical setting, [56] predicted improved outcomes with a hyperfractionated regimen (i.e., smaller doses more than once a day) compared to standard dosing and this was validated via experiments. At the clinical level, [7] explored intermittent radiation therapy (i.e., 6 Gy \times 1 day every 6 weeks) in comparison to hyperfractionated radiotherapy for recurring disease. [7] observed similar tumor control for both scenarios, but prolonged time to progression when additional intermittent doses of radiotherapy were delivered. In a future work, we plan to explore different clinical objectives and treatment controls, which could lead to other types of treatment regimens with further improvement in treatment efficacy.

5 Conclusions

We have developed a predictive digital twin methodology that enables end-to-end uncertainty quantification and optimization of personalized treatment regimens under uncertainty. The predictive digital twins employ a Bayesian perspective to account for uncertainty through the entirety of the clinical process to capture our knowledge (or lack thereof) of the individual dynamics. We illustrate the effectiveness of the predictive digital twin to optimize RT treatment plans for individual high-grade glioma patients. This subclass of tumors are aggressive and highly heterogeneous in nature, motivating the need for patient-specific modeling and treatment planning. We applied the predictive digital twin to a cohort of *in silico* patients with response and growth characteristics assigned from literature studies. We calibrated the digital twin with patient-specific data and used the calibrated digital twin to inform risk-aware clinical decision-making. We solve a multi-objective risk-based optimization under uncertainty problem to provide a suite of treatment plans balancing tumor control and toxicity. The risk-aware patient-specific optimization of the dose for the *in silico* study demonstrated that we were able to extend the progression-free survival for the standard total dose of 60 Gy relative to the standard-of-care. Additionally, we were able to identify potential reductions in the total dose delivered for similar tumor control response as the standard-of-care. These results represent a first step towards a practical digital twin for treatment optimization and future studies should expand this approach to data collected in the clinical setting.

Author Contributions

All authors contributed to conception and design of the study. AC, GP, MK, and KW contributed to the computational methodology. AC generated the results and plots. AC, GP, and DAH wrote the first draft of manuscript. All authors contributed to manuscript revision, read, and approved the submitted version.

Funding

A.C. and K.W. acknowledge support from DARPA grant number DE-AC05-76RL01830 under the Automating Scientific Knowledge Extraction and Modeling (ASKEM) program and Department of Energy (DOE) Grant DE-SC0021239. K.W. acknowledges support from AFOSR MURI grant FA9550-21-1-0084. G.P. acknowledges support from the U.S. Department of Energy Computational Science Graduate Fellowship under Award Number DE-SC0021110. G.L. acknowledges the European Union’s Horizon 2020 research

and innovation program under the Marie Skłodowska-Curie grant agreement No. 838786. T.E.Y. acknowledges support from the National Cancer Institute R01CA235800, U24CA226110, U01CA174706, and CPRIT RR160005. T.E.Y. is a CPRIT Scholar in Cancer Research. D.A.H. acknowledges support from CPRIT RP220225.

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. The views and opinions of authors expressed herein do not necessarily state or reflect those of the United States Government or any agency thereof.

References

- [1] H. AHMED AND L. DEVOTO, *The potential of a digital twin in surgery*, *Surgical Innovation*, 28 (2021), pp. 509–510.
- [2] J. C. L. ALFONSO, K. TALKENBERGER, M. SEIFERT, B. KLINK, A. HAWKINS-DAARUD, K. R. SWANSON, H. HATZIKIROU, AND A. DEUTSCH, *The biology and mathematical modelling of glioma invasion: a review*, *Journal of The Royal Society Interface*, 14 (2017), p. 20170490.
- [3] D. J. AUM, D. H. KIM, T. L. BEAUMONT, E. C. LEUTHARDT, G. P. DUNN, AND A. H. KIM, *Molecular and cellular heterogeneity: the hallmark of glioblastoma*, *Neurosurgical Focus*, 37 (2014), p. E11.
- [4] S. BENZEKRY, C. LAMONT, A. BEHESHTI, A. TRACZ, J. M. L. EBOS, L. HLATKY, AND P. HAHNFELDT, *Classical Mathematical Models for Description and Prediction of Experimental Tumor Growth*, *PLoS Comput Biol*, 10 (2014), p. e1003800.
- [5] G. BIROS, O. GHATTAS, M. HEINKENSCHLOSS, D. KEYES, B. MALLICK, L. TENORIO, B. VAN BLOEMEN WAANDERS, K. WILLCOX, Y. MARZOUK, AND L. BIEGLER, *Large-scale inverse problems and quantification of uncertainty*, John Wiley & Sons, 2011.
- [6] J. M. BLAND AND D. G. ALTMAN, *The logrank test*, *Bmj*, 328 (2004), p. 1073.
- [7] S. C. BRÜNINGK, J. PEACOCK, C. J. WHELAN, R. BRADY-NICHOLLS, M. Y. HSIANG-HSUAN, S. SAHEBJAM, AND H. ENDERLING, *Intermittent radiotherapy as alternative treatment for recurrent high grade glioma: A modeling study based on longitudinal tumor measurements*, *Sci. Rep.*, 11 (2021), p. 20219.
- [8] A. CHAUDHURI, B. KRAMER, M. NORTON, J. O. ROYSET, AND K. WILLCOX, *Certifiable risk-based engineering design optimization*, *AIAA Journal*, 60 (2022), p. 551–565.
- [9] A. C. COHEN, *Truncated and censored samples: theory and applications*, CRC press, 1991.

- [10] J. CORRAL-ACERO, F. MARGARA, M. MARCINIAK, C. RODERO, F. LONCARIC, Y. FENG, A. GILBERT, J. F. FERNANDES, H. A. BUKHARI, A. WAJDAN, M. V. MARTINEZ, M. S. SANTOS, M. SHAMOHAMMDI, H. LUO, P. WESTPHAL, P. LEESON, P. DIACHILLE, V. GUREV, M. MAYR, L. GERIS, P. PATHMANATHAN, T. MORRISON, R. CORNELUSSEN, F. PRINZEN, T. DELHAAS, A. DOLTRA, M. SITGES, E. J. VIGMOND, E. ZACUR, V. GRAU, B. RODRIGUEZ, E. W. REMME, S. NIEDERER, P. MORTIER, K. MCLEOD, M. POTSE, E. PUEYO, A. BUENO-OROVIO, AND P. LAMATA, *The "digital twin" to enable the vision of precision cardiology*, European Heart Journal, 41 (2020), pp. 4556–4564.
- [11] G. D'ARRIGO, D. LEONARDIS, S. ABD ELHAFEEZ, M. FUSARO, G. TRIPEPI, AND S. ROUMELIOTIS, *Methods to analyse time-to-event data: the kaplan-meier survival curve*, Oxidative Medicine and Cellular Longevity, 2021 (2021).
- [12] C. DAVIDSON-PILON, *lifelines: survival analysis in python*, Journal of Open Source Software, 4 (2019), p. 1317.
- [13] H. ENDERLING, J. C. L. ALFONSO, E. MOROS, J. J. CAUDELL, AND L. B. HARRISON, *Integrating Mathematical Modeling into the Roadmap for Personalized Adaptive Radiation Therapy*, Trends in Cancer, 5 (2019), pp. 467–474.
- [14] N. S. FLOYD, S. Y. WOO, B. S. TEH, C. PRADO, W.-Y. MAI, T. TRASK, P. L. GILDENBERG, P. HOLOYE, M. E. AUGSPURGER, L. CARPENTER, H. H. LU, J. CHIU, W. H. GRANT, AND E. BUTLER, *Hypofractionated intensity-modulated radiotherapy for primary glioblastoma multiforme*, International Journal of Radiation Oncology*Biophysics, 58 (2004), pp. 721–726.
- [15] M. GREENWOOD ET AL., *A report on the natural duration of cancer.*, A Report on the Natural Duration of Cancer., (1926).
- [16] P. H. GUTIN, F. M. IWAMOTO, K. BEAL, N. A. MOHILE, S. KARIMI, B. L. HOU, S. LYMBERIS, Y. YAMADA, J. CHANG, AND L. E. ABREY, *Safety and efficacy of bevacizumab with hypofractionated stereotactic irradiation for recurrent malignant gliomas*, International Journal of Radiation Oncology*Biophysics, 75 (2009), pp. 156–163.
- [17] Y. HAIMES, *On a bicriterion formulation of the problems of integrated system identification and system optimization*, IEEE transactions on systems, man, and cybernetics, (1971), pp. 296–297.
- [18] T. HERNANDEZ-BOUSSARD, P. MACKLIN, E. J. GREENSPAN, A. L. GRYSHUK, E. STAHLBERG, T. SYEDA-MAHMOOD, AND I. SHMULEVICH, *Digital twins for predictive oncology will be a paradigm shift for precision cancer care*, Nature medicine, 27 (2021), pp. 2065–2066.
- [19] R. P. HILL, R. G. BRISTOW, A. FYLES, M. KORITZINSKY, M. MILOSEVIC, AND B. G. WOUTERS, *Hypoxia and predicting radiation response*, Seminars in radiation oncology, 25 (2015), pp. 260–272.
- [20] M. HINGORANI, W. P. COLLEY, S. DIXIT, AND A. M. BEAVIS, *Hypofractionated radiotherapy for glioblastoma: strategy for poor-risk patients or hope for the future?*, The British Journal of Radiology, 85 (2012), pp. e770–e781. PMID: 22919020.
- [21] C. HORBINSKI, L. B. NABORS, J. PORTNOW, J. BAEHRING, A. BHATIA, O. BLOCH, S. BREM, N. BUTOWSKI, D. M. CANNON, S. CHAO, M. G. CHHEDA, A. J. FABIANO, P. FORSYTH, P. GIGILIO, J. HATTANGADI-GLUTH, M. HOLDHOFF, L. JUNCK, T. KALEY, R. MERRELL, M. M. MRUGALA,

- S. NAGPAL, L. A. NEDZI, K. NEVEL, P. L. NGHIEMPHU, I. PARNEY, T. R. PATEL, K. PETERS, V. K. PUDUVALLI, J. ROCKHILL, C. RUSTHOVEN, N. SHONKA, L. J. SWINNEN, S. WEISS, P. Y. WEN, N. E. WILLMARTH, M. A. BERGMAN, AND S. DARLOW, *Nccn guidelines® insights: Central nervous system cancers, version 2.2022: Featured updates to the nccn guidelines*, Journal of the National Comprehensive Cancer Network, 21 (2023), pp. 12 – 20.
- [22] D. A. HORMUTH, M. FARHAT, C. CHRISTENSON, B. CURL, C. CHAD QUARLES, C. CHUNG, AND T. E. YANKEELOV, *Opportunities for improving brain cancer treatment outcomes through imaging-based mathematical modeling of the delivery of radiotherapy and immunotherapy*, Advanced Drug Delivery Reviews, 187 (2022), p. 114367. Publisher: Elsevier B.V.
- [23] D. A. HORMUTH, A. M. JARRETT, G. LORENZO, E. A. LIMA, C. WU, C. CHUNG, D. PATT, AND T. E. YANKEELOV, *Math, magnets, and medicine: Enabling personalized oncology*, Expert review of precision medicine and drug development, 6 (2021), pp. 79–81.
- [24] D. A. HORMUTH, C. M. PHILLIPS, C. WU, E. A. B. F. LIMA, G. LORENZO, P. K. JHA, A. M. JARRETT, J. T. ODEN, AND T. E. YANKEELOV, *Biologically-based mathematical modeling of tumor vasculature and angiogenesis via time-resolved imaging data*, Cancers, 13 (2021).
- [25] D. A. HORMUTH II, K. A. A. FEGHALI, A. M. ELLIOTT, T. YANKEELOV, AND C. CHUNG, *Image-based personalization of computational models for predicting response of high-grade glioma to chemoradiation*, Scientific Reports, 11 (2021), pp. 1–14.
- [26] D. A. HORMUTH II, A. M. JARRETT, T. DAVIS, AND T. E. YANKEELOV, *Towards an Image-Informed Mathematical Model of In Vivo Response to Fractionated Radiation Therapy*, Cancers, (2021), pp. 1–23.
- [27] D. A. HORMUTH II, A. M. JARRETT, AND T. E. YANKEELOV, *Forecasting tumor and vasculature response dynamics to radiation therapy via image based mathematical modeling*, Radiation Oncology, 15 (2020), p. 4.
- [28] D. A. HORMUTH II, J. A. WEIS, S. L. BARNES, M. I. MIGA, E. C. RERICHA, V. QUARANTA, AND T. E. YANKEELOV, *Predicting in vivo glioma growth with the reaction diffusion equation constrained by quantitative magnetic resonance imaging data*, Physical Biology, 12 (2015), p. 46006.
- [29] C.-L. HWANG AND A. S. M. MASUD, *Multiple objective decision making—methods and applications: a state-of-the-art survey*, vol. 164, Springer Science & Business Media, 2012.
- [30] D. A. JAFFRAY, *Image-guided radiotherapy: from current concept to future perspectives.*, Nature reviews. Clinical oncology, 9 (2012), pp. 688–99.
- [31] A. M. JARRETT, E. A. B. F. LIMA, D. A. HORMUTH, M. T. MCKENNA, X. FENG, D. A. EKRUT, A. C. M. RESENDE, A. BROCK, AND T. E. YANKEELOV, *Mathematical models of tumor cell proliferation: A review of the literature*, Expert Review of Anticancer Therapy, 18 (2018), pp. 1271–1286. Publisher: Taylor & Francis.
- [32] E. L. KAPLAN AND P. MEIER, *Nonparametric estimation from incomplete observations*, Journal of the American statistical association, 53 (1958), pp. 457–481.
- [33] M. G. KAPTEYN, J. V. R. PRETORIUS, AND K. E. WILLCOX, *A probabilistic graphical model foundation for enabling predictive digital twins at scale*, Nature Computational Science, 1 (2021).

- [34] F.-M. KONG, R. K. TEN HAKEN, M. SCHIPPER, K. A. FREY, J. HAYMAN, M. GROSS, N. RAMNATH, K. A. HASSAN, M. MATUSZAK, T. RITTER, N. BI, W. WANG, M. ORRINGER, K. B. CEASE, T. S. LAWRENCE, AND G. P. KALEMKERIAN, *Effect of Midtreatment PET/CT-Adapted Radiation Therapy With Concurrent Chemotherapy in Patients With Locally Advanced Non-Small-Cell Lung Cancer: A Phase 2 Clinical Trial*, JAMA oncology, 3 (2017), pp. 1358–1365.
- [35] D. P. KOURI AND T. M. SUROWIEC, *Risk-averse pde-constrained optimization using the conditional value-at-risk*, SIAM Journal on Optimization, 26 (2016), pp. 365–396.
- [36] R. LAUBENBACHER, J. P. SLUKA, AND J. A. GLAZIER, *Using digital twins in viral infection*, Science, 371 (2021), pp. 1105–1106.
- [37] M. LE, H. DELINGETTE, J. KALPATHY-CRAMER, E. R. GERSTNER, T. BATCHELOR, J. UNKELBACH, AND N. AYACHE, *Personalized Radiotherapy Planning Based on a Computational Tumor Growth Model*, IEEE Transactions on Medical Imaging, PP (2016), p. 1.
- [38] K. LEDER, K. PITTER, Q. LAPLANT, D. HAMBARDZUMYAN, B. D. ROSS, T. A. CHAN, E. C. HOLLAND, AND F. MICHOR, *Mathematical Modeling of PDGF-Driven Glioblastoma Reveals Optimized Radiation Dosing Schedules*, Cell, 156 (2014), pp. 603–616.
- [39] E. LIMA, J. ODEN, B. WOHLMUTH, A. SHAHMORADI, D. HORMUTH, T. YANKEELOV, L. SCARABOSIO, AND T. HORGER, *Selection and validation of predictive models of radiation effects on tumor growth based on noninvasive imaging data*, Comput. Methods Appl. Mech. Eng., 327 (2017), pp. 277–305.
- [40] J. LIPKOVÁ, P. ANGELIKOPOULOS, S. WU, E. ALBERTS, B. WIESTLER, C. DIEHL, C. PREIBISCH, T. PYKA, S. E. COMBS, P. HADJIDOUKAS, K. VAN LEEMPUT, P. KOUMOUTSAKOS, J. LOWENGRUB, AND B. MENZE, *Personalized radiotherapy design for glioblastoma: Integrating mathematical tumor models, multimodal scans, and Bayesian inference*, IEEE Trans. Med. Imag., 38 (2019), pp. 1875–1884.
- [41] G. LORENZO, D. A. HORMUTH II, A. M. JARRETT, E. A. LIMA, S. SUBRAMANIAN, G. BIROS, J. T. ODEN, T. J. HUGHES, AND T. E. YANKEELOV, *Quantitative in vivo imaging to enable tumor forecasting and treatment optimization*, in Cancer, Complexity, Computation, I. Balaz and A. Adamatzky, eds., Springer, 2022, pp. 55–97.
- [42] S. MADHAVAN, R. A. BECKMAN, M. D. MCCOY, M. J. PISHVAIAN, J. R. BRODY, AND P. MACKLIN, *Envisioning the future of precision oncology trials*, Nature Cancer, 2 (2021), pp. 9–11.
- [43] G. P. MAZZARA, R. P. VELTHUIZEN, J. L. PEARLMAN, H. M. GREENBERG, AND H. WAGNER, *Brain tumor target volume determination for radiation treatment planning through automated mri segmentation*, International Journal of Radiation Oncology* Biology* Physics, 59 (2004), pp. 300–312.
- [44] S. J. MCMAHON, *The linear quadratic model: Usage, interpretation and challenges*, Physics in Medicine and Biology, 64 (2019).
- [45] S. A. NIEDERER, M. S. SACKS, M. GIROLAMI, AND K. WILLCOX, *Scaling digital twins from the artisanal to the industrial*, Nature Computational Science, 1 (2021), pp. 313–320.
- [46] J. NIJKAMP, F. J. POS, T. T. NUVER, R. DE JONG, P. REMEIJER, J.-J. SONKE, AND J. V. LEBESQUE, *Adaptive radiotherapy for prostate cancer using kilovoltage cone-beam computed tomography: First clinical results*, International Journal of Radiation Oncology* Biology* Physics, 70 (2008), pp. 75–82.

- [47] J. T. ODEN, I. BABUŠKA, AND D. FAGHIHI, *Predictive computational science: Computer predictions in the presence of uncertainty*, Encyclopedia of Computational Mechanics Second Edition, (2017), pp. 1–26.
- [48] 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), pp. 735–779.
- [49] A. OMURO AND L. DEANGELIS, *Glioblastoma and other malignant gliomas: A clinical review*, JAMA, 310 (2013), pp. 1842–1850.
- [50] M. PALDINO, K. HEDGES, K. RODRIGUES, AND D. BARBORIAK, *Repeatability of quantitative metrics derived from mr diffusion tractography in paediatric patients with epilepsy*, The British journal of radiology, 87 (2014), p. 20140095.
- [51] M. PEIRLINCK, F. COSTABAL, J. YAO, J. GUCCIONE, S. TRIPATHY, Y. WANG, D. OZTURK, P. SEGARS, T. MORRISON, S. LEVINE, AND E. KUHL, *Precision medicine in human heart modeling*, Biomechanics and modeling in mechanobiology, 20 (2021), p. 803–831.
- [52] J. POLESZCZUK, R. WALKER, E. G. MOROS, K. LATIFI, J. J. CAUDELL, AND H. ENDERLING, *Predicting Patient-Specific Radiotherapy Protocols Based on Mathematical Model Choice for Proliferation Saturation Index*, Bulletin of Mathematical Biology, 80 (2018), pp. 1195–1206.
- [53] M. J. POWELL, *A direct search optimization method that models the objective and constraint functions by linear interpolation*, in Advances in optimization and numerical analysis, Springer, 1994, pp. 51–67.
- [54] X. S. QI, C. J. SCHULTZ, AND X. A. LI, *An estimation of radiobiologic parameters from clinical outcomes for radiation treatment planning of brain tumor*, International Journal of Radiation Oncology*Biology*Physics, 64 (2006), p. 1570–1580.
- [55] B. W. RAAYMAKERS, J. LAGENDIJK, J. OVERWEG, J. KOK, A. RAALJMAKERS, E. KERKHOF, R. VAN DER PUT, I. MEIJSING, S. CRIJNS, F. BENEDESSO, M. VAN VULPEN, C. DE GRAAFF, J. ALLEN, AND K. BROWN, *Integrating a 1.5 t mri scanner with a 6 mv accelerator: proof of concept*, Physics in Medicine & Biology, 54 (2009), p. N229.
- [56] A. RANGLES, H.-G. WIRSCHING, J. A. DEAN, Y.-K. CHENG, S. EMERSON, S. S. PATTWELL, E. C. HOLLAND, AND F. MICHOR, *Computational modelling of perivascular-niche dynamics for the optimization of treatment schedules for glioblastoma*, Nature Biomedical Engineering, 5 (2021). Publisher: Springer US.
- [57] A. RASHEED, O. SAN, AND T. KVAMSDAL, *Digital twin: Values, challenges and enablers from a modeling perspective*, IEEE Access, 8 (2020), pp. 21980–22012.
- [58] E. REZNIK, A. W. SMITH, S. TAUBE, J. MANN, M. Z. YONDORF, B. PARASHAR, AND A. G. WERNICKE, *Radiation and Immunotherapy in High-grade Gliomas: Where Do We Stand?*, American Journal of Clinical Oncology: Cancer Clinical Trials, 41 (2018), pp. 197–212. ISBN: 0000000000000.
- [59] R. T. ROCKAFELLAR AND J. O. ROYSET, *Superquantiles and their applications to risk, random variables, and regression*, in Theory Driven by Influential Applications, Informs, 2013, pp. 151–167.
- [60] ———, *Engineering decisions under risk averseness*, ASCE-ASME Journal of Risk and Uncertainty in Engineering Systems, Part A: Civil Engineering, 1 (2015), p. 04015003.

- [61] R. T. ROCKAFELLAR, S. URYASEV, ET AL., *Optimization of conditional value-at-risk*, Journal of risk, 2 (2000), pp. 21–42.
- [62] R. ROCKNE, E. C. ALVORD, J. K. ROCKHILL, AND K. R. SWANSON, *A mathematical model for brain tumor response to radiation therapy.*, Journal of Mathematical Biology, 58 (2009), pp. 561–78.
- [63] E. SAAD AND A. KATZ, *Progression-free survival and time to progression as primary end points in advanced breast cancer: often used, sometimes loosely defined*, Annals of Oncology, 20 (2009), p. 460–464.
- [64] J. SALVATIER, T. V. WIECKI, AND C. FONNESBECK, *Probabilistic programming in python using PyMC3*, PeerJ Computer Science, 2 (2016), p. e55.
- [65] R. C. SMITH, *Uncertainty quantification: theory, implementation, and applications*, vol. 12, Siam, 2013.
- [66] J.-J. SONKE, M. AZNAR, AND C. RASCH, *Adaptive radiotherapy for anatomical changes*, Seminars in Radiation Oncology, 29 (2019), pp. 245–257. Adaptive Radiotherapy and Automation.
- [67] A. M. STUART, *Inverse problems: a bayesian perspective*, Acta numerica, 19 (2010), pp. 451–559.
- [68] R. STUPP, M. J. WARREN, MASONVAN DEN BENT, M. WELLER, B. FISHER, M. J. B. TAPHOORN, K. BELANGER, A. A. BRANDES, C. MAROSI, U. BOGDAHN, J. CURSCHMANN, R. C. JANZER, S. K. LUDWIN, T. GORLIA, A. ALLGEIER, D. LACOMBE, J. G. CAIRNCROSS, E. EISENHAUER, AND R. O. MIRIMANOFF, *Radiotherapy plus Concomitant and Adjuvant Temozolomide for Glioblastoma*, The New England Journal of Medicine, 352 (2005), pp. 987–996.
- [69] S. SUBRAMANIAN, K. SCHEUFELE, M. MEHL, AND G. BIROS, *Where did the tumor start? an inverse solver with sparse localization for tumor growth models*, Inverse problems, 36 (2020), p. 045006.
- [70] K. R. SWANSON, E. C. ALVORD, AND J. D. MURRAY, *A quantitative model for differential motility of gliomas in grey and white matter*, Cell Proliferation, 33 (2000), pp. 317–329.
- [71] J.-C. TRONE, A. VALLARD, S. SOTTON, M. BEN MRAD, O. JMOUR, N. MAGNÉ, B. POMMIER, S. LAPORTE, AND E. OLLIER, *Survival after hypofractionation in glioblastoma: a systematic review and meta-analysis*, Radiation Oncology, 15 (2020), pp. 1–10.
- [72] E. G. C. TROOST, D. THORWARTH, AND W. J. G. OYEN, *Imaging-Based Treatment Adaptation in Radiation Oncology.*, Journal of nuclear medicine : official publication, Society of Nuclear Medicine, 56 (2015), pp. 1922–1929.
- [73] J. UNKELBACH, B. MENZE, E. KONUKOGLU, F. DITTMANN, N. AYACHE, AND H. A. SHIH, *Radiotherapy planning for glioblastoma based on a tumor growth model: implications for spatial dose redistribution*, Physics in Medicine and Biology, 59 (2014), p. 747.
- [74] P. VIRTANEN, R. GOMMERS, T. E. OLIPHANT, M. HABERLAND, T. REDDY, D. COURNAPEAU, E. BUROVSKI, P. PETERSON, W. WECKESSER, J. BRIGHT, ET AL., *Scipy 1.0: fundamental algorithms for scientific computing in python*, Nature methods, 17 (2020), pp. 261–272.
- [75] D. J. WALES AND J. P. DOYE, *Global optimization by basin-hopping and the lowest energy structures of lennard-jones clusters containing up to 110 atoms*, The Journal of Physical Chemistry A, 101 (1997), pp. 5111–5116.

- [76] C. H. WANG, J. K. ROCKHILL, M. MRUGALA, D. L. PEACOCK, A. LAI, K. JUSENIUS, J. M. WARDLAW, T. CLOUGHESY, A. M. SPENCE, R. ROCKNE, J. ALVORD, ELLSWORTH C., AND K. R. SWANSON, *Prognostic significance of growth kinetics in newly diagnosed glioblastomas revealed by combining serial imaging with a novel biomathematical model*, *Cancer Research*, 69 (2009), p. 9133–9140.
- [77] R. T. WOODALL, D. A. HORMUTH II, C. WU, M. R. ABDELMALIK, W. T. PHILLIPS, A. BAO, T. J. HUGHES, A. J. BRENNER, AND T. E. YANKEELOV, *Patient specific, imaging-informed modeling of rhenium-186 nanoliposome delivery via convection-enhanced delivery in glioblastoma multiforme*, *Biomedical physics & engineering express*, 7 (2021), p. 045012.
- [78] 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, AND T. E. YANKEELOV, *Integrating mechanism-based modeling with biomedical imaging to build practical digital twins for clinical oncology*, *Biophysics Reviews*, 3 (2022), p. 021304.
- [79] M. U. ZAHID, A. S. MOHAMED, K. LATIFI, A. RISHI, L. B. HARRISON, C. D. FULLER, E. G. MOROS, J. J. CAUDELL, AND H. ENDERLING, *Proliferation Saturation Index to Characterize Response to Radiation Therapy and Evaluate Altered Fractionation in Head and Neck Cancer*, *Applied Radiation Oncology*, 10 (2021), pp. 32–39.