

IMAGE-DRIVEN BIOPHYSICAL TUMOR GROWTH MODEL CALIBRATION*

KLAUDIUS SCHEUFELE[†], SHASHANK SUBRAMANIAN[‡], ANDREAS MANG[§],
GEORGE BIROS[‡], AND MIRIAM MEHL[†]

Abstract. We present a novel formulation for the calibration of a biophysical tumor growth model from a single-time snapshot, multiparametric magnetic resonance imaging (MRI) scan of a glioblastoma patient. Tumor growth models are typically nonlinear parabolic partial differential equations (PDEs). Thus, we have to generate a second snapshot to be able to extract significant information from a single patient snapshot. We create this two-snapshot scenario as follows. We use an atlas (an average of several scans of healthy individuals) as a substitute for an earlier, pretumor, MRI scan of the patient. Then, using the patient scan and the atlas, we combine image-registration algorithms and parameter estimation algorithms to achieve a better estimate of the healthy patient scan and the tumor growth parameters that are consistent with the data. Our scheme is based on our recent work (Scheufele et al., *Comput. Methods Appl. Mech. Engrg.*, to appear), but we apply a different and novel scheme where the tumor growth simulation in contrast to the previous work is executed in the patient brain domain and not in the atlas domain yielding more meaningful patient-specific results. As a basis, we use a PDE-constrained optimization framework. We derive a modified Picard-iteration-type solution strategy in which we alternate between registration and tumor parameter estimation in a new way. In addition, we consider an ℓ_1 sparsity constraint on the initial condition for the tumor and integrate it with the new joint inversion scheme. We solve the subproblems with a reduced space, inexact Gauss–Newton–Krylov/quasi-Newton method. We present results using real brain data with synthetic tumor data that show that the new scheme reconstructs the tumor parameters in a more accurate and reliable way compared to our earlier scheme.

Key words. tumor progression inversion, biophysical model calibration, image registration, PDE-constrained optimization, Picard iteration

AMS subject classifications. 35K40, 49M15, 49M20, 65K10, 65N35, 65Y05, 92C50

DOI. 10.1137/19M1275280

1. Introduction. Glioblastoma multiforme (GBM) tumor is a terminal primary brain cancer—the most aggressive one. Biophysical models are increasingly used to help the analysis of GBM magnetic resonance imaging (MRI) scans for epidemiological studies and also for assisting clinical decision making [61, 69], survival estimation, diagnosis, and preoperative and treatment planning [20, 48, 54, 62]. The key step in integrating biophysical models with clinical information is to calibrate them with patient MRI scans. After calibration, we can either use the estimated parameters

*Submitted to the journal’s Computational Methods in Science and Engineering section July 16, 2019; accepted for publication (in revised form) December 3, 2019; published electronically May 6, 2020.

<https://doi.org/10.1137/19M1275280>

Funding: The work of the first author was supported by a DAAD scholarship. The work of the fifth author was supported by a Tinsley Oden Fellowship. This work was supported by AFOSR grant FA9550-17-1-0190; by Simons Foundation grant (Award 586055, AM); by NSF grant CCF-1337393; by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under Awards DE-SC0010518 and DE-SC0009286; by NIH grant 10042242; by DARPA grant W911NF-115-2-0121; and by the University of Stuttgart, Institute for Parallel and Distributed Systems.

[†]Institut für Parallel und Distributed Systems, Universität Stuttgart, Universitätsstraße 38, 70569, Stuttgart, Germany (kscheufele@austin.utexas.edu, miriam.mehl@ipvs.uni-stuttgart.de).

[‡]Oden Institute for Computational Engineering and Sciences, University of Austin, 201 E. 24th Street, Austin, TX 78712-1229 (shashank@ices.utexas.edu, biros@ices.utexas.edu).

[§]Department of Mathematics, University of Houston, 3551 Cullen Blvd., Houston, TX 77204-3008 (andreas@math.uh.edu).

as biomarkers, or we can evolve the calibrated PDE to estimate short-term tumor infiltration. Tumor growth models are typically nonlinear reaction-diffusion PDEs and their calibration is challenging. Not only do we need to estimate reaction and diffusion parameters, but also the tumor concentration, which is only partially (and implicitly) observed in the MRI scan. If we had an MRI scan of the patient without the tumor, we could solve an inverse problem for the initial condition, the reaction, and diffusivity parameters of the tumor. Unfortunately, a *healthy patient* MRI scan is rarely available. To resolve this conundrum, Davatzikos's group [48] pioneered the idea of using an MRI scan of another, *healthy*, individual as a proxy for the tumor-free patient scan. In practice, we use a standardized average brain of several individuals, also known as a *statistical atlas*. But it turns out that naively using such an atlas would result in erroneous results. The second key idea addresses this by simultaneously using image registration to deform the atlas towards the patient scan as far as possible. In this paper, we present a novel formulation of such a joint registration and inversion problem and a numerical solver scheme.

In the image registration problem, the goal is to estimate spatial point correspondences between a template image m_T (the atlas) and a reference image m_R (the patient scan). To solve the registration problem, we use an optimal control formulation, in which we seek a (stationary) velocity field $\mathbf{v}$ (parametrizing a deformation map $\mathbf{y}$), such that the transported template image intensities match the intensities in the reference image, i.e., $m_T \circ \mathbf{y} \approx m_R$. In the tumor inversion problem, we want to estimate tumor-growth parameters (such as tumor origin, infiltration, and proliferation rates) of our PDE model so that if we start growing the tumor at $t = 0$ we match the partial tumor observations at $t = 1$. (The astute reader is probably wondering how do we know the time horizon. We don't, but $t = 1$ is related to a nondimensional form of the tumor growth PDE.) We simultaneously solve for both $\mathbf{v}$ and the tumor parameters.

Let us try to explain the setting using a somewhat high-level notation. In our work we are not using the original MRI scan intensities, but assume that we have segmentation labels describing different brain tissue (white matter, gray matter, etc.) distributions for both atlas and patient images. Let $\mathbf{m}_A$ be the healthy atlas labels and $\mathbf{m}_D$ the patient labels (both vector functions in the unit cube in $\mathbb{R}^3$). Let $\mathcal{T}$ and $\mathcal{R}$, respectively, be abstract forward operators for a *tumor simulation* and *registration mapping* component. $\mathcal{T}(\mathbf{m}, \mathbf{p})$ takes a label image $\mathbf{m}$ without tumor and tumor model parameters $\mathbf{p}$ and creates a labeled image that now has tumor labels in addition to the original healthy tissue labels (e.g., gray matter, white matter, etc.). $\mathcal{R}(\mathbf{m}, \mathbf{v})$ takes a label image $\mathbf{m}$ and a velocity $\mathbf{v}$ that parametrizes the deformation and creates a deformed image.

Then, we can summarize the general idea of our previous *moving-patient* and the new *moving-atlas* schemes by defining two optimization problems:

$$(1.1) \quad \begin{aligned} \text{moving-patient (MP):} \quad & \min_{\mathbf{p}, \mathbf{v}} \|\mathcal{R}(\mathcal{T}(\mathbf{m}_A, \mathbf{p}), \mathbf{v}) - \mathbf{m}_D\|, \\ \text{moving-atlas (MA):} \quad & \min_{\mathbf{p}, \mathbf{v}} \|\mathcal{T}(\mathcal{R}(\mathbf{m}_A, \mathbf{v}), \mathbf{p}) - \mathbf{m}_D\|. \end{aligned}$$

Notice that the main difference is switching the order of the registration and the tumor operators; see Figure 1.1 for an illustration. This seemingly simple change has significant impact in the solution of this problem. In [55, 56], we used the MP scheme and only inverted for the tumor initial condition. Note that in MP, to evaluate the objective function, we first grow a tumor in the atlas and then deform the patient image to match the resulting tumor-bearing atlas. As long as the tumor-bearing atlas




FIG. 1.1. Schematic view of the MP and the (new) MA scheme. The MP coupling scheme performs the tumor simulation (operator $\mathcal{T}$) (and parameter estimation) in a healthy atlas brain and warps the patient anatomy and tumor labels towards the atlas space (operator $\mathcal{R}$) [55, 56]. For the new MA coupling scheme, we first compute an approximation of the actual healthy patient brain by warping the atlas anatomy labels towards the patient (operator $\mathcal{R}$) and perform the tumor simulation (operator $\mathcal{T}$) (and parameter estimation) in the resulting approximation of the healthy brain geometry. This is critical to obtain patient specific tumor parameter estimates (of, e.g., tumor origin, tumor cell migration rate, tumor cell proliferation rate) that may be of (future) clinical relevance in decision making and treatment. Brain images modified from [20].

is topologically similar to the tumor-bearing patient brain, image registration can always yield good tumor reconstruction, even with completely wrong tumor parameters. Thus this approach enables efficient nondiffeomorphic registration between an atlas and a tumor-bearing patient, but is limited in terms of meaningful tumor inversion.

In this paper, we propose the MA scheme. To evaluate the objective function given, we first deform the healthy atlas (using $\mathbf{v}$) so that it matches the patient. Then we grow a tumor in the transformed atlas. It turns out that the new formulation is more appropriate for the biophysical modeling since the tumor growth takes place in images that resemble the actual patient.

Contributions. The main contributions of this paper are as follows:

- (i) We present a new optimization problem formulation (MA) for tumor-growth model calibration based on patient individual single-snapshot data.
- (ii) The new formulation prevents fitting of patient input data to a possibly poor tumor reconstruction. The computed biophysical parameters “live” in the patient space, i.e., the tumor parameters are estimated assuming tumor growth in (an approximation of) the healthy patient brain (as opposed to the atlas brain for the former scheme). This renders the new scheme better for biophysical inversion.¹
- (iii) We derive a Picard-iteration-type solution scheme that alternates between the image registration and the inverse tumor-growth problem.²
- (iv) We enhance the tumor inversion component; the tumor solver used in [56] inverted for the tumor initial condition only, penalized with an ℓ_2 constraint. In this paper, we use an ℓ_1 constraint (which restricts the tumor to more plausi-

¹The presented results are to be seen as a proof of concept for the developed methodology. To enable predictive capabilities, we require a more complex tumor model.

²A modified objective function for the registration subcomponent allows us to fulfill the strongly coupled first order optimality conditions of the joint optimization problem formulation.

ble initial condition), and also invert for the diffusion parameter in the tumor growth PDE (modeling tumor infiltration).

- (v) We conduct numerical experiments to evaluate the new scheme using synthetically grown tumor data from real clinical brain imaging data. We demonstrate that the MA scheme yields better results in terms of the accuracy and robustness of tumor parameter reconstruction.³ For a synthetic example, we showcase a malfunction⁴ of the MP scheme, whereas the new MA scheme yields good results.

Limitations. The reaction-diffusion tumor-growth model we consider is purely phenomenological. Important biophysical, biomechanical, and biochemical effects and phenomena such as mitosis, apoptosis, chemotaxis, deformation of brain parenchyma (mass effect), and the modeling of edema, necrosis, and angiogenesis are neglected. It is, however, the most common model used for clinical analysis [9, 21, 23, 29, 60, 62]. The effect of tumor growth induced deformation of brain parenchyma is important and not accounted for in our current model. This is ongoing work. Note that the main work here lies in the enhancement of the tumor component. Integrating it into the coupled formulation and the Picard iteration is straightforward due to our modular approach. The inverse tumor-growth problem with inversion for all growth parameters is ill-posed. We have to introduce additional prior information. Currently, if we want to invert for the proliferation rate ρ and the diffusion rate, we run the solver for multiple values of ρ and compare the quality of reconstruction results. We do not account for uncertainties in input imaging data, employed algorithms and solution approaches, and model parameters (such as tumor initial condition and characteristic net cell migration into surrounding tissue). A Bayesian framework is subject to future work; we currently invert for the maximum a posteriori estimate of the parameters. Our approach does not directly process MRI data but requires prior segmentation of the imaging data. Note that reliably integrating models with MRI is an open problem, since one needs to estimate cell density from data or link model output to intensities. Last, we emphasize that within this work we present a proof-of-concept for a new, promising methodology; future work will include its application to real clinical data.

Related work. One important application of biophysical models is to enable non-diffeomorphic image registration. Classical image registration [46, 57] assumes that the input images (template and reference) are topologically equivalent. But registering a healthy atlas (no tumor) to a tumor-bearing patient image violates this assumption. We refer to this problem as the normal-to-abnormal registration. We refer to our previous work in [56] for an extensive review in nondiffeomorphic image registration.

Here, we are interested in fitting models to patient individual imaging data. Several groups have tackled this problem using derivative-free optimization approaches [8, 27, 29, 36, 43, 45, 67, 68], or address the parameter estimation problem within a derivative-free Bayesian framework [22, 31, 32, 33, 34, 44, 52]. There has been only limited work in the direction of PDE-constrained optimization for model-based image analysis [10, 16, 23, 27, 28, 35, 36, 43, 53]. Most of these approaches rely on longitudinal patient data. In particular, they require knowledge of the healthy patient brain before tumor occurrence. To allow for model-inversion based on single-snapshot data

³For all considered test cases the tumor-growth ground truth parameters are known, to allow for comparison of the respective reconstruction accuracy.

⁴The patient input data are fitted to a poor reconstruction of the tumor in the atlas (due to large anatomical differences); consequently, the target data are corrupted in the further course of the inversion.

only, deformable intersubject registration becomes necessary to artificially generate a second snapshot in time. Such a combination of registration and biophysical inversion has been previously targeted in [20, 23, 48, 71, 72, 73]. The work in [48, 71, 72, 73], however, uses a purely mechanical model for tumor progression and falls short in providing information about progression and infiltration of cancerous cells into surrounding healthy tissue. In [2, 20, 23, 30], the authors propose a framework for joint segmentation, registration, and tumor modeling, which is very similar to our approach. In our previous work [56], we significantly reduce the time-to-solution by employing second order derivative information in conjunction with highly scalable and efficient numerics, improved algorithms, and powerful preconditioners. We eliminate the need for manual tumor seeding, and improve intersubject registration performance.

Outline. The outline of the paper is as follows: In section 2 we present the mathematical formulation for the MA joint registration and biophysical inversion problem. We present the PDE-constrained optimization problem and derive the first order optimality conditions in section 2.1 and propose a Picard-iteration-type solution strategy, outlined in section 2.2. Details on continuation schemes, adaptivity, and convergence criteria are given in section 2.3. More details on the two main sub-components, diffeomorphic medical image registration and biophysical tumor-growth inversion are summarized in section 3. In section 4, we perform numerical experiments to analyze our scheme, and compare it to the MP scheme [56].

2. Formulation and Picard iteration for the MA coupled problem. To apply tumor inversion and image registration in a joint approach, we consider an optimal control formulation, which results in a PDE-constrained, nonlinear optimization problem. In this section, we present the new MA formulation along with an iterative fixed-point coupling scheme. This scheme allows us to solve the joint optimization problem based on two separate components: a tumor-growth inversion solver and a modified diffeomorphic image registration component. Before presenting the formulation and the iterative solver, we shortly introduce the notation we are going to use throughout the paper.

Notation. $m_X(\mathbf{x}, t) \in [0, 1]$, $(\mathbf{x}, t) \in \Omega \times [0, 1]$ with $X \in \{WM, GM, CSF\}$ represent probability maps of different brain tissue types, namely, white matter (WM), gray matter (GM), and cerebrospinal fluid (CSF), which includes ventricles. The probability maps for each tissue label are computed from (segmented) MRI data in a preprocessing step. In the following, we refer to $m_X(\mathbf{x}, t)$ as anatomy labels. $\Omega = [0, 2\pi]^3$ denotes the normalized spatial domain. We gather these probability maps in a space-time vector field

$$(2.1) \quad \mathbf{m}(\mathbf{x}, t) = (m_i(\mathbf{x}, t))_{i=WM, GM, CSF} \in \mathbf{R}^3.$$

The normalized tumor cell concentration $c: \Omega_B \times [0, 1] \rightarrow [0, 1]$, $(\mathbf{x}, t) \mapsto c(\mathbf{x}, t)$ (tumor map) is given as a fourth field and interpreted as a probability map for cancerous tissue. $\Omega_B \subset \Omega = [0, 2\pi]^3$ is embedded⁵ in the simulation domain Ω , and denotes the domain occupied by brain tissue; it is bounded by the skull. All anatomy and tumor labels evolve in space and time. To simplify notation, we introduce the space-time domain $U := \Omega \times (0, 1]$, and $\bar{U} := \Omega \times [0, 1]$. We omit the spatial dependency, and indicate temporal evolution by the pseudotime $t_R \in [0, 1]$ for the advection problem of

⁵We use the fictitious domain method [24, 63] and discretize the tumor equations in the extended cubic box $\Omega = [0, 2\pi]^3$. No-flux Neumann boundary conditions $\frac{\partial c}{\partial n} = 0$ on $\partial\Omega_B \times (0, 1)$, restricting tumor cell invasion beyond the boundary of the CSF edged skull of the brain, are (approximately) satisfied via a penalty approach [11].

TABLE 2.1

Notation for the MA joint tumor inversion and image registration formulation: $\mathbf{m}$ denotes the vector of anatomy labels (probability maps for WM, GM, and CSF) defining the brain geometry, and c denotes a probability map for the tumor. Subscripts A and P indicate variables in the atlas and patient spaces, respectively. Anatomy labels and tumor label with subscript D denote the patient input imaging data (after preprocessing). The fields evolve in time and space; anatomy labels $\mathbf{m}$ for brain tissue evolve along a pseudotime $t_R \in [0, 1]$ associated with registration, and tumor labels evolve along the (normalized) time $t_T \in [0, 1]$ of tumor growth. The integration of a simulated tumor map $c(1)$ into healthy brain tissue $\mathbf{m}$ is indicated by $\mathbf{m}'$, and modeled via the formula $\mathbf{m}' = \mathbf{m}(1 - c(1))$ (assuming that the probability of encountering WM, GM, CSF, or tumor at location $\mathbf{x}$ sum up to one). To judge the proximity of the predicted state of our mathematical model to the observed data $\mathbf{m}_D$, and c_D , we define the ℓ_2 -distance measures $\mathcal{D}_c$ and $\mathcal{D}_m$.

Healthy atlas brain (input)	$\mathbf{m}_A$	Atlas	
Patient brain with tumor (input; target data)	$\mathbf{m}_D$	$\mathbf{m}_A$	Patient
Patient tumor (input; target data)	c_D		$\mathbf{m}_D, c_D$
Healthy patient brain	$\mathbf{m}_P(1)$	$\mathbf{m}_P(0) \xrightarrow{\mathcal{R}} \mathbf{m}_P(1)$	$\mathbf{m}'_P(1)$
Tumor initial condition	$c(0)$	$c(0) \xrightarrow{\mathcal{T}} c(1)$	$c(1)$
Simulated tumor	$c(1)$		
Approximated patient brain	$\mathbf{m}'_P(1)$	$t_R = 0$	$t_R = 1$
ℓ_2 -misfit tumor	$\mathcal{D}_c[c_1, c_2]$	$t_T = 0$	$t_T = 1$
	$:= \frac{1}{2} \ c_1 - c_2\ _{L^2(\Omega)}^2$		
ℓ_2 -misfit anatomy labels	$\mathcal{D}_m[\mathbf{m}_1, \mathbf{m}_2]$		
	$:= \frac{1}{2} \ \mathbf{m}_1 - \mathbf{m}_2\ _{L^2(\Omega)^3}^2$		

the registration, and the normalized time $t_T \in [0, 1]$ for tumor growth; cf. Table 2.1. To judge the proximity of the predicted state of our mathematical model to the observed data $\mathbf{m}_D, c_D$, we define the ℓ_2 -distance measures $\mathcal{D}_c$ and $\mathcal{D}_m$ (see Table 2.1), which drive the optimization process.

2.1. MA coupled formulation. The MA coupled problem formulation is given by the minimization problem

$$\min_{\mathbf{v}, \mathbf{p}, k_f, w} \mathcal{D}_c[c(1), c_D] + \mathcal{D}_m[\mathbf{m}_P(1)(1 - c(1)), \mathbf{m}_D] + \mathcal{S}[\mathbf{p}, \mathbf{v}, w]$$

subject to

$$\begin{aligned} \mathcal{T}: \text{ (tumor forward op.)} \\ (2.2a) \quad & \partial_t c - \operatorname{div} \mathbf{k}(\mathbf{m}_P(1)) \nabla c - f(c, \rho(\mathbf{m}_P(1))) = 0 && \text{in } U, \\ (2.2b) \quad & c(0) = \Phi(\mathbf{m}_P(1)) \mathbf{p} && \text{in } \Omega, \\ \mathcal{R}: \text{ (registration advection op.)} \\ (2.2c) \quad & \partial_t \mathbf{m}_P + \nabla \mathbf{m}_P \cdot \mathbf{v} = \mathbf{0} && \text{in } \bar{U}, \\ (2.2d) \quad & \mathbf{m}_P(0) = \mathbf{m}_A && \text{in } \Omega, \\ (2.2e) \quad & \operatorname{div} \mathbf{v} = w && \text{in } U. \end{aligned}$$

That is, we seek parameters and space-time fields $(\mathbf{v}, \mathbf{p}, k_f, w)$, such that the predicted state of our model becomes similar to the observed data. The distance between data and predicted state is quantified by the two ℓ_2 -distance measures $\mathcal{D}_c[c(1), c_D]$ and $\mathcal{D}_m[\mathbf{m}_P(1)(1 - c(1)), \mathbf{m}_D]$. $\mathcal{D}_c[c(1), c_D]$ measures the discrepancy between the simulated tumor $c(1)$ computed from (2.2a)–(2.2b) in the *approximated patient space* and the input tumor data c_D . $\mathcal{D}_m[\mathbf{m}_P(1)(1 - c(1)), \mathbf{m}_D]$ quantifies the discrepancy

between the patient's anatomy labels and the *warped-to-patient* atlas anatomy labels in regions not affected by the tumor⁶ computed from (2.2c)–(2.2e). $\mathcal{S}[\mathbf{p}, \mathbf{v}, w] := \beta_{\mathbf{p}} \mathcal{S}_{\mathbf{p}}[\mathbf{p}] + \beta_{\mathbf{v}} \mathcal{S}_{\mathbf{v}}[\mathbf{v}] + \beta_w \mathcal{S}_w[w]$ is a regularization term.

To simulate *tumor growth* over time, we use a diffusion-reaction model with diffusion tensor $\mathbf{k}$ and a logistic reaction term $f(c, \rho) = \rho c(1 - c)$ with reaction rate ρ . For the tumor initial conditions, we use a parametrization $\Phi \mathbf{p}$, i.e., $c(0)$ lives in an $n_{\mathbf{p}}$ -dimensional space spanned by Gaussian basis functions ($\mathbf{p} \in \mathbf{R}^{n_{\mathbf{p}}}$, $\Phi \mathbf{p} := \sum_{i=1}^{n_{\mathbf{p}}} \Phi_i p_i$). Within CSF, the Gaussians are set to zero to prevent spurious diffusion of cancerous cells into CSF. We model the *diffusion* and *reaction* coefficient as

$$(2.3) \quad \mathbf{k}(\mathbf{m}) := k_f \sum_{i=1}^3 k_i m_i \mathbf{I} = \bar{\mathbf{k}} \mathbf{m} \mathbf{I} \quad \text{and} \quad \rho(\mathbf{m}) := \rho_f \sum_{i=1}^3 \rho_i m_i = \bar{\rho} \mathbf{m}$$

with $\bar{\mathbf{k}} = (k_1 k_f, k_2 k_f, k_3 k_f)^T$ and $\bar{\rho} = (\rho_1 \rho_f, \rho_2 \rho_f, \rho_3 \rho_f)^T$, i.e., both diffusion and reaction parameters vary in space depending on the value of the anatomy labels $\mathbf{m}$ describing the brain geometry. For simplicity, we only consider isotropic diffusion.⁷ To regularize the tumor growth initial condition parametrization, we use combinations of ℓ_1 - and ℓ_2 -regularization, i.e.,

$$(2.4) \quad \mathcal{S}_{\mathbf{p}}^1[\mathbf{p}] = \frac{1}{2} \|\mathbf{p}\|_{L^1(\Omega)}^2 \quad \text{and} \quad \mathcal{S}_{\mathbf{p}}^2[\mathbf{p}] = \frac{1}{2} \|\mathbf{p}\|_{\mathbf{W}(\Omega)}^2 := \sum_{i=1}^{p_n} w_i \cdot p_i^2,$$

where $\beta_{\mathbf{p}} > 0$ is the regularization parameter and $\|\cdot\|_{\mathbf{W}(\Omega)}$ is a weighted ℓ_2 -norm. For more details on how we switch between the two penalty terms in the solver algorithm, see section 2.3.

For *image registration*, we use an advection model with a stationary velocity field $\mathbf{v}$. We require $\mathbf{v}$ to be sufficiently smooth to ensure that the associated deformation map is a diffeomorphism [64, 65, 70]. These smoothness requirements are typically stipulated using a Sobolev norm for $\mathbf{v}$ [5]. We consider an H^1 regularization model for the velocity field

$$(2.5) \quad \mathcal{S}_{\mathbf{v}}[\mathbf{v}] = \frac{1}{2} \|\mathbf{v}\|_{H^1(\Omega)}^2 = \frac{1}{2} \int_{\Omega} \sum_{i=1}^3 |\nabla v^i(\mathbf{x})|^2 d\Omega$$

with an additional penalty on the divergence $w = \operatorname{div} \mathbf{v}$, which allows us to control volume change. We refer to [38, 42] for details. A similar model has been considered in [6]. We refer to [4, 5, 6, 7, 26, 64, 65, 70] for theoretical results for various regularity assumptions on $\mathbf{v}$ and the image data. We use periodic boundary conditions on $\partial\Omega$ for all involved physical fields of (2.2).

Optimality conditions. As a general optimization approach, we use the method of Lagrange multipliers. Taking variations yields the strong form of the *first order*

⁶The simulated pathology is embedded into the advected atlas anatomy labels when calculating the registration misfit. This can be seen as a “masking” of the tumor region, i.e., the registration does not alter the tumor concentration nor the discrepancy measure $\mathcal{D}_c$.

⁷Note that our approach can be generalized in a straightforward way to general diffusion tensors.

optimality conditions:

$$\begin{aligned}
 & \text{tum., state} \\
 (2.6a) \quad & \partial_t c - \operatorname{div} \mathbf{k}(\mathbf{m}_P(1)) \nabla c - f(c, \rho(\mathbf{m}_P(1))) = 0 \quad \text{in } U, \\
 (2.6b) \quad & c(0) - \Phi(\mathbf{m}_P(1)) \mathbf{p} = 0 \quad \text{in } \Omega, \\
 & \text{tum., adj.} \\
 (2.6c) \quad & -\partial_t \alpha - \operatorname{div} \mathbf{k}(\mathbf{m}_P(1)) \nabla \alpha + \partial_c f^*(c, \rho(\mathbf{m}_P(1))) \alpha = 0 \quad \text{in } \bar{U}, \\
 (2.6d) \quad & c_D - c(1) + (\mathbf{m}_P(1))^T (\mathbf{m}_P(1)(1 - c(1)) - \mathbf{m}_D) - \alpha(1) = 0 \quad \text{in } \Omega, \\
 & \text{reg., state} \\
 (2.6e) \quad & \partial_t \mathbf{m}_P + \nabla \mathbf{m}_P \mathbf{v} = \mathbf{0} \quad \text{in } \bar{U}, \\
 (2.6f) \quad & \mathbf{m}_P(0) - \mathbf{m}_A = \mathbf{0} \quad \text{in } \Omega, \\
 & \text{reg., adj.} \\
 (2.6g) \quad & -\partial_t \boldsymbol{\lambda} - \operatorname{div} (\boldsymbol{\lambda} \otimes \mathbf{v}) = \mathbf{0} \quad \text{in } \bar{U}, \\
 (2.6h) \quad & (c(1) - 1)(\mathbf{m}_P(1)(1 - c(1)) - \mathbf{m}_D) + \partial_{\mathbf{m}_P(1)} \Phi(\mathbf{p}) \alpha(0) \\
 & - \int_0^1 (\nabla c)^T \nabla \alpha \partial_{\mathbf{m}_P(1)} \mathbf{k}^*(\mathbf{m}_P(1)) + \partial_{\mathbf{m}_P(1)} f^*(c, \rho(\mathbf{m}_P(1))) \alpha \, dt - \boldsymbol{\lambda}(1) = \mathbf{0} \quad \text{in } \Omega, \\
 & \text{tum., inv.} \\
 (2.6i) \quad & \beta_{\mathbf{p}} \nabla_{\mathbf{p}} \mathcal{S}_{\mathbf{p}}[\mathbf{p}] - \Phi^T \alpha(0) = \mathbf{0} \quad \text{in } \Omega, \\
 (2.6j) \quad & \int_0^1 \int_{\Omega} \mathbf{m}_P(1) ((\nabla c)^T \nabla \alpha) \, dx \, dt = \mathbf{0} \quad \text{in } \Omega, \\
 & \text{reg., inv.} \\
 (2.6k) \quad & \beta_{\mathbf{v}} \nabla_{\mathbf{v}} \mathcal{S}_{\mathbf{v}}[\mathbf{v}] + \mathcal{K} \left[\int_0^1 (\nabla \mathbf{m}_P)^T \boldsymbol{\lambda} \, dt \right] = \mathbf{0} \quad \text{in } \Omega
 \end{aligned}$$

with adjoint variables α , $\boldsymbol{\lambda}$, ν associated with the state variables $c(1)$, $\mathbf{m}_P$, and $\mathbf{m}_P(1)$. We eliminate w from the optimization problem and use the resulting operator $\mathcal{K}$ in (2.6k). For $\operatorname{div} \mathbf{v} = w = 0$, $\mathcal{K}$ would be the Leray projection $\mathcal{K}(\mathbf{u}) := \mathbf{u} + \nabla \Delta^{-1} \operatorname{div} \mathbf{u}$; for a nonzero w , the projection operator becomes slightly more involved; we refer to [37, 38] for additional details. The gradients of the diffusion and reaction terms in (2.6h) with respect to $\mathbf{m}_P(1)$ can be derived from (2.3), and are given by

$$(2.7) \quad \partial_{\mathbf{m}_P(1)} \mathbf{k} = \bar{\mathbf{k}} = k_f(k_1, k_2, 0) \quad \text{and} \quad \partial_{\mathbf{m}_P(1)} \mathbf{f} = \bar{\rho} \mathbf{m}_P(1)(1 - 2c(1)).$$

Discussion. We summarize the advantages of the MA scheme over the previously used MP scheme [56]: (i) Image registration operates on the brain anatomy only such that the tumor solver is forced to actually produce a biophysically more meaningful⁸ tumor that is similar to the observed patient tumor. (ii) Tumor growth parameters, i.e., $\mathbf{p}$ and k_f , are estimated directly in an approximation of the patient brain such that their quality does not strongly deteriorate with differences between the atlas and the patient brain. This comes at a price, which we account for by employing suitable solver strategies presented in section 3: (a) The decomposition of the moving atlas formulation into tumor and registration components requires modifications of the registration formulation as shown in section 2.2. (b) The exploitation of the full

⁸We are assuming a biophysical tumor progression model which accurately reflects the physiological processes; the model used here does not fulfill this assumption.

potential of the MA scheme requires sparse localization of the tumor initial condition, encouraged by an ℓ_1 -regularization in the tumor component as outlined in section 3.1 (cf. also [55, 59]). The derivation of the optimality conditions of the proposed formulation is formal only. A rigorous analysis is beyond the scope of the present manuscript and remains subject to future work.

2.2. Picard iteration scheme. We solve the optimization problem by iterating over the registration velocity $\mathbf{v}$, the tumor growth initial guess parametrization $\mathbf{p}$ and (if switched on) the diffusion coefficient k_f . We use a superscript j to mark variables associated with the j th Picard iteration. Each Picard iteration is decomposed into two substeps:

- (1) Solve *inverse tumor problem*, i.e., given $\mathbf{v}^j$ solve for $(\mathbf{p}^{j+1}, k_f^{j+1})$ such that

$$(2.8) \quad (\mathbf{p}^{j+1}, k_f^{j+1}) = \arg \min_{\mathbf{p}, k_f} \underbrace{\mathcal{D}_c[c(1), c_D] + \mathcal{D}_m[\mathbf{m}_P(1)^j(1 - c(1)), \mathbf{m}_D] + \beta_p \mathcal{S}_p[\mathbf{p}]}_{=: \mathcal{J}_T}$$

subject to the tumor forward problem (2.2a) and (2.2b) for the healthy patient anatomy $\mathbf{m}_P(1) = \mathbf{m}_P(1)^j$ kept fixed (result of the previous Picard iteration or initial guess). This means, we fulfill the first order optimality conditions (2.6a)–(2.6d), (2.6i), and (2.6j) for $\mathbf{m}_P(1) = \mathbf{m}_P(1)^j$. Aside from new iterates for $\mathbf{p}$ and k_f , this step yields a new simulated tumor $c(1)^{j+1}$ and tumor adjoint α^{j+1} .

- (2) Solve *modified registration problem*, i.e., given $(\mathbf{p}^{j+1}, k_f^{j+1})$, solve for $\mathbf{v}^{j+1}$ such that

$$(2.9) \quad (\mathbf{v}^{j+1}, w^{j+1}) = \arg \min_{\mathbf{v}, w} \mathcal{D}_m[\mathbf{m}_P(1)(1 - c(1)^{j+1}), \mathbf{m}_D] + \int_{\Omega} \mathbf{q}^T \mathbf{m}_P(1) \, d\mathbf{x} + \beta_v \mathcal{S}_v[\mathbf{v}] + \beta_w \mathcal{S}_w[w]$$

subject to the image advection problem (2.2c)–(2.2e). $\mathbf{q} = (q_1, q_2, q_3)^T$ is defined as

$$\mathbf{q}(\mathbf{x}) = \int_0^1 \bar{k}((\nabla c(1)^{j+1})^T \nabla \alpha^{j+1}) + \bar{\rho} c(1)^{j+1} (1 - c(1)^{j+1}) \alpha^{j+1} \, dt.$$

- (3) Iterate until convergence.

The particular modification of (2.9) by introducing the term $\int_{\Omega} \mathbf{q}^T \mathbf{m}_P(1) \, d\mathbf{x}$ ensures that the registration problem reproduces (2.6e)–(2.6h) and (2.6k) as first order optimality conditions. Note that we freeze the variables $c(1)$, and α from the previous iteration, i.e., neglect the indirect impact of the registration velocity $\mathbf{v}$ on $c(1)$ and α via changes in $\mathbf{m}_P(1)$. Besides new iterates for $\mathbf{v}$, the registration step also provides a new approximation $\mathbf{m}_P(1)^{j+1}$ of the healthy patient brain. The converged solution of the respective Picard or fixed-point iteration fulfills all first order optimality conditions of (2.2) with the exception of (2.6h). For the latter, the term $\partial_{\mathbf{m}_P(1)} \Phi(\mathbf{p}) \alpha(0)$ is neglected and not enforced explicitly. We found experimentally that our scheme works well and produces satisfying results despite the simplification.

2.3. Adaptivity and convergence criteria. In our Picard scheme, we use several strategies to control convergence and to adapt the tumor and image registration regularization over the Picard iterations.

For the tumor solver, the regularization term $\beta_p \mathcal{S}_p[\mathbf{p}]$ is used to attain better conditioning, and also to encourage sparse localization of the initial condition. Thus, we use $\mathcal{S}_p[\mathbf{p}] = 1/2 \|\mathbf{p}\|_{L^1(\Omega)}^2$ to compute a sparse solution of the initial condition and

switch to a weighted ℓ_2 -regularization $\mathcal{S}_{\mathbf{p}}[\mathbf{p}] = \|\mathbf{p}\|_{\mathbf{W}(\Omega)}^2 = 1/2 \sum_{i=1}^{p_n} w_i \cdot p_i^2$ in a second phase of the tumor inversion solver, where the weights w_i are calculated according to

$$w_i = \begin{cases} w^{\text{large}} & \text{if } |p_i^{\ell_1}| < \tau_S \|\mathbf{p}\|_{L^\infty(\Omega)}, \\ w^{\text{small}} & \text{otherwise.} \end{cases}$$

Here, τ_S is a user defined tolerance. The ℓ_2 -phase uses $\mathbf{p}^{\ell_1}$ achieved in the ℓ_1 -phase as initial approximation. In the ℓ_1 -phase we invert for $\mathbf{p}$ only and keep the diffusivity k_f (estimated in the ℓ_2 -phase) fixed;⁹ initially, we start with $k_f = 0$. For more details on numerics of the tumor inversion solver, we refer to [56].

For the registration solver, convergence and robustness are improved by regularization. We perform a parameter continuation. That is, we start with a large value β_v^{init} and subsequently reduce β_v by a factor of ten after each Picard iteration until we reach the prescribed value β_v^{final} . Image registration calculates solutions and possibly increases β_v internally based on a check of lower admissible bounds for the determinant of the deformation gradient. The deviation of the determinant of the deformation gradient from one indicates local volume changes (expansion for a value larger than one and contraction for a value smaller than one; see [38, 56] for more details).

As a stopping criterion for the Picard iterations, we consider a mixed criterion that is fulfilled either if we reach β_v^{final} or if the image registration fails due to violation of the user defined lower bound on the determinant of the deformation gradient. We always execute two Picard iterations with the final β_v^{final} .

For both tumor inversion and image registration, we use fixed tolerances for the norm of the reduced gradients as given in (2.6i), (2.6j), and (2.6k), respectively. Note that we use norms relative to the gradient norm at the beginning of the first Picard iteration for all gradients. We combine these gradient tolerances with a prescribed maximal number of Newton iterations for the reduced space KKT systems (see also [56]).

3. Numerical methods for tumor inversion and image registration. In this section, we shortly describe the numerical and algorithmic features of the tumor and the image registration solvers. We give only a brief summary along with references to previous publications since these methods have not been changed from our previous work.

3.1. Regularization and nonlinear solver components for tumor inversion. The tumor inversion problem is solved based on an optimize-then-discretize approach. To solve the respective system of optimality conditions, we use different solvers for ℓ_1 or weighted ℓ_2 -regularization, respectively. For ℓ_1 -regularization, the general iterative shrinkage thresholding (GIST) [19] is applied. In every iteration, it determines a solution $\mathbf{p}^{k+1}$ close to the result $\mathbf{p}_{sd}^k$ of a gradient descent step in the direction $\nabla \mathcal{D}_c(c(1), c_D)$ with step length α^k :

$$(3.1a) \quad \mathbf{p}^{k+1} = \arg \min_{\mathbf{p}} \frac{1}{2} \|\mathbf{p} - \mathbf{p}_{sd}^k\|_{L^2(\Omega)}^2 + \alpha^k \frac{\beta_{\mathbf{p}}}{2} \|\mathbf{p}\|_{L^1(\Omega)}, \text{ and}$$

$$(3.1b) \quad \mathbf{p}_{sd}^k = \mathbf{p}^k - \alpha^k \nabla \mathcal{D}_c(c(1), c_D).$$

⁹We found experimentally, that diffusion inversion in the ℓ_1 -phase has vanishing impact on the proposed sparsity of the solution, and the ℓ_2 estimation for k_f is better.

To determine the step length α^k , we use an Armijo line search. For ℓ_1 -regularization, we find the analytical solution

$$(3.1c) \quad \mathbf{p}^{k+1} = \text{sign}(\mathbf{p}_{sd}^k) \max(0, |\mathbf{p}_{sd}^k| - \alpha^k \beta_{\mathbf{p}}).$$

The stopping criteria are given by

$$(3.2a) \quad |\mathcal{J}_T(\mathbf{p}^{k+1}) - \mathcal{J}_T(\mathbf{p}^k)| < \tau_J(1 + \mathcal{J}_T(\mathbf{p}^0)),$$

$$(3.2b) \quad \|\mathbf{p}^{k+1} - \mathbf{p}^k\|_{L^\infty(\Omega)} < \sqrt{\tau_J}(1 + \|\mathbf{p}^{k+1}\|_{L^\infty(\Omega)}),$$

where $\mathcal{J}_T$ is the objective function of the tumor inversion problem (2.8), and $\tau_J > 0$ is a user defined tolerance. We perform parameter continuation in $\beta_{\mathbf{p}}$ to estimate its value. We start with a large $\beta_{\mathbf{p}}$, bounded by $\|\nabla \mathcal{D}_c(c(1), c_D)\|_{L^\infty(\Omega)}$, and perform a binary search in subsequent GIST iterations based on the Hoyer sparsity measure [25], $\mathcal{H}_s(\mathbf{p}) \in [0, 1]$:

$$(3.3) \quad \mathcal{H}_s(\mathbf{p}) = \frac{n_{\mathbf{p}} - \|\mathbf{p}\|_{L^1(\Omega)} / \|\mathbf{p}\|_{L^\infty(\Omega)}}{n_{\mathbf{p}} - 1},$$

where $n_{\mathbf{p}}$ is the dimensionality of $\mathbf{p}$. Solutions with larger values of $\mathcal{H}_s$ feature higher sparsity, i.e., have fewer nonzero entries. We use a tolerance of 0.95 to identify sparse solutions. If the sparsity is large enough (larger than the tolerance), we search for a smaller regularization parameter and, alternatively, increase it if the sparsity is insufficient. GIST features slow first order convergence, which is, however, acceptable for our purpose as only a rough estimate for $c(0)$ is to be calculated with ℓ_1 -regularization.

In the ℓ_2 -phase, we employ a reduced space method. This means that we assume the state and adjoint equations to be fulfilled exactly; they can be eliminated from the KKT system. For the tumor inversion solver, the reduced gradients are calculated based on the current iterates for $\mathbf{p}$ and k_f as follows: (i) Solve the state equations (2.6a)–(2.6b); (ii) use the result to solve the adjoint equations (2.6c)–(2.6d); (iii) insert the result in (2.6i)–(2.6j) to calculate the reduced gradient with respect to $\mathbf{p}$ and k_f .

As a solver for the reduced KKT system, we use a quasi-Newton approach (LBFGS, [51, p. 135ff]) yielding a matrix-free approximation of the matrix-vector product of the gradient with the respective inverse Hessian and, thus, making an inner linear solver obsolete. We use the Moré–Thuente line search [49] for globalization.¹⁰ The ℓ_2 -phase is terminated when the relative change of the norm of the reduced gradient is below a user defined threshold $\text{opttol}_T > 0$. The reference gradient is the gradient obtained for the zero initial guess, $\mathbf{p}^0$.

3.2. Nonlinear solver components for image registration. We use an optimize-then-discretize approach for the image registration solver [42]. We employ a reduced space method: The current iterate for the velocity $\mathbf{v}$ is used to (i) solve the state equations (2.6e)–(2.6f) (forward advection), (ii) use the result to solve the adjoint equations (2.6g)–(2.6h), (iii) insert the result in (2.6k) to calculate the reduced gradient.

We use a Newton–Krylov approach with an inexact Gauss–Newton linearization and a matrix-free PCG method as inner linear solver (see [17, 37, 38, 39, 40] for more details) to compute the search direction.

For the convergence of the optimizer, we use a combination of the relative change of (i) the norm of the gradient in (2.6k), (ii) the objective in (2.9), and (iii) the

¹⁰Note that we do not use line search in the outer Picard iterations.

control variable $\mathbf{v}$, all controlled by a single parameter, $\text{opttol}_R > 0$, as a stopping criterion. More details on the stopping conditions can be found in [37, 47]; see [18, 305 ff.] for a discussion. For the inner PCG solver, we perform inexact solves [12, 13] with a tolerance that is proportional to the norm of the reduced gradient of our problem. This prevents oversolving of the Hessian system if far from the optimum (large gradient norm). See also [51, p. 165 ff.] for details. We use an Armijo line search for globalization.

For both tumor inversion and registration solver, we specify a maximum number ($\text{maxit}_{N,T}/\text{maxit}_{N,R}$) of Newton iterations, a maximum number ($\text{maxit}_{K,T}/\text{maxit}_{K,R}$) of Krylov iterations, and a lower bound of $1\text{E}-6$ for the absolute norm of the gradient as a safeguard against a prohibitively high number of iterations.

3.3. Discretization in space and time for the state and adjoint equations. For all state and adjoint equations in (2.6), we use the following numerical ingredients that have been published in more detail in [1, 16, 17, 37, 38, 39, 40]: (i) spectral elements for spatial discretization combined with a regular mesh; (ii) three-dimensional (3D) Fourier transforms to compute spatial derivatives; (iii) an unconditionally stable semi-Lagrangian time-stepping scheme to avoid stability issues and small time steps for image advection; (iv) an approximation of boundary conditions for the tumor equations at the brain surface based on a penalty approach combined with periodic boundary conditions at the boundary of the overall domain $\bar{\Omega} = [0, 2\pi]^3$; (v) an unconditionally stable, second order Strang-splitting approach [24, 58] with analytical solution of the reaction terms and an implicit Crank–Nicholson method for the diffusion terms in the tumor equations.

4. Numerical results. The experiments serve as a proof of concept for the newly developed joint registration and biophysical inversion strategy (MA) and demonstrate that the previously introduced MP scheme [56] can fail dramatically in terms of both tumor reconstruction and inversion for initial conditions and tumor growth parameters. We quantitatively and qualitatively compare the new MA strategy to the MP with an emphasis on model inversion and parameter estimation and address the scheme's eligibility to recover true model parameters. This is an important step towards reconstruction of meaningful biophysical characteristics. We consider the following classes of test cases:

SYN: synthetic brain. This first academic test case shows that the MP scheme can dramatically fail in terms of biophysical inversion. More precisely, we demonstrate that the MP scheme yields a largely wrong reconstruction of the tumor, the tumor initial condition, and the diffusion coefficient. We use two simple elliptical anatomies consisting of WM and GM only for atlas and healthy patient brain, respectively, and grow a tumor with known diffusion and reaction parameters in the patient anatomy. The resulting pathologic patient brain serves as target input data for our joint inversion approach. In order to achieve significant results, we create a case with large intersubject differences between atlas and patient anatomy.

STRV: synthetic tumor, real velocity. This second class of test cases is based on real brain geometries and serves as a means to examine the potential of the MA formulation for biophysical inversion in realistic geometries, i.e., to examine the quality of reconstruction for the known ground truth initial conditions, k_f and healthy patient geometry. We use two different real healthy brain geometries and grow a synthetic tumor in one of them, which we use as the patient brain. The second healthy brain serves as a normal brain template (atlas brain). In contrast to real patient tumors, in this case we know the ground truth for tumor initial conditions, diffusion coefficient

k_f , and healthy patient geometry. Note that we do *not* know the ground truth for the registration velocity $\mathbf{v}$.

In our test cases, we compare three different solution strategies:

I. *Tumor stand-alone (T)*. For the tumor stand-alone solution, the inversion is carried out assuming a statistical atlas brain geometry as the healthy patient brain. No intersubject registration is applied. The error is a function of the intersubject variability between atlas and patient brain.

II. *The joint registration and biophysical inversion MP scheme*. The MP scheme presented in [56] exhibits some inherent shortcomings with respect to the suitability for biophysical inversion and model calibration. The registration may fit the input data to a poor tumor reconstruction. We use the method described in [56] enhanced with a sparsity constraint for the tumor inversion solver section 3.1.

III. *The joint registration and biophysical inversion MA scheme*. The MA scheme presented in this work is designed to remedy some of the shortcomings of the MP scheme.

In the following, we shortly describe data and parameters for the experiments in section 4.1 as well as the hardware used and the general setup in section 4.2. We introduce performance measures used in the evaluation of our numerical results in section 4.3. In sections 4.4 and 4.5 we show results for the two classes of test cases.

4.1. Data and parameters. We use a common data basis for brain geometries as well as some fixed parameter settings for all test cases as described below.

Brain imaging data. For the STRV cases, we use normal brain MRI data obtained at the Perelman School of Medicine at the University of Pennsylvania. The 3D imaging data have an image resolution of 256^3 voxels. We use binary segmentations of MRI scans for WM, GM, and CSF. In a preprocessing step, these labels are smoothed and rescaled to probability maps, i.e., values between zero and one. To ensure *partition of unity* across all anatomy labels for each $\mathbf{x}$ in Ω , we introduce a fourth probability map for background based on

$$\forall \mathbf{x} \in \Omega : c^{(\cdot,1)}(\mathbf{x}) + \sum_{i=1}^4 m_i^{(\cdot,1)}(\mathbf{x}) = 1.$$

Common model and numerical parameters. The tumor model and solver parameters are summarized in Table 4.1. The solver parameters are based on experiments. We briefly discuss some of the more involved choices: For the parametrization of the tumor initial condition, we select a regular grid of n_p Gaussian basis functions with a variable standard deviation $\sigma \in \{2\pi/30, 2\pi/64\}$ and a grid spacing of $\delta = 1.5\sigma$ around the center of mass of the tumor. The number n_p is chosen a priori on a case-by-case basis, such that the pathological domain is covered sufficiently.

For all experiments, we consider tumor growth only in WM ($k_1 = 1; \rho_1 = 1$) and, consequently, set the characteristic diffusivity and net cell proliferation in GM and CSF to zero, i.e., $k_2 = k_3 = 0$ and $\rho_2 = \rho_3 = 0$. To grow the synthetic tumor, we choose $\rho_f^* = 15$ as cell proliferation rate (reaction coefficient), and $k_f^* = 1\text{E}-1$ as rate of cell migration into surrounding tissue (diffusion coefficient), respectively. For the tumor evolution simulation, we choose a time step of $\Delta t = 0.01$ and various time horizons $T \in \{0.16, 0.32, 0.44\}$. The regularization parameter β_p for the ℓ_2 -regularized inverse tumor problem has been determined experimentally from an L-curve study similar to [16], using $n_p = 125$ Gaussian basis functions and an image resolution of $n = 128^3$. We fix the relation $\sigma/\delta = 1.5$ for the spacing of the Gaussian

TABLE 4.1

Summary of common parameters used in all test cases. We report values of the following parameters: n denotes the image resolution with n_i , $i = 1, 2, 3$; n_p is the number of Gaussians for the parametrization of the tumor initial condition, σ is the standard deviation of the associated Gaussian basis functions (and varies throughout test cases), δ denotes the spacing in-between centers of adjacent Gaussians; k_1 , k_2 , and k_3 are the characteristic diffusion parameters for WM, GM, and CSF, k_f the overall scaling parameter for the isotropic part of the diffusion coefficient for net migration of cancerous cells into surrounding tissue; ρ_1 , ρ_2 , and ρ_3 are the characteristic reaction factors for WM, GM, and CSF, ρ_f is the overall reaction scaling factor; opttol_R , opttol_T are the convergence tolerances for registration and tumor inversion; $\text{maxit}_i = (\text{maxit}_{i,N}, \text{maxit}_{i,K})$ denotes the maximum number of Newton iterations and Krylov iterations (for the KKT system) for the tumor inversion ($i = T$) and registration ($i = R$), respectively; β_p is the regularization parameter for the tumor inversion; β_v^{init} and β_v^{final} are the initial and final values for the β -continuation scheme, applied in image registration; ε_∇ is the bound on the variation of the deformation gradient $\det(\nabla \mathbf{y})$ used in the continuation scheme.

Description	Parameter(s)	Value(s)
Image resolution	$n = n_1 \cdot n_2 \cdot n_3$	$n_i = 128$
Tumor initial condition parametrization	n_p, σ, δ	$\{125, 343\}, \{\frac{2\pi}{30}, \frac{2\pi}{64}\}, 1.5$
Tumor growth characteristic cell diffusivity	$\bar{k} = k_f^*(k_1, k_2, k_3)^T$	$0.1 \cdot (1, 0, 0)^T$
Tumor growth characteristic net cell proliferation	$\bar{\rho} = \rho_f^*(\rho_1, \rho_2, \rho_3)^T$	$15 \cdot (1, 0, 0)^T$
Optimizer tolerances	$\text{opttol}_R, \text{opttol}_T$	$1\text{E}-3, 1\text{E}-3$
Maximum Newton / Krylov iterations, tumor inversion	$\text{maxit}_T = (\text{maxit}_{T,N}, \text{maxit}_{T,K})$	$(50, -)$
Maximum Newton / Krylov iterations, registration	$\text{maxit}_R = (\text{maxit}_{R,N}, \text{maxit}_{R,K})$	$(50, 80)$
Tumor inversion regularization parameter	β_p	$1\text{E}-4$
Registration regularization parameter continuation	$\beta_v^{\text{init}}, \beta_v^{\text{final}}$	$1, 1\text{E}-4$
Bound on local volume change (variation $\det(\nabla \mathbf{y})$)	ε_∇	$1\text{E}-2$

basis functions.¹¹ The lower bound β_v^{final} on the regularization parameter β_v has been determined based on extensive numerical experiments for different synthetic and real brain data sets [38, 39]. By experimental analysis we found that reducing the regularization weight by a factor of 10 in every Picard iteration is sufficient.¹² For the optimizer, we use tolerances of $\text{opttol}_R = \text{opttol}_T = 1\text{E}-3$ for registration and tumor inversion, respectively.¹³ Further reduction of the gradient did not improve the final reconstruction quality, since the overall error is bound by an $\mathcal{O}(1\text{E} - 3)$ error introduced by the solution of the hyperbolic advection problem (cf. [56]). To prevent prohibitively large runtimes, we define upper bounds on the number of Newton (and Krylov) iterations.

4.2. Hardware and setup. All numerical experiments were executed on the supercomputer **HazelHen** at the High Performance Computing Center HLRS in Stuttgart (www.hlrs.de), a Cray XC40 system with a peak performance of 7.42 petaflops comprising 7,712 nodes with Xeon E5-2680 v3 processors and 24 cores on two sockets

¹¹This leads to an invariant condition number of $\Phi^T \Phi$ for the corresponding interpolation problem for varying standard deviations σ , i.e., width of the Gaussians.

¹²Using a smaller reduction or more Picard iterations per fixed regularization weight did not improve the overall result.

¹³Note that, in combination with the applied stopping conditions, this results in a required relative gradient of $1\text{E}-3$ for the tumor inversion, but a gradient reduction of only about two orders of magnitude for the registration.

and 128 GB memory per node. Our SIBIA framework is written in C++ and uses MPI for parallelism. It is compiled using the Intel 17 compiler. We use PETSc's implementations for linear algebra operations and PETSc's TAO package for the non-linear optimization [3, 50], AccFFT for Fourier transforms [15, 16], and PnetCDF for I/O [66]. We use 3 nodes with 64 MPI tasks for (down-sampled) data sizes of size $n_i = 128$, $i = 1, 2, 3$, for all runs reported in this study.

4.3. Performance metrics. The task of our numerical experiments is to assess (i) the convergence towards solutions with low mismatch between the prediction of our model in (2.2a)–(2.2e) and the observed patient input data, both in the brain geometry $\mathbf{m}$ and the tumor c (for all test cases), (ii) the reconstruction quality for the (in practice inaccessible) healthy patient brain¹⁴ as a direct output of the moving atlas formulation, (iii) the quality of inversion for the biophysical parameters, i.e., tumor growth initial conditions and diffusion coefficient, in terms of correctness of the parameters compared to the known ground truth.

We report all performance measures in the patient space, i.e., with respect to the patient anatomy and not the atlas anatomy, since the patient space is the relevant space from an applications point of view. We use the following metrics: We measure the relative mismatch/residual between patient anatomy and reconstructed anatomy, between the reconstructed healthy patient anatomy and the ground truth, and between patient tumor and simulated tumor in the patient domain (compare diagram on right of Table 2.1):

$$\begin{aligned}\mu_{B,L^2} &:= \frac{\|\mathbf{m}'_P(1) - \mathbf{m}_D\|_{L^2(\Omega)^3}}{\|\mathbf{m}_A - \mathbf{m}_D\|_{L^2(\Omega)^3}}, \\ \mu_{B_0,L^2} &:= \frac{\|\mathbf{m}_P(1) - \mathbf{m}_P^*(1)\|_{L^2(\Omega)^3}}{\|\mathbf{m}_A - \mathbf{m}_P^*(1)\|_{L^2(\Omega)^3}}, \\ \mu_{T,L^2} &:= \frac{\|c(1) - c_D\|_{L^2(\Omega)^3}}{\|c_D\|_{L^2(\Omega)^3}},\end{aligned}$$

where $\mathbf{m}_P^*(1)$ denotes the ground truth healthy patient brain. Based on the cardinality $|\cdot|$ of a set and a selection function H with threshold 0.5, i.e.,

$$H(\mathbf{u}) := \{u_i \geq 0.5\},$$

we calculate Dice coefficients for the individual label maps associated with the anatomy labels for $l \in \{WM, GM, CSF\}$ and their average across labels for the patient and the atlas anatomy

$$\text{DICE}_{l,B} := 2 \frac{|H(m_{P,l}(1)) \cap H(m_{D,l})|}{|H(m_{A,l})| + |H(m_{D,l})|}, \quad \text{DICE}_B = \sum_{l=1}^3 \text{DICE}_{l,B}/3.$$

Analogously, we report values for the Dice coefficient computed for the healthy patient geometry, denoted by DICE_{B_0} , and for the anatomy label of the tumor, denoted by DICE_T . To monitor convergence of the Picard iterations, we report the relative value of an *approximation* to the gradient for the coupled problem (see section 2.1) for the final iteration k :

$$\|\mathbf{g}\|_{\text{rel}} := \|\mathbf{g}^j\|_{L^2(\Omega)} / \|\mathbf{g}^0\|_{L^2(\Omega)},$$

¹⁴For our semisynthetic test case setting, where the patient target data are generated from synthetic tumor progression simulation in the healthy brain, the healthy patient brain is known and the approximation quality can be measured exactly.

where $\mathbf{g}^j$ is an approximation to the gradient of the coupled optimization problem (2.2) (given by (2.6i)–(2.6k)) after the i th Picard iteration and $\mathbf{g}^0$ the corresponding approximate gradient for the initial guess. The gradient is only approximate, since we neglect the term $\partial_{\mathbf{m}_{P(1)}}\Phi(\mathbf{p})\alpha(0)$ in the final condition (2.6h) of the adjoint equation for the registration problem. This term is guaranteed to be small if the norm of the remaining gradient is small as the latter requires $\alpha(0)$ and the change in the estimated patient geometry to be small. Finally, we calculate the relative ℓ_2 -error for the initial condition:

$$e_{c0,L^2} := \|\Phi\mathbf{p}^* - c(0)\|_{L^2(\Omega)} / \|\Phi\mathbf{p}^*\|_{L^2(\Omega)},$$

where $c(0) = \Phi\mathbf{p}$ and $\Phi\mathbf{p}^*$ denotes the parametrization of the ground truth tumor initial condition.

4.4. Results for SYN. This first, purely synthetic experiment uses two fairly different, very simplistic anatomies for the atlas and patient brain and an artificially grown tumor in one of them (called the patient brain).

Purpose. This case is designed to cause the MP scheme to fail due to large differences in the brain anatomy between patient and atlas in the tumor region. We demonstrate that the MA scheme retains a good reconstruction quality for the tumor anatomy label, the tumor initial condition, and the characteristic diffusivity of tumor cells, while the MP scheme falls short as it inverts tumor growth in the “wrong” (atlas) geometry.

Setup. For both, atlas and patient, we generate a purely synthetic brain anatomy, composed of an elliptical WM inclusion surrounded by GM. The WM inclusions for atlas and patient differ greatly in size. The target data c_D are generated synthetically from a tumor-progression simulation (with $\mathbf{p} = \mathbf{p}^*$) using our reaction-diffusion model with $\rho_f^* = 15$ and $k_f^* = 1\text{E-}2$ from an initial condition with two nearby Gaussians enabled in the patient geometry. The healthy atlas geometry is depicted in the first row, the tumor-bearing patient geometry in the last row of Figure 4.1. We use the same set of $n_{\mathbf{p}} = 343$ Gaussians with standard deviation $\sigma = 2\pi/30$ for the target data generation and inversion. For simplicity, tumor growth is enabled only in WM. We invert for the registration velocity $\mathbf{v}$, the parameters $\mathbf{p}$ for the tumor initial condition, and for the characteristic diffusivity k_f in WM.

Observation. In this case, the inclusions of WM in atlas and patient vary significantly in size (cf. Figure 4.1). Thus, registering the atlas to the patient requires strong local expansion of volume. Looking at Figure 4.1 (right tableau), in the first iteration, to prevent the registration from directly matching the target tumor unaided by the tumor solver, we choose a higher regularization parameter $\beta_{\mathbf{v}}$ for the registration problem. Therefore, warping the target tumor to the atlas space basically corresponds to a copy operation. (before $k = 2$). This warped-to-atlas tumor is used for tumor inversion in the atlas space. The subsequent registration step (with lower regularization) computes a velocity that registers the patient (small tumor in large WM blob) to the atlas (small tumor in small WM blob). This results in a compression of the patient WM and, consequently, yields a very small target tumor for the inversion in the atlas ($k = 3$).¹⁵ In subsequent iterations (with further reduced registration regularization), the warped-to-atlas target tumor is compressed further, ultimately causing the MP scheme to fail and returning a relative error of almost 100% for the

¹⁵Note, the displayed results are in the patient space, i.e., the inverse of the deformation is applied to display atlas labels in patient space, and the very small reconstructed tumor in the atlas is expanded when warped to the patient space.




FIG. 4.1. Results for SYN (ground truth proliferation rate $\rho_f = 15$ and diffusivity $k_f = 1\text{E}-2$ in WM, tumor growth disabled in GM). We compare the new MA Picard scheme (left tableau), and the MP scheme (right tableau). The figure shows anatomy and tumor labels (from left to right: WM, GM, CSF, and TU), the residual differences (if available), and a hard segmentation based on the given probabilities for the individual tissue classes. The synthetic healthy atlas and the patient are composed of an elliptic inclusion of WM in a GM rectangular brain (see text for details; axial slice 64 of a 3D volume). We show the initial configuration for the problem (top row: iteration $k = 1$), two intermediate results (second row: $k = 2$, and third row: $k = 3$), the final configuration after joint registration and tumor inversion (fourth row: iteration $k = 6$) and the target patient data (reference image: bottom row). All results are presented in patient space. In the table, we report the reconstructed diffusivity $\text{inv-}k_f$ and its relative error e_{k_f} with respect to the ground truth, the average mismatch for the anatomy labels of the tumor bearing brain μ_{B,L^2} , for the healthy brain tissue labels μ_{B_0,L^2} , and for the tumor μ_{T,L^2} , the mean Dice coefficients for brain tissue DICE_B , for the healthy patient brain tissue DICE_{B_0} , and for the tumor DICE_T . Furthermore, we measure the ℓ_2 -error e_{c0,L^2} of the reconstructed tumor initial condition with respect to the ground truth.

reconstruction of the tumor label, the reconstruction of the tumor initial condition, and the reconstruction of the characteristic diffusivity k_f in WM; cf. Figure 4.1). The new MA scheme, however, maintains a good reconstruction quality and results in only 3.3% relative errors for the tumor label, 18.4% relative errors for the reconstruction of the tumor initial condition, and 23.2% relative error in the reconstruction of the characteristic diffusivity rate k_f in WM.

4.5. Results for STRV. With the results for the STRV test cases, we show the general applicability of the new registration and tumor inversion coupling scheme as well as its improved properties in terms of correct reconstruction of initial tumor and healthy brain geometry. We evaluate the MA scheme for a real brain geometry with synthetically generated target tumor such that the true tumor model parameters are known and we have access to the true healthy patient brain.

Purpose. With the STRV test case experiments, we pursue the following main purposes: (i) We give experimental evidence for the convergence of the new MA scheme for several settings; (ii) we compare our coupled image registration and tumor inversion schemes MP and MA to stand-alone tumor inversion, replacing the missing second snapshot (healthy patient) by a statistical atlas without involving image

TABLE 4.2

STRV test case variations. Parameter choices for the STRV test case; ground truth ($\rho_f^* = 15$, $\rho_1^* = 1$, $\rho_1^* = 0$, $k_f^* = 1.00\text{E}-1$, $k_1^* = 1$, $k_1^* = 0$, $\mathbf{p} = \mathbf{p}^*$). For the initial condition parametrization, we use a regular grid of $n_{\mathbf{p}} = 125$ (large σ), and $n_{\mathbf{p}} = 343$ (small σ) Gaussian basis functions, respectively, with standard deviation σ as outlined below, and a spacing of $\delta = 1.5\sigma$. The grid is centered around the positions $\mathbf{x}_{c_1} = 2\pi/128(39, 63, 64)$ (STRV-C1) and $\mathbf{x}_{c_2} = 2\pi/128(61, 89, 64)$ (STRV-C2), respectively (cf. Figures 4.3 and 4.4 for the location of the tumor seed).


Test case	Setting #1	Setting #2	Setting #3
(STRV-C1)	$\sigma = 2\pi/30$, $T = 0.16$	$\sigma = 2\pi/64$, $T = 0.32$	$\sigma = 2\pi/64$, $T = 0.44$
(STRV-C2)	$\sigma = 2\pi/64$, $T = 0.16$	$\sigma = 2\pi/64$, $T = 0.32$	

registration; (iii) we examine the robustness of our joint inversion against noise in the data; (iv) we show the superiority of the new MA scheme over the previously presented MP scheme in terms of reconstruction quality of the (initially unknown) healthy patient brain anatomy and meaningful reconstruction of biophysical tumor-growth characteristics, in particular, shape and sparsity of the ground truth tumor initial condition and prediction of the grown tumor.

Setup. For both atlas and patient, we use real brain MRI data from (healthy) individuals. The target tumor c_D is generated synthetically from a tumor-progression simulation (with $\mathbf{p} = \mathbf{p}^*$) using our reaction-diffusion model. Accordingly, the ground truth velocity field is unknown (or may not even exist due to a possibly wrong model), but the true tumor model parameters are known. We consider two different tumor locations (STRV-C1) and (STRV-C2) and varying time horizons for the tumor-progression simulation. An overview is given in Table 4.2. For (STRV-C1) we also employ different settings of the parametrization of the tumor initial condition: For parameter setting #1, we use a set of $n_{\mathbf{p}} = 125$ Gaussian basis functions with standard deviation $\sigma = 2\pi/30$, whereas, for settings #2 and #3, we use smaller (but more; $n_{\mathbf{p}} = 343$) Gaussians with $\sigma = 2\pi/64$ to generate sparser ground truth initial conditions. The same set of Gaussian basis functions is used for target data generation and inversion. The target data are generated from a forward simulation using $n_t = 100$ time steps with parameters $\rho_f^* = 15$ and $k_f^* = 1\text{E}-1$, and an initial condition with two nearby Gaussians enabled. As suggested by recent investigations [14], tumor growth is enabled in WM only. We invert for the tumor initial condition and the characteristic diffusivity k_f in WM.

Misfit and gradient reduction of the MA scheme. To experimentally assess convergence of the MA scheme over the Picard iterations, we report the ℓ_2 -mismatch and Dice overlay coefficients for the reconstruction of the (pathologic) brain anatomy (μ_{B,L^2} and DICE_B), the healthy patient anatomy (μ_{B_0,L^2} and DICE_{B_0}), and the reconstructed (grown) tumor (μ_{T,L^2} and DICE_T) in Figure 4.2 for the (STRV-C1) test case. Furthermore, we monitor the ℓ_2 -error e_{c_0,L^2} for the reconstructed tumor initial condition with respect to the ground truth and the relative norm $\|\mathbf{g}\|_{\text{rel}}$ of the (approximated) reduced gradient of the coupled formulation (2.2).

Observations. Figure 4.2 shows a monotonic reduction of data-misfit values and the corresponding gain in tissue overlay Dice scores for the brain anatomy, the healthy patient anatomy, the grown tumor, and the initial tumor condition. Although we do not have a proof for the convergence of the MA Picard-iteration-type solution strategy (outlined in section 2.2), we monitor the relative norm $\|\mathbf{g}\|_{\text{rel}}$ of the (approximated) reduced gradient for the MA formulation (2.2), indicating convergence to an optimum. The norm of the approximated gradient continually decreases throughout our Picard-iteration-type solution strategy. Our solution scheme, and hence the gradient, neglect



It	$\beta_{\mathbf{v}}$	μ_{B,L^2}	DICE_B	μ_{B_0,L^2}	DICE_{B_0}	μ_{T,L^2}	DICE_T	$\ \mathbf{g}\ _{\text{rel}}$	e_{c0,L^2}	T^{it} [s]
1	1E0	1.00E0	5.45E-1	6.77E-1	5.51E-1	1.00E0	0.00E0	1.00E0	1.00E0	5.64E2
2	1E-1	8.79E-1	5.80E-1	6.07E-1	5.93E-1	2.90E-1	7.92E-1	2.04E-2	4.25E-1	9.45E2
3	1E-2	6.54E-1	7.29E-1	4.51E-1	7.38E-1	2.04E-1	8.24E-1	1.45E-2	3.51E-1	8.44E2
4	1E-3	4.06E-1	8.66E-1	2.80E-1	8.72E-1	1.16E-1	8.80E-1	1.36E-2	2.08E-1	1.63E3
5	1E-4	2.62E-1	9.32E-1	1.81E-1	9.37E-1	7.57E-2	9.48E-1	7.01E-3	9.55E-2	8.46E2
6	1E-4	2.36E-1	9.42E-1	1.64E-1	9.46E-1	6.74E-2	9.51E-1	4.91E-3	5.75E-2	8.58E2
7	1E-4	2.33E-1	9.43E-1	1.62E-1	9.47E-1	6.49E-2	9.50E-1	4.15E-3	7.13E-2	5.69E3

FIG. 4.2. Misfit reduction of the new MA Picard iteration scheme. We show qualitative and quantitative results for the (STRV-C1) test case with parameter setting #1 from Table 4.2 over the course of the MA Picard-iteration-type solution scheme with a sparsity constraint on the initial condition for the tumor inversion solver. We use the ground truth values $\rho_f = \rho_f^* = 15$ and $k_f = k_f^* = 1E-1$ in WM for the inversion. The figure shows anatomy labels of the healthy atlas brain and the patient brain with an (artificially) grown tumor (see text for details; axial slice 64 of a 3D volume). We show the initial configuration for the problem (top row: iteration $k = 1$), an intermediate result (second row: $k = 3$), the final configuration after joint registration and tumor inversion (third row: iteration $k = 7$), and the target patient data (bottom row: reference image). Each row contains (from left to right) the anatomy labels for WM, GM, CSF, and TU, the residual differences (if available) between the anatomy labels, and a hard segmentation based on the given probabilities for the individual tissue classes. In the table, we report the average mismatch for the anatomy labels for the pathologic brain tissue labels μ_{B,L^2} , the healthy brain tissue labels μ_{B_0,L^2} , and the tumor μ_{T,L^2} , the mean Dice coefficient for brain tissue DICE_B , healthy patient brain tissue DICE_{B_0} , and tumor DICE_T over the Picard iterations. Furthermore, we measure the ℓ_2 -error e_{c0,L^2} of the reconstructed tumor initial condition to the ground truth, and the relative norm $\|\mathbf{g}\|_{\text{rel}}$ of the approximated reduced gradient of the coupled formulation. Timings per iteration are given in seconds for parallel execution using 64 MPI tasks on three nodes of HazelHen. The last row shows the final values for mismatch and Dice as well as the time-to-solution (accumulated runtime) in seconds.

the term $\partial_{\mathbf{m}_{P(1)}} \Phi(\mathbf{p}) \alpha(0)$ in (2.6h). This term, however, becomes very small as the misfit and, thus, $\alpha(0)$ decrease.

We note, that within seven iterations of our Picard scheme, both $\mathbf{p}$ and $\mathbf{v}$ have essentially converged, and the respective updates become very small (i.e., for (STRV-C1) we get $\|\Delta \mathbf{p}\|_r := \|\mathbf{p}^k - \mathbf{p}^{k-1}\|_2 / \|\mathbf{p}^{k-1}\|_2 \approx \mathcal{O}(8E-4)$ and $\|\Delta \mathbf{v}\|_r := \|\mathbf{v}^k - \mathbf{v}^{k-1}\|_2 / \|\mathbf{v}^{k-1}\|_2 \approx \mathcal{O}(1E-4)$). These small updates have little to no effect on the model prediction,

and are not reflected in significant changes of the objective function value (after an additional 23 iterations, $\|\Delta \mathbf{v}\|_r$ drops from $\mathcal{O}(1E - 12)$ to zero, while $\|\Delta \mathbf{p}\|_r$ remains on the order of $\mathcal{O}(1E - 4)$). Since after completion of the β_v -continuation scheme, little to no progress is made, we terminate the solver. We explain such deterioration in convergence, and justify termination of our solver by a few remarks: (i) We employ an optimize-then-discretize approach, which together with the semi-Lagrangian time stepping for the pure advection steps, results in inconsistencies between the discretized objective function and the discretized gradient and, therefore, ultimately results in a stagnation of convergence (in other words, the associated numerical errors in the forward and adjoint operators are on the order of the norm of the gradient). (ii) In each subsolver we are reducing the gradient by three orders of magnitude, which has been confirmed to be adequate if real data input images, presenting many image features and details, are considered [37, 39, 41]. Our empirical observations on numerous datasets suggest that further reducing the gradient does, in general, not result in an improved data mismatch. For practitioners in the field data mismatch is the main benchmark (aside from the requirement that the deformation map is a diffeomorphism).

Stand-alone tumor inversion versus MP versus MA. To compare the different inversion schemes, we show results obtained using the newly introduced MA strategy to the results for the MP solution scheme and those for a tumor stand-alone (T) solver without intersubject registration, respectively, in Table 4.3 for both test cases (STRV-C1) and (STRV-C2).

Observations. Solving the single-snapshot tumor inversion problem without intersubject registration using a tumor stand-alone solver for some arbitrary normal brain anatomy (atlas) (cf. [16]) exhibits poor reconstruction quality. Our joint registration and biophysical inversion MA approach (runs #3 and #6 in Table 4.3) results in significantly higher reconstruction quality and outperforms the other strategies (MP, runs #2 and #5, and tumor stand-alone, runs #1 and #4) by a large margin, both in terms of data misfit and Dice coefficient, as well as in reconstructing the true initial condition.

Comparing to the tumor stand-alone approach, we improve the reconstruction quality of the input data (grown tumor) from a relative data misfit of $2.90E-1$ (run #1) to $6.49E-2$ (run #3) using the new MA joint inversion scheme, and likewise improve the reconstruction quality for the tumor initial condition (compared to the ground truth) from a relative ℓ_2 -error of $4.25E-1$ for tumor stand-alone (run #1) to $7.13E-2$ (run #3) for the MA scheme. Furthermore, the estimation of the characteristic diffusivity in WM k_f is significantly improved for the joint inversion scheme (going from a relative error of $4.80E-1$ (tumor stand-alone; run #4) to a relative error of $7.70E-2$ (MA; run #6)).

MP versus MA. For a more detailed comparison of the MA and the MP schemes, we show qualitative results of the reconstructed tumor initial condition and the grown tumor compared to the target data and ground truth initial condition in Figure 4.3 (for STRV-C1) and Figure 4.4 for (STRV-C2). In Figure 4.5, we outline the initial and final residuals for the different brain tissue labels (WM, GM, CSF, grown tumor ($TU(t = 1)$), and tumor initial condition $TU(t = 0)$; for axial slice 64 of a 3D volume) before and after the joint registration and biophysical inversion for the test case variants (STRV-C1) and (STRV-C2). Furthermore, we display structural differences between the utilized normal brain anatomies for both inversion schemes. We show qualitative results for a longer time horizon of tumor evolution for the (STRV-C1) test case (with a very sparse initial condition) in Figure 4.6.

TABLE 4.3
Comparison of the tumor inversion stand-alone (T) solver only (without intersubject registration), the MP Picard iteration scheme, and the new MA Picard iteration scheme. We show quantitative results for the (STRV-C1) and (STRV-C2) test cases with centers $\mathbf{x}_{c_1}$ and $\mathbf{x}_{c_2}$, and parameter setting #1 from Table 4.2, respectively. We invert for the registration velocity $\mathbf{v}$, the parameters $\mathbf{p}$ of the tumor initial condition, and for the characteristic diffusivity k_f in WM. We set the proliferation rate to the true value $\rho_f = 15$. We report the reconstructed diffusivity $\text{inv-}k_f$ and its relative error e_{k_f} with respect to the ground truth, the average mismatch for the anatomy labels for the pathologic brain tissue labels μ_{B,L^2} , the healthy brain tissue labels μ_{B_0,L^2} , and the tumor μ_{T,L^2} , the mean Dice coefficients for brain tissue DICE_B , for the healthy patient brain tissue DICE_{B_0} and for the tumor DICE_T . We also measure the ℓ_2 -error e_{c_0,L^2} of the reconstructed tumor initial condition to the ground truth. Timings are given for parallel execution using 64 MPI tasks on three nodes of HazelHen.

ID	Center	Setting	Solver	$\text{inv-}k_f$	e_{k_f}	μ_{B,L^2}	DICE_B	μ_{B_0,L^2}	DICE_{B_0}	μ_{T,L^2}	DICE_T	e_{c_0,L^2}	T^{tt} [s]	T^{tu}_{inv} [s]	$T^{\text{reg}}_{\text{inv}}$ [s]
#1	#2 C1	#1	<i>T</i>	1.15E-1	1.50E-1	N/A	N/A	N/A	N/A	2.90E-1	7.92E-1	4.25E-1	N/A	-	N/A
#2			<i>MP</i>	9.22E-2	7.80E-2	3.72E-1	8.72E-1	2.56E-1	8.79E-1	1.44E-1	9.25E-1	3.25E-1	5.72E3	1.99E3	7.88E2
#3			<i>MA</i>	9.05E-2	9.50E-2	2.33E-1	9.43E-1	1.62E-1	9.47E-1	6.49E-1	9.50E-1	7.13E-2	5.69E3	1.59E3	1.14E3
#4	#5 C2	#1	<i>T</i>	1.48E-1	4.80E-1	N/A	N/A	N/A	N/A	2.93E-1	5.54E-1	4.08E-1	N/A	-	N/A
#5			<i>MP</i>	8.86E-2	1.14E-1	3.76E-1	8.79E-1	2.55E-1	8.80E-1	1.90E-1	7.96E-1	4.70E-1	6.05E3	2.50E3	7.49E2
#6			<i>MA</i>	9.23E-2	7.70E-2	2.37E-1	9.48E-1	1.60E-1	9.48E-1	6.07E-1	9.70E-1	6.33E-2	7.89E3	4.40E3	9.26E2




FIG. 4.3. Comparison of the MP and the new MA schemes. We show qualitative results for the (STRV-C1) test case with parameter setting #1 from Table 4.2 (left), and parameter setting #2 from Table 4.2 (right). The image shows the reconstructed grown tumor and tumor initial condition for the MA and MP solution schemes, featuring different time horizons and sparsity of the initial condition. We show parts of the patient brain anatomy with the respective reconstructed tumor initial condition (magenta wireframe/volume) as compared to the ground truth initial condition (cyan volume). The grown tumor is indicated as white wireframe (3D cut image), and as blue semitransparent volume compared to the target data given as red wireframe (close-up image). The top row shows the test case target data (grown tumor; white wireframe) and initial condition ground truth (cyan volume) for each set of parameters. The close-up image shows a volume rendering of the grown tumor, overlaid with a wireframe indicating the boundary of the tumor resulting from the segmentation (used for visualization purposes). The light blue/white area indicates the ventricles with CSF.

We also study the sensitivity of the MA and MP solution strategies with respect to perturbations in the tumor model parameters. In Table 4.4, we vary the value of the characteristic diffusivity in WM from the ground truth, and monitor the obtained reconstruction performance.

Observations. The new MA scheme outperforms the MP counterpart [56] in various ways. Table 4.3 shows that, for the reconstruction of the brain anatomy (model prediction for anatomy labels of brain tissue labels WM, GM, and CSF), we reach a relative data misfit (ℓ_2 -error) of $2.33\text{E}-1$ (run #3) compared to $3.72\text{E}-1$ (run #2) for the MP scheme (this translates to a Dice score of $9.43\text{E}-1$ (averaged over all labels) for the MA compared to a Dice score of $8.72\text{E}-1$ for the old scheme). Similarly, for the approximation of the actual healthy patient anatomy, we improve from a Dice score of $8.79\text{E}-1$ (relative ℓ_2 -error of 25.6%; run #3) obtained from the MP solution to a Dice score of $9.47\text{E}-1$ (relative ℓ_2 -error of only 16.2%; run #2) using the new scheme. Figure 4.5 illustrates the structural differences between the atlas and patient brain, the initial residuals, and the residuals after joint inversion using the MA and the MP schemes, respectively, for an exemplary slice (axial slice 64) of the 3D volume. We observe smaller errors for the MA solution. The improved anatomy reconstruction is characteristic for the MA strategy; we observe similar trends for various solver and




FIG. 4.4. Comparison of the MP and the new MA schemes. We show qualitative results for the (STRV-C2) test case with parameter setting #1 from Table 4.2 (left image), and parameter setting #2 from Table 4.2 (right image). The image shows the reconstructed grown tumor and the tumor initial condition for the MA and MP solution schemes featuring different time horizons for the tumor evolution. Illustrated are parts of the (approximated) patient brain anatomy with the respective reconstructed tumor initial condition (magenta wireframe/volume) as compared to the ground truth initial condition (cyan volume). The grown tumor is indicated as white wireframe (bottom images), and as blue semitransparent volume compared to the target data given as red wireframe (top images). The first row shows the test case target data (grown tumor; white wireframe (left) and volume rendering (right)) and initial condition ground truth (cyan volume) for each set of parameters.

parameter configurations (compare also Table 4.4 with inversion under perturbed or wrong tumor model parameters).

The MA scheme furthermore results in improved reconstruction quality of the predicted grown tumor and higher similarity to the target data. For (STRV-C1), the MA scheme results in a relative error of 6.4% (Dice score of $9.50\text{E}-1$; run #3) compared to a relative error of 14.4% (Dice score of $9.25\text{E}-1$; run #2) for the MP solution. The numbers for (STRV-C2) are very similar with 6.0% relative error (Dice score of $9.70\text{E}-1$; run #6) versus 19.0% relative error (Dice score of $7.96\text{E}-1$; run #5).

The primary objective for the derivation of the MA scheme was to allow for a more informative inversion for biophysical model parameters. We primarily consider the estimation of the tumor initial condition, but also invert for the characteristic diffusivity k_f in WM. The relative errors¹⁶ e_{c0,L^2} for the inversion of the ground truth tumor initial condition in Table 4.3 obtained for each scheme clearly attest the MA scheme to be more reliable and sound in recovering biophysical parameters.¹⁷ Us-

¹⁶Note that for the MA scheme, the reconstructed initial condition naturally lives in the patient space; for the MP solution strategy, we invert for the tumor initial condition in the atlas space. For a fair comparison, the reconstructed initial condition for the MP scheme is advected to the patient space before computing the relative error to the ground truth.

¹⁷Our analysis assumes that tumor growth is perfectly described by a reaction-diffusion model, which is quite certainly not true. When we say “biophysically meaningful” this is to be seen under the aforementioned assumption. Note furthermore, that the tumor model can easily be exchanged in the modular setting of our Picard iteration approach.




FIG. 4.5. Comparison of the MP Picard iteration scheme, and the new MA Picard iteration scheme. We show initial and final residuals for the (STRV-C1) test case (top) and the (STRV-C2) test case (bottom) with parameter setting #1 from Table 4.2, respectively. The tableaux show the initial and final residual for the reconstruction of the pathologic and healthy patient for the MA and MP solution schemes, respectively for parameter setting #1 from Table 4.2 for both tumor seeds $\mathbf{x}_{c_1}$ and $\mathbf{x}_{c_2}$. For each case, the top row shows the initial residual between patient input data and the healthy atlas anatomy; the healthy anatomies are illustrated on the right, where gray indicates WM, blue indicates GM, and black indicates CSF. The final residual for the reconstruction of the patient input data and the healthy patient anatomy are shown in the second and third row, for joint inversion using the MP (run #2 for (STRV-C1) and run #5 for (STRV-C2) in Table 4.3) and the MA scheme (run #3 for (STRV-C1) and run #6 for (STRV-C2) in Table 4.3), respectively. All images show axial slice 64 (of 128) of the 3D volume.




FIG. 4.6. Comparison of the MP, and the new MA schemes. We show qualitative results for the (STRV-C1) test case with parameter setting #3 from Table 4.2. The image shows the reconstructed grown tumor and the tumor initial condition for the MA and MP solution schemes. The images show parts of the patient brain anatomy and illustrate the respective reconstructed tumor initial condition (magenta wireframe/volume) as compared to the ground truth initial condition (cyan volume). The grown tumor is indicated as white wireframe (bottom row), and as blue semitransparent volume compared to the target data given as red wireframe (top row). The light blue/white area indicates the ventricles with CSF. The top row shows the test case target data (grown tumor; white wireframe) and initial condition ground truth (cyan volume). The close-up image shows a volume rendering of the grown tumor, overlaid with a wireframe indicating the boundary of the tumor resulting from the segmentation (used for visualization purposes). We observe that the MP reconstruction yields a wrong position for the initial condition and inferior reconstruction of the grown tumor.

ing the improved MA scheme, we achieve excellent reconstruction of the true initial condition with a relative error of only 7.1% (for (STRV-C1); run #3) and 6.3% (for (STRV-C2); run #6), as opposed to a relative error of 32.5% (for (STRV-C1); run #2) and 42.4% (for (STRV-C2); run #5), respectively, when using the MP scheme instead. An illustration of the error is given in Figure 4.5 (for an axial cut at slice 64). We explain the improved inversion properties of the MA scheme by its general idea to seek for a good approximation of the healthy patient brain anatomy first and, thus, to carry out the inversion in the “correct” space. For the MP scheme, the inversion in the wrong anatomy (atlas space) has the potential to induce large errors for the estimation of biophysical parameters as shown in section 4.4 since the intersubject deformation map can aid matching the pathologic brains. In particular, there is no implication that this deformation map produces reasonable results when applied to the tumor initial condition (or other quantities that describe the time point of tumor genesis) in order to translate the latter into the individual patient application space. Visual inspection of our simulation results for longer time horizons in Figure 4.6 support this point: The MP solution recovers a wrong position of the tumor initial condition (compare the green (ground truth) and purple (reconstruction) isovolume representations of the initial tumor condition). A wrong position, wrong shape, or wrong sparsity of the initial condition can, in a second step, also cause a wrong estimation of the characteristic proliferation rate or tumor cell infiltration rate. For our experiments, the difference for the estimation of the characteristic diffusivity k_f is, however, not signif-

TABLE 4.4

Comparison of the MP, and the new MA schemes. We report results for the (STRV-C1) and (STRV-C2) test cases with centers $\mathbf{x}_{c_1}$ and $\mathbf{x}_{c_2}$, and parameter setting #1 from Table 4.2, respectively. Here we are not inverting for k_f , but only invert for the tumor initial condition and registration velocity. We analyze the sensitivity of the joint inversion with respect to perturbations in the characteristic diffusivity k_f by inverting with different choices of k_f ; bold numbers correspond to inversion with the ground truth value ($k_f^* = 1\text{E-}1$). We always use the ground truth proliferation rate $\rho_f = 15$. We report the average mismatch for the anatomy labels for the pathologic brain tissue labels μ_{B,L^2} , the healthy brain tissue labels μ_{B_0,L^2} , and the tumor μ_{T,L^2} , the mean Dice coefficient for brain tissue $DICE_B$, for the healthy patient brain tissue $DICE_{B_0}$, and for the tumor $DICE_T$ (in cases without Dice score for the tumor reconstruction, the tumor probability map has values below 0.5 everywhere). Furthermore, we measure the ℓ_2 -error e_{c0,L^2} of the reconstructed tumor initial condition to the ground truth.

ID	k_f	μ_{B,L^2}	$DICE_B$	μ_{B_0,L^2}	$DICE_{B_0}$	μ_{T,L^2}	$DICE_T$	e_{c0,L^2}
#7	5E-1	3.80E-1	8.62E-1	2.58E-1	8.79E-1	3.96E-1	—	4.99E-1
#8	3E-1	3.76E-1	8.64E-1	2.57E-1	8.79E-1	2.88E-1	7.76E-1	8.68E-1
#9	$k_f^* = 1\text{E-}1$	3.72E-1	8.72E-1	2.56E-1	8.79E-1	1.45E-1	9.21E-1	3.23E-1
#10	1E-3	3.73E-1	8.75E-1	2.59E-1	8.73E-1	2.49E-1	8.75E-1	5.39E-1
#11	5E-1	2.40E-1	9.46E-1	1.62E-1	9.47E-1	5.65E-1	—	1.37E0
#12	3E-1	2.40E-1	9.37E-1	1.63E-1	9.48E-1	2.86E-1	3.66E-1	9.41E-1
#13	$k_f^* = 1\text{E-}1$	2.35E-1	9.43E-1	1.62E-1	9.48E-1	6.77E-2	9.56E-1	8.02E-2
#14	1E-3	2.37E-1	9.48E-1	1.61E-1	9.48E-1	5.40E-1	—	7.04E-1
#15	3E-1	3.76E-1	8.79E-1	2.55E-1	8.80E-1	3.66E-1	—	4.96E-1
#16	2E-1	3.76E-1	8.79E-1	2.55E-1	8.80E-1	2.64E-1	—	5.25E-1
#17	$k_f^* = 1\text{E-}1$	3.76E-1	8.80E-1	2.55E-1	8.80E-1	1.93E-1	8.34E-1	4.24E-1
#18	1E-2	3.82E-1	8.75E-1	2.59E-1	8.75E-1	2.57E-1	6.86E-1	6.94E-1
#19	3E-1	2.41E-1	9.46E-1	1.63E-1	9.47E-1	4.20E-1	—	5.99E-1
#20	2E-1	2.40E-1	9.46E-1	1.62E-1	9.47E-1	2.54E-1	—	7.96E-1
#21	$k_f^* = 1\text{E-}1$	2.37E-1	9.48E-1	1.60E-1	9.48E-1	6.07E-2	9.70E-1	6.33E-2
#22	1E-2	2.40E-1	9.46E-1	1.62E-1	9.47E-1	3.39E-1	6.21E-1	6.78E-1

icant; both schemes yield comparable results. One possible reason for this could be that the diffusion part of the gradient initially is large but then flattens out compared to the gradient component for the initial condition parametrization. The theoretical and numerical analysis of this effect is ongoing work.

Robustness of the joint inversion against perturbed target data. We examine the sensitivity of our joint inversion scheme to the synthetic data generation process by considering low-frequency noise in the target data c_D , and varying image resolutions (i.e., mesh sizes) for target data generation and inversion. We investigate two variations for the (STRV-C1) test case: (i) We perturb the synthesized target data by adding 1%, 5%, and 10% noise. Rather than using uncorrelated high frequency (white) noise, we study the effects of a more realistic low-frequency noise model. To this end, we modify the low-frequency spectrum of the target data c_D by adding uniform random noise scaled by the inverse squared frequency to the respective frequency components of the data. Thereafter, the inverse Fourier transform of this noisy spectrum is applied to obtain the noise-corrupted data. This results in tumors with a slightly different shape compared to the one generated by plain forward simulation. (ii) We consider the higher resolution $n = 256^3$ patient brain to synthesize the target data c_D , and resample the latter to $n = 128^3$ for application of the joint inversion scheme. Note, that the original MRI data have a resolution of $n = 256^3$, exposing a higher fidelity of brain structures and imaging features, which, consequently, results in the generation of more detailed target data. Table 4.5 shows numerical results for both variations of the (STRV-C1) test case, comparing the MA joint inversion against the MP schemes.

TABLE 4.5
Robustness of joint inversion. We report reconstruction errors for the (STRV-C1), parameter setting #1 test case from Table 4.3, and sensitivity to (i) influence of low frequency noise in the target data, and (ii) the process of synthesizing c_D , using the MP and MA Picard iteration schemes.

Solver	Perturbation	inv- k_f	e_{k_f}	μ_{B,L^2}	DICE $_B$	μ_{B_0,L^2}	DICE $_{B_0}$	μ_{T,L^2}	DICE $_T$	e_{c_0,L^2}
MA	0%	9.05E-2	9.50E-2	2.33E-1	9.43E-1	1.62E-1	9.47E-1	6.49E-2	9.50E-1	7.13E-2
	1%	9.04E-2	9.61E-2	2.33E-1	9.43E-1	1.62E-1	9.46E-1	6.59E-2	9.53E-1	7.86E-2
	5%	8.98E-2	1.02E-1	2.34E-1	9.43E-1	1.63E-1	9.46E-1	6.97E-2	9.56E-1	1.29E-1
	10%	9.02E-2	9.82E-2	2.30E-1	9.43E-1	1.61E-1	9.47E-1	8.40E-2	9.57E-1	1.63E-1
	Mesh $256^3 \rightarrow 128^3$	6.54E-2	3.46E-1	2.34E-1	9.43E-1	1.62E-1	9.46E-1	1.58E-1	9.11E-1	1.62E-1
MP	0%	9.22E-2	7.80E-2	3.72E-1	8.72E-1	2.56E-1	8.79E-1	1.44E-1	9.25E-1	3.25E-1
	1%	9.30E-2	7.04E-2	3.72E-1	8.75E-1	2.56E-1	8.79E-1	1.44E-1	9.29E-1	3.22E-1
	5%	9.10E-2	8.99E-2	3.71E-1	8.74E-1	2.56E-1	8.78E-1	1.46E-1	9.27E-1	3.18E-1
	10%	9.27E-2	7.27E-2	3.70E-1	8.74E-1	2.56E-1	8.78E-1	1.53E-1	9.28E-1	3.26E-1
	Mesh $256^3 \rightarrow 128^3$	5.98E-2	4.02E-1	3.72E-1	8.75E-1	2.56E-1	8.78E-1	2.07E-1	8.95E-1	3.82E-1

Observations. We observe that the MA scheme exhibits a slight deterioration in the reconstruction accuracy of the tumor initial condition for increasingly high noise level in the observed data. Nonetheless, even with using target data with 10% noise corruption, the new MA scheme outperforms the MP scheme by a large margin with respect to reconstruction of the initial and grown tumor, as well as the approximation of the healthy patient brain. Overall, both joint inversion schemes are quite robust against the employed low-frequency noise model.

Similarly, when using target data obtained from a resampled high resolution, high fidelity forward simulation, we observe a deterioration in reconstruction accuracy for both MP and MA schemes. Despite this fact, the MA scheme remains superior, and produces higher accuracy reconstructions of the initial and grown tumor, and the healthy patient brain.

From this experiment, we conclude that the presented small reconstruction errors μ_{T,L^2} , and e_{c0,L^2} for the grown tumor and its initial condition are partly attributed to the synthetic data generation process, and deterioration of reconstruction quality is expected when moving towards clinical data. Notwithstanding the above, the MA scheme results in higher quality estimates due to a better (and direct) approximation of the healthy patient anatomy.

5. Conclusion. We present a new inverse problem formulation to solve the patient-specific tumor inversion problem given a single data snapshot in time only. We remedy the lack of healthy patient imaging data via registration to a healthy atlas brain, resulting in a joint intersubject registration and biophysical inversion scheme. We simultaneously solve for an estimation of the intersubject deformation map and (a subset of) the biophysical model parameters for tumor progression. More specifically, we solve for an estimation of a sparse tumor initial condition and for the characteristic diffusivity rate of tumor cells in WM. The presented scheme conceptually improves on our earlier coupled solution approach [56]. The latter uses the registration to warp patient target data towards the atlas space, and inverts for tumor model parameters in the atlas brain (proxy for the unknown healthy patient brain). This may introduce large errors due to unwanted data fitting and model inversion in a wrong anatomy. We demonstrate this for synthetic cases. The new scheme computes an approximation of the healthy patient and performs the tumor inversion in the patient space. We furthermore encourage a sparse localization of the initial tumor (presented in our work [59]) to allow for a more reliable reconstruction of the model parameters. This biophysically motivated constraint can help to understand tumor genesis and its role in the subsequent tumor evolution. We derive a Picard-iteration-type solution strategy by dividing the strongly coupled set of first order optimality conditions into two solvers.

For our new method, we observe (i) improved similarity of the reconstructed probability maps of brain tissue labels (WM, GM, CSF) with respect to the patient data, (ii) significantly smaller errors for the reconstruction of the grown tumor, (iii) a better approximation of the actual healthy patient brain anatomy, and, (iv) more reliable reconstruction of biophysical parameters, such as the tumor initial condition.

The reconstruction result obtained from the MA scheme seems to be more “natural” as we invert in the correct brain anatomy of the (healthy) patient, and the registration does not act on the tumor probability map (as is the case for the MP strategy). In particular, using the MA scheme, we were able to improve the Dice score for the brain anatomy from $8.72\text{E}-1$ for the MP strategy to $9.43\text{E}-1$; the Dice score

for the approximation of the healthy patient brain anatomy from $8.79\text{E}-1$ for MP to $9.47\text{E}-1$; and the Dice coefficient for the tumor reconstruction from $8.34\text{E}-1$ for MP to $9.70\text{E}-1$. We observe an excellent reconstruction of the true initial condition with a relative error of only 7.1% (STRV-C1) and 6.3% (STRV-C2) (compared to 32.3% and 42.4% for MP) using the more sophisticated MA scheme. We further conclude, that a sparse localization of the initial condition is essential if targeting biophysical parameter estimation from patient MRI. The previously used ℓ_2 -regularization results in rich initial conditions. Last, the MA solution scheme seems to be more useful in identifying wrong parameters (such as cell proliferation rate and cell migration rate of brain tissue) than the MP counterpart.

In this study, we focused on the differences between two registration approaches and, thus, used a single healthy brain as “atlas.” Typically, the latter refers to a representative image that provides label maps for, e.g., functional areas or to a statistical template, generated from a population of images. The presented approach works for either dataset. If we are interested in normalizing patient individual results to a common template [20], using a statistical atlas, would be a good choice. For image-driven biophysical model calibration, an atlas brain may be selected from a database of healthy individuals based on a measure of closeness to the patient brain (e.g., mutual information, sum of squared differences, or cross correlation).

Acknowledgments. Any opinions, findings, and conclusions or recommendations expressed herein are those of the authors and do not necessarily reflect the views of the AFOSR, DOE, NIH, DARPA, and NSF. Computing time on the High-Performance Computing Centers (HLRS) Hazel Hen system (Stuttgart, Germany) was provided by an allocation of the federal project application ACID-44104.

REFERENCES

- [1] A. GHOLAMI, J. HILL, D. MALHOTRA, AND G. BIROS, *Accfft: A Library for Distributed-Memory FFT on CPU and GPU Architectures*, preprint, arXiv:1506.07933, 2015.
- [2] S. BAKAS, Z. ZENG, A. SOTIRAS, S. RATHORE, H. AKBARI, B. GAONKAR, M. ROZYCKI, S. PATI, AND C. DAVATZIKOS, *GLISTRboost: Combining multimodal MRI segmentation, registration, and biophysical tumor growth modeling with gradient boosting machines for glioma segmentation*, *Brainlesion*, 9556 (2015), pp. 144–155.
- [3] S. BALAY, S. ABHYANKAR, M. F. ADAMS, J. BROWN, P. BRUNE, K. BUSCHELMAN, L. DALCIN, V. ELJKHOUT, W. D. GROPP, D. KAUSHIK, M. G. KNEPLEY, L. C. MCINNES, K. RUPP, B. F. SMITH, S. ZAMPINI, H. ZHANG, AND H. ZHANG, *PETSc Users Manual*, Technical report ANL-95/11 - Revision 3.7, Argonne National Laboratory, Argonne, IL, 2016.
- [4] V. BARBU AND G. MARINOSCHI, *An optimal control approach to the optical flow problem*, *Systems Control Lett.*, 87 (2016), pp. 1–9.
- [5] M. F. BEG, M. I. MILLER, A. TROUVÉ, AND L. YOUNES, *Computing large deformation metric mappings via geodesic flows of diffeomorphisms*, *Int. J. Comput. Vis.*, 61 (2005), pp. 139–157.
- [6] A. BORZI, K. ITO, AND K. KUNISCH, *An optimal control approach to optical flow computation*, *Internat. J. Numer. Methods Fluids*, 40 (2002), pp. 231–240.
- [7] K. CHEN AND D. A. LORENZ, *Image sequence interpolation based on optical flow, segmentation and optimal control*, *IEEE Trans. Image Process.*, 21 (2012), pp. 1020–1030.
- [8] X. CHEN, R. M. SUMMERS, AND J. YOA, *Kidney tumor growth prediction by coupling reaction-diffusion and biomechanical model*, *IEEE Trans. Biomed. Eng.*, 60 (2012), pp. 169–173.
- [9] O. CLATZ, M. SERMESANT, P. Y. BONDIAU, H. DELINGETTE, S. K. WARFIELD, G. MALANDAIN, AND N. AYACHE, *Realistic simulation of the 3D growth of brain tumors in MR images coupling diffusion with biomechanical deformation*, *IEEE Trans. Med. Imaging*, 24 (2005), pp. 1334–1346.
- [10] T. COLIN, A. IOLLO, J.-B. LAGAERT, AND O. SAUT, *An inverse problem for the recovery of the vascularization of a tumor*, *J. Inverse Ill-Posed Probl.*, 22 (2014), pp. 759–786.

- [11] S. DEL PINO AND O. PIRONNEAU, *A fictitious domain based general PDE solver*, in Numerical Methods for Scientific Computing: Variational Problems and Applications, Barcelona, 2003, CIMNE, Barcelona, 2003.
- [12] R. S. DEMBO AND T. STEIHAUG, *Truncated-Newton algorithms for large-scale unconstrained optimization*, Math. Program., 26 (1983), pp. 190–212.
- [13] S. C. EISENSTAT AND H. F. WALKER, *Choosing the forcing terms in an inexact Newton method*, SIAM J. Sci. Comput., 17 (1996), pp. 16–32.
- [14] M. ESMAEILI, A. L. STENSJØEN, E. M. BERNTSEN, O. SOLHEIM, AND I. REINERTSEN, *The direction of tumour growth in glioblastoma patients*, Sci. Rep., 8 (2018), 1199, <https://doi.org/10.1038/s41598-018-19420-z>.
- [15] A. GHOLAMI AND G. BIROS, *AccFFT*, 2017, <http://www.accfft.org>.
- [16] A. GHOLAMI, A. MANG, AND G. BIROS, *An inverse problem formulation for parameter estimation of a reaction-diffusion model of low grade gliomas*, J. Math. Biol., 72 (2016), pp. 409–433.
- [17] A. GHOLAMI, A. MANG, K. SCHEUFELE, C. DAVATZIKOS, M. MEHL, AND G. BIROS, *A framework for scalable biophysics-based image analysis*, in Proceedings of the ACM/IEEE Conference on Supercomputing, ACM, New York, 2017, 19.
- [18] P. E. GILL, W. MURRAY, AND M. H. WRIGHT, *Practical Optimization*, Academic Press, Waltham, MA, 1981.
- [19] P. GONG, C. ZHANG, Z. LU, J. Z. HUANG, AND J. YE, *A general iterative shrinkage and thresholding algorithm for non-convex regularized optimization problems*, in Proceedings of the 30th International Conference on International Conference on Machine Learning, JMLR Workshop Conf. Proc. 28, 2013, pp. 37–45.
- [20] A. GOOYA, K. M. POHL, M. BILELLO, L. CIRILLO, G. BIROS, E. R. MELHEM, AND C. DAVATZIKOS, *GLISTR: Glioma image segmentation and registration*, IEEE Trans. Med. Imaging, 31 (2013), pp. 1941–1954.
- [21] H. L. P. HARPOLD, E. C. ALVORD, AND K. R. SWANSON, *The evolution of mathematical modeling of glioma proliferation and invasion*, J. Neuropathol. Exper. Neurol., 66 (2007), pp. 1–9.
- [22] A. HAWKINS-DAARUD, S. PRUDHOMME, K. G. VAN DER ZEE, AND J. T. ODEN, *Bayesian calibration, validation, and uncertainty quantification of diffuse interface models of tumor growth*, J. Math. Biol., 67 (2013), pp. 1457–1485.
- [23] C. HOGEA, C. DAVATZIKOS, AND G. BIROS, *Brain-tumor interaction biophysical models for medical image registration*, SIAM J. Comput. Sci., 30 (2008), pp. 3050–3072.
- [24] C. HOGEA, C. DAVATZIKOS, AND G. BIROS, *An image-driven parameter estimation problem for a reaction-diffusion glioma growth model with mass effects*, J. Math. Biol., 56 (2008), pp. 793–825.
- [25] N. P. HURLEY AND S. T. RICKARD, *Comparing measures of sparsity*, IEEE Trans. Interm. Theory, 55 (2009), pp. 4723–4741.
- [26] P. P. JARDE AND M. ULBRICH, *Existence of Minimizers for Optical Flow Based Optimal Control Problems under Mild Regularity Assumptions*, manuscript.
- [27] D. KNOPOFF, D. R. FERNÁNDEZ, G. A. TORRES, AND C. V. TURNER, *A mathematical method for parameter estimation in a tumor growth model*, Comput. Appl. Math., 36 (2017), pp. 733–748.
- [28] D. A. KNOPOFF, D. R. FERNÁNDEZ, G. A. TORRES, AND C. V. TURNER, *Adjoint method for a tumor growth PDE-constrained optimization problem*, Comput. Math. Appl., 66 (2013), pp. 1104–1119.
- [29] E. KONUKOGLU, O. CLATZ, P. Y. BONDIAU, H. DELINGETTE, AND N. AYACHE, *Extrapolating glioma invasion margin in brain magnetic resonance images: Suggesting new irradiation margins*, Med. Image Anal., 14 (2010), pp. 111–125.
- [30] D. KWON, M. NIETHAMMER, H. AKBARI, M. BILELLO, C. DAVATZIKOS, AND K. M. POHL, *PORTR: Pre-operative and post-recurrence brain tumor registration*, IEEE Trans. Med. Imaging, 33 (2014), pp. 651–667.
- [31] M. LÊ, H. DELINGETTE, J. KALPATHY-CRAMER, E. R. GERSTNER, T. BATCHELOR, J. UNKELBACH, AND N. AYACHE, *Bayesian personalization of brain tumor growth model*, in Medical Image Computing and Computer-Assisted Intervention, Springer, Cham, Switzerland, 2015, pp. 424–432.
- [32] M. LÊ, 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 Trans. Med. Imaging, 36 (2017), pp. 815–825.
- [33] E. LIMA, J. ODEN, B. WOHLMUTH, A. SHAHMORADI, D. HORMUTH II, T. YANKEELOV, L. SCARABOSIO, AND T. HORGER, *Selection and validation of predictive models of radia-*

- tion effects on tumor growth based on noninvasive imaging data, *Comput. Methods Appl. Mech. Engrg.*, 327 (2017), pp. 277–305.
- [34] E. A. B. F. LIMA, J. T. ODEN, D. A. HORMUTH, T. E. YANKEELOV, AND R. C. ALMEIDA, *Selection, calibration, and validation of models of tumor growth*, *Math. Models Methods Appl. Sci.*, 26 (2016), pp. 2341–2368.
 - [35] Y. LIU, S. M. SADOWKI, A. B. WEISBROD, E. KEBEBEW, R. M. SUMMERS, AND J. YAO, *Patient specific tumor growth prediction using multimodal images*, *Med. Image Anal.*, 18 (2014), pp. 555–566.
 - [36] A. MANG, *Methoden zur numerischen Simulation der Progression von Gliomen: Modellentwicklung, Numerik und Parameteridentifikation [Methods for the Numerical Simulation of the Progression of Glioma: Model Design, Numerics and Parameter Identification]*, Springer Fachmedien, Wiesbaden, Germany, 2014 (in German).
 - [37] A. MANG AND G. BIROS, *An inexact Newton–Krylov algorithm for constrained diffeomorphic image registration*, *SIAM J. Imaging Sci.*, 8 (2015), pp. 1030–1069.
 - [38] A. MANG AND G. BIROS, *Constrained H^1 -regularization schemes for diffeomorphic image registration*, *SIAM J. Imaging Sci.*, 9 (2016), pp. 1154–1194.
 - [39] A. MANG AND G. BIROS, *A semi-Lagrangian two-level preconditioned Newton–Krylov solver for constrained diffeomorphic image registration*, *SIAM J. Sci. Comput.*, 39 (2017), pp. B1064–B1101.
 - [40] A. MANG, A. GHOLAMI, AND G. BIROS, *Distributed-memory large-deformation diffeomorphic 3D image registration*, in *Proceedings of the ACM/IEEE Conference on Supercomputing*, ACM, New York, 2016.
 - [41] A. MANG, A. GHOLAMI, C. DAVATZIKOS, AND G. BIROS, *PDE-constrained optimization in medical image analysis*, *Optim. Eng.*, 19 (2018), pp. 765–812.
 - [42] A. MANG, A. GHOLAMI, C. DAVATZIKOS, AND G. BIROS, *CLAIRE: A distributed-memory solver for constrained large deformation diffeomorphic image registration*, *SIAM J. Sci. Comput.*, 41 (2019), pp. C548–C584.
 - [43] A. MANG, A. TOMA, T. A. SCHUETZ, S. BECKER, T. ECKEY, C. MOHR, D. PETERSEN, AND T. M. BUZUG, *Biophysical modeling of brain tumor progression: From unconditionally stable explicit time integration to an inverse problem with parabolic PDE constraints for model calibration*, *Med. Phys.*, 39 (2012), pp. 4444–4459.
 - [44] B. H. MENZE, K. V. LEEMPUT, E. KONUKOGLU, A. HONKELA, M.-A. WEBER, N. AYACHE, AND P. GOLLAND, *A generative approach for image-based modeling of tumor growth*, in *Proceedings of Information Processing in Medical Imaging*, Springer, Berlin, 2011, pp. 735–747.
 - [45] H. MI, C. PETITJEAN, B. DUBRAY, P. VERA, AND S. RUAN, *Prediction of lung tumor evolution during radiotherapy in individual patients with PET*, *IEEE Trans. Med. Imaging*, 33 (2014), pp. 995–1003.
 - [46] J. MODERSITZKI, *Numerical Methods for Image Registration*, Oxford University Press, New York, 2004.
 - [47] J. MODERSITZKI, *FAIR: Flexible Algorithms for Image Registration*, *Fundam. Algorithms* 6, SIAM, Philadelphia, 2009.
 - [48] A. MOHAMED, D. SHEN, AND C. DAVATZIKOS, *Deformable registration of brain tumor images via a statistical model of tumor-induced deformation*, *Med. Image Anal.*, 10 (2006), pp. 752–763.
 - [49] J. J. MORÉ AND D. J. THUENTE, *Line search algorithms with guaranteed sufficient decrease*, *ACM Trans. Math. Software*, 20 (1994), pp. 286–307, <https://doi.org/10.1145/192115.192132>.
 - [50] T. MUNSON, J. SARICH, S. WILD, S. BENSON, AND L. C. MCINNES, *TAO 3.6 Users Manual*, Argonne National Laboratory, Mathematics and Computer Science Division, Argonne, IL, 2015.
 - [51] J. NOCEDAL AND S. J. WRIGHT, *Numerical Optimization*, Springer, New York, 2006.
 - [52] J. T. ODEN, E. E. PRUDENCIO, AND A. HAWKINS-DAARUD, *Selection and assessment of phenomenological models of tumor growth*, *Math. Models Methods Appl. Sci.*, 23 (2013), pp. 1309–1338.
 - [53] A. A. I. QUIROGA, D. FERNÁNDEZ, G. A. TORRES, AND C. V. TURNER, *Adjoint method for a tumor invasion PDE-constrained optimization problem in 2D using adaptive finite element method*, *Appl. Math. Comput.*, 270 (2015), pp. 358–368.
 - [54] S. JBABDI, E. MANDONNET, H. DUFFAU, L. CAPELLE, K. SWANSON, M. PÉLÉGRINI-ISSAC, R. GUILLEVIN, AND H. BENALI, *Simulation of anisotropic growth of low-grade gliomas using diffusion tensor imaging*, *Magn. Reson. Med.*, 54 (2005), pp. 616–624.

- [55] K. SCHEUFELE, *Coupling Schemes and Inexact Newton for Multi-Physics and Coupled Optimization Problems*, Ph.D. thesis, University of Stuttgart, Stuttgart, Germany, 2019.
- [56] K. SCHEUFELE, A. MANG, A. GHOLAMI, C. DAVATZIKOS, G. BIROS, AND M. MEHL, *Coupling brain-tumor biophysical models and diffeomorphic image registration*, Comput. Methods Appl. Mech. Engrg., pp. 533–567, <https://doi.org/10.1016/j.cma.2018.12.008>.
- [57] A. SOTIRAS, C. DAVATZIKOS, AND N. PARAGIOS, *Deformable medical image registration: A survey*, IEEE Trans. Med. Imaging, 32 (2013), pp. 1153–1190.
- [58] G. STRANG, *On the construction and comparison of difference schemes*, SIAM J. Numer. Anal., 5 (1968), pp. 506–517.
- [59] 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, to appear.
- [60] K. R. SWANSON, E. C. ALVORD, AND J. D. MURRAY, *A quantitative model for differential motility of gliomas in grey and white matter*, Cell Prolif., 33 (2000), pp. 317–330.
- [61] K. R. SWANSON, E. C. ALVORD, AND J. D. MURRAY, *Virtual brain tumours (gliomas) enhance the reality of medical imaging and highlight inadequacies of current therapy*, British J. Cancer, 86 (2002), pp. 14–18.
- [62] K. R. SWANSON, R. C. ROSTOMILY, AND E. C. ALVORD, *A mathematical modelling tool for predicting survival of individual patients following resection of glioblastoma: A proof of principle*, British J. Cancer, 98 (2008), pp. 113–119.
- [63] P. TRACQUI, G. CRUYWAGEN, D. WOODWARD, G. BARTOO, J. MURRAY, AND E. ALVORD, *A mathematical model of glioma growth: The effect of chemotherapy on spatio-temporal growth*, Cell Prolif., 28 (1995), pp. 17–31.
- [64] A. TROUVE, *A infinite dimensional Group Approach for Physics Based Models in Pattern Recognition*, Technical report, Laboratoire d'Analyse Numerique CNRS URA, Universite Paris, Paris, 1995.
- [65] A. TROUVE, *Diffeomorphism groups and pattern matching in image analysis*, Int. J. Comput. Vis., 28 (1998), pp. 213–221.
- [66] *PnetCDF: A Parallel I/O Library for NetCDF File Access*, <https://trac.mcs.anl.gov/projects/parallel-netcdf>.
- [67] K. C. L. WONG, R. M. SUMMERS, E. KEBEBEW, AND J. YAO, *Pancreatic tumor growth prediction with multiplicative growth and image-derived motion*, in Information Processing in Medical Imaging, Lecture Notes in Comput. Sci., 9123, Springer, Cham, Switzerland, 2015, pp. 501–513.
- [68] K. C. L. WONG, R. M. SUMMERS, E. KEBEBEW, AND J. YOA, *Pancreatic tumor growth prediction with elastic-growth decomposition, image-derived motion, and FDM-FEM coupling*, IEEE Trans. Med. Imaging, 36 (2017), pp. 111–123.
- [69] T. YANKEELOV, N. ATUEGWU, D. HORMUTH, J. WEIS, S. BARNES, M. MIGA, E. RERICH, AND V. QUARANTA, *Clinically relevant modeling of tumor growth and treatment response*, Sci. Transl. Med., 5 (2013), 187ps9.
- [70] L. YOUNES, *Shapes and Diffeomorphisms*, Springer, Heidelberg, Germany, 2010.
- [71] E. I. ZACHARAKI, C. S. HOGEE, G. BIROS, AND C. DAVATZIKOS, *A comparative study of biomechanical simulators in deformable registration of brain tumor images*, IEEE Trans. Biomed. Eng., 55 (2008), pp. 1233–1236.
- [72] E. I. ZACHARAKI, C. S. HOGEE, D. SHEN, G. BIROS, AND C. DAVATZIKOS, *Parallel optimization of tumor model parameters for fast registration of brain tumor images*, in Proceedings of SPIE Medical Imaging: Image Processing, SPIE, Bellingham, WA, 2008, pp. 69140K1–69140K10.
- [73] E. I. ZACHARAKI, C. S. HOGEE, D. SHEN, G. BIROS, AND C. DAVATZIKOS, *Non-diffeomorphic registration of brain tumor images by simulating tissue loss and tumor growth*, NeuroImage, 46 (2009), pp. 762–774.