Published in final edited form as:
Comput Methods Appl Mech Eng. 2019 April 15; 347: 533–567. doi:10.1016/j.cma.2018.12.008.

Coupling brain-tumor biophysical models and diffeomorphic image registration

Klaudius Scheufele[a,*], **Andreas Mang**[b], **Amir Gholami**[d], **Christos Davatzikos**[e], **George Biros**[c], **Miriam Mehl**[a]

Andreas Mang: andreas@math.uh.edu; Amir Gholami: amirgh@berkeley.edu; Christos Davatzikos: christos.davatzikos@uphs.upenn.edu; George Biros: biros@ices.utexas.edu; Miriam Mehl: miriam.mehl@ipvs.uni-stuttgart.de

[a]University of Stuttgart, IPVS, Universitätstraße 38, 70569 Stuttgart, Germany

[b]University of Houston, Department of Mathematics, 3551 Cullen Blvd., Houston, TX 77204-3008, USA

[c]University of Texas, ICES, 201 East 24th St, Austin, TX 78712-1229, USA

[d]University of California Berkeley, EECS, Berkeley, CA 94720-1776, USA

[e]Department of Radiology, University of Pennsylvania School of Medicine, 3700 Hamilton Walk, Philadelphia, PA 19104, USA

Abstract

We present SIBIA (Scalable Integrated Biophysics-based Image Analysis), a framework for joint image registration and biophysical inversion and we apply it to analyze MR images of glioblastomas (primary brain tumors). We have two applications in mind. The first one is normal-to-abnormal image registration in the presence of tumor-induced topology differences. The second one is biophysical inversion based on single-time patient data. The underlying optimization problem is highly non-linear and non-convex and has not been solved before with a gradient-based approach.

Given the segmentation of a normal brain MRI and the segmentation of a cancer patient MRI, we determine tumor growth parameters and a registration map so that if we "grow a tumor" (using our tumor model) in the normal brain and then register it to the patient image, then the registration mismatch is as small as possible. This "*coupled problem*" two-way couples the biophysical inversion and the registration problem. In the image registration step we solve a large-deformation diffeomorphic registration problem parameterized by an Eulerian velocity field. In the biophysical inversion step we estimate parameters in a reaction–diffusion tumor growth model that is formulated as a partial differential equation (PDE). In SIBIA, we couple these two sub-components in an iterative manner. We first presented the components of SIBIA in *"Gholami et al., Framework for Scalable Biophysics-based Image Analysis, IEEE/ACM Proceedings of the SC2017"*, in which we derived parallel distributed memory algorithms and software modules for the *decoupled* registration and biophysical inverse problems. In this paper, our contributions are the introduction of a PDE-constrained optimization formulation of the coupled problem, and the

---

[*]Corresponding author. klaudius.scheufele@ipvs.uni-stuttgart.de (K. Scheufele).

derivation of a Picard iterative solution scheme. We perform extensive tests to experimentally assess the performance of our method on synthetic and clinical datasets. We demonstrate the convergence of the SIBIA optimization solver in different usage scenarios. We demonstrate that using SIBIA, we can accurately solve the coupled problem in three dimensions (256^3 resolution) in a few minutes using 11 dual-x86 nodes.

Keywords

49K20; 49M15; 35K57; 65K10; 68W10; Biophysically constrained diffeomorphic image registration; Tumor growth; Atlas registration; Adjoint-based methods; Parallel algorithms

1. Introduction

We present the SIBIA framework that comprises a mathematical formulation, algorithms, and software for joint image registration and biophysical inversion. Given a volumetric segmentation of a magnetic resonance imaging (MRI) dataset of a glioma (primary brain tumor) patient, SIBIA simultaneously registers this image to a segmented MRI dataset of a normal brain (an atlas) and fits a biophysical tumor growth model. Such a joint registration–biophysical inversion approach is motivated by two practical problems: (1) Often, we need to calibrate complex macroscopic biophysical PDE models using a single-time snapshot dataset. Since the biophysical model is typically a dynamical system, it is unclear how to fit model parameters using just a single-time dataset since we do not have the initial state. Replacing the missing second patient snapshot by a healthy atlas and coupling registration between patient and atlas with the tumor growth model attempts to resolve this problem. (2) We want to register normal MR images to MR images with abnormalities. This normal-to-abnormal registration finds applications in surgical planning (e.g., mapping structural and functional information to a patient image), longitudinal studies, followups, or population studies.

We focus on applying SIBIA to the registration of segmented normal images to segmented MR images of glioma patients. Many image analysis workflows for gliomas involve fitting biophysical models for tumor characterization and prognosis [1,2], for image registration [3–5], and for image segmentation [6–8]. In many methods for these image analysis and biophysical modeling problems, an essential step is the solution of an inverse tumor-growth problem, i.e., for instance the calculation of the initial conditions for tumor growth; this process can be combined with registration and segmentation as in [7].

1.1. Statement of the problem and summary of the approach

The specific scenario we are interested in is summarized in Fig. 1. We are given a single snapshot of the multimodal MRI of a glioma patient. Let us assume that we have obtained a segmentation of the MRI in healthy tissue and abnormal tissue, i.e., we have labels for white matter, gray matter, cerebrospinal fluid (CSF), cerebellum, ventricles, and tumor.[1] Since the

---

[1] In reality, the tumor has subclasses like edema, necrotic tumor, enhancing tumor, and others. Since we are using a very simple tumor model, we do not consider these subclasses here.

tumor-growth problem is typically a time dependent problem, we need at least two time snapshots to invert for parameters. Ideally we would like to have both initial and final conditions of the tumor in order to infer parameters. That is, we would like an image of the healthy brain of the patient, the white, gray, and CSF structure right before the tumor appeared, which of course we do not have. One way for partially tackling this problem is to use a pre-segmented normal healthy brain (atlas). This is the step that requires registration. The best way to describe this is to consider the *"forward problem"*: Assume we are given the (healthy) atlas, a tumor-growth model, and a deformation map (represented as advection with a given velocity field). Then, we first apply the tumor growth model to the healthy image and produce a new segmentation that has *both* healthy tissue and tumor. In a second step, we warp this image using the given deformation map. This image is the tumor-plus-deformation warped atlas. In the *"inverse problem"*, we seek to compute the tumor-growth parameters and deformation parameters so that the tumor-plus-deformation warped atlas image matches the patient image.

Thus, SIBIA combines the following steps: (i) a forward tumor solver growing a virtual (or synthetic) tumor in the segmented healthy atlas brain, (ii) the inverse tumor solver determining the initial condition for tumor growth given an observation at a later time, (iii) a forward linear advection solver that implicitly computes the warped image given the velocity field, and (iv) a large-deformation diffeomorphic image registration solver. We couple tumor inversion and image registration (both are optimization problems) to an outer, *coupled* optimization problem. Because both tumor models and registration formulations are always in flux and in varying complexity, tightly coupling them using a Newton method would require continuous significant software engineering effort. Our goal is a scheme that allows leveraging of the underlying optimization solvers and allows more flexibility. Motivated by this need for model flexibility, we consider the viability of a Picard (or fixed point) iteration.

1.2. Contributions

i. First, we formulate the coupled tumor growth and registration optimization problem in Section 3 and derive its first order optimality conditions and a *reduced gradient* of the objective function with respect to the tumor parameters and the velocity. This reduced gradient is used in our solver scheme as a convergence and correctness check. To our knowledge this is a new result.

ii. We propose a Picard iteration including a parameter-continuation scheme for solving the system based on modular tumor and image registration components. This allows great flexibility in both the type of tumor models and registration but also in the underlying solvers for the two subproblems.

iii. We conduct a large number of numerical experiments on synthetic and real data that demonstrate the validity and efficiency of our approach, in particular, we show that the Picard iteration efficiently reduces the gradient and converges to a local minimum. We conduct a mesh-dependence analysis in the convergence rate in which we show that our scheme is not sensitive to the mesh size (for both velocity and tumor parameters). We consider six clinical images that include multifocal tumors and we demonstrate the effectiveness of the method.

iv. We examine the sensitivity of our solver on the choice of the tumor model variant. We consider three variants of the tumor model: a reaction model, a reaction–diffusion model, and a simple radial basis approximation (no time evolution; the tumor is simply approximated at the observation time). We conclude on the usefulness of each of these variants depending on the goal of the analysis.

1.3. Limitations and open issues

Several limitations and unresolved issues remain. One limitation is that we do not provide a general theoretical proof that the Picard scheme converges. We discuss this further in Section 4.1. A second limitation is the fact that we are not using a Newton method. Indeed, we do not claim that our algorithm is algorithmically optimal. It is quite possible that a Newton method for the coupled problem would perform better. However, we present numerical evidence that our Picard is efficient and converges to acceptable accuracy in a mesh independent way (see Section 5). Another important limitation is that our tumor model does not include mass-effect, instead we use a simple phenomenological reaction–diffusion model [9–11]. Our simple model is quantitatively useful for image analysis and tumor characterization [1,12–14], but has limited predictive capabilities. We are currently working on integrating a mass-effect, which is considered critical for characterizing tumor aggressiveness [3,15,16] as well as more complex reaction–diffusion models. SIBIA can naturally be extended to these more complex models although of course more tests will be necessary to demonstrate convergence of the Picard scheme.

1.4. Literature review

We presented fast algorithms for the individual components of SIBIA in [17] where we introduced the key computational kernels (reaction–diffusion, advection, and inverse solution based on FFT and particle-in-mesh methods) and demonstrated their scalability on very large images and distributed-memory architectures. Our problem formulation results in a mixed-type PDE-constrained optimization problem that poses significant numerical challenges. We refer to [18–21] for a general introduction into PDE-constrained optimization and to [22–24] for reviews on its application to medical image analysis. We will limit ourselves to the work most closely related to ours in the following.

The first component of SIBIA is image registration. We refer to [25,26] for a general overview of medical image registration. Registering the atlas to the patient requires finding correspondences between two topologically different images — one with tumor and one without tumor. The key issue here is the ill-defined correspondence arising from the presence of a tumor in only one of the images to be registered. A simple strategy to deal with this issue is to consider the tumor area as non-informative and mask it from the optimization [27–29] or to relax the registration in the area affected by the tumor [30]. This, however, yields poor registration quality for tumors with severe tissue deformation due to mass effect. Another strategy is to simultaneously invert for the deformation map and a drift in intensity representing the imaging abnormality associated with the tumor [31,32], which can also be applied to other registration problems with topological differences (for instance, pre- and postoperative image registration [4]). While it may produce acceptable results for the purpose of atlas-based segmentation and registration, it cannot be used in the context of model prediction—which is our ultimate goal. Our prior work on diffeomorphic image registration [33–37] forms the basis for the proposed methodology based on the pioneering work in [38–40] (see [33] for additional references).

The second component of SIBIA is the calibration of brain tumor models to medical imaging data [16,17,41–43]. The associated PDE operator is a parabolic, non-linear, reaction–diffusion equation [10,11]. Despite its phenomenological character, this model is capable of generating simulations that are in good agreement with observations of abnormalities in standard MR imaging data and has been used by other groups besides ours [1,12,14,44–51]. Our implementation features inversion operators for the initial tumor concentration, for diffusivity parameters, or for the growth rate [41]. Related optimal control formulations can, e.g., be found in [12,52–59]. Unlike most existing approaches (with the exception of [60]), our parametrization of the problem [41,42] allows us to invert for multifocal tumors. Numerical methods for solving the tumor calibration problem are based either on point estimate or Bayesian inference methods [14,61]. Bayesian inference characterizes the uncertainty in the model parameters. In this paper, we just focus on methods for point estimates using an adjoint formulation. Our approach can be extended to the Bayesian setting [62].

The integration of biophysical brain tumor simulations with deformable image registration is not new [3,5,7,63–65]. In [3,5,64,65], a purely mechanical model for tumor progression was used. The two key limitations of the work in [3,5,64,65] are (i) that the model is oversimplified; it did not provide the capabilities to generate tumors with complex shapes (ii) that these models do not provide information about the progression and infiltration of cancerous cells into surrounding healthy tissue. The proposed formulation does not share these limitations.

The work that is closest to ours is [4,6,7,63]. In those papers the authors presented a framework for joint segmentation, registration, and tumor modeling. What sets our work apart are (i) the solver (we do not iterate on both control variables simultaneously; we perform a block elimination and iterate resulting in an interleaved optimization on the controls using globalized Newton–Krylov solvers as opposed to derivative free-optimization [3,5,7,59,63–65]); (ii) an efficient parallel implementation [17,35]; (iii) the derivation of the optimality systems for the fully coupled problem; (iv) the parametrization for the initial condition of the tumor, which not only allows us to represent multifocal tumors but also to significantly simplify the PDE constraint without losing segmentation accuracy; (v) and finally the integration with a state-of-the-art algorithm for constrained large deformation diffeomorphic image registration [33–35,37].

1.5. Outline

After introducing our notation in Section 2, we present the coupled formulation and the individual subcomponents (tumor growth) and registration and derive the first order optimality conditions in Section 3. In Section 4, we discuss our main contribution besides the formulation of the coupled problem and its decomposition—a modular iterative Picard solver scheme along with further algorithmic details. We summarize the setup for the experiments in Section 5.1 and present numerical results on both synthetic and clinical data in Section 5. We provide additional details on the numerical experiments for the clinical dataset in the Appendix.

2. Notation

Before presenting the models used for tumor growth and image registration, we summarize the notation used throughout the manuscript.

Segmentation labels and probability maps.

For each healthy brain tissue type, i.e., gray matter (GM), white matter (WM), and cerebrospinal fluid (CSF; which includes the ventricles), we use a separate probability map. We represent these probability maps as a space–time vector field

$$\boldsymbol{m}(\boldsymbol{x}, t) = (m_i(\boldsymbol{x}, t))_{i\,=\,1,\,\ldots,\,3} \in \mathbb{R}^3 \quad \text{with} \quad m_1 = m_{GM}, m_2 = m_{WM}, m_3 = m_{CSF} \tag{1}$$

with $m_i(\boldsymbol{x}, t) \in [0, 1]$ defined on the space–time interval $\Omega \times [0, 1]^3$ with boundary $\partial\Omega$. The domain occupied by brain tissue (healthy or unhealthy) is denoted by $\Omega_B \subset \Omega$. The fourth probability map is $c(\boldsymbol{x}, t) \in [0\ 1]$, also defined in Ω_B; c is the output of the tumor forward simulation and represents the probability to encounter cancerous cells at a location $\boldsymbol{x}$ at time t. For convenience, we also define the space–time domains $U = \Omega \times (0, 1]$ and $\bar{U} = \Omega \times [0, 1)$.

Labels, i.e., characteristic functions calculated based on given threshold values for tumor and healthy tissue types are used only in the evaluation of our results to compute Dice coefficients.

Patient and atlas data.

Data of the actual patient image are marked with a subscript T (template, $\boldsymbol{m}_T$, c_T), data of patient images advected to the atlas domain with a subscript P ($\boldsymbol{m}_P$, c_P), and those of the atlas brain with a subscript A ($\boldsymbol{m}_A$, c_A).

Dependency on time and space.

In most formulations, we do not explicitly include the dependency on the spatial position $\boldsymbol{x}$, but only the time dependency. For instance, $c_A(0)$ denotes the initial tumor probability map defined in the atlas image, whereas $c_A(1)$ is the tumor at time $t = 1$ (solution of the tumor forward problem). This is the point in time associated with the patient image. Probability maps in the atlas image evolve in the non-dimensional (tumor growth) time interval [0, 1]. For diffeomorphic image registration, we introduce a *pseudo-time variable* $t \in [0, 1]$ and invert for a stationary velocity field $\boldsymbol{v}(\boldsymbol{x})$; t does not have a physical meaning; we associate $t = 0$ with the undeformed patient image (image to be registered; template image) and $t = 1$ with its deformed representation. We make more explicit definitions below.

Vector notation.

Given a vector field $\boldsymbol{m} \in \mathbb{R}^3$, we compute $\nabla\boldsymbol{m}\ (\partial_j \boldsymbol{m}_i)_{i,\,j=1,2,3,} \in \mathbb{R}^{3,3}$. That is, given a velocity field $\boldsymbol{v} \in \mathbb{R}^3$, $\nabla\boldsymbol{m}\,\boldsymbol{v} \in \mathbb{R}^3$ indicates a matrix–vector multiplication. The standard scalar product in $\mathbb{R}^3$ is denoted by "·" and the outer product between two vector fields will be denoted by "⊗". In addition, we define the following inner products:

$$\langle \boldsymbol{m}, \widetilde{\mathbf{m}} \rangle_{L^2(\Omega)^3} := \sum_{i=1}^{3} \left\langle m^i, \widetilde{m}^i \right\rangle_{L^2(\Omega)}, \qquad \left\| \boldsymbol{m} \right\|^2_{L^2(\Omega)^3} := \sum_{i=1}^{3} \left\| m^i \right\|^2_{L^2(\Omega)}. \tag{2}$$

3. Formulation

3.1 Formulation of the coupled problem

Our formulation is a coupled PDE-constrained optimization problem. The input data are the (1) tumor-free atlas brain image $\boldsymbol{m}_A(\boldsymbol{x}, 0)$ (which we call *"reference"* image or *"healthy atlas"*) and (2) the tumor-bearing patient image (which we term *"template image"* or *"patient with tumor image"*). The latter includes the brain geometry $\boldsymbol{m}_T(\boldsymbol{x})$ and the patient's tumor $c_T(\boldsymbol{x})$. These images are vector images, i.e., they are probability maps for every tissue type class (white matter, gray matter, cerebrospinal fluid, and tumor). We invert for a stationary velocity field $\boldsymbol{v}(\boldsymbol{x}) := (v^1(\boldsymbol{x}), v^2(\boldsymbol{x}), v^3(x))^\top$ (establishes the spatial correspondence between the patient and atlas image), a parameter vector $\boldsymbol{p}$ (parametrizes the simulated tumor in the tumor-free atlas image), and a mass-source map $\mathrm{w}(\boldsymbol{x})$ (controls the computed deformation pattern) as follows:

$$\min_{\boldsymbol{v}, \boldsymbol{p}, w} \quad \mathcal{J}(v, p, w) \qquad \text{with} \quad \mathcal{J}(v, p, w) := \mathcal{D}_c(c_A, c_p) + \mathcal{D}_{\boldsymbol{m}}(\boldsymbol{m}_A, \boldsymbol{m}_p) + \beta_{\boldsymbol{p}} \mathcal{S}_{\boldsymbol{p}} + \beta_{\boldsymbol{v}} \mathcal{S}_{\boldsymbol{v}} + \beta_w \mathcal{S}_w \tag{3a}$$

subject to

$$\partial_t c_A - \operatorname{div} k \nabla c_A - f(c_A) = 0 \qquad \text{in } U, \tag{3b}$$

$$c_A(0) = \Phi \boldsymbol{p} \qquad \text{in } \Omega, \tag{3c}$$

$$\boldsymbol{m}_A(1) = \boldsymbol{m}_A(0)(1 - c_A(1)) \qquad \text{in } \Omega, \tag{3d}$$

$$\partial_t \boldsymbol{m}_P + \nabla \boldsymbol{m}_P \boldsymbol{v} = 0 \qquad \text{in } U, \tag{3e}$$

$$\boldsymbol{m}_P(0) = \boldsymbol{m}_T \qquad \text{in } \Omega, \tag{3f}$$

$$\partial_t c_P + \nabla c_p \cdot \boldsymbol{v} = 0 \qquad \text{in } U, \tag{3G}$$

$$c_P(0) = c_T \qquad \text{in } \Omega, \tag{3h}$$

$$\operatorname{div} \quad \boldsymbol{v} = w \qquad\qquad \text{in } U, \tag{3i}$$

with periodic boundary conditions on $\partial\Omega$.

Here (3b)–(3d) represent the forward tumor model; and (3e)–(3i) represent the forward registration problem. The objective functional $\mathcal{J}$ in (3b) consists of the following building blocks:

i. the two driving L^2-distance measures

$$\mathcal{D}_c(c_A, c_p) := \frac{1}{2}\left\|c_A(1) - c_P(1)\right\|^2_{L^2(\Omega)}, \qquad \text{and} \qquad \mathcal{D}_{\boldsymbol{m}}(\boldsymbol{m}_A, \boldsymbol{m}_P) := \frac{1}{2}\left\|\boldsymbol{m}_A(1) - \boldsymbol{m}_P(1)\right\|^2_{L^2(\Omega)^3}$$

that measure the discrepancy between the simulated tumor in atlas space $c_A(1)$ (solution of the tumor forward problem (3b) with initial condition (3c) parameterized by $\boldsymbol{p}$) and the transported probability map $c_P(1)$ of cancerous cells for the patient data (solution of the registration forward problem (3g) with initial condition (3h)) and the discrepancy between the healthy tissue probability maps in atlas space with tumor (calculated according to (3d)) and the transported probability maps of healthy tissue for the patient data (solution of the registration forward problem (3e) with initial condition (3f));

ii. three regularization operators balanced against the discrepancy measures $\mathcal{D}_c$ and $\mathcal{D}_{\boldsymbol{m}}$ based on regularization weights $\beta_j > 0$, $j \in \{v, w, p\}$ involving

- **a.** a regularization operator $\mathcal{S}_{\boldsymbol{p}}$ for $\boldsymbol{\Phi p}$ (an L^2-norm),
- **b.** a regularization operator $\mathcal{S}_{\boldsymbol{v}}$ for $\boldsymbol{v}$ (an H^1 Sobolev norm; defined in (10b)), and
- **c.** a regularization operator $\mathcal{S}_w$ for w (an H^1 Sobolev norm; defined in (10b)).

Next we discuss in detail the two forward problems (tumor growth and image advection).

Forward Tumor Problem.—We model the tumor growth based on the population density $c_A(\boldsymbol{x}, t)$, which is obtained by solving (3b)–(3c). Two main phenomena are included: ***proliferation*** of cancerous cells and the ***net migration*** of cancerous cells into surrounding healthy tissue [11].[2] The proliferation model is a logistic growth function $f(c_A) := \rho\, c_A(1 - c_A)$ with reaction coefficient $\rho(\boldsymbol{x}) := \rho_f \rho_0(\boldsymbol{x})$. ρ_f denotes the scaling of the spatially variable characteristic growth rate parameters defined by white and gray matter, i.e., $\rho_0(\boldsymbol{x}) := \rho_w m_{WM}(\boldsymbol{x})\, \rho_g\, m_{GM}(\boldsymbol{x})$. The migration model is based on an inhomogeneous (potentially anisotropic) diffusion process with diffusion coefficient $\boldsymbol{k}(\boldsymbol{x}) := k_f k_0(\boldsymbol{x})\boldsymbol{I} + k_a \boldsymbol{T}(\boldsymbol{x})$. Here, k_f

---

[2] Displacement of healthy tissue as a result of tumor growth (mass effect) is currently neglected in our model.

and k_a are the scaling factors for the isotropic and anisotropic parts of the diffusion tensor, $\boldsymbol{T}(\boldsymbol{x})$ is a weighted diffusion tensor modeling anisotropy. In our test cases in Section 5, we always use isotropic diffusion, i.e., $k_a = 0$. The isotropic part is $k_0(\boldsymbol{x}) := k_w m_{WM}(\boldsymbol{x})\ k_g m_{GM}(\boldsymbol{x})$. This yields a *non-linear* parabolic PDE with non-constant coefficients for the tumor concentration c given by (3b)–(3c).

We use a parametrization $\boldsymbol{\Phi p}$ for the tumor initial condition $c_A(0)$ as originally in (3c) in an $n_{\boldsymbol{p}}$-dimensional space spanned by Gaussian basis functions, i.e., $\boldsymbol{p} \in \mathbb{R}^{n_{\boldsymbol{p}}}, \boldsymbol{\Phi p} := \sum_{i=1}^{n_{\boldsymbol{p}}} \Phi_i p_i$ with Gaussian basis functions $\Phi_i : \Omega_B \to \mathbb{R}$. We set the Gaussians in CSF to zero to prevent a spurious diffusion of cancerous cells into the area associated with CSF. Once the tumor density has been computed we can adjust the probability of the other tissue labels by Eq. (3d).

For notational convenience, we represent the process of solving (3b)–(3d) with the operator

$$\mathscr{T}^{\text{fwd}}(\boldsymbol{p}, \boldsymbol{m}_A(0)) := (\boldsymbol{m}_A(1), c_A(1)), \tag{4}$$

which maps $\boldsymbol{p}$ in (3c) to the tissue labels of the atlas image at time $t = 1$. This simple model is by no means predictive on its own, but is the de-facto standard approach when it comes to modeling tumor progression as seen in medical imaging [1,9,44,46,47,63]. Some results available in the literature have suggested that this model can offer (to some extent) predictive capabilities when integrated with medical imaging information [9,66]. Its usefulness is in segmentation and registration algorithms that use normal atlas information and in producing features (e.g., tumor parameters) to augment image-based features for tumor staging and prognosis. Note that usually the stand-alone tumor model is stated in the actual brain domain Ω_B with Neumann boundary conditions, which is the actual biophysical problem. In our numerical implementation, however, we extend $\boldsymbol{k}$ by a small parameter, set $\rho = 0$ in $\Omega\backslash\Omega_B$, discretize in Ω using periodic boundary conditions, and use a penalty approach to approximate the boundary conditions at $\partial\Omega_B$. One can show that this "fictitious domain method" approximates the correct boundary conditions and as we refine the discretization it converges to the correct solution (compare [41,63]).

Advective Image Transformation (Forward Problem).—The image registration velocity $\boldsymbol{v}$ is used to transport both the brain geometry $\boldsymbol{m}_T$ and the tumor concentration c_T from patient to atlas space. The input for our formulation is not image intensities [33,34,36,37] but the probability maps for tissue classes (see (1); WM, GM, CSF, and tumor). The formulation we propose here is suited for general problems that involve the registration of vector fields. The *template image* (image to be registered) $\boldsymbol{m}_T : \overline{\Omega} \to \mathbb{R}^3$ is given by the probability maps of the patient's healthy anatomy in all areas except the part hidden by the tumor. We treat the probability map for the tumor, $c_T(\boldsymbol{x})$, as an individual entity to make the coupled formulation in (3) more accessible. We register the three probability maps for healthy tissue and the tumor concentrations. Given some template image $\boldsymbol{m}_T(\boldsymbol{x})$, some tumor concentration $c_T(\boldsymbol{x})$, and a *stationary velocity field* $\boldsymbol{v}(\boldsymbol{x})$, the forward problem describes the advective transformation of $\boldsymbol{m}_T$ and c_T in a pseudo-time interval [0, 1) as given in (3e) –(3i).

Solving (3e) –(3i) defines the implicitly given operator

$$\mathscr{R}^{\mathrm{fwd}}(\boldsymbol{v}, \boldsymbol{m}_T, c_T) := (\boldsymbol{m}_P(1), c_P(1)), \tag{5}$$

which maps the template images $\boldsymbol{m}_T$ and c_T to images $\boldsymbol{m}_P$ and c_P defined at pseudo-time t = 1. For simplicity, we will later slightly abuse our notation and use the operator $\mathscr{R}^{\mathrm{fwd}}$ also for the advection of only $\boldsymbol{m}$, c or one of the components of $\boldsymbol{m}$, i.e., m_{WM}, m_{GM}, or m_{CSF}.

3.2. Optimality conditions and fully coupled gradient

We use the method of Lagrange multipliers to transform the constrained problem (3) into an unconstrained one.

The Lagrangian of (3) reads

$$\begin{aligned}
&\mathscr{L}_G(c_A, c_P, \boldsymbol{m}_P, \boldsymbol{m}_A(1), \alpha, \lambda_c, \boldsymbol{\lambda}_{\boldsymbol{m}}, \nu, \boldsymbol{\xi}, \boldsymbol{p}, \boldsymbol{v}, w) = \mathscr{J}(\boldsymbol{p}, \boldsymbol{v}, w) \\
&+ \left\langle \alpha(0), (c_A(0) - \Phi_{\boldsymbol{p}}) \right\rangle_{L^2(\Omega)} + \int_0^1 \left\langle \alpha, \partial_t c_A - \nabla \cdot (\boldsymbol{k} \nabla c_A) - f(c_A) \right\rangle_{L^2(\Omega)} \mathrm{d}t \\
&+ \int_0^1 \left\langle \lambda_{\boldsymbol{m}}, \partial_t \boldsymbol{m}_P + \nabla \boldsymbol{m}_P \boldsymbol{v} \right\rangle_{L^2(\Omega)^3} \mathrm{d}t + \left\langle \lambda_{\boldsymbol{m}}(0), \boldsymbol{m}_P(0) - \boldsymbol{m}_T \right\rangle_{L^2(\Omega)^3} \\
&+ \int_0^1 \left\langle \lambda_c, \partial_t c_P + \nabla c_P \cdot \boldsymbol{v} \right\rangle_{L^2(\Omega)} \mathrm{d}t + \left\langle \lambda_c(0), c_P(0) - c_T \right\rangle_{L^2(\Omega)} \\
&+ \int_0^1 \left\langle \nu, \operatorname{div} \boldsymbol{v} - w \right\rangle_{L^2(\Omega)} \mathrm{d}t + \left\langle \xi, \boldsymbol{m}_A(1) - (1 - c_A(1)) \boldsymbol{m}_A(0) \right\rangle_{L^2(\Omega)^3},
\end{aligned} \tag{6}$$

with the state fields c_A, c_P, $\boldsymbol{m}_P$, and $\boldsymbol{m}_A(1)$, the adjoint fields α, λ_c, $\boldsymbol{\lambda}_{\boldsymbol{m}}$, ν, and $\boldsymbol{\xi}$, and the inversion fields $\boldsymbol{p}$, $\boldsymbol{v}$, and w. The strong form of the ***first-order optimality conditions*** for Eq. (6) is given by the following equations:

$$\mathit{tumor\ forward} \qquad \delta_\alpha \mathscr{L}_G = 0: \qquad \text{given by } (3\text{b}) – (3\text{c}) \tag{7a}$$

$$\delta_\xi \mathscr{L}_G = 0: \qquad \text{given by } (3\text{d}) \tag{7b}$$

$$\mathit{tumor\ ad}\text{joint} \quad \delta_{c_A} \mathscr{L}_G = 0: \qquad -\partial_t \alpha - \nabla \cdot \boldsymbol{k} \nabla \alpha - \alpha\rho + 2\alpha\rho c_A = 0 \quad \text{in } \bar{U}, \tag{7c}$$

$$\delta_{c_A(1)} \mathscr{L}_G = 0: \qquad c_P(1) - c_A(1) - \boldsymbol{\xi} \boldsymbol{m}_A(0) - \alpha(1) = 0 \qquad \text{in } \Omega\,. \tag{7d}$$

$$\delta_{\boldsymbol{m}A(1)}\mathscr{L}_G = 0: \qquad \boldsymbol{m}_P(1) - \boldsymbol{m}_A(1) - \boldsymbol{\xi} = \boldsymbol{0} \qquad \text{in } \Omega, \tag{7e}$$

$$\textit{registration forward} \qquad \delta_{\lambda}\mathscr{L}_G = 0: \qquad \text{given by } \big(\text{3e}\big)\text{–}(\text{ 3i}\big) \tag{7f}$$

$$\textit{registration adjoint} \qquad \delta_{\boldsymbol{m}P}\mathscr{L}_G = 0: \qquad -\partial_t\boldsymbol{\lambda}_{\boldsymbol{m}} - \operatorname{div}(\boldsymbol{\lambda}_{\boldsymbol{m}} \otimes \boldsymbol{v}) = \boldsymbol{0} \text{in } \bar{U}, \tag{7g}$$

$$\delta_{\boldsymbol{m}P(1)}\mathscr{L}_G = 0: \qquad \boldsymbol{m}_A(1) - \boldsymbol{m}_P(1) - \boldsymbol{\lambda}_{\boldsymbol{m}}(1) = \boldsymbol{0} \qquad \text{in } \Omega, \tag{7h}$$

$$\delta_{c_P}\mathscr{L}_G = 0: \qquad -\partial_t\lambda_c - \operatorname{div}(\lambda_c\boldsymbol{v}) = 0 \quad \text{in } \bar{U}, \tag{7i}$$

$$\delta_{cP(1)}\mathscr{L}_G = 0: \qquad c_A(1) - c_P(1) - \lambda_c(1) = 0 \quad \text{in } \Omega, \tag{7j}$$

$$\textit{tumor inversion} \qquad \delta_{\boldsymbol{p}}\mathscr{L}_G = 0: \qquad \nabla_{\boldsymbol{p}}\mathcal{S}_{\boldsymbol{P}}(\boldsymbol{P}) - \Phi^{\mathrm{T}}\alpha(0) = 0 \quad \text{in } \Omega, \tag{7k}$$

$$\textit{registration inversion} \qquad \delta_{\boldsymbol{v}}\mathscr{L}_G = 0: \qquad \nabla_{\boldsymbol{v}}\mathcal{S}_{\boldsymbol{v}}(\boldsymbol{v}) + \mathscr{K}\Big[\int_0^1 (\nabla\boldsymbol{m}_P)^T\boldsymbol{\lambda}_{\boldsymbol{m}} + \nabla c_P\lambda_c \mathrm{d}t\Big] = \boldsymbol{0}\,. \tag{7l}$$

Note, that $(\nabla\boldsymbol{m})^T\boldsymbol{\lambda} = \sum_i^3 \lambda_i\big(\partial_1 m_i, \partial_2 m_i, \partial_3 m_i\big)^T$ and that we have not given an equation for w. The operator $\mathscr{K}$ in (7l) is a pseudo-differential operator that is derived by eliminating w in (3i). We discuss more details below.

In summary, the first-order optimality conditions (7) define the gradient of the Lagrangian and comprise a system of non-linear partial differential equations, which is quite formidable since it has 11 fields in addition to the tumor initial condition parameters $\boldsymbol{p}$. Though we later decompose our fully coupled problem into sub-components, the gradient is required to assess convergence of our modular Picard scheme. Instead of using the full gradient of the Lagrangian given by (7), we use an elimination method (or also called a *"reduced-space method"* [19]) to solve these equations. That is, we solve *forward* and *adjoint equations* exactly and evaluate the gradient of the Lagrangian with respect to the inversion variables $\boldsymbol{v}$ and $\boldsymbol{p}$[3]: Given $\boldsymbol{p}$ and $\boldsymbol{v}$, we wish to compute the gradient of $\mathcal{J}$, i.e., $\boldsymbol{g}_p$ and $\boldsymbol{g}_v$. This ***gradient computation*** involves the following steps:

---

[3] Note that, in our algorithm, w is eliminated.

1. Solve the forward tumor growth and registration equations (3c)–(3d) for the state variables $c_A(t)$, $c_P(t)$, $\boldsymbol{m}_P(t)$, $\boldsymbol{m}_A(1)$.
2. Compute the coupling adjoint variable ξ from Eq. (7e).
3. Solve the adjoint tumor equations (7c) and (7d) for $\alpha(t)$.
4. Solve the adjoint registration equations (7g)–(7j) for $\boldsymbol{\lambda}_{\boldsymbol{m}}(t)$ and $\lambda_c(t)$.
5. Evaluate the gradients using the inversion equations (7k) and (7l) at $\boldsymbol{v}$ and $\boldsymbol{p}$:

$$\boldsymbol{g}_{\boldsymbol{v}} = \beta_{\boldsymbol{v}} \nabla_{\boldsymbol{v}} \mathcal{S}_{\boldsymbol{v}}(\boldsymbol{v}) + \mathscr{K}[\int_0^1 (\nabla \boldsymbol{m}_P)^T \boldsymbol{\lambda}_{\boldsymbol{m}} + \nabla c_P \lambda_c \mathrm{d}t], \qquad \boldsymbol{g}_{\boldsymbol{p}} = \beta_{\boldsymbol{p}} \nabla_{\boldsymbol{p}} \mathcal{S}_{\boldsymbol{p}}(\boldsymbol{p}) - \Phi^{\mathrm{T}} \alpha(0)\,. \tag{8}$$

At a stationary point $\boldsymbol{g} = (\boldsymbol{g}_{\boldsymbol{v}}, \boldsymbol{g}_{\boldsymbol{p}})^T$ vanishes.

3.3. Decomposition into subcomponents — inverse tumor and image registration

In the following, we formulate separate inverse problems for tumor growth and image registration that are used as high-level components in our coupling algorithm presented in 4.1. We show that the optimality conditions of the sub-problems together yield the optimality conditions of the coupled problem. This observation is the basis for our Picard iteration using tumor and registration as modular components instead of a gradient descent or Newton method with line search for the coupled problem (3).

Inverse Tumor Problem.—In the inverse tumor problem, we seek an initial condition for the forward tumor problem that recovers $c_P(1)$ and $\boldsymbol{m}_P(1)$ at time $t = 1$ as well as possible, i.e., we solve the minimization problem

$$\min_{\boldsymbol{p}} \mathcal{J}_T(\boldsymbol{p}) \quad \text{with} \quad \mathcal{J}_{\mathrm{T}}(\boldsymbol{p}) :=\mathcal{D}_c(c_A, c_P) + \mathcal{D}_m(\boldsymbol{m}_A, \boldsymbol{m}_P) + \frac{\beta_p}{2}\|\Phi \boldsymbol{p}\|_2^2 \tag{9a}$$

subject to the constraints (3b)–(3d). Eq. (9) defines a concrete instance for the tumor regularization operator $\mathcal{S}_{\boldsymbol{p}}$ in (3).

The Lagrangian of (9) reads

$$\mathcal{L}_T(c_A, \boldsymbol{p}, \alpha) = \mathcal{J}_T(\boldsymbol{p}) + \left\langle \alpha(0), c_A(0) - \Phi \boldsymbol{p} \right\rangle_{L^2(\Omega_{\mathrm{B}})} + \int_0^1 \left\langle \alpha, \partial_t c_A - \nabla \cdot \boldsymbol{k} \nabla c_A - \rho c_A (1 - c_A) \right\rangle_{L^2(\Omega_{\mathrm{B}})} \mathrm{d}t$$

The first-order optimality condition of the tumor problem are given by (3b) –(3d), (7c)–(7e), and (7k).

Image Registration (Inverse Problem).—We seek a velocity $\boldsymbol{v}$ that advects the given template (patient) images to images that are as close as possible to the corresponding images in the atlas brain (denoted with a subscript A). That is, we solve the following minimization problem:

$$\min_{\boldsymbol{v}} \mathscr{J}_R(\boldsymbol{v}) \quad \text{with} \quad \mathscr{J}_{\mathrm{R}}(\boldsymbol{v}) := \mathscr{D}_{\boldsymbol{m}}(\boldsymbol{m}_A, \boldsymbol{m}_P) + \mathscr{D}_c(c_A, c_P) + \beta_{\boldsymbol{v}} \mathscr{S}_{\boldsymbol{v}}(\boldsymbol{v}) + \beta_w \mathscr{S}_w(w) \tag{10a}$$

subject to (3e)–(3i). The regularization in (10a) is a *smoother* for $\boldsymbol{v}$ and w and given by an H^1-seminorm for $\boldsymbol{v}$, and an H^1-norm for w (i.e., for div $\boldsymbol{v}$ according to (3i)), respectively:

$$\mathscr{S}_{\boldsymbol{v}}(\boldsymbol{v}) = \frac{1}{2} \int_{\Omega} \sum_{i=1}^{3} \left| \nabla v^i(\boldsymbol{x}) \right|^2 \mathrm{d}\Omega, \qquad \mathscr{S}_w(w) = \frac{1}{2} \int_{\Omega} |\nabla w(\boldsymbol{x})|^2 \mathrm{d}\Omega . \tag{10b}$$

The *Lagrangian* for the image registration problem reads

$$\begin{aligned} \mathscr{L}_R(c_P, \boldsymbol{m}_P, \boldsymbol{v}, \boldsymbol{\lambda}_{\boldsymbol{m}}, \lambda_c) = \mathscr{J}_R(\boldsymbol{v}) & \\ & + \int_0^1 \left\langle \boldsymbol{\lambda}_{\boldsymbol{m}}, \partial_t \boldsymbol{m}_P + \nabla \boldsymbol{m}_P \boldsymbol{v} \right\rangle_{L^2(\Omega)^3} \mathrm{d}t + \left\langle \boldsymbol{\lambda}_{\boldsymbol{m}}(0), \boldsymbol{m}_P(0) - \boldsymbol{m}_T \right\rangle_{L^2(\Omega)^3} \\ & + \int_0^1 \left\langle \lambda_c, \partial_t c_P + \nabla c_P \cdot \boldsymbol{v} \right\rangle_{L^2(\Omega)} \mathrm{d}t + \left\langle \lambda_c(0), c_P(0) - c_T \right\rangle_{L^2(\Omega)} \\ & + \int_0^1 \langle \nu, \operatorname{div} \ \boldsymbol{v} - w \rangle_{L^2(\Omega)} \mathrm{d}t . \end{aligned}$$

Taking variations, we obtain the first-order optimality conditions [34] for the registration problem as (3e)–(3i), (7g)–(7j), and (7l). As mentioned before, the operator $\mathscr{K}$ in (7l) is a pseudo-differential operator that is derived by eliminating w in (3i). For the case of exact incompressibility ($w = 0$), it is the Leray projection $\mathscr{K}(u) := u + \nabla \Delta^{-1}$ div $\boldsymbol{u}$; for a non-zero w, the projection operator becomes slightly more involved; we refer to [33,34] for additional details.

Next, we discuss the Picard iteration and the numerical discretization of the forward and adjoint problems.

4. Numerical ingredients — Picard and sub-components

The focus of the present work is on the formulation of the coupled problem and the numerical scheme used to solve the associated optimization problem. The algorithms for the individual subblocks have been published in [17,33–35,37,41,67]. The main ingredients can be described as follows: (i) All PDEs are spatially discretized in $\Omega = [0, 2\boldsymbol{\pi}]^3$. (ii) All spatial derivatives (∇, div , and higher derivatives) are computed using 3D Fourier transforms. (iii) Although in the formulation we present the derivatives in the strong form, in our implementation we use a discretize-then-optimize approach for the tumor equations and an optimize-then-discretize approach for the registration. (iv) The solution of pure advection equations is done using a semi-Lagrangian time-stepping scheme to avoid stability issues and small time-steps.[4] (v) We use Krylov and matrix-free Newton methods for linear and nonlinear solvers. (vi) We use a Picard iteration scheme for the coupled optimization problem, without line search.[5]

In the following subsections, we give more details on the proposed Picard iteration scheme, and then give a short overview of the numerical methods used to solve the tumor and image registration forward and inverse problems.

4.1. The Picard iteration algorithm

Here we discuss the solution of (3). As we discussed in the introduction, we use a modular approach in which we combine tumor growth inversion and diffeomorphic registration models in an (interleaved) Picard iteration scheme. This scheme iteratively improves both the tumor initial condition's parametrization $\boldsymbol{p}$ and the registration velocity $\boldsymbol{v}$. This allows us to establish a coupling of both components in an easy and efficient way; the solver subcomponents can be exchanged as required. Our results presented in Section 5 demonstrate that the Picard iteration is a powerful approach that reduces the gradient of the fully coupled formulation as given in (8).

The outline of our Picard scheme is simple.[6] It is convenient to restate (3) in terms of the compact notation for the forward operators given by (4) and (5). With this notation the optimization problem is given by

$$\min_{\boldsymbol{v},\boldsymbol{p}} \mathscr{D}_c(c_A(1), c_P(1)) + \mathscr{D}_{\boldsymbol{m}}(\boldsymbol{m}_A(1), \boldsymbol{m}_P(1)) + \beta_{\boldsymbol{p}}\mathscr{S}_{\boldsymbol{p}} + \beta_{\boldsymbol{v}}\mathscr{S}_{\boldsymbol{v}}$$

$$\text{s.t.:} \left(\boldsymbol{m}_A(1), c_A(1)\right) = \mathscr{T}^{\text{fwd}}(\boldsymbol{p}, \boldsymbol{m}_A(0)), \quad (\boldsymbol{m}_P, c_P(1)) = \mathscr{R}^{\text{fwd}}(\boldsymbol{v}, \boldsymbol{m}_T, c_T).$$

One step in our Picard iteration is summarized as follows.

1. Given $\boldsymbol{m}_p$, c_P, solve the inverse tumor problem for $\boldsymbol{p}$, $\boldsymbol{m}_A$, and c_A:

$$min\, \mathscr{D}_c\left(c_A(1), c_P(1)\right) + \mathscr{D}_{\boldsymbol{m}}\left(\boldsymbol{m}_A(1), \boldsymbol{m}_P(1)\right) + \mathscr{S}_{\boldsymbol{p}}, \quad s.t.:(\boldsymbol{m}_A(1), c_A(1)) = \tau^{\text{fwd}}\left(\boldsymbol{p}, \boldsymbol{m}_A(0)\right).$$

2. Given c_A, $\boldsymbol{m}_A$, solve the image registration problem for $\boldsymbol{v}$, $\boldsymbol{m}_P$ and c_P:

$$min\, \mathscr{D}_c\left(c_A(1), c_P(1)\right) + \mathscr{D}_{\boldsymbol{m}}\left(\boldsymbol{m}_A(1), \boldsymbol{m}_P(1)\right) + \mathscr{S}_{\boldsymbol{v}}, \quad s.t.:\left(\boldsymbol{m}_P(1), c_P(1)\right) = \mathscr{R}^{\text{fwd}}(\boldsymbol{v}, \boldsymbol{m}_T, c_T).$$

This iteration defines new iterates for $\boldsymbol{p}$ and $\boldsymbol{v}$. Here we just described the two inverse problems for the tumor initial condition and registration velocity in a compact form. The precise details are given in Section 3.3. Since the individual optimization problems are

---

[4] This accelerates the solver significantly, but the adjoint and forward solvers are not consistent, that is the discretized gradient is not the gradient of the discretized forward problem.

[5] The Newton-type iterations for the sub-problems registration and tumor inversion are enhanced with a globalizing line search method.

[6] Complications arise due to parameter continuation and inexact solves. We discuss these later.

subcomponents of the original optimization and $\mathcal{S}_{\boldsymbol{p}}$ and $\mathcal{S}_{\boldsymbol{v}}$ are independent from the respective other subsolver's result, each Picard iteration is guaranteed to reduce the objective function of (3). However, this does not imply convergence since the Picard reductions may stagnate. Our numerical experiments indicate however that this is not the case and the scheme is convergent. Our Picard scheme can be interpreted (and eventually analyzed) as a block Gauss–Seidel type iteration over the optimality system (7) of the coupled problem. The reduced gradient's norm of (3) can be interpreted as a residual norm after such an iteration over $\boldsymbol{p}$ and $\boldsymbol{v}$.

Now, let us give more details on the initialization of the Picard iteration. We start with an initial guess for $\boldsymbol{p} \in \mathbb{R}^{n_p}$ and $\boldsymbol{v}$ (both zero), compute the center of mass of the given patient tumor c_T and place the Gaussian basis functions for the tumor initial condition parametrization as a regular grid around the center of mass of the tumor in patient space. In our current algorithm, the grid of Gaussian basis functions is fixed throughout the Picard iterations, i.e., we do not re-compute the center of mass of the tumor or advect the Gaussian basis functions with the velocity $\boldsymbol{v}$.

Let us also mention some additional implementation details that significantly affect the performance of the method: (i) We use the solution $\boldsymbol{v}$ from the previous Picard iteration as an initial guess for the new registration iteration to reduce the run time (warm start). (ii) We perform a continuation in the regularization parameter $\beta_{\boldsymbol{v}}$ (we start with a large value and successively reduce it). We explain this below. (iii) We do not have a proof for the convergence of the proposed Picard scheme to a minimizer of the fully-coupled optimization problem in (3). Such a proof is beyond the scope of the present paper and remains for future work. We provide numerical evidence that shows that our scheme reduces the gradient of the fully-coupled problem.

Let us remark that we have implemented the gradient of the coupled problem in Section 3 in order to verify that our Picard iteration actually reduces the gradient of the global optimization problem. An implementation of a scheme that iterates simultaneously on both control variables (i.e., solves the global problem), requires more work and will be addressed in a follow-up paper. We show experimentally in various settings that our algorithm is effective and generates registration results that are in excellent agreement with the patient data. Next, we give additional details on the numerical methods used in SIBIA.

Parameter continuation.—Eq. (3) is a nonlinear, non-convex, and ill-conditioned optimization problem that involves a system of reaction–diffusion and pure advection vector equations. The joint forward problem alone has eight unknowns per grid point. The joint inversion problem has twenty unknowns in four dimensions. The computational costs can be quite significant. Furthermore, due to the ill-conditioning and non-convexity, the solver can stagnate. To accelerate convergence, we employ a parameter-continuation scheme to convexify the problem and reduce the chance of getting trapped in local minima. Extensive experimental studies showed that it is sufficient to use parameter continuation for the registration regularization parameter $\beta_{\boldsymbol{v}}$. We use a Tikhonov-type regularization (additional details can be found in [33]) and start with β_v^0 in the first Picard iteration. In each iteration k, the candidate β_v^k is reduced by one order of magnitude until the specified lower bound for the determinant of the deformation gradient is breached for a candidate β_v^k. The registration solver disregards the associated solution, sets the candidate value for β_v^k to $\beta_v^k \leftarrow \beta_v^{k-1} - \left(\beta_v^{k-1} - \beta_v^k\right)/2$ and restarts the inversion with the velocity of the former Picard iteration as initial guess. This process is repeated until the lower bound for the determinant of the deformation gradient is no longer violated. If no violation of the determinant of the deformation gradient was detected during the registration solve, we finalize the current Picard iteration and proceed with the next iteration until we reach the specified smallest admissible value for $\beta_{\boldsymbol{v}}$. Besides globalization, this continuation scheme stabilizes the registration problem and automatically identifies an adequate regularization parameter $\beta_{\boldsymbol{v}}$ for the H^1-Sobolev norm for $\boldsymbol{v}$ in every Picard iteration.

Stopping conditions.—Since we can compute the reduced gradient and the objective function we could, in principle, use standard optimization stopping criteria [68]. However, for practical reasons we use a simpler scheme. Since we use parameter continuation, we finalize our Picard iteration either when $\beta_{\boldsymbol{v}}$ reaches the prescribed minimum or when we reach a user-defined lower bound for the determinant of the deformation gradient. The lower bound for $\beta_{\boldsymbol{v}}$ is a safeguard against numerical instabilities that can occur if $\boldsymbol{v}$ becomes highly irregular (see [37] for a discussion). In both cases we execute two additional Picard iterations with the final value for $\beta_{\boldsymbol{v}}$, instead of iterating until convergence. This scheme leads to more predictable behavior in terms of wall-clock time. Note that the reduced gradient in (7) is computed using an optimize-then-discretize approach and is consistent with the objective function only up to discretization errors.

4.2. Numerics for the tumor inversion and registration sub-blocks

The optimality conditions of both (9a) and (10a) are complex, multi-component, non-linear operators for the state, adjoint, and control fields. We employ an inexact, globalized, preconditioned Gauss–Newton–Krylov method for both problems. In reduced space methods, the reduced space Hessian system has a very similar structure as the first order optimality conditions; we can employ the same numerical strategies we use for the solution of the PDE operators in the first order optimality conditions (see below). We refer to [24,33,37,41] for details on the reduced space Hessian system and its discretization.

Numerical solution and discretization of the PDE operators.—To discretize forward and adjoint tumor and registration problems in space and time, we use regular grids consisting of $N_0 \times N_1 \times N_2, \times N_i \in \mathbb{N}$ grid points. For all spatial differential operators, we use a spectral projection scheme as described in [17,35]. The mapping between the spatial and spectral coefficients is done using FFTs (the implementation of the parallel FFT library is presented in [67]). Corresponding to the spectral collocation scheme, we assume that the functions in our formulation (including images) are periodic and continuously differentiable and apply filtering operations and periodically extensions of the discrete data to meet these requirements.

To solve the forward and adjoint tumor problem, we use an unconditionally stable, second-order Strang-splitting method, where we split the right-hand side of (3b) and (7c) into diffusion and reaction (see [41,63] for details). The diffusion sub-steps are solved using an implicit Crank–Nicolson method. The solver is a preconditioned conjugate gradient method with a fixed tolerance of 1E−6. The reaction sub-steps are solved analytically. We enforce a positivity constraint on the initial tumor concentration before we apply the forward operator to make sure that $c(\boldsymbol{x}, t)$ is in [0, 1] for all $\boldsymbol{x} \in \Omega$ and $t \in [0, 1]$. This thresholding operation is necessary since the parametrization allows for negative concentrations for the initial condition $c_A(0)$ (the coefficient vector $\boldsymbol{p}$ can have negative entries) and the logistic growth function would amplify these negative values in time.

In the image registration module, we solve the hyperbolic transport equations in (3g), (7g), and (7i) based on a semi-Lagrangian scheme [35,37]. Semi-Lagrangian schemes are a hybrid between Lagrangian and Eulerian schemes. They are unconditionally stable and, like Lagrangian schemes (see [36] for an example in the context of diffeomorphic registration), require the evaluation of the space–time fields that appear in our optimality conditions at off-grid locations defined by the characteristic associated with $\boldsymbol{v}$. We compute the values of these space–time fields at the off-grid locations with a cubic Lagrange polynomial interpolation model. The time integration for computing the characteristic and for the solution of ordinary differential equations along it (if necessary) is based on a second order accurate Runge–Kutta scheme. We refer to [35,37] for a precise definition of these computations.

Newton–Krylov solver.—We initialize our solvers for both reduced gradient systems (7k) and (7l) with a zero initial guess and use an inexact, globalized, preconditioned Gauss–Newton–Krylov method. We terminate the tumor inversion if the relative change of the norm of the reduced gradient in (7k) is below a user defined threshold $\text{opttol}_T > 0$. The reference gradient is the gradient obtained for the zero initial guess for $\boldsymbol{p}$ in the first Picard iteration. For the registration problem, we use a combination of the relative change of (i) the norm of the gradient in (7l), (ii) the objective in (10a) and (iii) the control variable $\boldsymbol{v}$, all controlled by a single parameter $\text{opttol}_R > 0$, as a stopping criterion. We also specify a maximal number of Newton iterations ($\text{maxit}_{N,T}$ and $\text{maxit}_{N,R}$) and a lower bound of 1E−6 for the absolute norm of the gradient as a safeguard against a prohibitively high number of iterations. Details for the stopping conditions can be found in [33,69]; see [70, 305 ff.] for a discussion.

We invert the inner linear KKT systems using a matrix-free PCG method. We terminate the PCG method when we either reach a predefined tolerance for the relative residual norm of the PCG method or exceed the maximum number of iterations ($\text{maxit}_{K,T}$ and $\text{maxit}_{K,R}$). We perform inexact solves [71,72] with a tolerance that is proportional to the norm of the reduced gradient of our problem. The key idea here is to not invert the Hessian accurately if we are far from an optimum (large gradient norm). Details about this approach can be found in [68, p. 165ff.].

4.3. Parallel algorithms and computational kernels

Our parallel implementation is described in detail in [17,35]. We use a 2D pencil decomposition for 3D FFTs [73,74] to distribute the data among processors. We denote the number of MPI tasks by $P = P_0 P_1$. Each MPI task gets a total of $N_0/ P_0 \times N_1/ P_1 \times N_2$ values. We use the open source library AccFFT [67], which supports parallel FFT on CPU and GPU for both single and double precision computations. For parallel linear algebra operations, we use PETSC as well as its optimization interface TAO [75,76]. The other computational kernel besides FFTs is the cubic Lagrange polynomial interpolation model used for the semi-Lagrangian time integration of hyperbolic transport equations. We refer to [17,34] for additional details on the parallel implementation of this interpolation kernel. We provide an algorithmic complexity analysis for the registration in [35] and for the FFT in [41].

5. Numerical experiments

We test the performance of our method on synthetically generated and clinical datasets. Note that all experiments carried out are full 3D simulations. Following standard practice, we use 2D cross-sections of 3D data to depict our results.

SIN.

In our first problem, we consider a synthetic problem in which all the data are periodic, band-limited functions. Our goal with this example is to conduct a mesh refinement to investigate the convergence of the Picard solver. To construct the images, first we grow a tumor in the atlas and then, we advect it with a smooth, band-limited velocity to create a patient image. For this smooth data there are no issues with respect discretization errors and we can focus on the performance of the optimization solver. For a Newton method we expect the number of optimization iterations for the outer solver to be mesh independent. We conduct this experiment to test our Picard scheme in a similar scenario.

ATAV.

This first class of real brain synthetic test cases is used as a proof of concept to assess the convergence of our Picard scheme, its robustness with respect to the tumor model, and the sensitivity of tumor reconstruction quality in terms of the tumor model and its parameters. It uses analytic tumor and analytic velocity (fully synthetic; ground truth for tumor parameters and velocity is known); see Section 5.3 for details. It comes in three different flavors: (i) ***ATAV-REAC*** with reaction only and diffusion disabled, (ii) ***ATAV-DIF*** with reaction–diffusion model, and (i) ***ATAV-LD*** with reaction only and diffusion disabled and a small number $n_{\boldsymbol{p}}$ of Gaussian basis functions for the initial condition.[7]

---

[7] LD stands for low-dimensional. That is, we use only very few active Gaussian basis functions with a limited support in a small neighborhood surrounding the center of mass of the patient's tumor. For the other test cases, we allow the grid for the Gaussian to cover the entire area identified as tumor in the patient data.

RTRV.

This second class of real brain test cases uses real clinical images used in the study described in [7]. It uses real patient data, diffusion disabled and enabled; see Section 5.4 for details.

5.1. General setup

Common Parameters.—We list all model and numerical parameters that are common to all test cases in Table 1 and those that are specific for the test cases in Table 2. In the following, we shortly motivate some of the more involved choices:

β_p has been determined experimentally for a purely synthetic test case with image resolution $N_i = 128$ and $n_p = 125$ Gaussian basis functions. For variations of N_i and n_p in our test cases, we observed that the inversion is not sensitive with respect to β_p. Accordingly, we fixed it for all test cases. Smaller values did not further reduce the tumor mismatch in our Picard iteration scheme. To initialize the Gaussian basis, the sub-domain for the grid of basis functions is chosen a priori for the synthetic test cases and determined automatically for the real patient data (cases RTRV) to cover the actual tumor volume in an optimal way: We set the center of the grid for the parametrization to the center of mass of the tumor. The support of the domain covered by the basis functions is the ℓ^∞-ball that covers the entire patient's tumor in $\mathbb{R}^3$. We adjust the standard deviation σ of the Gaussian functions and the distance of their centers automatically. The values are chosen such that their quotient remains constant, and equal to the value used in the synthetic cases. This allows us to ensure a similar conditioning for Φ and to use a fixed β_p.

The choice for n_p depends on the appearance of the patient's tumor (shape and size). We need to be able to parameterize sufficiently complex initial conditions to recover tumors with a complex appearance. Accordingly, we invert for a varying number of parameters n_p on a case-by-case basis ($n_p \in \{8, 125, 343\}$) (see Table 2). For the class of synthetic test cases (ATAV), we use an image resolution of 128^3 to be able to perform more experiments in limited time, for the real patient data, we use 256^3 to ensure high accuracy. Two additional parameters that control the subcomponent solvers are opttol_R, opttol_T as described in Section 4.

Data.—For the generation of the synthetic cases, we use normal brain imaging data obtained at the University of Pennsylvania. For the test case on real imaging data, we use the data available after the first iteration of GLISTR [6,7,77]. The data for these results are the patient data used in the study presented in [6,7]. We consider six datasets from this repository (patient IDs: AAMH, AAAN, AAAC, AAMP, AAQD and AAWI). The original datasets have more labels than we use in our Picard iterations. In particular, they contain background (BG), white matter (WM), gray matter (GM), cerebellum (CB), cerebrospinal fluid (CSF), ventricles (VE), edema (ED), enhancing tumor (ENH), and necrotic tumor (NEC). We construct the labels (WM), (GM), (CSF), (BG) and (TU) by integrating (i) (CB) into (BG), (ii) (VE) into (CSF), and (iii) (ENH), (NEC) and (ED) into (TU).

For all brains, we use the labels white matter (WM), gray matter (GM), cerebrospinal fluid (CSF), and background (BG). The motivation for introducing an additional label BG is technical. We have to ensure the partition of unity across all probability maps for each $\boldsymbol{x}$ in Ω, i.e., all labels have to sum up to one. For example, for the atlas data at $t = 1$ we have

$$\forall \boldsymbol{x} \in \Omega : c_A(\boldsymbol{x}, 1) + \sum_{i=1}^{4} m_{A,i}(x, 1) \overset{!}{=} 1$$

Note that m_{BG} is not used as a label in the image registration formulation (10a). Glial matter is integrated in BG.

Performance Measures.—We perform the registration from the patient space to the atlas space. However, we report all performance measures in the patient space, since the patient space is the relevant space from an applications point of view (for atlas based segmentation).[8] An important point to note here is that velocity based image registration offers an immediate access to the inverse of the deformation applied to an image; we can essentially solve the forward problem with a negative velocity to obtain the action of the inverse deformation map. We use the following measures to assess the performance of our approach: The relative mismatch/residual μ_{B,L^2} between patient anatomy and atlas anatomy after registration ($i = 1$: GM; $i = 2$: WM; $i = 3$: CSF), and the relative mismatch/residual μ_{T,L^2} between patient tumor and atlas tumor after registration:

$$\mu_{B,L^2} := \frac{\sum_{i=1}^{3} \left\| \mathscr{R}^{\text{fwd}}(-\boldsymbol{v}, m_{A,i}(1)) - m_{T,i} \right\|_2}{\sum_{i=1}^{3} \left\| m_{A,i}(0) - m_{T,i} \right\|_2} \quad \mu_{T,L^2} := \frac{\left\| \mathscr{R}^{\text{fwd}}(-\boldsymbol{v}, c_A(1)) - c_T \right\|_2}{\left\| c_T(1) \right\|_2}$$

The Dice coefficient $\text{DICE}_{l,B}$ for the individual label maps (generated by thresholding; see below) associated with the probability maps for $l = WM$, $l = GM$, and $l = CSF$, for the patient and atlas anatomy, as well as the average Dice coefficient DICE_B across all labels:

$$\text{DICE}_{l,B^c} := 2\frac{\left| H(\mathscr{R}^{\text{fwd}}(-\boldsymbol{v}, m_{A,l}(1))) \cap H(m_{T,l}) \right|}{\left| H\left(\mathscr{R}^{\text{fwd}}(-\boldsymbol{v}, m_{A,l}(1))\right) \right| + \left| H(m_{T,l}) \right|}, \quad \text{DICE}_B = \sum_{l=1}^{3} \text{DICE}_{\text{l,B}}/3$$

where $|\cdot|$ is the cardinality of the set and H is a characteristic function of a label with threshold 0.5, i.e.,

$$H(u(x)) := \begin{cases} 1 & \text{for } \mathrm{u}(x) \geq 0.5, \\ 0 & \text{else,} \end{cases}$$

---

[8] If we perform cohort studies, this is different (see [7] for an example).

for all $\boldsymbol{x} \in \Omega$. We also report values for the Dice coefficient computed for the probability maps of the tumor, denoted by $\text{DICE}_{\mathcal{T}}$. Further, we monitor the relative change of the gradient for the coupled problem (see Section 3) for the final iteration k:

$$\|\boldsymbol{g}\|_{\text{rel}} := \|\boldsymbol{g}^k\|_2 / \|\boldsymbol{g}^0\|_2 .$$

where $\boldsymbol{g}^k$ is the gradient of the coupled optimization problem (3) after the kth Picard iteration and $\boldsymbol{g}^0$ the gradient for the initial guess. Finally, the relative ℓ^2-error for the computed velocity and the initial condition (under the assumption that we know the true velocity $\boldsymbol{v}^\star$ and the true tumor parameters $\boldsymbol{p}^\star$; synthetic test problem) at the final (kth Picard) iteration is calculated according to

$$e_{\boldsymbol{v},L^2} := \|\boldsymbol{v}^\bigstar - \boldsymbol{v}^k\|_2 / \|\boldsymbol{v}^\bigstar\|_2 \qquad e_{c0,L^2} := \left\|c_A^\bigstar(0) - c_A^k(0)\right\|_2 / \left\|c_A^\bigstar(0)\right\|_2$$

$$e_{\boldsymbol{v},L^2,rel} := \|\boldsymbol{v}^{k-1} - \boldsymbol{v}^k\|_2 / \|\boldsymbol{v}^{k-1}\|_2 \qquad e_{c0,L^2,rel} := \left\|c_A^{k-1}(0) - c_A^k(0)\right\|_2 / \left\|c_A^{k-1}(0)\right\|_2$$

Hardware.—The runs for all test cases were executed on the Tier-1 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 7712 nodes with Xeon E5–2680 v3 processors and 24 cores on two sockets per node. The nodes are connected via an Aries interconnect. For data sizes of $N_i = 128$, $i = 0, 1, 2$, we use 3 nodes with 64 MPI tasks and for $N_i = 256$, $i = 0, 1, 2$, we use 11 nodes and 256 MPI tasks.

5.2. Test case SIN

Purpose.—The purpose of this experiment is to verify the convergence of our scheme and check the dependence of the convergence rate on the mesh size. For smooth problems, Newton's method is known to converge quadratically near the solution and is expected to have a convergence rate that is mesh independent [21]. What about our Picard scheme? We will use this numerical experiment to informally answer this question.

Results.—We report results in Table 3 and Fig. 2. We report the relative (to the initial value) reduction in the objective function value $\|\mathcal{J}\|_{\text{rel}}$, the relative (to the initial value) norm of the reduced gradient $\|g\|_{\text{rel}}$, as well as the relative norm $e_{c0,L^2,rel}$ and $e_{\boldsymbol{v},L^2,rel}$, respectively of the update for the inversion variables is reported over the first 23 Picard iterations. We report the relative norm of the reduced gradient (8) of the coupled formulation (3) in Fig. 2 for three different mesh resolutions.

Setup.—We consider a fully synthetic, smooth test case using sinusoidal functions for white matter and gray matter, defining the artificial brain. We use white matter for the atlas configuration as $m_{WM} := \frac{1}{2}cos(x_1)sin(x_2) + \frac{1}{2}$ and gray matter as $m_{GM} := 1 - m_{WM}$. No

CSF is used here ($m_{CSF} := 0$). Based on a smooth, sinusoidal velocity $\boldsymbol{v}^{\star}(\boldsymbol{x}) = \left(v_1^{\star}(\boldsymbol{x}), v_2^{\star}(\boldsymbol{x}), v_3^{\star}(\boldsymbol{x})\right)^T$ with $\boldsymbol{v}_1^{\star}(\boldsymbol{x}) = sin(x_3)cos(x_2)sin(x_2), v_2^{\star}(\boldsymbol{x}) = sin(x_3)cos(x_3)sin(x_3)$, and $v_3^{\star}(\boldsymbol{x}) = sin(x_2)cos(x_1)sin(x_1)$, we generate a patient configuration by growing a tumor in the atlas configuration and advecting the resulting abnormal brain with $\boldsymbol{v}^{\star}$ to the patient space. Different from all other test cases and Table 1, we set the tolerance for the relative gradient in the registration and in the tumor inversion solver to opttol$_R$ = opttol$_T$=1E−7 and allow maxit$_R$ = (100, 100) and maxit$_T$ = (100, 100) Newton and Krylov iterations, respectively. We solve the subcomponent problems accurately to isolate the coupled solver convergence from issues related to early termination in the component solvers. The target value for the regularization parameter β_v in the continuation scheme is set to 1E−2. This also differs from the value chosen for all other test cases and Table 1. The smooth data used in SIN provide less geometric structures to the image registration problem and, thus, requires higher regularization for $\boldsymbol{v}$. We consider three mesh resolutions in which we increase the resolution for both $\boldsymbol{v}$ and $\boldsymbol{p}$. In particular we consider grids with 64, 128, and 256 points per dimension and Gaussian grids of size 8, 64, and 512 for the tumor parametrization with a slightly larger spacing of 2σ than in the other test cases.

Observations.—This experiment indicates that our method is mesh-independent. For example, in Table 3 we see that within seven iterations for all resolutions both $\boldsymbol{p}$ and $\boldsymbol{v}$ have essentially converged. We also achieve five orders of magnitude reduction in the relative reduced gradient $\|g\|_{\mathrm{rel}}$, given by (8). Once the reduction in the gradient reaches 1e-5 or so, our method slows down. A Newton method would converge much faster assuming we used a gradient consistent to the objective function. Recall that we use an optimize-then-discretize approach. Due to the semi-Lagrangian time-stepping for the pure advection steps, this gradient is different than the discretize-then-optimize gradient that is needed for the second-order convergence of a Newton method.

Conclusion.—We conclude that although our modular iteration scheme does not perform as well as Newton's method, it features mesh-independent linear convergence with good convergence rates until we reach gradient values that are significantly lower than what is required and feasible (due to discretization and modeling errors) in real clinical cases. Indeed, as we will see next, the Dice coefficients for the different tissue types converge within a few iterations to their final value for all synthetic and clinical datasets. So the behavior of our scheme appears to be robust and independent of the particular dataset.

5.3. Test case ATAV

Purpose.—This experiment is a proof of concept. We test the numerical accuracy of our scheme, identify the inversion accuracies we can ideally expect (i.e., the errors we get if we use the forward operators to generate the observations for our coupled inversion), and study the convergence of our solver.[9] With ATAV-REAC and ATAV-DIF, we in addition test the

---

[9] In general, considering real data, our forward model is – at best – a crude approximation of the actual observations. That is, first of all, we cannot expect our registration algorithm to recover a bio-physically meaningful deformation between the brain anatomies of different subjects (such a model does not even exist). Second, our forward tumor model is far from being realistic (see Section 3.1 for a discussion).

sensitivity of our approach with respect to perturbations in the model and model parameters. In ATAV-LD, we restrict the admissible initial condition parametrization for tumor growth to a very low dimensionality in order to exclude fully grown tumors as initial conditions. We examine, whether this increases the sensitivity of the reconstruction quality for the tumor with respect to correct model parameters.

Setup.—ATAV is based on real brain geometries, but uses a tumor grown with our forward solver in a synthetically generated patient image. We use a resolution of $N_i = 128$. We choose $\boldsymbol{p} = \boldsymbol{p}^{\star}$, which defines $c_A(0) = c_A^{\star}(0)$, grow a tumor in the tumor-free atlas, which gives $c_A^{\star}(1)$, choose $\boldsymbol{v} = \boldsymbol{v}^{\star}$, deform the atlas with the grown tumor by advecting the probability maps with the negative velocity $\boldsymbol{v}^{\star}$. This gives us $\boldsymbol{m}_T$ and c_T (patient image with tumor). The velocity $\boldsymbol{v}^{\star}$ is generated by registering two tumor-free images of two different individuals ($\beta_{\boldsymbol{v}}^{\star} = 1E-4$). The center of mass for the synthetic tumor is set to $(x^1, x^2, x^3) = 2\boldsymbol{\pi} \cdot (0.285, 0.36, 0.5)$. As initial condition for the artificial tumor generation, we enable two of the Gaussians at the center of the grid of Gaussians. See Section 5.1 for further details on the parameters. As a baseline, we also report results for the sole registration of the healthy anatomy (i.e., neglecting the tumor forward solve to generate the data). In addition to that, we quantify the numerical error of our scheme for solving the transport equations (forward registration). This is done by solving the forward problem twice, once with the original and once with the reverted (negative) velocity. The associated error is given by

$$\frac{\left\| c_A(1)^{\star} - \mathscr{R}^{\text{fwd}}(-\boldsymbol{v}^{\star}, \mathscr{R}^{\text{fwd}}(\boldsymbol{v}^{\star}, c_A(1)^{\star})) \right\|_2}{\left\| c_A(1)^{\star} \right\|_2} = 9.37\text{E} - 2 \tag{11}$$

Whereas we disable diffusion ($k_f = 0$) in ATAV-REAC, we use the full tumor model including diffusion ($k_f 0$) for ATAV-DIF. In ATAV-REAC, the same growth rates ρ_w, ρ_g with scaling ρ_f and k_w, k_g with scaling k_f are used for growing the tumor and for the inversion to reconstruct the initial condition. In ATAV-DIF, we use values for the reaction and diffusion coefficients for the inversion in the Picard iterations, that are either the same or differ from those used for the generation of the tumor. In ATAV-LD, we use the full reaction–diffusion tumor model and enforce the initial tumor to be small, i.e., only invert for $n_{\boldsymbol{p}} = 8$ parameters, which results in a grid of 2 × 2 × 2 Gaussians that cover the true initial condition of the artificially grown tumor. We expect this to increase the sensitivity of our inversion with respect to the tumor parameters (we cannot fully represent the whole patient tumor purely based on a linear combination of the basis functions). For the inversion, we again consider a variety of models and model parameter combinations, which includes the use of the “correct” (ground truth) tumor parameters. We also consider the case in which we completely neglect the tumor model (i.e., $\rho_f = k_f = 0$) and just invert for the basis.

Results.—We report results for the registration of anatomy (without tumor) in Table 4. Results for the inversion in ATAV-REAC are presented in Fig. 3 and assess the reconstruction quality (Dice and residuals) and the reduction of the reduced gradient with respect to the iteration index. We also report the error between the ground truth $c_A(0)^{\star} = \boldsymbol{\Phi}$

$\boldsymbol{p}^{\star}$ and $\boldsymbol{v}^{\star}$ and the estimated iterates as well as the relative norm of the gradient of the fully coupled problem in (3).

For ATAV-DIF, quantitative results for the inversion are shown in Table 5, qualitative results can be found in the supplementary material in Fig. A.10. In addition to reconstruction quality and gradient reduction, we list the run time for the Picard scheme per iteration and the percentage spent in each individual solver (tumor and registration), respectively. Note that tumor and registration run times do not add up to 100% as further parts of the code such as the calculation of the reduced gradient and the steering of the Picard iteration are not included in the measurements.

For ATAV-LD, we show simulation results in Fig. 5 (for sagittal and coronal slices, Fig. A. 11) and report quantitative results in Table 6. We plot the trend of the relative tumor mismatch with respect to the regularization parameter β_v for the Sobolev norm for the registration velocity (and by that the Picard iteration index) in Fig. 4.

Observations.—The most important observations are (i) that the reconstructed data (tumor and registered anatomy) are in excellent agreement with the patient data, and (ii) we are able to reduce the reduced gradient (8) by two orders of magnitude in seven iterations of our Picard scheme for ATAV-REAC and significantly more than a factor of 50 for all experiments in ATAV-DIF and ATAV-LD.

The numerical error for the advection in (11) is 9.37E−2. The relative mismatch for the anatomy obtained for our iterative Picard scheme is in the order of the advection error for the forward image registration problem. We can also see that the reconstruction of $c_A(0)^{\star}$ and $c_A(1)^{\star}$ seems to be bounded by this error. In fact, due to the advection error that leads to a mismatch in this order in the atlas domain, this is the best we can expect without over-fitting the data. Similar observations can be made if we compare the inversion results with the results obtained for the registration for healthy brains (neglecting the tumor simulations) reported in Table 4. Hence, the quality of tumor reconstruction is comparable to the quality of pure image registration between the healthy geometries. This is an excellent result that clearly demonstrates the potential of our approach. The obtained Dice coefficient for the brain anatomy is in the order of what we see for the sole registration of healthy anatomies.[10] Note that the mismatch between the true velocity $\boldsymbol{v}^{\star}$ and the recovered velocity $\boldsymbol{v}^{k}$ reported for ATAV-REAC is due to the fact that image registration is an inherently ill-posed problem: the velocity can only be reconstructed exactly in image areas with non-zero gradients and if there are only non-zero intensity differences between the images to be registered in areas that do correspond to one another.[11] In addition, we ask for the reconstruction of a vector field from scalar data.

The Dice coefficient for the brain anatomy increases from 7.10E−1 to 9.47E−1 in ATAV-REAC, where we obtain a final Dice coefficient of 9.73E−1 for the tumor. The results for

---

[10] In fact, it is even slightly better for ATAV-REAC. This slight increase might be a consequence of numerical inaccuracies in our scheme, discrepancies in the number of Newton-steps and Krylov iterations taken.

[11] As in any formulation based on an L^2-distance functional, it is the mismatch between the reference and template image and the gradient of the deformed template image that drive the optimization (see (7l)).

ATAV-DIF and ATAV-LD show that the used model and the dimensionality of the initial conditions do not have a significant impact on the quality of the inversion (Dice and mismatch).

We can furthermore see that we can significantly reduce the norm of the reduced gradient (8) to 1.21E−2. We can also see that once we have reached the target regularization parameter β_v = 1E−4 we do not make any more progress. The update for the velocity tends to zero, the changes in the reduced gradient are small and the error measures (residual and Dice) do no longer change significantly.

For ATAV-DIF, we see that we obtain a slightly better mismatch (5.63E−2) and Dice coefficient (9.76E−1) for the tumor probability map if we use the correct ρ_f. These results suggest that we can identify the correct ρ_f if we run multiple inversions for the initial condition. However, we note that these differences are small (the mismatch is between 5.95E −2 and 5.63E−2 and the Dice is between 9.71E−1 and 9.76E−1, respectively, for the tumor). Overall, we obtain an excellent agreement between patient data and atlas data irrespective of the model choice. These observations are confirmed by careful visual inspection of Fig. A. 10 (supplementary material; shows only results for the “correct” tumor parameters). This can be explained by the parametrization of the initial condition. We can reconstruct complex tumor shapes even with a simple model. Overall, this indicates that using a reaction-only model might be sufficient for pure diffeomorphic image registration,[12] where a comparison of the total run time for the last row in Table 5 to the total run times attained when enabling diffusion shows that we can save a factor of 10 in run time by disabling diffusion. Another interesting behavior within our scheme is that, during the first few iterations, most time is spent in the tumor inversion, whereas, as we reduce β_v, the registration does most of the work. This is to be expected, since the run time of our scheme (more precisely, the condition number of the Hessian) for diffeomorphic registration, is not independent of β_v [33,36].

The most important observation for ATAV-LD is that we can identify the correct pair of diffusion and reaction rates that have been used to generate the synthetic test based on slightly more pronounced differences in the overall mismatch and Dice scores than in ATAV-DIF. The sensitivity with respect to changes in the diffusion parameter k_f is larger than with respect to changes in the growth rate ρ_f. We expect this much more pronounced dependence on the diffusion model if we parameterize the initial condition on a grid with a smaller support, since we need the cancerous cells to diffuse to areas more distant to the tumor center in order to be able to reconstruct the whole tumor. The trend of the relative mismatch for the tumor in Fig. 4 highlights this effect. The curves almost cluster in terms of the different choices for the diffusion coefficient; we overall achieve significantly better mismatch if we choose the correct diffusion coefficient. We can also see that the curves plateau much earlier in cases where we use the wrong parameter combinations. The tumor mismatch takes values between 3.37E−1 and 4.68E−2 with a Dice score of 4.71E−2 to 9.65E−1.

---

[12] This is clearly not the case when targeting the design of computational framework to aid clinical decision making by generating model based prediction of a patient’s future tumor state. In this case, we quite certainly have to use more complex, high-fidelity models. We discuss this in more detail in Section 6.

Conclusion: We conclude that our Picard scheme is efficient and converges to a valid (local) minimum in the search space (we reduce the relative global gradient, the distance measure, and significantly increase the Dice coefficient). We attain an excellent agreement between the data (patient tumor and geometry) and the predicted state (transported atlas geometry and predicted tumor) with a final Dice coefficient of 9.47E−1 and 9.73E−1 for the labels of the anatomy and tumor in ATAV-REAC and only slightly worse values in ATAV-DIF and ATAV-LD.

The integration of a diffusion model into our inversion is very costly, at least for our current implementation. Designing a more efficient forward solver for the diffusion operator requires more work. We could demonstrate that our parametrization of the initial condition allows us to generate high-fidelity registration results, especially for the healthy anatomy, irrespective of the model that has been used to generate the data. These are clearly preliminary results, but they provide some evidence that reaction-only models might be sufficient for pure image registration (something that is quite certainly not the case if we target parameter identification and tumor growth prediction) [63].

We can also observe that there are subtle differences in the reconstruction quality of the tumor if we use the “correct” growth rate for the inversion for ATAV-DIF, that are more pronounced in ATAV-LD. Overall, we conclude that we (i) can neglect the diffusion model in the context of diffeomorphic registration and compensate the resulting loss in accuracy by a higher-dimensional Gaussian basis, (ii) might be able to identify appropriate growth rates if we run multiple inversions for different parameters with a low-dimensional Gaussian basis.

5.4. Test case RTRV

Purpose.—We test our approach on real data of patients diagnosed with glioma tumors and study the registration quality for a variety of parameter choices for the tumor growth model.

Setup.—This test scenario consists of real patient brains with real tumors for which we do not know any parameters. The patient datasets are the first proposal for a patient segmentation produced in the first iteration of GLISTR [7]. We provide additional details about this database in Section 5.1. We have to identify the support of the domain spanned by the Gaussian basis functions for the tumor initial condition parametrization as well as their spacing d, and the standard deviation σ for any unseen patient. This is done automatically. As some of the real tumors are multifocal, we use $n_{\boldsymbol{p}} = 343$ Gaussians in our first set of experiments on all six brains for the initial condition parametrization. In contrast to the synthetic test cases ATAV-REAC, ATAV-DIF and ATAV-LD, we allow the tumor to grow also in gray matter instead of in white matter only, but with a reaction parameter that is five times smaller than in white matter (see Table 2 for details). We use a variety of models and parameter settings in our Picard scheme for the two patients AAMH and AAAN to not only assess the performance of our method but also study its sensitivity towards parameter changes and model complexity. We vary the reaction parameter ρ_f between 0 and 15 and choose the diffusion coefficient k_f either as 0 or 1.00E−2. See Section 5.1 for additional details on the setup of the test case and the parameters.

Results.—Figs. 6–8 show the healthy atlas ($k = 1$) in the top row and the corresponding patient image with tumor in the bottom row for axial, saggital, and coronal orientations. The hard segmentations for the results computed with the proposed approach are shown in the middle row using the same orientations. In Table 7, we summarize the results for all patient datasets. We report the initial and final values for the mismatch and Dice coefficients associated with the probability maps for the tumor and the brain anatomy, as well as the relative norm of the gradient of the coupled problem in (3). We also report timings for the entire inversion. Fig. 9 shows more detailed images of the probability maps for the patient AAMH (complex and large tumor). More detailed visual results with information on more iterations of our inversion algorithm for all patients are given in Appendix (Figs. A.12 through A.17). Results for varying reaction and diffusion for AAMH and AAAN are listed in Table 8. Note that all parameter choices refer to the model used for tumor reconstruction in the Picard scheme. The true growth parameters of the tumors are unknown.

Observations.—The most important observation is that we obtain very good registration results – qualitatively and quantitatively – in what is an extremely challenging registration problem. From visual inspection of the data alone (Fig. 6 through Fig. 8) we can immediately see that there are significant anatomical differences between the atlas image and the patient images, and across patients. The tumors vary significantly in shape and size. Overall, this poses considerable challenges to our framework. The results reported in Fig. 6 through Fig. 8 and in Fig. 9 clearly demonstrate that the deformed atlases are in very good agreement with the patient data for all six subjects. We reach Dice coefficients between 7.87E−1 to 8.44E−1 and 8.74E−1 to 9.75E−1 for the probability maps associated with the anatomy and the tumor. These results are slightly worse than those obtained for the artificially grown tumors in the former sections, but still competitive. We also note that the initial Dice coefficients for the anatomy range between 2.92E−1 and 4.74E−1 for these data, which is drastically worse than what we have seen in our synthetic test cases. The run times are comparable to our former experiment. We again achieve a reduction of the relative norm of the gradient for the coupled problem in (3) of about two orders of magnitude (slightly less than what we saw before). The results of our study with varying tumor model parameters presented in Table 8 show that there are slight variances in the results depending on the parameter choices, but the Picard iteration scheme is successful in all cases. Further enhancements of our environment are required to enable the identification of the ‘correct’ tumor growth parametrization.

Conclusions: We have tested our formulation on real data that pose significant challenges due to large inter-subject anatomical variability and a strong variation in the appearance of the tumor, in shape, size, location and growth behavior. We use a very simple model that only accounts for logistic growth. This, in combination with a flexible, high-dimensional parametrization of the initial condition allows us to overcome these challenges. We achieve extremely promising registration accuracies with a Dice score of up to 8.44E−1 and 9.75E−1 for the label maps associated with the probability maps of the brain anatomy and the tumor in what is an extremely challenging problem. Runs with variation of the growth parameter ρ_f show, on the one hand, that it is important to identify the correct parameters to achieve optimal quality of tumor reconstruction. We, therefore, believe that it will be possible to identify physical tumor growth parameters from our coupled solver if the tumor model is enhanced (anisotropic diffusion, mass effect, . . .) and we consider ‘correct’ time horizons and/or restrict the initial conditions of tumor to point seeds (see first steps following this idea in ATAV-LD). On the other hand, the results for varying model parameters show the robustness of our Picard scheme with respect to the model and parameter choice. Furthermore, the reconstruction results for the AAAN patient data show that we can reconstruct multifocal tumors with comparable quality and computational costs.

6. Conclusion

We have presented a new method for the registration of images of patients diagnosed with mono- or multifocal brain tumors to a common reference atlas. Application scenarios are biophysical model calibration and image registration. Our method combines stand-alone forward and inverse tumor [41] and diffeomorphic registration [33–35,37] by means of an efficient coupling scheme based on a Picard iteration. It allows us to exploit available, tailored implementations for the solution of the subproblems [17,35], while monitoring convergence using the coupled gradient information. We invert simultaneously for the control variables of both problems—a parametrization of the initial condition for the tumor model, and a smooth velocity field to capture the inter-subject variability of brain anatomy. Here is what we have learned from our experiments on synthetic and real data:

i. Despite the fact that our scheme neglects coupling terms that appear in our coupled optimization problem, we could experimentally show that it efficiently reduces the coupled gradient. A convergence proof of the Picard scheme to a local minimum is beyond the scope of this paper and remains subject to future work.

ii. We could demonstrate that our parametrization of the initial condition allows us to generate high-fidelity registrations irrespective of the complexity of the data or the model used for the tumor simulations. If we reduce the number of basis functions used for initial condition parametrization, this is no longer true.

iii. In studies with various models from a pure interpolation with the basis functions that parameterize the initial condition (zero reaction and diffusion coefficient) over a reaction-only, to a full reaction–diffusion model, we could show in our synthetic cases, that we get the highest accuracy in tumor reconstruction, if we use the correct model. This implies that our framework could eventually serve as a powerful tool for model selection. A rigorous verification of this claim requires significantly more work and remains for the future.

iv. Overall, our numerical study, which includes real brain images with real tumors, shows that we can achieve high-fidelity results with an overall low mismatch and high Dice score in particular for the simulated and observed tumor, with Dice coefficients ranging from 93% for real tumors above a critical size, and up to 97% for artificially grown tumors in a real brain geometry.

v. We are able to solve the joint registration and biophysical inversion problem on hardware and wall-clock time that fits clinical requirements. We used 128 × 86

> cores for our simulations. This configuration is not exotic and it is found in many hospitals that already use such clusters for bioinformatics and image analysis. On such a machine, our real data simulations would require about two hours. Considering that standard MRI imaging already takes 30–45 min, the simulation run time is acceptable.

Let us emphasize that our tumor model is currently not sophisticated enough to allow prediction, but the tumor-registration coupling approach that we present in this work and for which we can show good computational performance and high accuracy for various real brain data test cases lays the basis for further developments with improved tumor solvers and, finally, a parameter identification and growth prediction tool. The next steps for SIBIA are to improve the biophysical tumor-growth model by first adding mass-effect and second higher-fidelity tumor growth models.

Acknowledgments

This material is based upon work supported by AFOSR grant FA9550-17-1-0190; by NSF grants CCF-1725743 and CCF-1337393; by the U.S. Department of Energy, Office of Science, Office of Advanced Scientific Computing Research, Applied Mathematics program under Award Number DE-SC0019393; by NIH grant 10042242; by resources at the Texas Advanced Computing Center (TACC); and by the University of Stuttgart, Institute for Parallel and Distributed Systems. 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, 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. Computing time on the Texas Advanced Computing Centers Stampede system was provided by an allocation from TACC.

Appendix.: Supplementary material

A.1. Supplementary material for the ATAV testcase

We present additional qualitative data for the ATAV-DIF testcase and the ATAV-LD testcase.

A.2. Supplementary material for the RTRV testcase

We show detailed results for the RTRV real data test cases for six patient datasets (AAAC, AAAN, AAMP, AAMH, AAQD, AAWI; first proposal for a patient segmentation produced in the first iteration of GLISTR [7]). Fig. A.12 through Fig. A.17 show axial slices of the evolution of the probability maps of the atlas brain tissue labels (in patient space) and the reconstructed tumor (in patient space) throughout our inversion algorithm. Brain tissue probability maps, tumor probability map, mismatch for all labels as well as the hard segmentation image are given for Picard iterations $k = \{1, 2, 4, 6\}$.

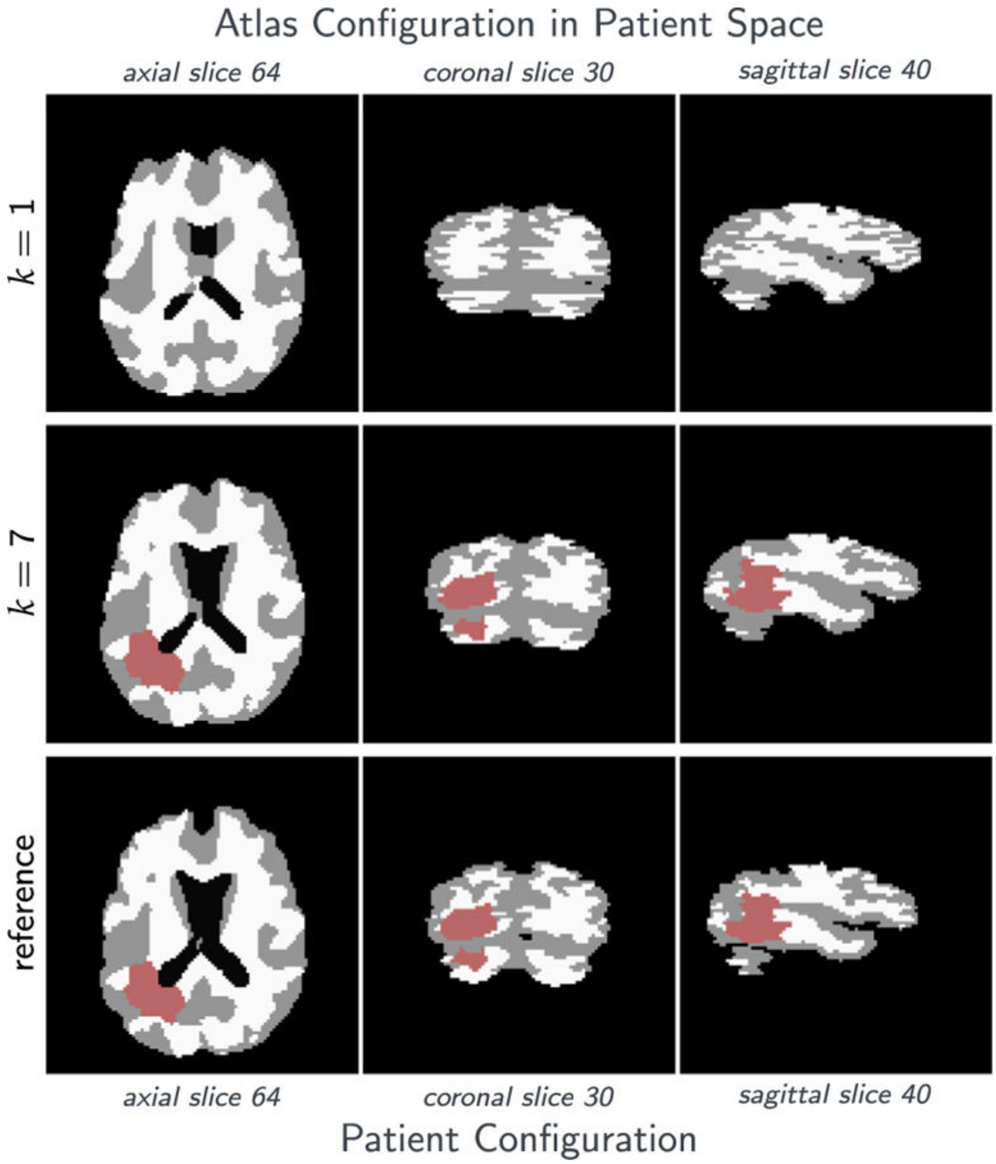

Fig. A.10.

Tumor and brain labels (obtained by thresholding probability maps) for the **analytic tumor with analytic velocity with non-zero diffusion (ATAV-DIF)** test case; ground truth: ($\rho_f = 10$, $\rho_w = 1$, $\rho_g = 0$, $k_f = 1.00\text{E}{-2}$, $k_w = 1$, $k_g = 0$, $\boldsymbol{p} = \boldsymbol{p}^\star$, $\boldsymbol{v} = -\boldsymbol{v}^\star$). We show results for the inversion (velocity and initial condition) if we use the true (“correct”) tumor parameters used to generate the test case (i.e., $\rho_f = 10$ and $k_f = 1.00\text{E}{-2}$). The top row shows the initial label maps for the atlas image. The middle row shows the label maps for the atlas image after registration (transported to the patient space) and the bottom row shows the label maps for the patient data. We can see that the results are qualitatively in excellent agreement. This is confirmed by the values for the mismatch and Dice coefficients for the labels and probability maps for the individual tissue classes reported in Table 5.

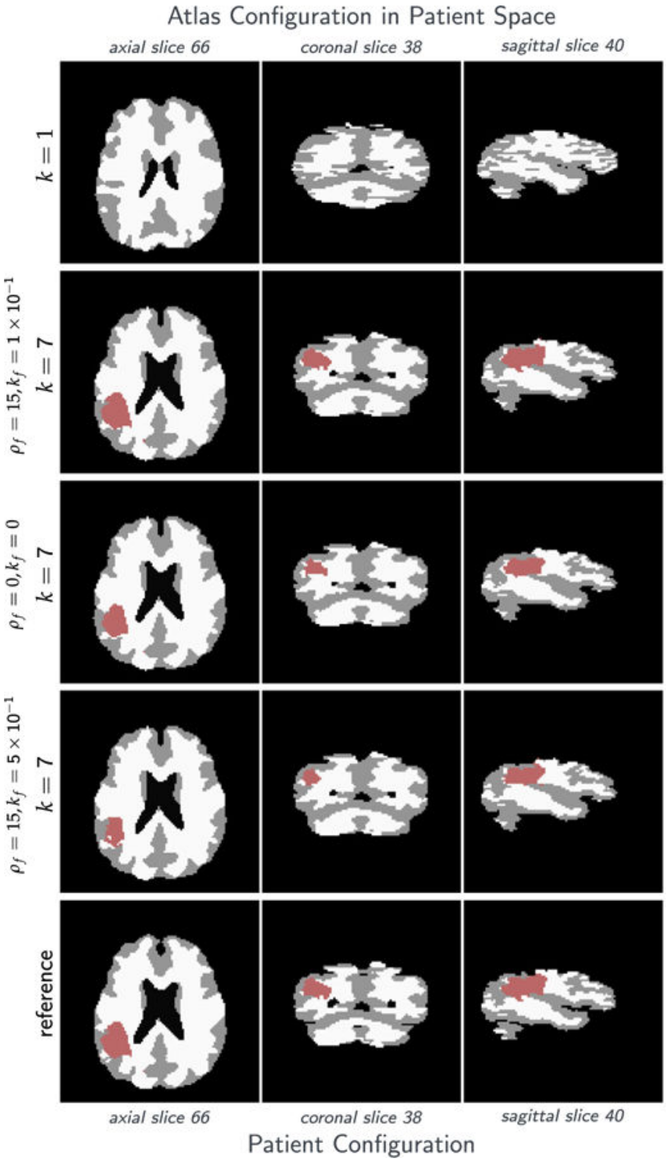

Fig. A.11.

Tumor and brain labels for the **analytic tumor with analytic velocity with non-zero diffusion and low-dimensional initial condition (ATAV-LD)** test case; ground truth: ($\rho_f = 15$, $\rho_w = 1$, $\rho_g = 0$, $k_f = 1.00\text{E}{-1}$, $k_w = 1$, $k_g = 0$, $\boldsymbol{p} = \boldsymbol{p}^\star$ (in patient domain), $\boldsymbol{v}$ N/A); for the 'correct' tumor parameters $\rho_f = 15$, $k_f = 1.00\text{E}{-1}$, and the two different settings $\rho_f = 15$, $k_f = 5.00\text{E}{-1}$ and $\rho_f = k_f = 0$.

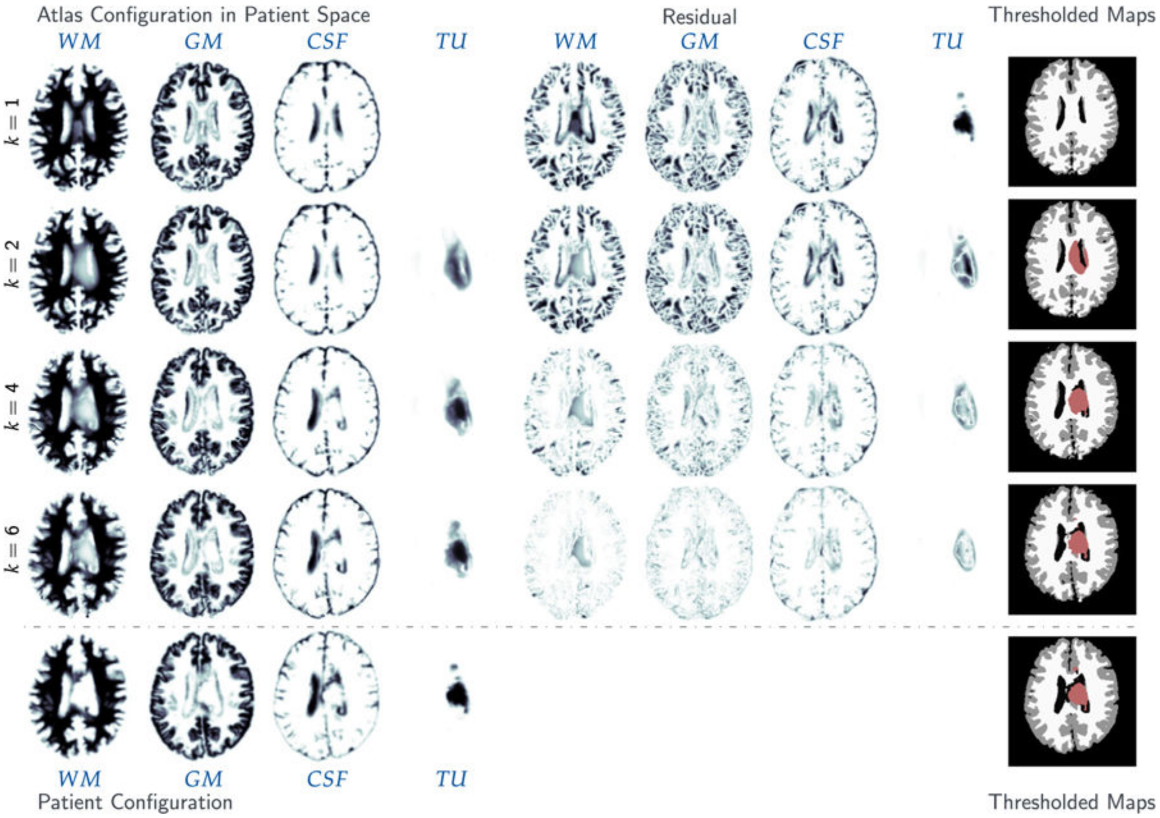

Fig. A.12.

Results for the **real tumor/ real velocity (RTRV)** test case, ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A) and the **AAAC** patient. The figure shows probability maps for the labels of the healthy atlas brain ($k = 1$; top row) and the AAAN patient (target) brain probability maps with tumor (bottom row), along with the reconstructed probability maps throughout the Picard iterations ($k = 2, 4, 6$) (axial-slice 132).

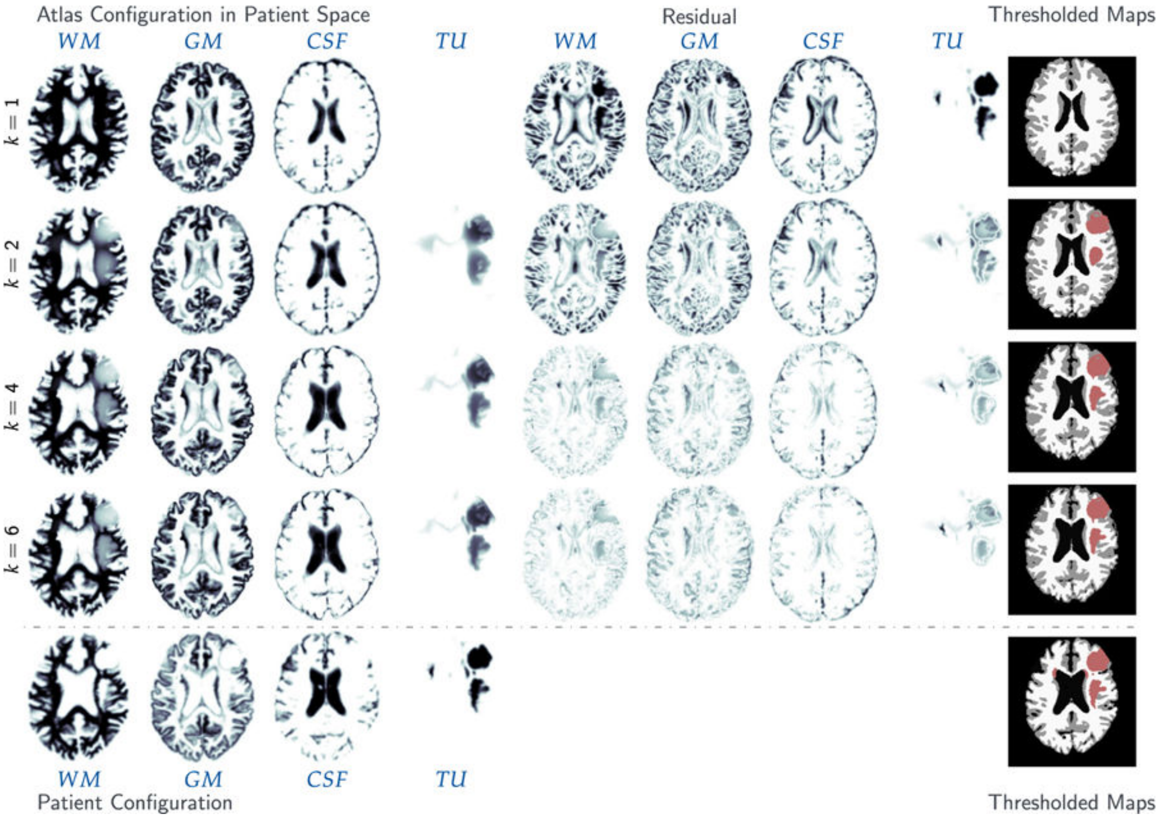

Fig. A.13.

Results for the **real tumor/ real velocity (RTRV)** test case, ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A) and the **AAAN** patient. The figure shows probability maps for the labels of the healthy atlas brain ($k = 1$; top row) and the AAAN patient (target) brain probability maps with tumor (bottom row), along with the reconstructed probability maps throughout the Picard iterations ($k = 2, 4, 6$) (axial-slice 132).

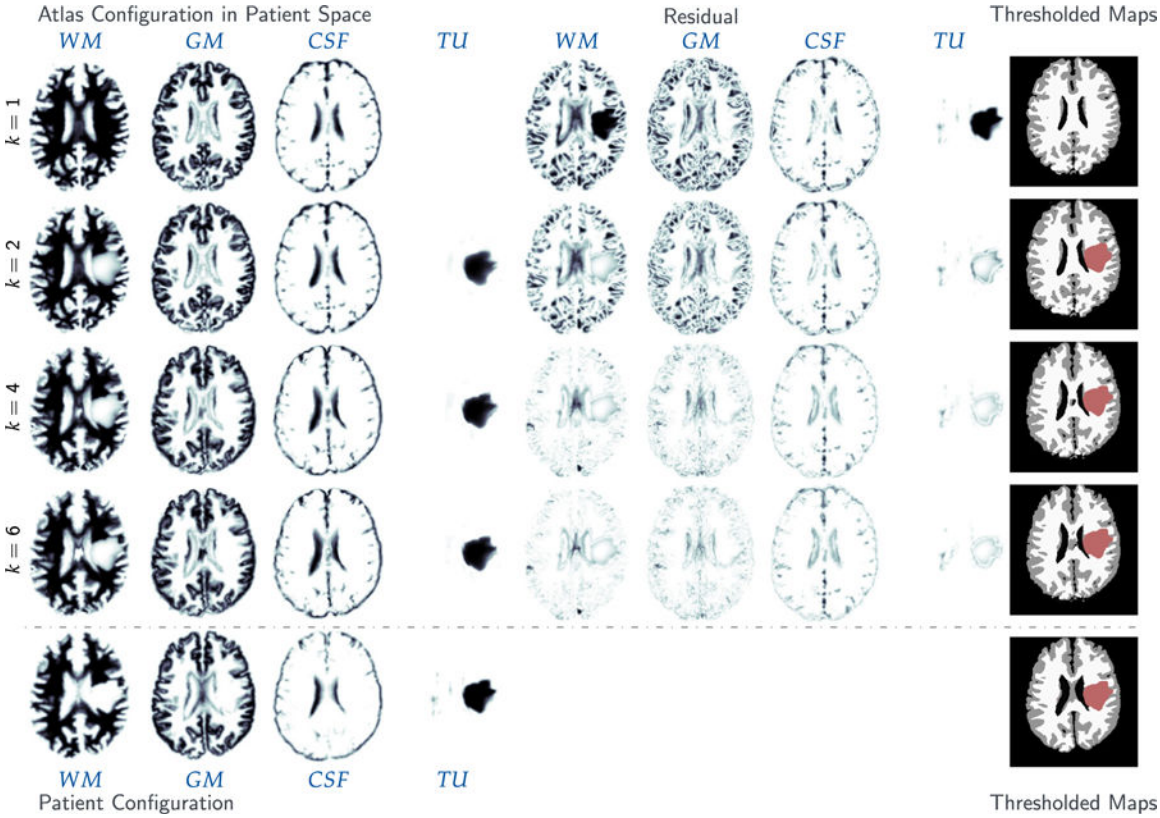

Fig. A.14.

Results for the **real tumor/ real velocity (RTRV)** test case, ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A) and the **AAMP** patient. The figure shows probability maps for the labels of the healthy atlas brain ($k = 1$; top row) and the AAAN patient (target) brain probability maps with tumor (bottom row), along with the reconstructed probability maps throughout the Picard iterations ($k = 2, 4, 6$) (axial-slice 132).

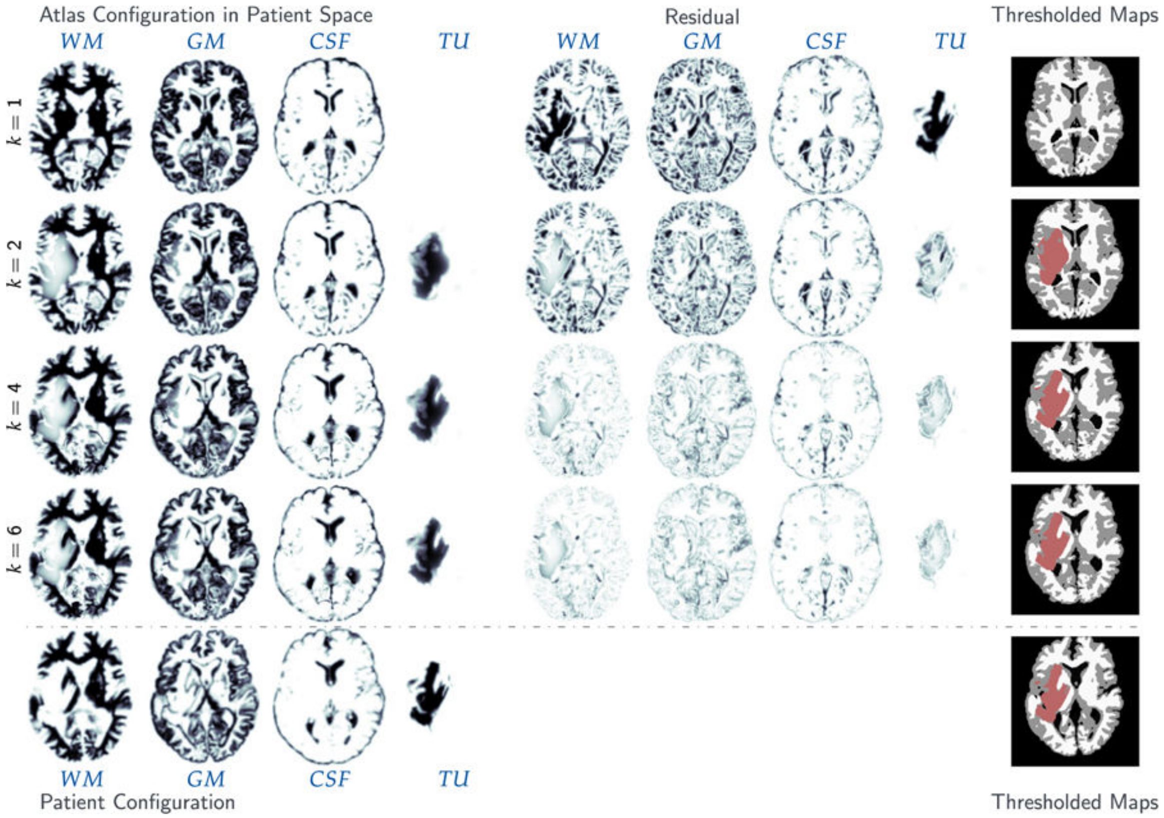

Fig. A.15.

Results for the **real tumor/ real velocity (RTRV)** test case, ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A) and the **AAMH** patient. The figure shows probability maps for the labels of the healthy atlas brain ($k = 1$; top row) and the AAAN patient (target) brain probability maps with tumor (bottom row), along with the reconstructed probability maps throughout the Picard iterations ($k = 2, 4, 6$) (axial-slice 120).

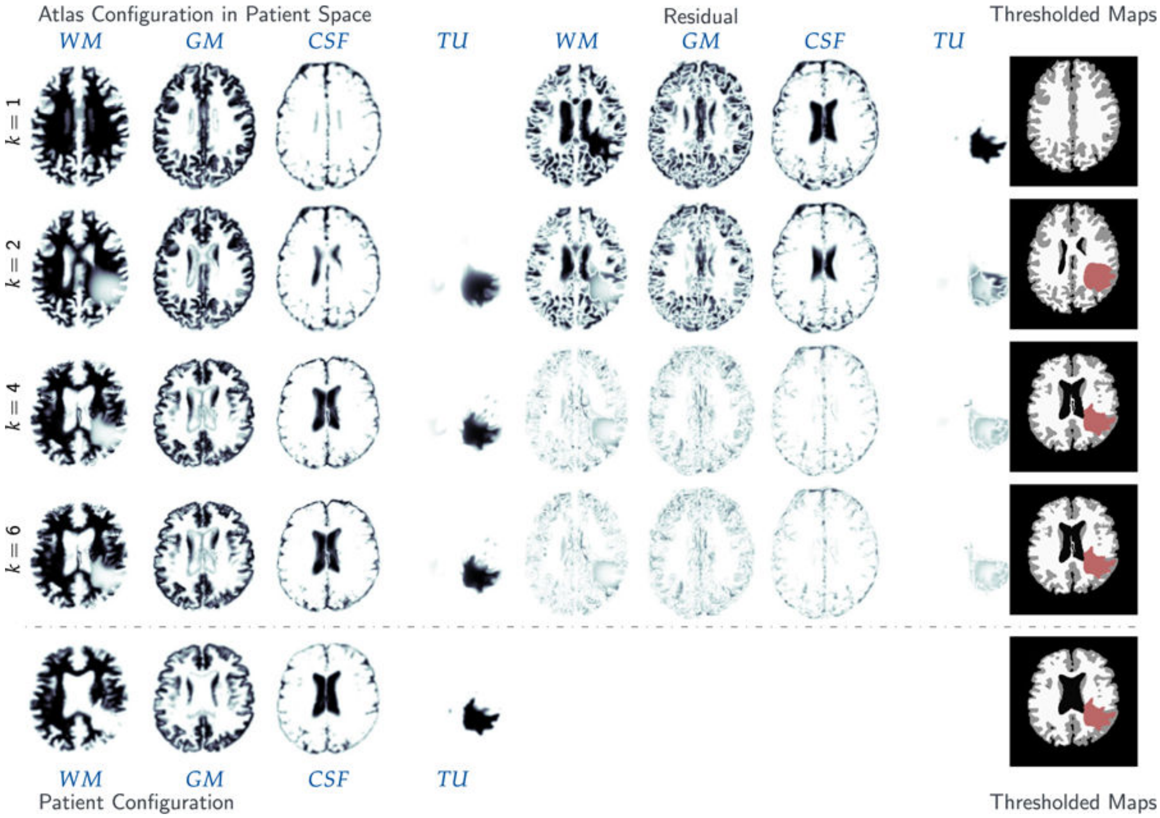

Fig. A.16.

Results for the **real tumor/ real velocity (RTRV)** test case, ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A) and the **AAQD** patient. The figure shows probability maps for the labels of the healthy atlas brain ($k = 1$; top row) and the AAAN patient (target) brain probability maps with tumor (bottom row), along with the reconstructed probability maps throughout the Picard iterations ($k = 2, 4, 6$) (axial-slice 136).

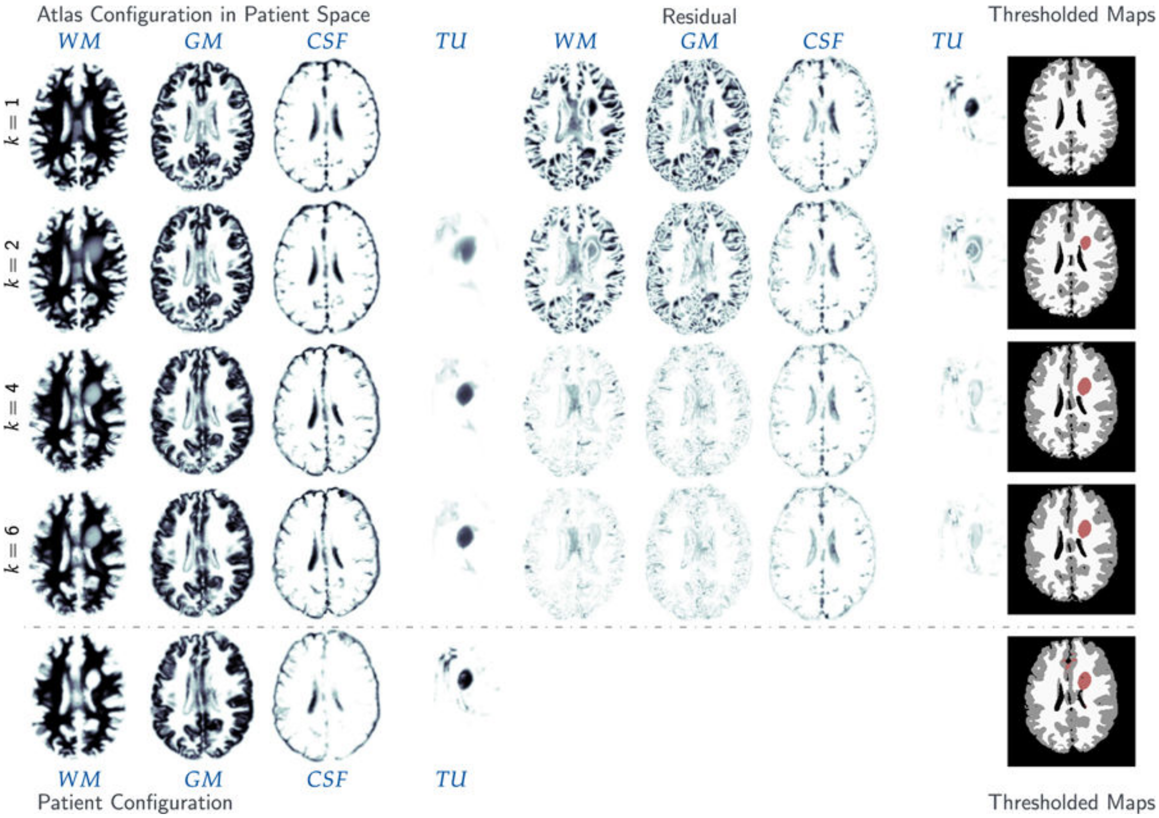

Fig. A.17.

Results for the **real tumor/ real velocity (RTRV)** test case, ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A) and the **AAWI** patient. The figure shows probability maps for the labels of the healthy atlas brain ($k = 1$; top row) and the AAAN patient (target) brain probability maps with tumor (bottom row), along with the reconstructed probability maps throughout the Picard iterations ($k = 2, 4, 6$) (axial-slice 132).

References

[1]. Swanson KR, Rostomily RC, Alvord EC, A mathematical modelling tool for predicting survival of individual patients following resection of glioblastoma: A proof of principle, Br. J. Cancer 98 (1) (2008) 113–119. [PubMed: 18059395]

[2]. Rahman MM, Feng Y, Yankeelov TE, Oden JT, A fully coupled space-time multiscale modeling framework for predicting tumor growth, Comput. Methods Appl. Mech. Engrg 320 (2017) 261–286.

[3]. Mohamed A, Shen D, Davatzikos C, Deformable registration of brain tumor images via a statistical model of tumor-induced deformation, Med Image. Anal 10 (2006) 752–763. [PubMed: 16860588]

[4]. Kwon D, Niethammer M, Akbari H, Bilello M, Davatzikos C, Pohl KM, PORTR: Pre-operative and post-recurrence brain tumor registration, IEEE Trans. Med. Imag 33 (3) (2014) 651–667.

[5]. Zacharaki EI, Hogea CS, Shen D, Biros G, Davatzikos C, Non-diffeomorphic registration of brain tumor images by simulating tissue loss and tumor growth, NeuroImage 46 (3) (2009) 762–774. [PubMed: 19408350]

[6]. Bakas S, Zeng Z, Sotiras A, Rathore S, Akbari H, Gaonkar B, Rozycki M, Pati S, Davatzikos C, GLISTRboost: Combining multimodal MRI segmentation, registration, and biophysical tumor

growth modeling with gradient boosting machines for glioma segmentation, Brain Lesion 9556 (2015) 144–155.

[7]. Gooya A, Pohl KM, Bilello M, Cirillo L, Biros G, Melhem ER, Davatzikos C, GLISTR: Glioma image segmentation and registration, IEEE Trans. Med. Imag 31 (10) (2013) 1941–1954.

[8]. Prastawa M, Buillitt E, Gerig G, Simulation of brain tumors in MR images for evaluation of segmentation efficacy, Med. Image Anal 13 (2009) 297–311. [PubMed: 19119055]

[9]. Swanson KR, Alvord EC, Murray JD, A quantitative model for differential motility of gliomas in grey and white matter, Cell Proliferation 33 (5) (2000) 317–330. [PubMed: 11063134]

[10]. Swanson KR, Alvord EC, Murray JD, Virtual brain tumours (gliomas) enhance the reality of medical imaging and highlight inadequacies of current therapy, Br. J. Cancer 86 (1) (2002) 14–18. [PubMed: 11857005]

[11]. Murray JD, Mathematical Biology, Springer-Verlag, New York, 1989.

[12]. Mang A, Toma A, Schuetz TA, Becker S, Eckey T, Mohr C, Petersen D, Buzug TM, 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 (7) (2012) 4444–4459, 10.1118/1.4722749. [PubMed: 22830777]

[13]. Jackson PR, Juliano J, Hawkins-Daarud A, Rockne RC, Swanson KR, Patient-specific mathematical neuro-oncology: Using a simple proliferation and invasion tumor model to inform clinical practice, Bull Math Biol 77 (5) (2015) 846–856, 10.1007/s11538-015-0067-7. [PubMed: 25795318]

[14]. Lima EABF, Oden JT, Hormuth DA, Yankeelov TE, Almeida RC, Selection, calibration, and validation of models of tumor growth, Math. Models Methods Appl. Sci 26 (12) (2016) 2341–2368. [PubMed: 28827890]

[15]. Kyriacou S, Davatzikos C, et al., Nonlinear elastic registration of brain images with tumor pathology using a biomechanical model, IEEE Trans. Med. Imaging 18 (1999) 580–592. [PubMed: 10504092]

[16]. Yankeelov TE, Atuegwu N, Hormuth D, Weis JA, Barnes SL, Miga MI, Rericha EC, Quaranta V, Clinically relevant modeling of tumor growth and treatment response, Sci. Translational Med 5 (187) (2013) 187ps9–187ps9.

[17]. Gholami A, Scheufele K, Mang A, Mehl M, Biros G, A framework for scalable biophysics-based image analysis, in: Proc ACM/IEEE Conference on Supercomputing, ACM, New York, NY, USA, 2017, pp. 19:1–19:13, URL http://doi.acm.org/10.1145/3126908.3126930.

[18]. Biegler LT, Ghattas O, Heinkenschloss M, van Bloemen Waanders B, Large-Scale PDE-Constrained Optimization, Springer, 2003.

[19]. Borzì A, Schulz V, Computational Optimization of Systems Governed by Partial Differential Equations, SIAM, Philadelphia, Pennsylvania, US, 2012.

[20]. Herzog R, Kunisch K, Algorithms for PDE-constrained optimization, GAMM-Mitt 33 (2) (2010) 163–176.

[21]. Hinze M, Pinnau R, Ulbrich M, Ulbrich S, Optimization with PDE Constraints, Springer, Berlin, DE, 2009.

[22]. Angelini ED, Clatz O, Mandonnet E, Konukoglu E, Capelle L, Duffau H, Glioma dynamics and computational models: A review of segmentation, registation, in silico growth algorithms and their clinical applications, Curr. Med. Imaging Rev 3 (4) (2007) 262–276.

[23]. Bauer S, Wiest R, Nolte L-P, Reyes M, A survey of MRI-based medical image analysis for brain tumor studies, Phys. Med. Biol 58 (13) (2013) R97. [PubMed: 23743802]

[24]. Mang A, Gholami A, Davatzikos C, Biros G, PDE constrained optimization in medical image analysis, Opt. Eng (2017) in review.

[25]. Modersitzki J, Numerical Methods for Image Registration, Oxford University Press, New York, 2004.

[26]. Sotiras A, Davatzikos C, Paragios N, Deformable medical image registration: A survey, IEEE Trans. Med. Imaging 32 (7) (2013) 1153–1190. [PubMed: 23739795]

[27]. Henn S, Hömke L, Witsch K, Lesion preserving image registration with application to human brains, in: Lect Notes Comput Sc, in: Lect Notes Comput Sc, vol. 3175, 2004, pp. 496–503.

[28]. Stefanescu R, Commowick O, Maladain G, Bondiau PY, Ayache N, Pennec X, Non-rigid atlas to subject registration with pathologies for conformal brain radiotherapy, in: Lect Notes Comput Sc, Vol. 3216, 2004, pp. 704–711.

[29]. Brett M, Leff AP, Rorden C, Ashburner J, Spatial normalization of brain images with focal lesions using cost function masking, NeuroImage 14 (2) (2001) 486–500. [PubMed: 11467921]

[30]. Parisot S, Wells W, Chemouny S, Duffau H, Paragios N, Concurrent tumor segmentation and registration with uncertainty-based sparse non-uniform graphs, Med. Image Anal 18 (4) (2014) 647–659. [PubMed: 24717540]

[31]. Li X, Registration of Images with Varying Topology using Embedded Maps (Ph.D. thesis), Department of Electrical & Computer Engineering, Virginia Polytechnic Institute and State University, Blacksburg, Virginia, US, 2010.

[32]. Li X, Long X, Laurienti P, Wyatt C, Registration of images with varying topology using embedded maps, IEEE Trans. Med. Imaging 31 (3) (2012) 749–765. [PubMed: 22194239]

[33]. Mang A, Biros G, An inexact Newton–Krylov algorithm for constrained diffeomorphic image registration, SIAM J. Imaging Sci 8 (2) (2015) 1030–1069, 10.1137/140984002. [PubMed: 27617052]

[34]. Mang A, Biros G, Constrained H1-regularization schemes for diffeomorphic image registration, SIAM J. Imaging Sci 9 (3) (2016) 1154–1194, 10.1137/15M1010919. [PubMed: 29075361]

[35]. Mang A, Gholami A, Biros G, Distributed-memory large-deformation diffeomorphic 3D image registration, in: Proc ACM/IEEE Conference on Supercomputing, 2016, 10.1109/SC.2016.71.

[36]. Mang A, Ruthotto L, A Lagrangian Gauss–Newton–Krylov solver for mass- and intensity-preserving diffeomorphic image registration, SIAM J. Sci. Comput 39 (5) (2017) B860–B885, 10.1137/17M1114132. [PubMed: 29097881]

[37]. Mang A, Biros G, A Semi-Lagrangian two-level preconditioned Newton–Krylov solver for constrained diffeomorphic image registration, SIAM J. Sci. Comput 39 (6) (2017) B1064–B1101, 10.1137/16M1070475. [PubMed: 29255342]

[38]. Christensen GE, Rabbitt RD, Miller MI, Deformable templates using large deformation kinematics, IEEE Trans. Image Process 5 (10) (1996) 1435–1447. [PubMed: 18290061]

[39]. Trouvé A, Diffeomorphism groups and pattern matching in image analysis, Int. J. Comput. Vis 28 (3) (1998) 213–221.

[40]. Beg MF, Miller MI, Trouvé A, Younes L, Computing large deformation metric mappings via geodesic flows of diffeomorphisms, Int. J. Comput. Vis 61 (2) (2005) 139–157.

[41]. Gholami A, Mang A, Biros G, An inverse problem formulation for parameter estimation of a reaction-diffusion model of low grade gliomas, J. Math. Biol 72 (1) (2016) 409–433, 10.1007/s00285-015-0888-x. [PubMed: 25963601]

[42]. Gholami A, Fast Algorithms for Biophysically-Constrained Inverse Problems in Medical imaging (Ph.D. thesis), The University of Texas at Austin, 2017.

[43]. Miga M, Paulsen K, et al., Initial in-vivo analysis of 3D heterogeneous brain computations for model-updated image-guided neurosurgery, LNCS 1496 (1998) 743–752.

[44]. Clatz O, Sermesant M, Bondiau PY, Delingette H, Warfield SK, Malandain G, Ayache N, Realistic simulation of the 3D growth of brain tumors in MR images coupling diffusion with biomechanical deformation, IEEE Trans. Med. Imaging 24 (10) (2005) 1334–1346. [PubMed: 16229419]

[45]. Hogea CS, Biros G, Abraham F, Davatzikos C, A robust framework for soft tissue simulations with application to modeling brain tumor mass effect in 3D MR images, Phys. Med. Biol 52 (23) (2007) 6893. [PubMed: 18029982]

[46]. Harpold HL, Alvord EC, Swanson KR, The evolution of mathematical modeling of glioma proliferation and invasion, J Neuropathol. Exp. Neurol 66 (2007) 1–9. [PubMed: 17204931]

[47]. Konukoglu E, Clatz O, Menze BH, Weber MA, Stieltjes B, Mandonnet E, Delingette H, Ayache N, Image guided personalization of reaction-diffusion type tumor growth models using modified anisotropic eikonal equations, IEEE Trans. Med. Imaging 29 (1) (2010) 77–95. [PubMed: 19605320]

[48]. Konukoglu E, Clatz O, Bondiau PY, Delingette H, Ayache N, Extrapolating glioma invasion margin in brain magnetic resonance images: Suggesting new irradiation margins, Med. Image Anal 14 (2) (2010) 111–125. [PubMed: 20042359]

[49]. Le M, Delingette H, Kalapathy-Cramer J, Gerstner ER, Batchelor T, Unkelbach J, Ayache N, MRI based Bayesian personalization of a tumor growth model, IEEE Trans. Med. Imaging 35 (10) (2016) 2329–2339. [PubMed: 27164582]

[50]. Mosayebi P, Bobzas D, Murtha A, Jagersand M, Tumor invasion margin on the Riemannian space of brain fibers, Med. Image Anal 16 (2) (2012) 361–373. [PubMed: 22154876]

[51]. Menze sB.H, Leemput KV, Konukoglu E, Honkela A, Weber M-A, Ayache N, Golland P, A generative approach for image-based modeling of tumor growth, in: Proc Information Processing in Medical Imaging, Vol. 2, 2011, pp. 735–747.

[52]. Colin T, Iollo A, Lagaert J-B, Saut O, An inverse problem for the recovery of the vascularization of a tumor, J. Inverse Ill-posed Probl 22 (6) (2014) 759–786.

[53]. Hogea C, Davatzikos C, Biros G, An image-driven parameter estimation problem for a reaction-diffusion glioma growth model with mass effect, J. Math. Biol 56 (2008) 793–825. [PubMed: 18026731]

[54]. Knopoff DA, Fernández DR, Torres GA, Turner CV, Adjoint method for a tumor growth PDE-constrained optimization problem, Comput. Math. Appl 66 (6) (2013) 1104–1119.

[55]. Knopoff D, Fernández DR, Torres GA, Turner CV, A mathematical method for parameter estimation in a tumor growth model, Comput. Appl. Math 36 (1) (2017) 733–748.

[56]. Liu Y, Sadowki SM, Weisbrod AB, Kebebew E, Summers RM, Yao J, Patient specific tumor growth prediction using multimodal images, Med. Image Anal 18 (3) (2014) 555–566. [PubMed: 24607911]

[57]. Quiroga AAI, Fernández D, Torres GA, Turner CV, Adjoint method for a tumor invasion PDE-constrained optimization problem in 2D using adaptive finite element method, Appl. Math. Comput 270 (2015) 358–368.

[58]. Quiroga AAI, Torres GA, Fernández DR, Turner CV, Nonlinear optimization for a tumor invasion PDE model, Comput. Appl. Math (2016) 1–15.

[59]. Wong KCL, Summers RM, Kebebew E, Yoa J, Tumor growth prediction wit reaction-diffusion and hyperelastic biomechanical model by physiological data fusion, Med. Image Anal 25 (1) (2015) 72–85. [PubMed: 25962846]

[60]. Kwon D, Shinohara RT, Akbari H, Davatzikos C, Combining generative models for multifocal glioma segmentation and registration, in: Proc Medical Image Computing and Computer-Assisted Intervention, in: LNCS 8673, 2014.

[61]. Lima E, Oden J, Wohlmuth B, Shahmoradi A, Hormuth D II, Yankeelov T, Scarabosio L, Horger T, Selection and validation of predictive models of radiation effects on tumor growth based on noninvasive imaging data, Comput. Methods Appl. Mech. Engrg 327 (2017) 277–305.

[62]. Flath H, Wilcox L, Akçelik V, Hill J, van Bloemen Waanders B, Ghattas O, Fast algorithms for Bayesian uncertainty quantification in large-scale linear inverse problems based on low-rank partial Hessian approximations, SIAM J. Sci. Comput 33 (1) (2011) 407–432.

[63]. Hogea C, Davatzikos C, Biros G, Brain-tumor interaction biophysical models for medical image registration, SIAM J. Imag. Sci 30 (6) (2008) 3050–3072.

[64]. Zacharaki EI, Hogea CS, Shen D, Biros G, Davatzikos C, Parallel optimization of tumor model parameters for fast registration of brain tumor images, in: Proc SPIE Medical Imaging: Image Processing, 2008, pp. 69140K1–69140K10.

[65]. Zacharaki EI, Hogea CS, Biros G, Davatzikos C, A comparative study of biomechanical simulators in deformable registration of brain tumor images, IEEE Trans. Biomed. Eng 55 (3) (2008) 1233–1236. [PubMed: 18334420]

[66]. Tomer R, Ye L, Hsueh B, Deisseroth K, Advanced CLARITY for rapid and high-resolution imaging of intact tissues, Nat. Protoc 9 (7) (2014) 1682–1697. [PubMed: 24945384]

[67]. Gholami A, Hill J, Malhotra D, Biros G, AccFFT: A library for distributed-memory FFT on CPU and GPU architectures, 2016, arXiv e-prints, in review (arXiv preprint: http://arxiv.org/abs/1506.07933);.

[68]. Nocedal J, Wright SJ, Numerical Optimization, Springer, New York, New York, US, 2006.

[69]. Modersitzki J, FAIR: Flexible Algorithms for Image Registration, SIAM, Philadelphia, Pennsylvania, US, 2009.

[70]. Gill PE, Murray W, Wright MH, Practical Optimization, Academic Press, Waltham, Massachusetts, US, 1981.

[71]. Dembo RS, Steihaug T, Truncated-Newton algorithms for large-scale unconstrained optimization, Math. Program 26 (2) (1983) 190–212.

[72]. Eisentat SC, Walker HF, Choosing the forcing terms in an inexact Newton method, SIAM J. Sci. Comput 17 (1) (1996) 16–32.

[73]. Grama A, Gupta A, Karypis G, Kumar V, An Introduction to Parallel Computing: design and Analysis of Algorithms, second ed., Addison Wesley, 2003.

[74]. Czechowski K, Battaglino C, McClanahan C, Iyer K, Yeung P-K, Vuduc R, On the communication complexity of 3D FFTs and its implications for exascale, in: Proc ACM/IEEE Conference on Supercomputing, 2012, pp. 205–214.

[75]. Balay S, Abhyankar S, Adams MF, Brown J, Brune P, Buschelman K, Dalcin L, Eijkhout V, Gropp WD, Kaushik D, Knepley MG, McInnes LC, Rupp K, Smith BF, Zampini S, Zhang H, PETSc users manual, Tech. Rep ANL-95/11 - Revision 3.7, Argonne National Laboratory, 2016.

[76]. Munson T, Sarich J, Wild S, Benson S, McInnes LC, TAO 3.6 Users Manual, Argonne National Laboratory, Mathematics and Computer Science Division, 2015.

[77]. Gooya A, Pohl K, Bilello M, Biros G, Davatzikos C, Joint Segmentation and Deformable Registration of Brain Scans Guided by a Tumor Growth Model, in: Medical Image Computing and Computer-Assisted Intervention – MICCAI 2011, in: Lecture Notes in Computer Science, vol. 6892, Springer Berlin Heidelberg, 2011, pp. 532–540.

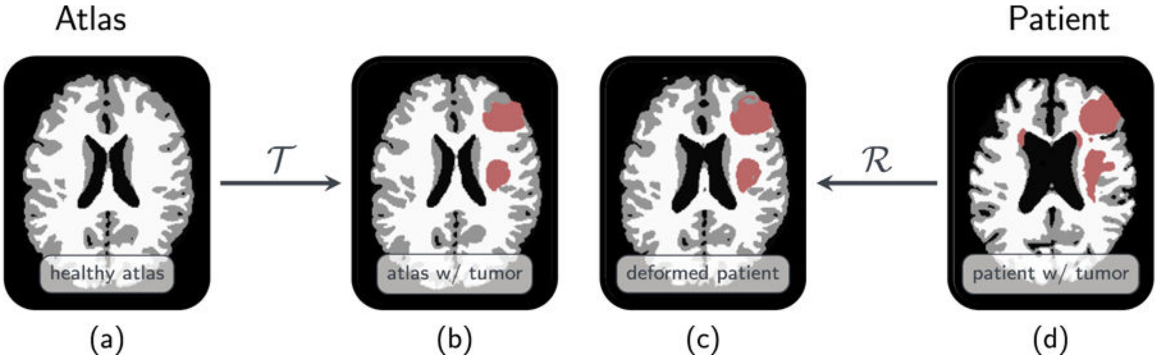

Fig. 1.

Here we summarize the joint registration and biophysical inversion in SIBIA. (a) shows the segmented healthy brain (the atlas), (b) tumor-bearing atlas brain generated by biophysical simulation, (c) patient tumor registered to atlas (in atlas space), (d) tumor-bearing patient brain. The inputs are the images (a) and (d). The outputs are the tumor growth parameters, the registration parameters, and images (b) and (c). Here is τ the forward tumor map, which given image (a) grows a tumor and generates image (b). Similarly, $\mathscr{R}$ is the forward registration map, which given image (d) generates image (c). In SIBIA, we compute tumor-growth parameters and registration parameters so that images (b) and (c) are as similar as possible (in the L^2 norm). In this example, the biophysical tumor growth parameters are the initial conditions for a reaction–diffusion equation. The registration is parameterized using an Eulerian framework and a velocity field that is used to advect image (d) to image (c).

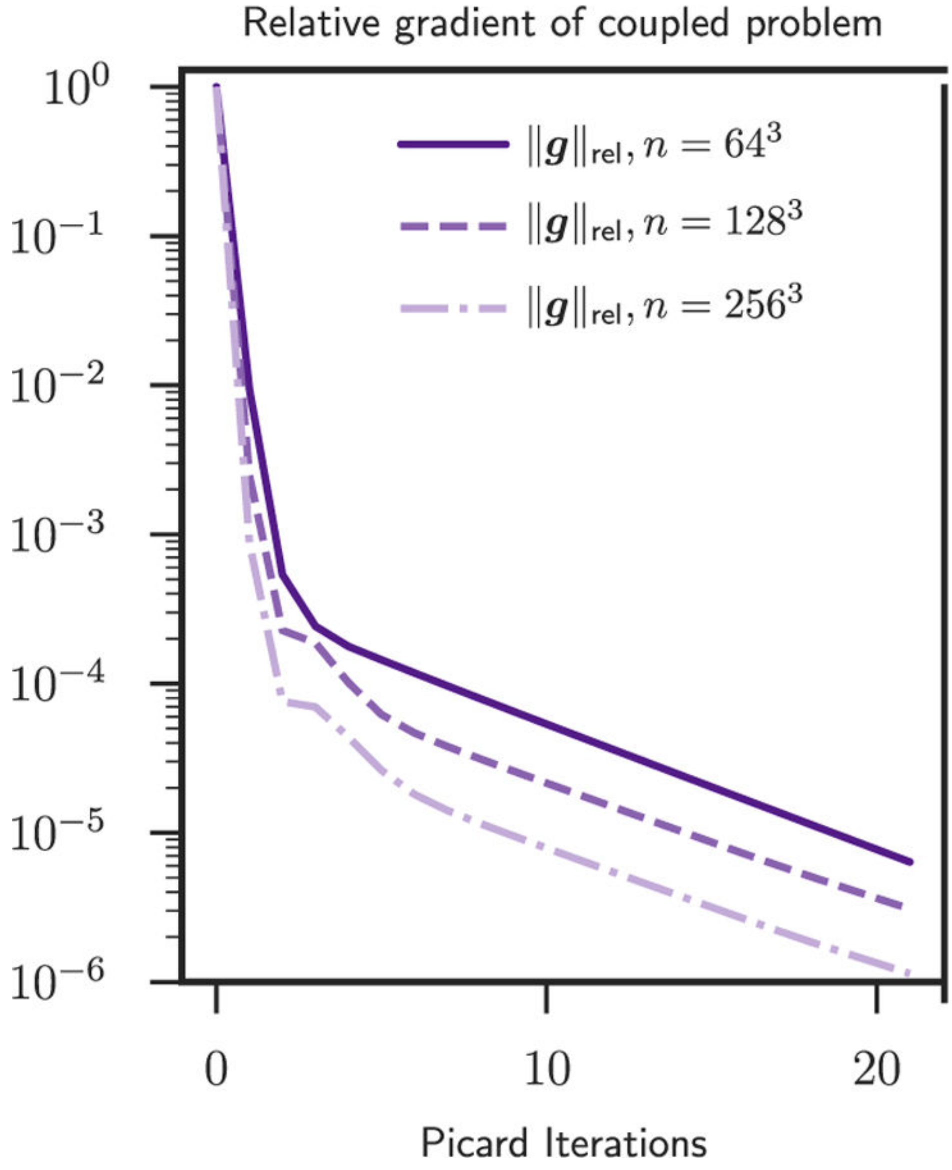

Fig. 2.

Here we summarize the norm of the reduced gradient of (3) as a function of the Picard iteration number for the **sinusoidal analytic tumor / analytic velocity (SIN)** test case. The exact numbers are reported in Table 3. The tree different curves correspond to the three different spatial resolutions for the velocity and tumor parameters. We use $N_i \in \{64, 128, 256\}$ points per dimension for the velocity and analogously $\boldsymbol{p} \in \{8, 64, 512\}$. We perform 20 Picard iterations after termination of the parameter-continuation scheme for the regularization parameter of the registration solver. We have two observations. First, the convergence of our scheme does not deteriorate as we increase the mesh resolution. As we refine the mesh we can reduce the gradient further and the number of iterations does not deteriorate (as it would be the case for a steepest descent method). Second, the convergence of our scheme eventually slows down for all mesh sizes.

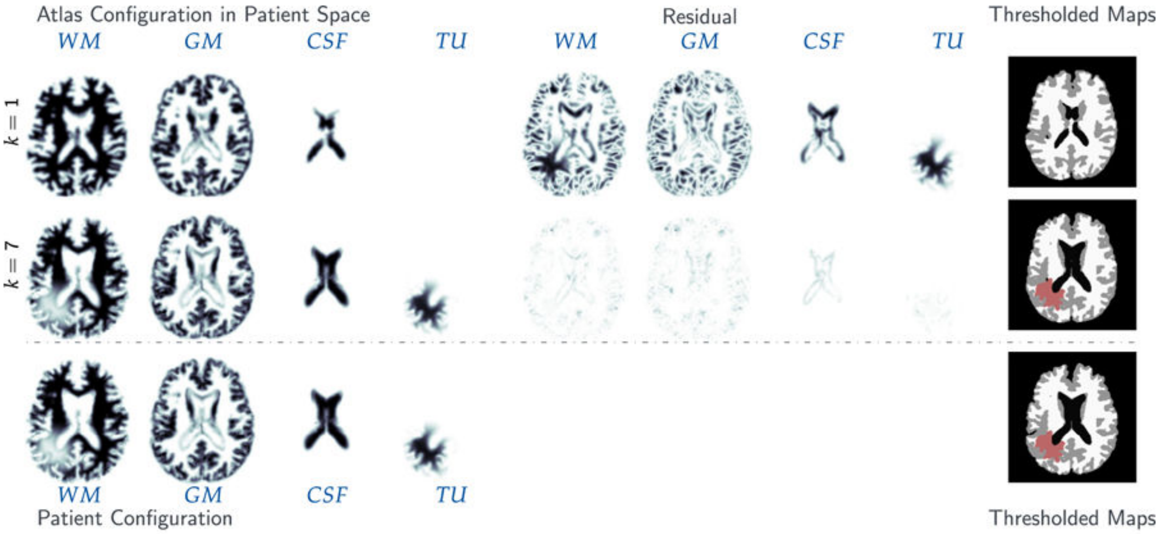

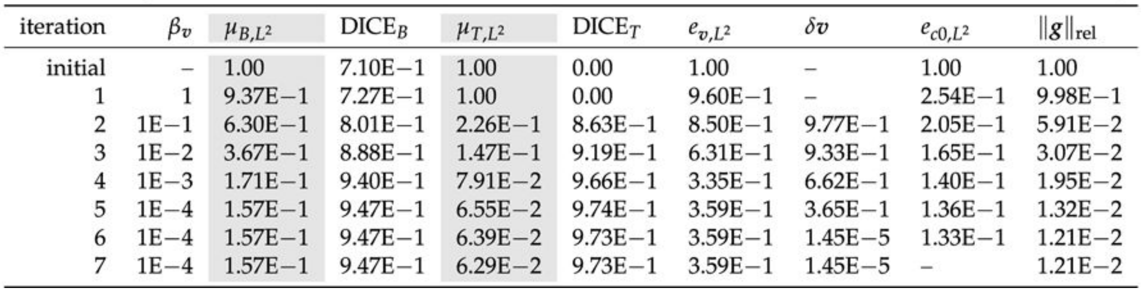

iteration	β_v	μ_{B,L^2}	DICE$_B$	μ_{T,L^2}	DICE$_T$	e_{v,L^2}	δv	e_{c0,L^2}	$\|g\|_{\text{rel}}$
initial	–	1.00	7.10E−1	1.00	0.00	1.00	–	1.00	1.00
1	1	9.37E−1	7.27E−1	1.00	0.00	9.60E−1	–	2.54E−1	9.98E−1
2	1E−1	6.30E−1	8.01E−1	2.26E−1	8.63E−1	8.50E−1	9.77E−1	2.05E−1	5.91E−2
3	1E−2	3.67E−1	8.88E−1	1.47E−1	9.19E−1	6.31E−1	9.33E−1	1.65E−1	3.07E−2
4	1E−3	1.71E−1	9.40E−1	7.91E−2	9.66E−1	3.35E−1	6.62E−1	1.40E−1	1.95E−2
5	1E−4	1.57E−1	9.47E−1	6.55E−2	9.74E−1	3.59E−1	3.65E−1	1.36E−1	1.32E−2
6	1E−4	1.57E−1	9.47E−1	6.39E−2	9.73E−1	3.59E−1	1.45E−5	1.33E−1	1.21E−2
7	1E−4	1.57E−1	9.47E−1	6.29E−2	9.73E−1	3.59E−1	1.45E−5	–	1.21E−2

Fig. 3.

Results for the **analytic tumor / analytic velocity reaction-only (ATAV-REAC)** test case; ground truth: ($\rho_f = 15$, $\rho_w = 1$, $\rho_g = 0$, $k_f = 0$, $\boldsymbol{p} = \boldsymbol{p}^\star$, $\boldsymbol{v} = -\boldsymbol{v}^\star$). The figure shows probability maps for the labels of the healthy atlas brain and the patient brain with tumor generated from a tumor grown in the atlas and known atlas to patient advection velocity (see text for details; axial-slice 64). We show the initial configuration for the problem (top row; iteration $k = 1$), the final configuration after joint registration and tumor inversion (middle row; iteration $k = 7$; the atlas image probability maps are transported to the patient space), and the target patient data (reference image; bottom row). Each row contains (from left to right) the probability maps for WM, GM, CSF, and TU, the residual differences (if available) between the probability maps, and a hard segmentation based on the given probabilities for the individual tissue classes. The table on the bottom provides quantitative results for the inversion. We report the average mismatch for the probability maps for the brain tissue labels μ_{B,L^2} and the tumor μ_{T,L^2}, the mean DICE coefficient for brain tissue DICE$_B$ and tumor DICE$_T$, respectively. The reconstruction quality is given in terms of convergence of $\boldsymbol{v}^k$ and $c_A^k(0)$ towards the ground truth $\boldsymbol{v}^\star$ and $c_A^k(0)$, respectively ($e_{\boldsymbol{v},L^2}$ and e_{c0,L^2}). We cannot expect this error to go to zero for several reasons. First, we lose information when we construct the test case (zero gradients in the intensity of the image), second our numerical solver introduces errors (in particular, the solver for the transport equations). We in addition to that report the change in update in the velocity $\boldsymbol{v}$ across successive iterations

$\delta v = \| v^k - v^{k-1} \| / \| v^{k-1} \|$. Finally, we also list the relative norm of the gradient for the coupled problem in (3) $\left(\|g\|_{\text{rel}}\right)$.

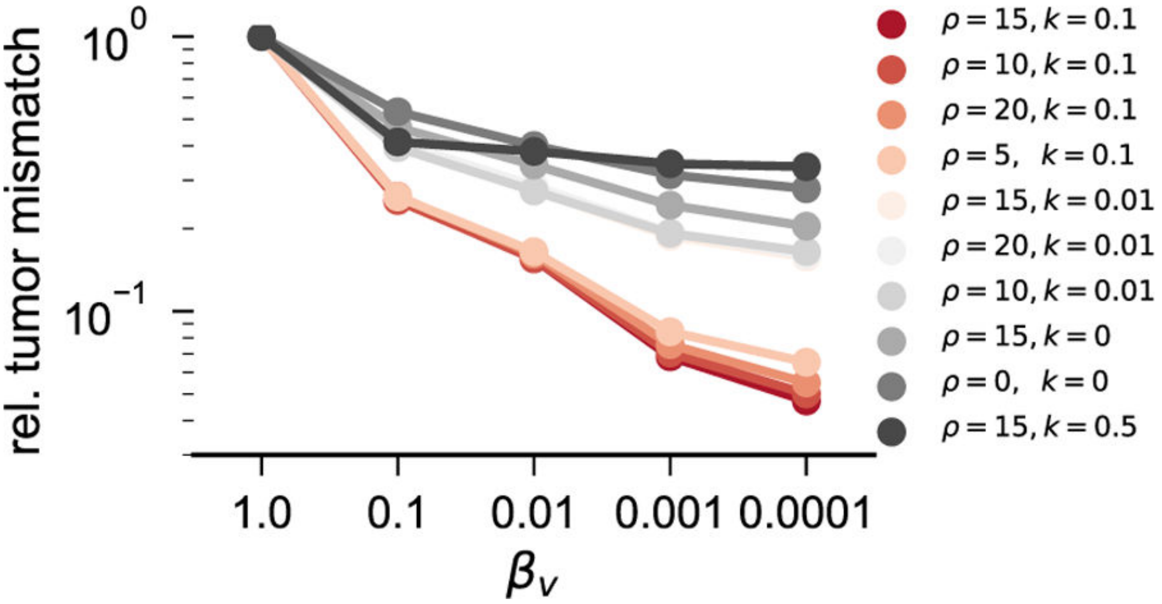

Fig. 4.

Mismatch reduction as a function of the regularization parameter β_v for the **analytic tumor with analytic velocity, diffusion and low-dimensional initial condition (ATAV-LD)** test case; ground truth: ($\rho_f = 15$, $\rho_w = 1$, $\rho_g = 0$, $k_f = 1.00\text{E}{-1}$, $k_w = 1$, $k_g = 0$, $\boldsymbol{p} = \boldsymbol{p}^\star$ (in patient domain), $\boldsymbol{v}$ N/A). We use an initial condition parameterized with only $n_{\boldsymbol{p}} = 8$ Gaussians ($\sigma = \pi/15$) and an analytic tumor with non-zero diffusion. The plot shows the relative mismatch for the tumor probability map μ_{T,L^2} in patient space. Note that we reduce β_v by a factor of ten in each Picard iteration; the plot shows the mismatch reduction over the Picard iterations if read from right to left.

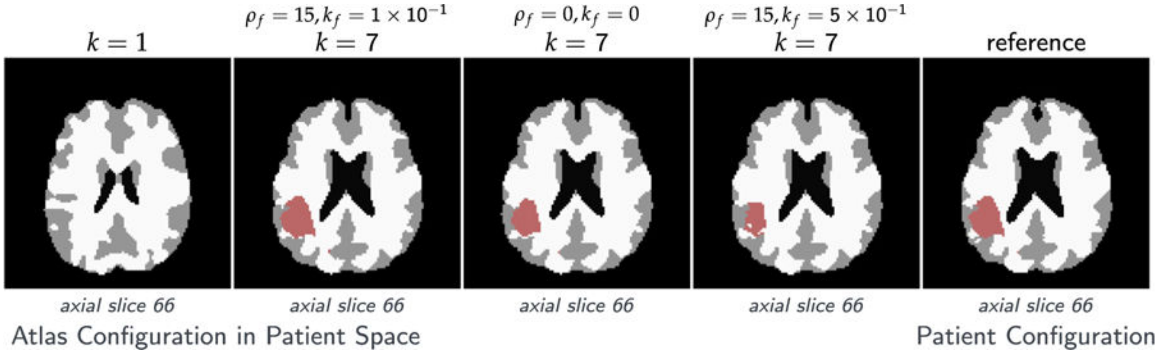

Fig. 5.

Tumor and brain labels for the **analytic tumor with analytic velocity with non-zero diffusion and low-dimensional initial condition (ATAV-LD)** test case; ground truth: ($\rho_f = 15$, $\rho_w = 1$, $\rho_g = 0$, $k_f = 1.00\text{E}{-1}$, $k_w = 1$, $k_g = 0$, $\mathbf{p} = \mathbf{p}^\star$ (in patient domain), $\mathbf{v}$ N/A); for the ‘correct’ tumor parameters $\rho_f = 15$, $k_f = 1.00\text{E}{-1}$, and the two different settings $\rho_f = 15$, $k_f = 5.00\text{E}{-1}$ and $\rho_f = k_f = 0$.

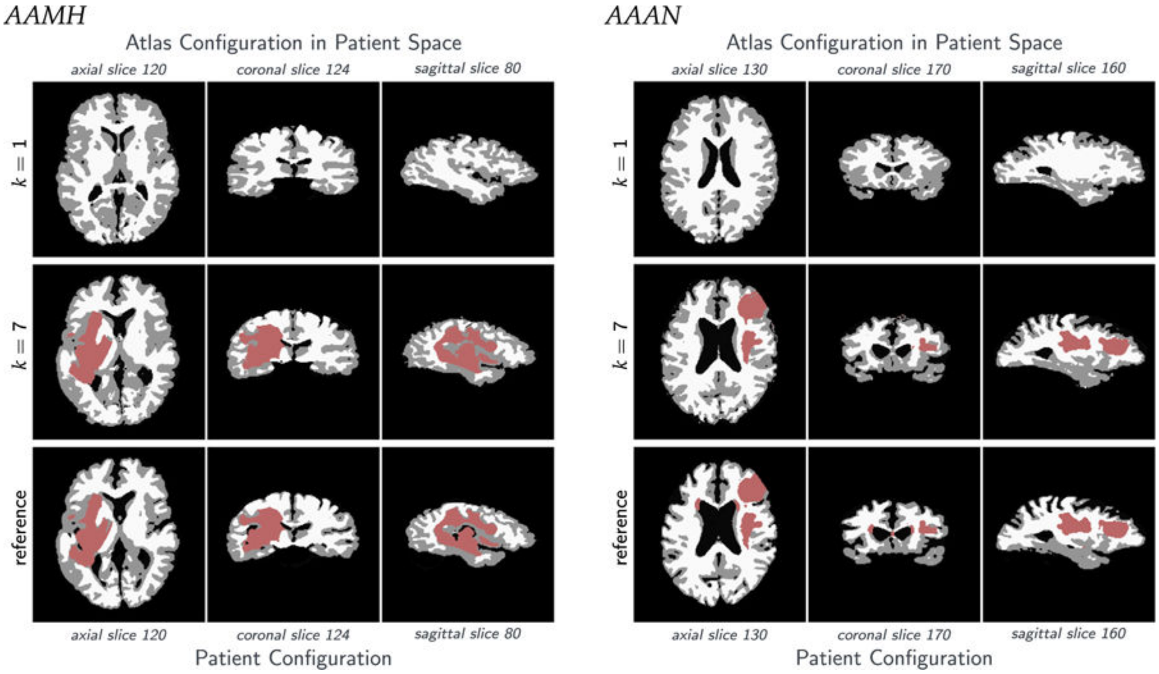

Fig. 6.

Tumor and brain labels for the **real tumor with real velocity (RTRV)** test case; ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A); patients **AAMH, AAAN**. We set the parameters for the tumor solver to $\rho_f = 15$, $k_f = 0$ (reaction-only). We use $n_{\boldsymbol{p}} = 343$ Gaussians for the inversion. The top row shows the original atlas image. The bottom row shows the patient image. The row in the middle shows the solution for our coupled scheme.

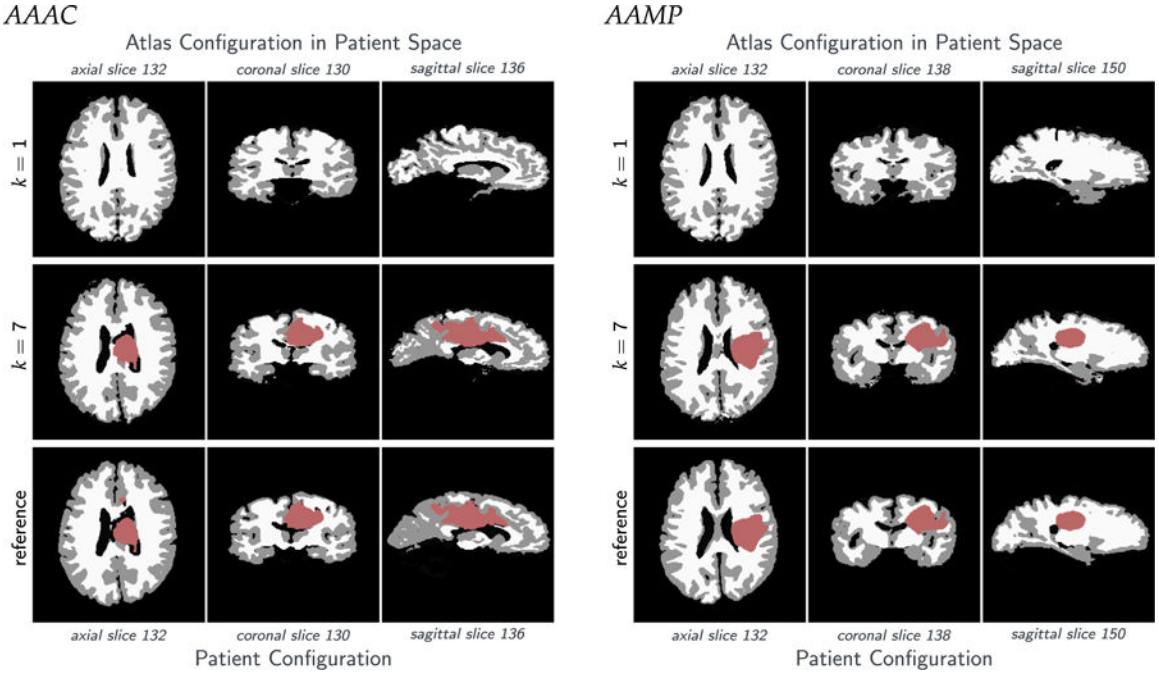

Fig. 7.

Tumor and brain geometry for the **real tumor with real velocity (RTRV)** test case; ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A); patients **AAAC, AAMP**. We set the parameters for the tumor solver to $\rho_f = 15$, $k_f = 0$ (reaction-only). We use $n_{\boldsymbol{p}} = 343$ Gaussians for the inversion. The top row shows the original atlas image. The bottom row shows the patient image. The row in the middle shows the solution for our coupled scheme.

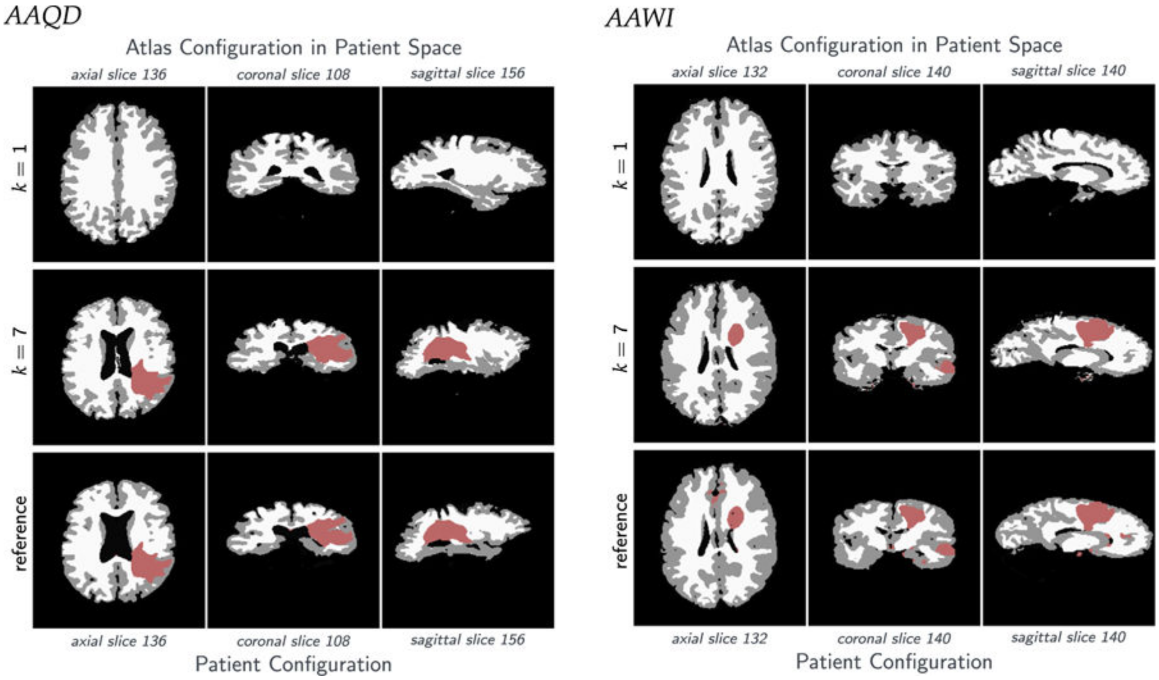

Fig. 8.

Tumor and brain geometry for the **real tumor with real velocity (RTRV)** test case; ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A); patients **AAQD, AAWI**. We set the parameters for the tumor solver to $\rho_f = 15$, $k_f = 0$ (reaction-only). We use $n_{\boldsymbol{p}} = 343$ Gaussians for the inversion. The top row shows the original atlas image. The bottom row shows the patient image. The row in the middle shows the solution for our coupled scheme.

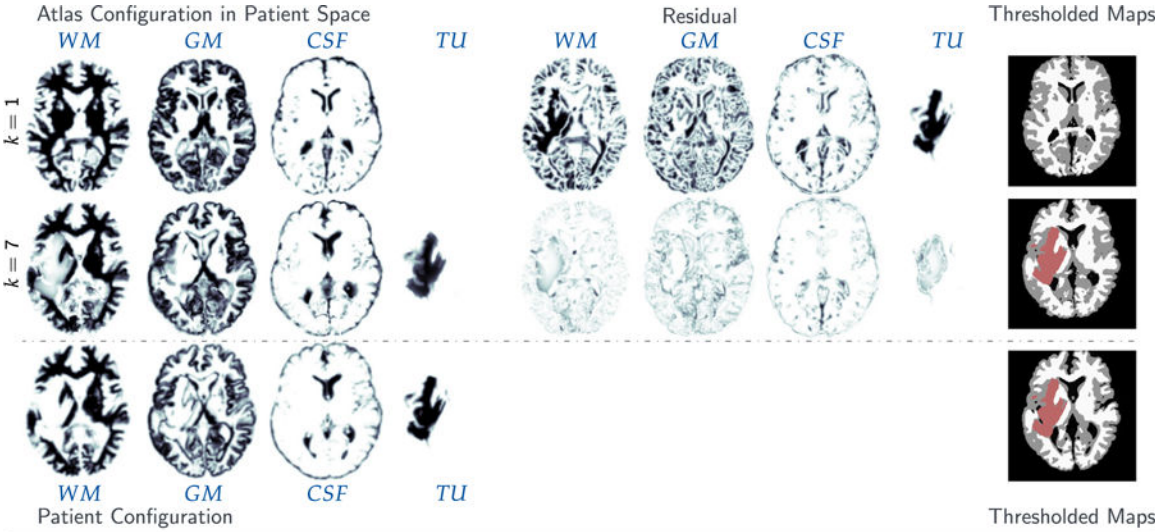

iterations	β_v	μ_{B,L^2}	DICE$_B$	μ_{T,L^2}	DICE$_T$	$\|g\|_{\text{rel}}$	T^{it} [s]	T^{tu}_{inv} [%]	T^{reg}_{inv} [%]
initial	–	1.00	4.36E−1	1.00	0.00	1.00	–	–	–
1	1	9.59E−1	4.62E−1	1.00	0.00	9.99E−1	9.80E+1	74.6	15.0
2	1E−1	7.42E−1	5.60E−1	3.90E−1	8.43E−1	9.06E−2	4.55E+1	51.2	47.7
3	1E−2	5.05E−1	7.01E−1	3.13E−1	9.03E−1	4.61E−2	9.59E+1	25.8	73.7
4	1E−3	3.91E−1	7.66E−1	2.52E−1	9.39E−1	3.69E−2	1.18E+2	21.5	78.0
5	1E−4	3.60E−1	7.84E−1	2.20E−1	9.51E−1	3.62E−2	1.15E+2	21.2	78.3
6	1E−4	3.51E−1	7.89E−1	2.05E−1	9.55E−1	3.84E−2	8.99E+1	27.0	72.4
7	1E−4	3.45E−1	7.92E−1	1.95E−1	9.57E−1	3.92E−2	6.63E+1	0.0	99.3
final	1E−4	3.45E−1	7.92E−1	1.95E−1	9.57E−1	3.92E−2	6.28E+2	1.95E+2	4.20E+2

Fig. 9.

Results for the **real tumor/ real velocity (RTRV)** test case, ground truth ($\boldsymbol{\rho}$ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A) and the **AAMH** patient. The figure shows probability maps for the labels of the healthy atlas brain ($k = 1$; top row) and the AAMH patient (target) brain probability maps with tumor (bottom row), along with the reconstructed probability maps, i.e., the final result of our inversion algorithm ($k = 7$; middle row) (axial-slice 120). In the table, we report the (summed) mismatch for the brain tissue probability maps (μ_{B,L^2}) and tumor probability map (μ_{T,L^2}) in patient space, the mean Dice coefficient for the hard segmentation corresponding to the brain tissue (DICE$_B$) and the tumor DICE$_T$, respectively, as well as the relative norm of the gradient for the coupled problem (3) ($\|g\|_{\text{rel}}$). We also report the run time per iteration in seconds (T^{it}), and the percentages (T^{tu}_{inv}) and (T^{reg}_{inv}) of this run time spent in the tumor solver and the image registration solver, respectively. Note that the latter sums up to less than 100% as we do not explicitly measure time spent in additional coupling functionality and in the forward solvers. The last row shows the final state and summed absolute timings for the respective solvers in seconds.

Table 1

Common parameters used in all brain data based test cases: opttol_R, opttol_T are the convergence tolerances for registration and tumor inversion ((7l) and (7k)); $\boldsymbol{\beta_p}$ is the regularization parameter for the tumor inversion (see (9a)); d is the spacing between the Gaussian basis function, σ their standard deviation; $\beta_{\boldsymbol{v}}^0$ and $\beta_{\boldsymbol{v}}^{\min}$ are the initial and final values for the β-continuation scheme as described in 4.1 applied in image registration, determined based on which values have been shown to yield best results in numerical tests for the RTRV test problem.

$\mathbf{opttol}_R$	$\mathbf{opttol}_T$	$\boldsymbol{\beta_p}$	$\boldsymbol{d}$	β_v^0	$\beta_{\boldsymbol{v}}^{\min}$
1E−3	1E−3	2.50E−4	1.5σ	1	1E−3

Table 2

Summary of the parameters for the generation of the synthetic test cases and the inversion. We report values for the following parameters: N_i , $i = 1, 2, 3$ denotes the grid size used for the discretization of the problem; ρ_w and ρ_g are the characteristic reaction factors for white and gray matter, ρ_f is the overall reaction scaling factor; k_w and k_g are the characteristic diffusion parameters for white and gray matter, k_f the overall scaling parameter for the isotropic part of the inhomogeneous diffusion coefficient for net migration of cancerous cells into surrounding tissue; $n_{\boldsymbol{p}}$ is the number of Gaussian basis functions used for the parametrization of the tumor initial condition, σ is the standard deviation of the associated Gaussian basis functions, and $\text{maxit}_i = (\text{maxit}_{i,N}, \text{maxit}_{i,K})$ 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,.

	SIN	ATAV-REAC	ATAV-DIF	ATAV-LD	RTRV
N_i	{64, 128, 256}	128	128	128	256
ρ_w	1	1	1	1	1
ρ_g	0	0	0	0	0.2
ρ_f	15	15	{5, 10, 15}	{5, 10, 15}	{0, 5, 10, 15}
k_w	1	1	1	1	1
k_g	0	0	0	0	0.1
k_f	0	0	{0, 1E−2}	{1E−2, 1E−1}	{0,1E−2}
$n_{\boldsymbol{p}}$	{8, 64, 512}	125	125	8	{125, 343}
σ	$\pi/10$	$\pi/10$	$\pi/10$	$\pi/15$	Auto
maxit_T	(100, 100)	(50, 100)	(30, 60)	(30, 60)	(30, 30)
maxit_R	(100, 100)	(50, 80)	(50, 80)	(50, 80)	(10, 20)

Table 3

Results for the **sinusoidal analytic tumor / analytic velocity (SIN)** test case; ground truth: ($\rho_f = 15$, $\rho_w = 1$, $\rho_g = 0$, $k_f = 0$, $\boldsymbol{p} = \boldsymbol{p}^{\star}$, $\boldsymbol{v} = -\boldsymbol{v}^{\star}$). We report convergence of the Picard iteration scheme for the SIN test case for three runs with increasingly refined mesh. For the registration, we increase the number of inversion variables by a factor of eight starting from $N_i = 64$ points per dimension, to $N_i = 128$ and $N_i = 256$; for the tumor parameters, we choose $n_{\boldsymbol{p}}$ to be 8,64, and 512 respectively. To study the convergence of our scheme, we perform 20 Picard iterations after termination of the parameter-continuation scheme for the regularization parameter of the registration solver. The relative objective function value $\|\mathcal{J}\|_{\text{rel}}$ and the relative gradient norm $\|g\|_{\text{rel}}$ of the coupled problem in (3) are reported for every iteration. Further, for each iteration we report the relative norm of the update $e_{c_0,L^2,rel}$ and $e_{\boldsymbol{v},L^2,rel}$ of the inversion variables $\boldsymbol{p}$ and $\boldsymbol{v}$ respectively.

		$N_i = 64, n_p = 8$				**$N_i = 128, n_p = 64$**				**$N_i = 256, n_p = 512$**			
It	β_v	$\|\mathcal{J}\|_{\text{rel}}$	$\|g\|_{\text{rel}}$	$e_{c0,L^2,rel}$	$e_{v,L^2,rel}$	$\|\mathcal{J}\|_{\text{rel}}$	$\|g\|_{\text{rel}}$	$e_{c0,L^2,rel}$	$e_{v,L^2,rel}$	$\|\mathcal{J}\|_{\text{rel}}$	$\|g\|_{\text{rel}}$	$e_{c0,L^2,rel}$	$e_{v,L^2,rel}$
init	-	1.00	1.00	–	–	1.00	1.00	–	–	1.00	1.00	–	–
1	1	6.67E−1	9.42E−3	–	–	6.73E−1	2.48E−3	–	–	6.67E−1	8.76E−4	–	–
2	1E−1	5.84E−1	5.37E−4	–	–	5.88E−1	2.27E−4	–	–	5.84E−1	7.58E−5	–	–
3	1E−2	3.08E−1	2.42E−4	8.08E−3	1.01E+1	3.08E−1	1.88E−4	2.66E−2	6.02	3.08E−1	6.97E−5	2.28E−2	4.97
4	1E−2	3.03E−1	1.77E−4	2.48E−3	1.51	3.04E−1	1.01E−4	2.35E−2	2.48	3.03E−1	4.37E−5	2.55E−2	2.94
5	1E−2	3.02E−1	1.44E−4	1.34E−3	2.60E−2	3.00E−1	6.20E−5	7.16E−3	1.33E−1	3.02E−1	2.62E−5	1.00E−2	1.48E−1
6	1E−2	3.01E−1	1.18E−4	8.72E−4	1.43E−2	3.02E−1	4.66E−5	2.51E−3	7.70E−2	3.01E−1	1.81E−5	3.61E−3	9.79E−2
7	1E−2	3.01E−1	9.64E−5	5.61E−4	7.25E−3	3.02E−1	3.78E−5	1.29E−3	4.61E−2	3.01E−1	1.41E−5	1.63E−3	6.43E−2
8	1E−2	3.00E−1	7.91E−5	3.62E−4	4.37E−3	3.02E−1	3.12E−5	8.04E−4	2.70E−2	3.00E−1	1.15E−5	9.29E−4	3.96E−2
9	1E−2	3.00E−1	6.50E−5	2.38E−4	3.30E−3	3.02E−1	2.59E−5	5.29E−4	1.61E−2	3.00E−1	9.51E−6	5.89E−4	2.48E−2
10	1E−2	3.00E−1	5.35E−5	1.57E−4	2.71E−3	3.02E−1	2.15E−5	3.54E−4	9.98E−3	3.00E−1	7.89E−6	3.89E−4	1.58E−2
11	1E−2	3.00E−1	4.40E−5	1.05E−4	2.31E−3	3.02E−1	1.79E−5	2.39E−4	6.42E−3	3.00E−1	6.56E−6	2.62E−4	1.02E−2
12	1E−2	3.00E−1	3.63E−5	6.99E−5	1.90E−3	3.02E−1	1.49E−5	1.63E−4	4.22E−3	3.00E−1	5.47E−6	1.78E−4	6.75E−3
13	1E−2	3.00E−1	2.99E−5	4.69E−5	1.58E−3	3.02E−1	1.24E−5	1.12E−4	2.09E−3	3.00E−1	4.55E−6	1.21E−4	4.54E−3
14	1E−2	3.00E−1	2.46E−5	3.16E−5	1.30E−3	3.02E−1	1.04E−5	7.63E−5	3.09E−3	3.00E−1	3.50E−6	8.33E−5	3.10E−3
15	1E−2	3.00E−1	2.03E−5	2.13E−5	1.07E−3	3.02E−1	8.64E−6	5.26E−5	1.35E−3	3.00E−1	3.17E−6	5.74E−5	2.19E−3
16	1E−2	3.00E−1	1.67E−5	1.44E−5	8.85E−4	3.02E−1	7.21E−6	3.63E−5	1.71E−3	3.00E−1	2.64E−6	3.96E−5	1.60E−3

		$N_i = 64, n_p = 8$				**$N_i = 128, n_p = 64$**				**$N_i = 256, n_p = 512$**			
It	β_v	$\|\mathcal{J}\|_{\mathrm{rel}}$	$\|g\|_{\mathrm{rel}}$	$e_{c0,L^2,rel}$	$e_{v,L^2,rel}$	$\|\mathcal{J}\|_{\mathrm{rel}}$	$\|g\|_{\mathrm{rel}}$	$e_{c0,L^2,rel}$	$e_{v,L^2,rel}$	$\|\mathcal{J}\|_{\mathrm{rel}}$	$\|g\|_{\mathrm{rel}}$	$e_{c0,L^2,rel}$	$e_{v,L^2,rel}$
17	1E−2	3.00E−1	1.38E−5	9.73E−6	7.29E−4	3.02E−1	6.05E−6	2.51E−5	9.31E−4	3.00E−1	2.21E−6	2.74E−5	1.20E−3
18	1E−2	3.00E−1	1.14E−5	6.59E−6	6.01E−4	3.02E−1	5.08E−6	1.59E−5	7.94E−4	3.00E−1	1.86E−6	1.89E−5	8.04E−4
19	1E−2	3.00E−1	9.37E−6	4.47E−6	4.95E−4	3.02E−1	4.30E−6	1.12E−5	6.62E−4	3.00E−1	1.57E−6	1.16E−5	6.50E−4
20	1E−2	3.00E−1	7.73E−6	3.03E−6	4.08E−4	3.02E−1	3.64E−6	7.61E−6	5.22E−4	3.00E−1	1.34E−6	8.04E−6	5.29E−4
21	1E−2	3.00E−1	6.37E−6	2.06E−6	3.36E−4	3.02E−1	3.09E−6	5.23E−6	4.27E−4	3.00E−1	1.14E−6	5.69E−6	4.42E−4
22	1E−2	3.00E−1	5.26E−6	1.40E−6	2.77E−4	3.02E−1	2.62E−6	3.71E−6	3.59E−4	3.00E−1	9.63E−7	4.08E−6	3.74E−4
23	1E−2	3.00E−1	4.35E−6	6.45E−7	2.29E−4	3.02E−1	2.23E−6	1.91E−6	3.05E−4	3.00E−1	8.17E−7	2.10E−6	3.17E−4

Table 4

Reference results for geometry registration only between healthy atlas and healthy patient. The table shows values for the relative mismatch for the geometry (μ_{B,L^2}) and the associated Dice coefficient DICE_B as well as the relative ℓ^2-error for the reconstruction of the velocity field $e_{\boldsymbol{v},L^2}$ with respect to the ground truth $\boldsymbol{v}^\star$.

maxit$_R$	μ_{B,L^2}	**DICE$_B$**	$e_{\boldsymbol{v},L2}$
(50, 80)	1.78E−1	9.37E−1	3.59E−1
(10, 20)	1.68E−1	9.36E−1	3.14E−1

Table 5

Results for the **analytic tumor/analytic velocity with non-zero diffusion (ATAV-DIF)** test case; ground truth: ($\rho_f = 10$, $\rho_w = 1$, $\rho_g = 0$, $k_f = 1.00\text{E}{-2}$, $k_w = 1$, $k_g = 0$, $\mathbf{p} = \mathbf{p}^\star$, $\mathbf{v} = -\mathbf{v}^\star$). We report values for the (summed) norm of the residual between the respective probability maps for the different brain tissue classes μ_{B, L^2} and tumor μ_{T, L^2} in patient space, the mean Dice coefficient for brain tissue DICE$_B$ and tumor DICE$_T$, respectively, as well as the relative norm of the gradient $\|g\|_{\text{rel}}$ for the global coupled problem (3). We report results for different values of $\rho_f \in \{5, 10, 15\}$ used in the inversion ($\rho_f = 10$ is the ground truth). We report the time spent per iteration (in seconds; top run) or in total (in seconds; bottom runs) for the entire Picard inversion, and the amount of that time spent in the tumor inversion and image registration (in percent (top run); in seconds (bottom runs)), respectively. Note that the latter sums up to less than 100% as we do not explicitly measure time spent in additional coupling functionality and forward solvers. These runs are performed using 64 MPI tasks on three nodes of *HazelHen* (see Section 5.1 for details). The top block shows the course of the inversion with respect to the Picard iteration index for the correct parameters (ground truth) for ρ_f and k_f. The four rows on the bottom show the final result for our Picard scheme for different parameter and model combinations.

iterations	β_v	μ_{B, L^2}	**DICE$_B$**	μ_{T, L^2}	**DICE$_T$**	$\|g\|_{\text{rel}}$	$T^{it}[s]$	$T^{tu}_{inv}[\%]$	$T^{reg}_{inv}[\%]$
non-zero diffusion with ground truth parameters $\rho_f = 10, k_f = 1\text{E}{-2}$									
ref	−	1.00	7.14E−1	1.00	0.00	1.00			
1	1	9.32E−1	7.32E−1	1.00	0.00	9.94E−1	2.11E+3	99.7	0.2
2	1E−1	6.52E−1	8.01E−1	2.08E−1	8.56E−1	7.45E−2	2.88E+2	96.0	2.1
3	1E−2	3.79E−1	8.88E−1	1.47E−1	9.12E−1	3.87E−2	3.94E+2	91.8	6.9
4	1E−3	1.75E−1	9.39E−1	8.08E−2	9.63E−I	2.45E−2	3.98E+2	72.7	26.1
5	1E−4	1.60E−1	9.48E−1	5.93E−2	9.78E−1	1.77E−2	2.00E+2	35.4	62.1
6	1E−4	1.60E−1	9.48E−1	5.77E−2	9.77E−1	1.72E−2	1.34E+2	46.6	49.6
7	1E−4	1.60E−1	9.47E−1	5.63E−2	9.76E−1	1.73E−2	7.10E+1	0.0	93.0
varying $\rho_f \in \{5,10,15\}, k_f = 1\text{E}{-2}$, *ground truth* $\rho_f = 10, k_f = 1\text{E}{-2}$							T ^{otal} [s]	$T^{tu}_{inv}[s]$	$T^{reg}_{inv}[s]$
$\rho_f = 5$	1E−4	1.61E−1	9.46E−1	6.48E−2	9.67E−1	2.21E−2	4.37E+3	3.95E+3	3.97E+2
$\rho_f = 10$	1E−4	1.60E−1	9.47E−1	5.63E−2	9.76E−1	1.73E−2	3.60E+3	3.17E+3	4.01E+2
$\rho_f = 15$	1E−4	1.61E−1	9.48E−1	6.39E−2	9.72E−1	1.41E−2	3.96E+3	3.53E+3	4.00E+2
$\rho_f = 10$, $k_f = 0$, *ground truth* $p_f = 10, k_f = 1\text{E}{-2}$							T^{total} [s]	$T^{tu}_{inv}[s]$	$T^{reg}_{inv}[s]$
$\rho_f = 10$	1E−4	1.60E−1	9.48E−1	5.95E−2	9.71E−1	1.74E−2	4.25E+2	2.25E+1	4.01E+2

Table 6

Results for the **analytic tumor with analytic velocity with diffusion and low-dimensional initial condition (ATAV-LD)** test case; ground truth: ($\rho_f = 15$, $\rho_w = 1$, $\rho_g = 0$, $k_f = 1.00\mathrm{E}{-1}$, $k_w = 1$, $k_g = 0$, $\boldsymbol{p} = \boldsymbol{p}^\star$ (in patient domain), $\boldsymbol{v}$ N/A). We use a low-dimensional parametrization for the initial condition with $n_{\boldsymbol{p}} = 8$ Gaussians ($\sigma = \pi/15$). We report the (summed) mismatch for the probability maps for the brain tissue μ_{B,L^2} and tumor μ_{T,L^2} in patient space, the mean Dice coefficient for hard segmentations of brain tissue (DICE$_B$) and tumor (DICE$_T$), respectively, and the relative norm of the gradient $\left(\|g\|_{\mathrm{rel}}\right)$ for the coupled problem (3). We assess the final state of the reconstruction using different values for $\rho_f \in \{5, 10, 15, 20\}$ and $k_f \in \{0, 1\mathrm{E}{-1}\}$. Absolute timings are given for the tumor inversion and image registration, respectively using 64 MPI tasks on three nodes of HazelHen. Note that the latter sums up to less than the reported total run time as we do not explicitly measure time spent in additional coupling functionality and forward solvers. We always reach the target value of $\beta_v = 1\mathrm{E}{-4}$.

$\boldsymbol{\rho}$	$\boldsymbol{k}$	**It**	μ_{B,L^2}	**DICE$_B$**	μ_{T,L^2}	**DICE$_T$**	$\|g\|_{\mathrm{rel}}$	T^{total}[s]	T^{tu}_{inv}[s]	T^{reg}_{inv}[s]
		initial	1.00	7.18E−1	1.00	0.00	1.00	–	–	–
0	0	final	1.96E−1	9.45E−1	2.80E−1	8.48E−1	4.16E−2	3.78E+2	3.33	3.77E+2
15	0		1.87E−1	9.46E−1	2.04E−1	9.16E−1	1.70E−2	3.78E+2	2.19	3.78E+2
15	IE−1		**1.69**E−**1**	**9.48**E−**1**	**4.68**E−**2**	**9.65**E−**1**	1.73E−2	4.72E+3	4.17E+3	4.59E+2
15	5E−1		1.90E−1	9.41E−1	3.37E−1	4.71E−2	1.78E−2	6.67E+3	5.54E+3	4.29E+2
15	IE−2		1.79E−1	9.47E−1	1.57E−1	9.31E−1	1.51E−2	2.23E+3	1.50E+3	4.31E+2
5	IE−1		1.71E−1	9.46E−1	6.50E−2	9.57E−1	2.45E−2	1.96E+3	1.50E+3	3.87E+2
10	IE−1		1.70E−1	9.47E−1	4.98E−2	9.62E−1	2.01E−2	3.90E+3	2.90E+3	3.73E+2
10	IE−2		1.79E−1	9.47E−1	1.65E−1	9.26E−1	1.91E−2	2.17E+3	1.63E+3	3.66E+2
20	IE−1		1.70E−1	9.48E−1	5.45E−2	9.56E−1	2.02E−2	3.07E+3	2.62E+3	3.88E+2
20	IE−2		1.80E−1	9.47E−1	1.60E−1	9.18E−1	1.20E−2	2.03E+3	4.51E+2	3.78E+2

Table 7

Summary of results for the **real tumor with real velocity (RTRV)** test case, ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A); based on real clinical data (taken from [7]). We set the tumor parameters to $\rho_f = 15$, $k_f = 0$ (reaction-only). We use a parametrization of tumor initial conditions with $n_{\boldsymbol{p}} = 343$ Gaussians. We report the (summed) mismatch for the brain tissue probability maps (μ_{B,L^2}) and tumor probability map (μ_{T,L^2}) in patient space, the mean Dice coefficient for the hard segmentation corresponding to the brain tissue (DICE$_B$) and the tumor (DICE$_T$), respectively, as well as the relative norm of the gradient for the coupled problem (3) $\left(\|g\|_{\text{rel}}\right)$. We also report the total run time in seconds (T_{total}), and the run time of the individual components of our Picard scheme, respectively (also in seconds; tumor solver:T_{inv}^{tu}; image registration: T_{inv}^{reg}). Note that the latter sums up to less than the reported total run time as we do not explicitly measure time spent in additional coupling functionality and forward solvers. We execute our code in parallel on 11 nodes using 256 MPI tasks of HazelHen.

Patient		β_v	μ_{B,L^2}	**DICE$_B$**	μ_{T,L^2}	**Dice$_T$**	$\|g\|_{\text{rel}}$	T^{total}[s]	T_{inv}^{tu}[s]	T_{inv}^{reg}[s]
AAMH	Initial	1	1.00	4.36E−1	1.00	0.00	1.00	–	–	–
	final	1E−4	3.45E−1	7.92E−1	1.95E−1	9.57E−1	3.92E−2	6.28E+2	1.95E+2	4.35E+2
AAAN	initial	1	1.00	4.37E−1	1.00	0.00	1.00	–	–	–
	final	1E−4	3.54E−1	8.35E−1	3.77E−1	8.91E−1	1.06E−1	6.34E+2	2.15E+2	4.17E+2
AAAC	initial	1	1.00	4.82E−1	1.00	0.00	1.00	–	–	–
	final	1E−4	3.36E−1	8.44E−1	2.45E−1	9.55E−1	4.38E−2	4.92E+2	1.59E+2	3.32E+2
AAMP	initial	1	1.00	4.74E−1	1.00	0.00	1.00	–		–
	final	1E−4	3.32E−1	8.05E−1	1.34E−1	9.75E−1	3.31E−2	6.48E+2	2.36E+2	4.13E+2
AAQD	initial	1.00	1.00	2.92E−1	1.00	0.00	1.00	–	–	–
	final	1E−4	3.24E−1	8.06E−1	2.55E−1	9.31E−1	9.88E−2	1.10E+3	2.19E+2	8.81E+2
AAWI	initial	1	1.00	4.53E−1	1.00	0.00	1.00	–	–	–
	final	1E−4	3.45E−1	7.87E−1	3.92E−1	8.74E−1	5.26E−2	6.84E+2	2.92E+2	3.90E+2

Table 8

Results for the **AAMH** and the **AAAN** patient **real tumor/ real velocity (RTRV)** test case, ground truth (ρ N/A, $\boldsymbol{k}$ N/A, $\boldsymbol{p}$ N/A, $\boldsymbol{v}$ N/A) using real clinical input data. The table shows the (summed) mismatch for the brain tissue probability maps (μ_{B,L^2}) and tumor probability map (μ_{T,L^2}) in patient space, the mean Dice coefficient for the hard segmentation corresponding to the brain tissue (DICE$_B$) and the tumor (DICE$_T$), respectively for different values for reaction scaling parameter ρ_f and of the diffusion coefficient scaling parameter k_f., as well as the relative norm of the gradient for the coupled problem (3) $\left(\|g\|_{\text{rel}}\right)$. We also report the total run time in seconds (T_{total}), and the run time of the individual components of our Picard scheme, respectively (also in seconds; tumor solver:T_{inv}^{tu}; image registration:T_{inv}^{reg}). Note that the latter sums up to less than the reported total run time as we do not explicitly measure time spent in additional coupling functionality and forward solvers. We execute our code in parallel on 11 nodes using 256 MPI tasks of HazelHen.

	ρ_f	k_f	n_p	μ_{B,L^2}	**DICE$_B$**	μ_{T,L^2}	**DICE$_T$**	$\|g\|_{\text{rel}}$	T^{total}[s]	T_{inv}^{tu}[s]	T_{inv}^{reg}[s]
ID AAMH	*initial canfig*			1.00	4.36E−1	1.00	0.00	1.00	–	–	–
	0	0	343	3.40E−1	7.92E−1	l.47E−1	9.69E−1	l.19E−1	6.09E+2	2.13E+2	3.63E+2
	0	0	125	3.38E−1	7.95E−1	1.73E−1	9.62E−1	1.31E−1	5.09E+2	7.86E+l	4.04E+2
	15	0	343	3.45E−1	7.92E−1	1.95E−1	9.57E−1	3.92E−2	6.28E+2	l.95E+2	4.35E+2
	15	0	125	3.50E−1	7.91E−1	2.36E−1	9.53E−1	4.83E−2	4.47E+2	7.20E+l	3.48E+2
	5	lE−2	343	3.44E−1	7.92E−1	l.88E−1	9.64E−1	7.68E−2	l.23E+4	l.18E+4	3.41E+2
	10	lE−2	343	3.46E−1	7.92E−1	2.0 7E−1	9.62E−1	5.16E−2	l.33E+4	l.28E+4	3.44E+2
	15	lE−2	343	3.48E−1	7.93E−1	2.35E−1	9.59E−1	3.47E−2	1.18E+4	l.13E+4	3.63E+2
ID AAMN	*initial canfig.*			1.00	4.37E− 1	1.00	0.00	1.00	-	-	-
	0	0	343	3.42E−1	8.39E−1	2.56E−1	9.36E−1	2.19E−1	6.58E+2	2.02E+2	4.23E+2
	0	0	125	3.56E−1	8.33E−1	4.40E−1	6.68E−1	2.36E−1	4.54E+2	8.89E+l	3.38E+2
	15	0	343	3.54E−1	8.35E−1	3.77E−1	8.91E−1	l.06E−1	6.34E+2	2.15E+2	4.17E+2
	15	0	125	3.67E−1	8.30E−1	5.71E−1	3.44E− 1	1.28E−1	4.67E+2	8.38E+l	3.57E+2
	5	lE−2	343	3.50E−1	8.35E−1	3.17E−1	9.06E−1	1.79E−1	1.82E+4	1.77E+4	3.33E+2
	10	lE−2	343	3.52E−1	8.35E−1	3.60E−1	8.85E−1	l.41E−1	l.88E+4	l.82E+4	3.41E+2
	15	lE−2	343	3.S5E−1	8.34E−1	3.98E−1	8.62E−1	l.13E−1	1.85E+4	l.80E+4	3.28E+2