arXiv:2111.04090v2 [physics.med-ph] 18 Jul 2022

Learn-Morph-Infer: a new way of solving the inverse problem for brain tumor modeling

Ivan Ezhov[a,b,1], Kevin Scibilia[a], Katharina Franitza[c], Felix Steinbauer[a], Suprosanna Shit[a,b], Lucas Zimmer[d,b], Jana Lipkova[e,f,g], Florian Kofler[a,b,i], Johannes Paetzold[a,b], Luca Canalini[h], Diana Waldmannstetter[a], Martin J. Menten[a,b], Marie Metz[i,b], Benedikt Wiestler[b,i,2], Bjoern Menze[d,2]

[a] *Department of Informatics, TUM, Munich, Germany*
[b] *TranslaTUM - Central Institute for Translational Cancer Research, TUM, Munich, Germany*
[c] *Department of Mechanical Engineering, TUM, Munich, Germany*
[d] *Department of Quantitative Biomedicine, UZH, Zurich, Switzerland*
[e] *Department of Pathology, Brigham and Women's Hospital, Harvard Medical School, Boston, MA*
[f] *Broad Institute of Harvard and MIT, Cambridge, MA*
[g] *Data Science Program, Dana-Farber Cancer Institute, Boston, MA*
[h] *Fraunhofer MEVIS, Bremen, Germany*
[i] *Neuroradiology Department of Klinikum Rechts der Isar, TUM, Munich, Germany*

Abstract

Current treatment planning of patients diagnosed with a brain tumor, such as glioma, could significantly benefit by accessing the spatial distribution of tumor cell concentration. Existing diagnostic modalities, e.g. magnetic resonance imaging (MRI), contrast sufficiently well areas of high cell density. In gliomas, however, they do not portray areas of low cell concentration, which can often serve as a source for the secondary appearance of the tumor after treatment. To estimate tumor cell densities beyond the visible boundaries of the lesion, numerical simulations of tumor growth could complement imaging information by providing estimates of full spatial distributions of tumor cells. Over recent years a corpus of literature on medical image-based tumor modeling was published. It includes different mathematical formalisms describing the forward tumor growth model. Alongside, various parametric inference schemes were de-

Email address: ivan.ezhov@tum.de (Ivan Ezhov)
[1] Corresponding author.
[2] Contributed equally as senior authors.

veloped to perform an efficient tumor model personalization, i.e. solving the inverse problem. However, the unifying drawback of all existing approaches is the time complexity of the model personalization which prohibits a potential integration of the modeling into clinical settings. In this work, we introduce a deep learning based methodology for inferring the patient-specific spatial distribution of brain tumors from T1Gd and FLAIR MRI medical scans. Coined as *Learn-Morph-Infer* the method achieves real-time performance in the order of minutes on widely available hardware and the compute time is stable across tumor models of different complexity, such as reaction-diffusion and reaction-advection-diffusion models. We believe the proposed inverse solution approach not only bridges the way for clinical translation of brain tumor personalization but can also be adopted to other scientific and engineering domains.

Keywords: Inverse modeling, physics-based deep learning, glioma, model calibration, tumor modeling, MRI

---

1. Introduction

Glioblastoma (GBM) is the most aggressive brain tumor, characterized by varying and unknown infiltration into the surrounding tissue. After resection of the tumor mass visible in MRI scans, current treatment includes radiotherapy targeting tissue around the visible lesion to account for residual tumor cells. Tumor recurrence is however present in most cases, possibly due to patient-specific and non-uniform distribution of residual tumor cells. Personalization of the clinical (irradiation) target volume, could spare more healthy tissue and increase progression-free survival by potentially avoiding recurrence (Stupp et al., 2014, Harpold et al., 2007, Jackson et al., 2015, Lipkova et al., 2019).

Current computational approaches for personalizing radiotherapy planning often rely on solving an inverse problem for GBM growth models (Hogea et al., 2008, Konukoglu et al., 2010b, Geremia et al., 2012, Menze et al., 2011, Le et al., 2017, Lipkova et al., 2019, Scheufele et al., 2020, Subramanian et al., 2020a, Hormuth et al., 2021, Lorenzo and et al, 2021). In this context, the

growth (forward) models are based on partial differential equations (PDEs) that describe the evolution of tumor cell density in the brain anatomy. The inverse model aims to identify free parameters of the forward model that best match the observation, e.g. tumor outlines from medical imaging modalities. To identify such parameters, the inverse problem can be cast as constrained optimization (Hogea et al., 2008, Mang et al., 2012, Scheufele et al., 2019) or Bayesian inference formulations (Menze et al., 2011, Lipkova et al., 2019, Ezhov et al., 2020, 2019).

Predominantly, existing forward GBM models view tumor progression at the macroscopic level by describing gross biomechanical phenomena. These include diffusive motion and proliferation of tumor cells (under simplistic reaction-diffusion PDEs) (Menze et al., 2011), interaction between the tumor and surrounding tissue (i.e. mass-effect) (Subramanian et al., 2020a), necrotic core formation (Patel and Hathout, 2017), etc. Despite methodological advances in computing the inverse model, the total time for model personalization is still large amounting to multiple hours using such simplistic forward models (Subramanian et al., 2020a, Hormuth et al., 2018, Scheufele et al., 2019). For example, in (Scheufele et al., 2019) the authors exploit a highly efficient quasi-Newton optimization scheme to infer parameters of the reaction-diffusion model. The inference converges after $\sim$5 hours of compute on 11 dual-x86 CPU nodes for 256^3 resolution. In (Subramanian et al., 2020a), the mass-effect model is solved using an analogous optimization scheme but implemented on GPU leading to the same order of compute time for the 256^3 grid (and up to 1 hour for 128^3 resolution). Bayesian methods providing uncertainty estimate of the parametric inference (Lipkova et al., 2019) can take an even longer time (up to days) of computing on specialized CPU clusters.

Recently, machine learning solutions entered the field of PDEs. Learnable solutions for both forward (Raissi et al., 2019, Sitzmann et al., 2020, Stevens and Colonius, 2020, Thuerey et al., 2020, Kasim et al., 2020, Kim et al., 2018) and inverse (Papamakarios and Murray, 2016, Lueckmann et al., 2017, Dax et al., 2021) models were developed. In (Papamakarios and Murray, 2016, Lueckmann

et al., 2017), the authors proposed a Bayesian framework that allows bypassing a compute of intractable likelihoods necessary for the Bayesian inference. This notably speeds up the inference, however, it still requires thousands of simulations to be generated by the forward model, which in turn can still result in many hours of compute for the inverse problem. In concurrent to our work (Dax et al., 2021), a solution to the inverse problem providing point estimates was proposed. This is achieved via learning a mapping from the physical model simulations directly to the model parameters by leveraging access to large simulated data. Unfortunately, this method is not capable to deal with the varying geometry of the simulation domain. Moreover, as our paper shows, direct mapping to physical model parameters can result in a deterioration of prediction accuracy since the inverse mapping between simulations and parameters is not bijective.

Few machine learning approaches (Ezhov et al., 2020, Pati et al., 2020) also appeared in the brain tumor modeling context. However, (Ezhov et al., 2020) requires a vast amount of forward model evaluations for convergence of parametric estimation under Bayesian settings for each new patient. In turn, (Pati et al., 2020) requires access to a dataset of inferred model parameters that can become prohibitively expensive to collect with the growing complexity of the tumor model.

Potential integration of brain tumor modeling into clinical practice would require access to a large cohort of longitudinal clinical data, allowing to estimate the clinical value of a patient outcome's forecast by the tumor models (Yankeelov et al., 2013). Integration of current macroscopic models would also require a thorough analysis of forecast consistency between the macroscopic description and higher complexity microscopic models encompassing subcellular biophysics. For this, it is paramount to bring the computational cost of the inverse modeling to a reasonable time. Here, we propose a neural network based methodology for predicting a patient-specific spatial distribution of GBM (from single time-point medical scans, namely T1Gd and FLAIR MRI) that requires neither sampling nor optimization. As such, it is potentially well suited for rapid

model personalization in a clinical workflow, but also for scaling up preceding feasibility studies to large patient cohorts.

Our contribution is a) a pipeline allowing to perform tumor model personalization in a fixed for all patients space, and b) a special learning setup for the network (the main part of the pipeline) performing inverse model inference. The method achieves lightning-fast performance in the order of minutes on widely available hardware and the compute time is stable across tumor models of different complexity. This in turn opens the possibility of rapid testing various biophysical models on a large dataset of patients and hence bridging the way for clinical translation.

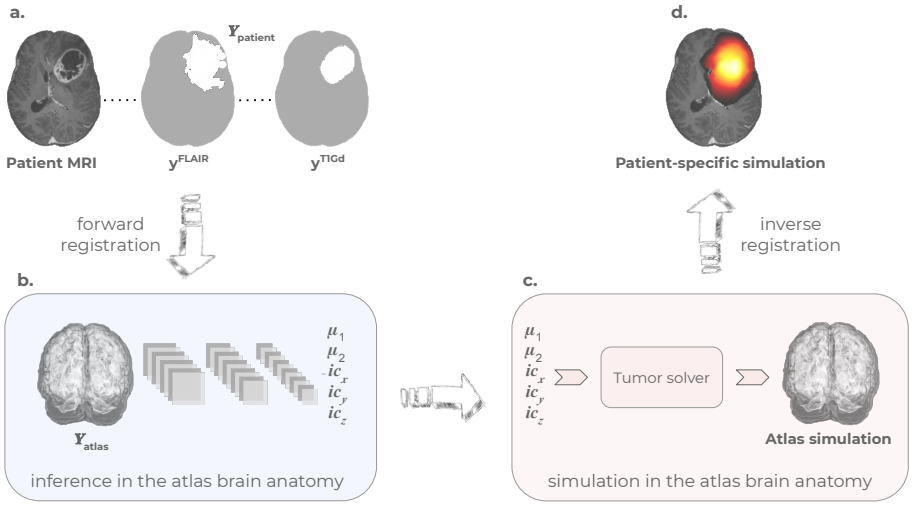

Figure 1: A sketch of the inference procedure of the *Learn-Morph-Infer* pipeline. First, the patient's brain tumor segmentations $\{\mathbf{y}^{T1Gd}, \mathbf{y}^{FLAIR}\}$ are morphed to the brain atlas space (a). A network trained on synthetic data inputs the morphed segmentations and outputs corresponding tumor model parameters $\{\mu_1, \mu_2, ic_x, ic_y, ic_z\}$ (b). The inferred parameters are used to simulate a tumor in the atlas space (c). Finally, the simulated tumor is morphed back to the patient space (d).

2. Method

2.0.1. Learn-Morph-Infer Pipeline

In contrast to a few works like (Hormuth et al., 2019, Ezhov et al., 2019, Tunç et al., 2021), we solve the inverse problem by relying only on single time-point MRI observations of the brain tumor $Y = \{\mathbf{y}^{T1Gd}, \mathbf{y}^{FLAIR}\}$. This is the most realistic clinical scenario as normally treatment follows immediately after the first MRI scanning. Given a tumor observation Y, our goal is to calibrate a set of personalized parameters θ_c of the forward tumor growth model that infers the underlying patient-specific tumor cell density $c_{patient}$ in the patient anatomy[3]. For a given patient, this is achieved via the proposed *Learn-Morph-Infer* pipeline:

- We register a patient MRI image to the brain atlas (Rohlfing et al., 2010) and obtain a transformation matrix.
- The transformation matrix is used to morph scans based on the patient anatomy $Y_{patient}$ to scans in the atlas anatomy Y_{atlas} as illustrated in Fig. 1a-b.
- A neural network, that has learned to solve the inverse problem $Y_{atlas} \rightarrow \theta_c$ through prior training on simulated data Y_{sim} (Fig. 1), predicts θ_c during inference, Fig. 1b.

[3]Let us explain how we see the potential clinical utility of glioma modeling. In clinical practice, for glioma patients normally only a single MRI scan is taken before treatment. It is not feasible to capture the dynamics of the pathology by calibrating a tumor model from a single scan - multiple sets of model parameters can result in the same simulated tumor profile. Thus, one cannot reliably predict tumor evolution over time as different sets of parameters will result in tumor trajectories diverging over time. Instead, we are only interested in identifying a simulated tumor whose morphology (and, hence, tumor cell distribution) best matches the morphology of the tumor we see on a single MRI scan. The best matching simulated profile will then serve to predict tumor infiltration beyond contrast-enhancing tumor areas and thereby define the target volume for radiotherapy (as opposed to the currently used simple EORTC criteria - 2cm border surrounding the cavity (Young et al., 1999)).

- These parameters θ_c are used as input in the tumor growth model (forward solver) to infer a tumor cell density c_{atlas} in the atlas space, Fig. 1c.
- The tumor volume c_{atlas} is transformed back to the patient space with the inverse transformation matrix, yielding $c_{patient}$ as displayed in Fig. 1d.

2.0.2. Forward tumor model

We independently probe two types of non-linear PDEs describing tumor growth:

Reaction-diffusion equation. First, we consider the Fisher-Kolmogorov PDE describing the evolution of the tumor cell concentration c by considering cell diffusion and proliferation,

$$\frac{\partial c}{\partial t} = \nabla \cdot (\mathbf{D} \nabla c) + \rho c(1 - c) \tag{1}$$

$$\nabla c \cdot \mathbf{n} = 0 \quad \textit{boundary condition} \tag{2}$$

Here, ρ denotes the tumor proliferation rate while the infiltrative behaviour of the tumor is modelled by the diffusion tensor $\mathbf{D} = D \cdot \mathbb{I}$. The equation is solved in a three dimensional atlas brain anatomy segmented into white matter (WM), grey matter (GM) and cerebrospinal fluid (CSF). The diffusion coefficient D is computed for each voxel $\mathbf{i}$ with location (i_x, i_y, i_z) as $D_i = p_{w_i} D_w + p_{g_i} D_g$, where p_w, p_g describe percentages and D_w, D_g diffusion coefficients of WM and GM respectively, and a relation $D_w = 10 \cdot D_g$ is assumed (Lipkova et al., 2019). No cell diffusion into CSF is feasible according to the model. The solver based on this growth model takes $\theta_c = \{D_w, \rho, T, ic_x, ic_y, ic_z\}$ as input and returns a tumor cell density c. The parameters $\mathbf{x} = (ic_x, ic_y, ic_z)$ define the initial condition where the tumor is initialized at time t=0 as a point seed. The tumor is simulated until the time of detection T.

Reaction-diffusion-advection equation. The second type is a non-linear reaction-diffusion-advection PDE analogous to (Subramanian et al., 2020a). In the following, the brain tissue will be represented by $\mathbf{m} = (m_{WM}(\mathbf{i}, t),\ m_{GM}(\mathbf{i}, t),\ m_{CSF}(\mathbf{i}, t))$ for each voxel $\mathbf{i}$ and time t. The normalized tumor cell density $c = c(\mathbf{i}, t)$ can be modelled by the following equations:

$$\frac{\partial c}{\partial t} = \nabla \cdot (\mathbf{D}\nabla c) - \nabla(c\mathbf{v}) + \rho c(1 - c) \tag{3}$$

$$\frac{\partial \mathbf{m}}{\partial t} + \nabla \cdot (\mathbf{m} \otimes \mathbf{v}) = 0 \tag{4}$$

$$\nabla \cdot (\lambda \nabla \mathbf{u} + \mu(\nabla \mathbf{u} + \nabla \mathbf{u}^{\top})) = \gamma \nabla c \tag{5}$$

$$\frac{\partial \mathbf{u}}{\partial t} = \mathbf{v} \tag{6}$$

$$\nabla c \cdot \mathbf{n} = 0 \quad \textit{boundary condition} \tag{7}$$

$$\mathbf{m} = 0 \quad \textit{boundary condition} \tag{8}$$

$$\mathbf{u} = 0 \quad \textit{boundary condition} \tag{9}$$

$$\mathbf{v} = 0 \quad \textit{boundary condition} \tag{10}$$

Coupling the Eqn. 3 to a linear elasticity model, Eqn. 5, allows considering deformation in the anatomy due to a mass effect induced by tumor growth (Subramanian et al., 2019). The linear elasticity model is defined by the Lamè coefficients λ and μ as specified in Eqn. 5. The displacement $\mathbf{u}$ is represented in the advection term of Eqn. 3. The degree of the mass effect depends on the selection of the mass effect parameter γ.

Parameter symbol	Parameter meaning
c	Tumor cell density
$\mathbf{D}$	Diffusion tensor
D	Diffusion coefficient
$\mathbb{I}$	Identity tensor
ρ	Proliferation rate
$\mathbf{n}$	Unit vector normal to boundary
$\mathbf{u}$	Advection displacement
$\mathbf{v}$	Advection velocity
$\mathbf{m}$	Brain tissue maps
λ, μ	Lamè coefficients
γ	Mass effect parameter

Table 1: Description of parameters driving the reaction-diffusion and reaction-advection-diffusion PDEs.

2.0.3. Linking cell density with MRI signal

MRI modalities capture structural information about the brain tumor. T1Gd contrasts the tumor core, whereas FLAIR informs about the area of the edema in addition to the tumor core. It is established practice (Lê et al., 2016, Lipkova et al., 2019, Subramanian et al., 2020a, Tunc et al., 2021, Konukoglu et al., 2010a, Menze et al., 2011) to consider binary segmentations corresponding to the MRI scans to inform biophysical models. The binary masks contain zeros in the area of healthy tissues and are non-zero in the tumor-related area. In order to relate the segmentations $Y = \{\mathbf{y}^{T1Gd}, \mathbf{y}^{FLAIR}\}$ to a simulated tumor cell density c, we threshold the density at randomly sampled levels c_t^{T1Gd} and c_t^{FLAIR} ($c_t^{T1Gd} > c_t^{FLAIR}$) to obtain $Y_{sim} = \{\mathbf{y}_{sim}^{T1Gd}, \mathbf{y}_{sim}^{FLAIR}\}$ reproducing the real segmentations [4].

[4] We want to clarify here the thresholding step in more details. The output of our tumor solver is a continuous tumor cell density distribution. The experimental observation to which we want to fit our tumor model simulations comes in a form of binary tumor segmentation. In order to identify a simulation best fitting the binary segmentation, we threshold our simulated continuous tumor cell distribution to obtain a binary volume, which in turn can already be nicely compared with the real patient segmentation. Oftentimes, in literature people resort to fixed values of cell density for the threshold levels. However, there is no understanding

2.0.4. Learning the inverse model in atlas space

As discussed in the previous section, the key step of the *Learn-Morph-Infer* pipeline is to learn the inverse tumor model using a neural network that can infer a set of personalized parameters θ_c from corresponding tumor observations Y_{atlas}. In order to create a dataset for the network training, we generate 100,000 tumors in the atlas space by randomly sampling tumor model parameters $\{D_w, \rho, T, ic_x, ic_y, ic_z, c_t^{t1Gd}, c_t^{FLAIR}\}$ within physiologically plausible ranges (Swanson et al., 2000). To form the neural network input, the simulated MRI segmentations are combined into one volume $\mathbf{y}_{sim}^{MRI} = 0.5 \cdot \mathbf{y}_{sim}^{T1Gd} + 0.5 \cdot \mathbf{y}_{sim}^{FLAIR}$.

Reformulation of the inverse problem. Now, the question is what should be used as a network prediction? It is tempting to try to predict the tumor model parameters directly. However, it is well known that the inverse problem is highly ill-posed, i.e. numerous sets of dynamic parameters $\{D_w, \rho, T\}$ correspond to the same simulated cell density profile. Thus, we have two sources of prediction error: a) coming from the fact that we learn a mapping from one to many. Imagine we have two sets of dynamic parameters $\{D_w, \rho, T\}$ and $\{D_w^*, \rho^*, T^*\}$ that result in the same tumor profile $\mathbf{y}_{sim}^{MRI}$. If we train a network in a supervised fashion, every time the network predicts $\{D_w^*, \rho^*, T^*\}$ for a $\{D_w, \rho, T\}$-$\mathbf{y}_{sim}^{MRI}$ pair (and vice-versa), it will be falsely penalized; b) the actual error that comes from limited network capacity to accurately learn the mapping.

(neither agreement) in the community of what levels would best correspond to a real scenario. Different works used different numbers in their studies (Le et al., 2017). To consider the most general case, we randomly sampled the values within the ranges reported in the literature.

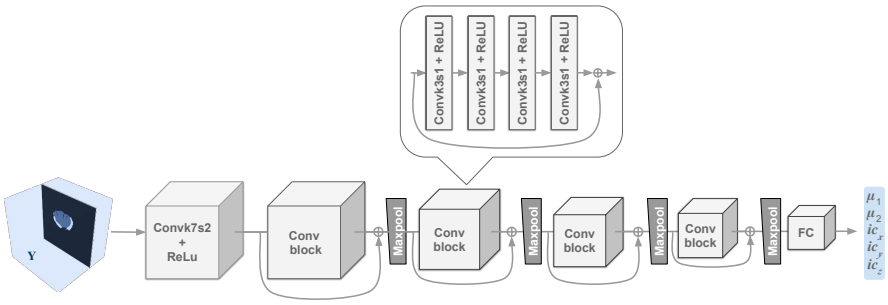

Figure 2: A design of the inverse model network - one of the key elements of the *Learn-Morph-Infer* pipeline. The network design represents a ResNet type architecture. It takes as an input the binary brain tumor segmentations $\{\mathbf{y}^{T1Gd}, \mathbf{y}^{FLAIR}\}$ and outputs $\{\mu_1, \mu_2, ic_x, ic_y, ic_z\}$. Crucially, the network predicts $\{\mu_1, \mu_2\}$ which are not the exact forward model parameters but time-independent combinations thereof.

Clearly, the error of type (a) should negatively affect the learning performance, as the network may be penalized for making a sensible prediction. To circumvent this, we do not predict the parameters directly. As evident from Eqn. 1, normalization of the time parameter T in Eqn. 1 is equivalent to re-scaling of the proliferation ρ and diffusion D coefficients (Subramanian et al., 2020b). This means that for sets $\{D_w, \rho, T\}$ and $\{D_w^*, \rho^*, T^*\}$ corresponding to the same simulated tumor, combinations of time-independent parameters $\mu_1 = \sqrt{D_w T}$, $\mu_2 = \sqrt{T\rho}$ stay constant ($\sqrt{D_w T}$=$\sqrt{D_w^* T^*}$, $\sqrt{T\rho}$=$\sqrt{T^*\rho^*}$) (Konukoglu et al., 2010b). Hence, predicting these combinations of time-independent parameters μ_1 and μ_2 relaxes the error type (a). In order to calculate back the $\{D_w, \rho, T\}$, we introduce a third combination as $v = 2\sqrt{D_w \rho}$. As it is not possible to infer the velocity v from a single time-point observation, we set the velocity equal to mean velocity of the used sampling range, 200 $mm/year$ (note also that for our purpose the choice of v is irrelevant as any tumor simulation can be obtained with arbitrary v (Menze et al., 2011)). Given $\{\mu_1, \mu_2, v\}$, we can calculate $\{D_w = \frac{\mu_1 v}{2\mu_2}, \rho = \frac{\mu_2 v}{2\mu_1}, T = \frac{2\mu_2 \mu_1}{v}\}$.

Following this reasoning, we make the network to predict five parameters: $\{\mu_1, \mu_2\}$ and $\{ic_x, ic_y, ic_z\}$. Note that we do not predict $\{c_t^{t1Gd}, c_t^{FLAIR}\}$, as

we do not need the threshold parameters at further steps of the *Learn-Morph-Infer* pipeline.

The inverse network design. The network we chose to learn the inverse model represents a convolutional architecture, depicted in Fig. 2. Every convolution in the network is followed by a rectified linear unit (ReLU) nonlinearity. The input is passed through an initial convolution of kernel size 7, stride 2 and 64 filters (k7s2f64), downsampling the volume to 64^3 and increasing the number of channels to 64. This volume of size 64^3 x 64 is input through four convolutional blocks, where every convolutional block contains four convolutions of kernel size 3, stride 1, and 64 filters (i.e. number of filters is kept constant throughout the network). A convolutional block uses a skip connection to learn a residual mapping (He et al., 2016), with the input being added element-wise to the output of the four convolutions. The first three convolutional blocks are followed by a MaxPool3D layer (with parameters k2s2) to downsample the 3D volumes by two. The last convolutional block is followed by a global average pooling layer, shrinking the 64 3D volumes to 64 neurons that can be linked through a fully connected (FC) layer to the output. These outputs are linearly interpolated into the [-1, 1] range for training.

3. Results

3.1. Data and implementation details

Synthetic data. The simulated tumors used for training the network have resolution of 128^3. The simulations were generated by randomly sampling patient-specific parameters from the following ranges: $D_w \in [0.0002\,\frac{\mathrm{cm}^2}{\mathrm{d}}, 0.015\,\frac{\mathrm{cm}^2}{\mathrm{d}}]$, $\rho \in [0.002\,\frac{1}{\mathrm{d}}, 0.2\,\frac{1}{\mathrm{d}}]$, $T \in [50\mathrm{d}, 150\mathrm{d}]$, $x \in [0.15, 0.7]$, $y \in [0.2, 0.8]$, $z \in [0.15, 0.7]$, $c^{T1Gd} \in [0.5, 0.85]$, $c^{FLAIR} \in [0.05, 0.5]$. The elasticity model parameters λ, μ, γ were taken from (Subramanian et al., 2019). Tumors that are unrealistically small or large have been discarded (based on minimum and maximum tumor sizes of real tumors from BraTS dataset (Menze et al., 2014)). The simulations

dataset is then divided into a training set (80000 tumors), validation set (8000 tumors) and test set (12000 tumors).

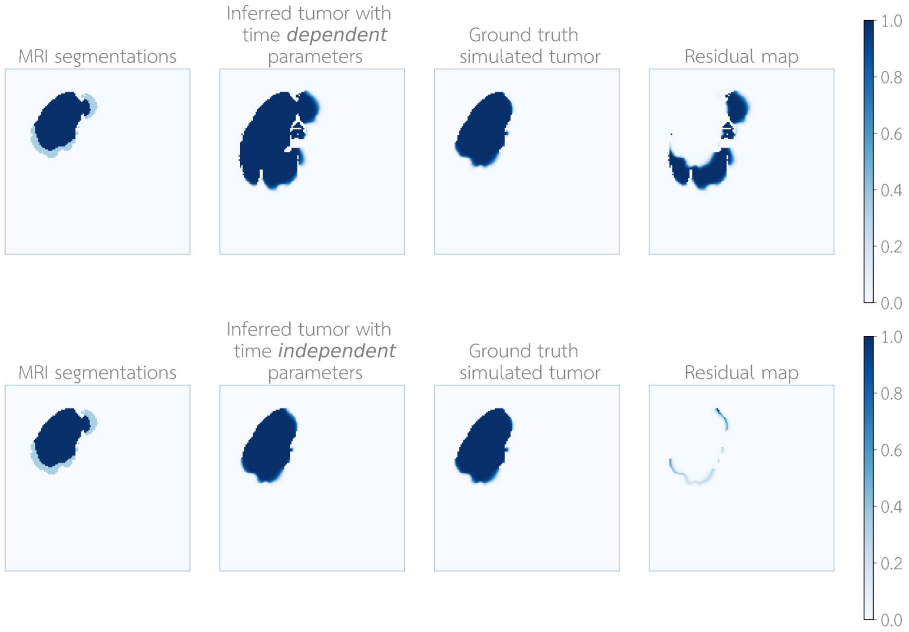

Figure 3: Qualitative comparison between the tumors inferred using the *Learn-Morph-Infer* pipeline with two different network training strategies: with time-dependent $(D_w, \rho, T, \mathbf{x})$ parameters as network output, and time-independent $(\mu_1, \mu_2, \mathbf{x})$ parameters as output. The examples correspond to the Fisher-Kolmogorov tumor model with MAE equal to 0.495 (within WM) and 0.496 (within GM) for time-dependent inference, and 0.048 (within WM) and 0.048 (within GM) for time-independent inference. The first columns "MRI segmentations" correspond to the FLAIR+T1Gd segmentation, the second one "Inferred tumor" corresponds to the simulation inferred by the proposed *Learn-Morph-Infer* pipeline, the third column "Ground truth simulated tumor" corresponds to the simulated tumor which we used for the sensitivity analysis (see the "Synthetic data" section for details how these test data were formed), and the fourth "Residual map" column is the difference between the 2nd and 3rd columns' images.

Analogously to how we created the training data, we form synthetic test data by thresholding a simulated tumor c_{gt} at two levels. Then we pass the obtained

thresholded volumes through the pipeline to infer the tumor cell distribution c_{pred}. Finally, we quantitatively compare the difference between c_{gt} and c_{pred}.

Real data. The data set of real MRI images consists of an in-house cohort of 80 patients with a newly diagnosed glioblastoma, IDH wild type as per the 2016 WHO classification of brain tumors. All patients gave written informed consent to be part of an observational cohort. Preoperative MR images were converted to NIfTI format. For image preprocessing (co-registration, skull stripping) and automated tumor segmentation, we used BraTS Toolkit (Kofler et al., 2020), a tool we developed locally and which is freely available. The average age of the patient cohort is 60 years with a minimum of 26 and a maximum of 79 years. The cohort is equally represented by gender.

Implementation details. For the registration between the patient and atlas brain MRI scans, we use the Advanced Normalization Tools (ANTs) (Avants et al., 2009). We choose a deformable SyN registration that ensures providing both forward and inverse transformation with step-size 0.25, weight 1, and region radius for cross-correlation computation r=4. Cross-correlation is used as a similarity metric. The optimization is performed over two resolutions with a maximum of 50 iterations at the coarsest level, and 20 at the final level. The tumor area on the patient scan was masked for the registration. We use a Gaussian regularizer with a sigma of 3 operating on the similarity gradient. These settings provided high morphing quality at relatively fast compute ($\sim$2 minutes).

The network is initialized with He initialization as in the original ResNet architecture (He et al., 2016), and is trained with the AdamW (Loshchilov and Hutter, 2019) optimizer, which is a variant of the Adam optimizer (Kingma and Ba, 2015) using decoupled weight decay. We use an initial learning rate of 6 x 10^{-5} which is decayed exponentially after every batch by a factor of 0.999997. Weight decay with a factor of 0.05 is used as a regularization technique. Furthermore, we train the network with a batch size of 32 and the Mean Squared Error (MSE) loss function. All training and testing runs were executed on an NVIDIA Quadro RTX 6000 with the PyTorch framework.

3.2. Experiments

We perform two sets of experiments: a) on synthetic data to estimate the accuracy of the learnable inverse model, and b) on patient MRI scans to qualitatively probe the transferability of the method to real data.

Experiment			MAE			DICE		
ID	Tumor Model	Predicted Parameters	GM	WM	CSF	c_t=0.01	c_t=0.1	c_t=0.8
1		$\{D_w, \rho, T\}, \{\mathbf{x}\}$	0.461	0.463	0.0	0.607	0.558	0.423
2	FK	$\{\mu_1, \mu_2\}, \{\mathbf{x}\}$	0.059	0.059	0.0	0.940	0.928	0.855
3		$\{\mu_1, \mu_2, \mathbf{x}\}$	**0.057**	**0.057**	**0.0**	**0.943**	**0.932**	**0.861**
4		$\{D_w, \rho, T\}, \{\mathbf{x}\}$	0.089	0.087	0.059	0.846	0.807	0.737
5	ME	$\{\mu_1, \mu_2\}, \{\mathbf{x}\}$	0.059	0.057	0.054	0.877	0.841	0.772
6		$\{\mu_1, \mu_2, \mathbf{x}\}$	**0.055**	**0.054**	**0.054**	**0.886**	**0.850**	**0.783**

Table 2: Ablation analysis on the test set (12k samples). In total, we perform 6 experiments. First three experiments are performed for the Fisher-Kolmogorov (FK) tumor model: 1) Two separate neural networks predicting $\{D_w, \rho, T\}$ and $\{\mathbf{x} = (ic_x, ic_y, ic_z)\}$, 2) Two separate neural networks for prediction of the growth $\{\mu_1, \mu_2\}$ and location $\{\mathbf{x}\}$ parameters, 3) Single neural network predicting $\{\mu_1, \mu_2, \mathbf{x}\}$. The last experiments 4-6 are analogous but performed for the mass-effect (ME) model. Mean absolute error (MAE) $||c_{sim} - c_{gt}||$ in white and gray matter and DICE score for the simulated tumor cell densities thresholded at different c_t is used for all experiments. The non-zero error in the CSF area for ME model comes from the fact that the simulated tumor is allowed to displace the healthy tissue including the CSF. For both FK and ME models the usage of single network predicting time-independent parameters results in a notable increase in accuracy compared to the time-dependent counterpart.

3.3. Synthetic test

In Tab.1, for two different tumor models, we show results of the ablation analysis, wherein we perform multiple experiments varying the neural network input and output configurations. First, we provide empirical proof that a network predicting D_w, ρ, T instead of time-independent parameter combinations

μ_1, μ_2 cannot be trained reliably. Mean absolute error, as well as DICE score, improve significantly when the network predicts μ_1, μ_2 (for the ME model the improvement is less pronounced that we attribute to a higher numerical error in our forward solver implementation for this model). Second, we tested whether the performance can be improved by learning two separate networks predicting growth $\{\mu_1, \mu_2\}$ and initial location parameters $\{\mathbf{x} = (ic_x, ic_y, ic_z)\}$, respectively. This test did not reveal an improvement compared to a single network predicting all parameters.

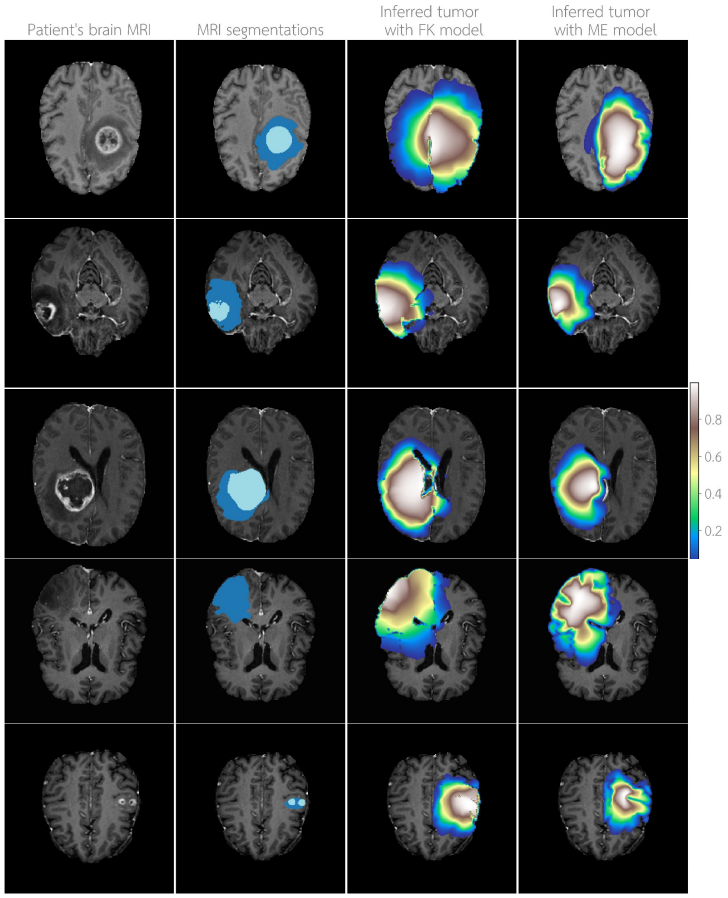

Figure 4: Examples of patient-specific simulations produced using the *Learn-Morph-Infer* method. The inverse model network was trained on samples from the Fisher-Kolmogorov (3rd column) and mass effect (4th column) forward models.

Fig. 3 qualitatively showcases the accuracy of inference using the *Learn-Morph-Infer* pipeline. As discussed before, depending on the network training strategy (either predicting time-dependent $D_w, \rho, T, \mathbf{x}$ parameters, or time-independent $\mu_1, \mu_2, \mathbf{x}$ parameters as network output), the accuracy of the final simulated tumor notably differs. If the proposed learning with time-independent combinations of parameters provides close to the ground truth tumor profile, then the learning with time-dependent combinations makes the inference hardly useful.

Finally, we also analyzed the robustness of the proposed inverse network against the wide range of model parameters. Fig. 5 demonstrates the distribution of the mean absolute error over the range of values for the parameters μ_1, μ_2 pertaining to the test set. From this scatter plots, we conclude that the network performance is stable across the ranges used for the test set.

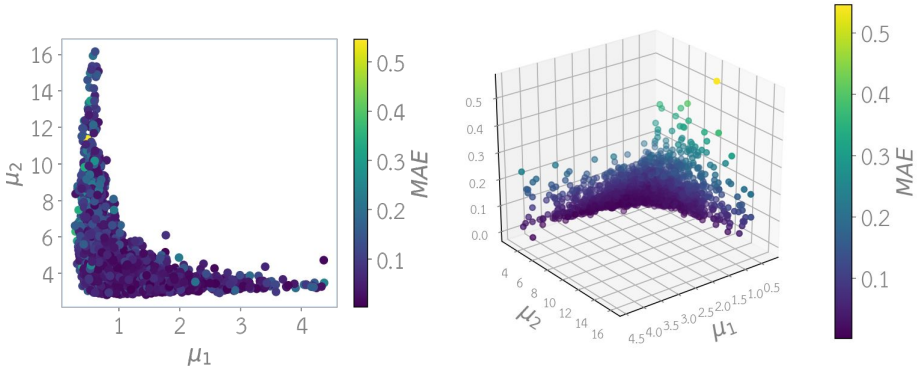

Figure 5: Distribution of the mean absolute error over the range of values for the parameters μ_1, μ_2, in 2D (left) and 3D (rights) views.

3.4. Real MRI patient data

We performed qualitative validation of the method on a large cohort of brain tumor patients who underwent MRI testing. Binary segmentation corresponding to the T1Gd and FLAIR modalities were used as input to the *Learn-Morph-Infer* method. Fig. 4 showcases examples of inferred tumor simulations for

various tumor grades and locations in patients' brains. Out of 80 cases, there were less than 10% cases in which the pipeline outputted wrong results: tumor occupying the whole brain or absence of tumor.

To quantitatively estimate the accuracy of the proposed method, we compared it with the publicly available glioma solver from (Lipkova et al., 2019). Due to the computationally costly inference of the latter, we evaluated the performance on 10 randomly chosen glioma patients. The results of the comparison are provided in Fig. 6. For most of the cases, the overlap in terms of DICE

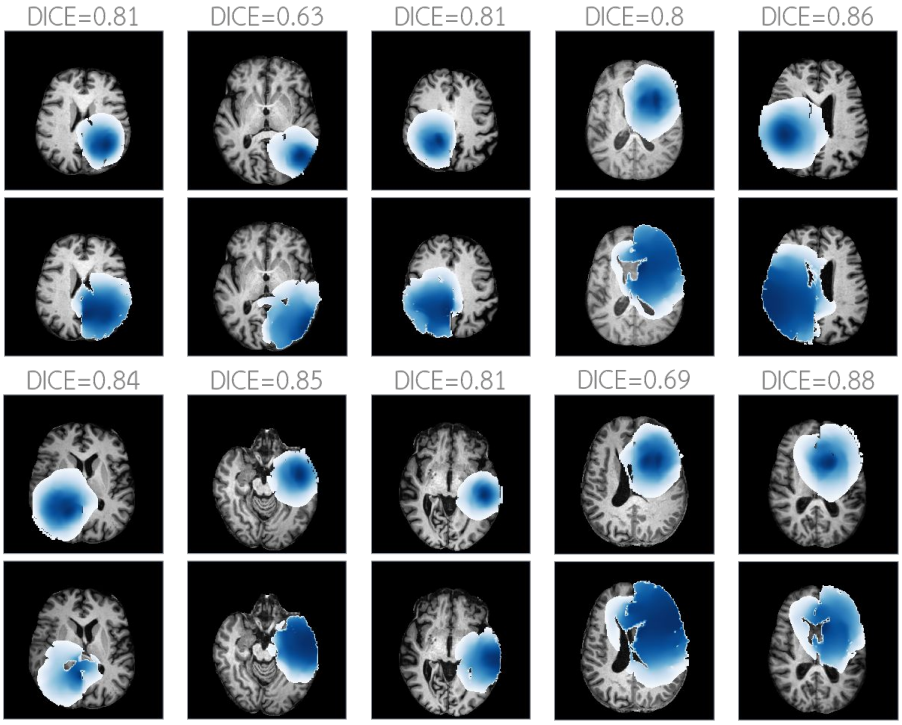

Figure 6: Examples of patient-specific simulations produced using the *Learn-Morph-Infer* method (above) and MCMC-based glioma solver from (Lipkova et al., 2019) (below) for the FK model. Despite of the different nature of the methods and inevitable errors coming from the network misprediction, the DICE score for most of the cases is around 0.8 (the DICE was computed after thresholding the tumor cell concentration at $c_t = 0.01$).

is around 0.8. Note that this error includes not only the contribution from network misprediction and registration back and forth to the brain atlas space but likely also from just different nature of the methods (our proposed method provides point estimates, while the glioma inverse solver by (Lipkova et al., 2019) is a probabilistic method). A natural extension of our work to reduce the error would be to run the sampling-based (or optimization) method in patient space after performing inference with the proposed *Learn-Morph-Infer* method but within narrow parametric ranges centered around estimates predicted by our learnable method. Or even a simpler extension - instead of the steps "C" and "D" in Fig.1, one can run the tumor solver directly in patient space with model parameters inferred at step "B" (such procedure would mitigate the error coming from the mismatch between brain atlas and patient anatomy).

3.5. Computing time

The total time including registration, morphing to atlas space, inference, tumor simulation, and morphing back to patient space is 4-7 minutes. The time for the inverse model network's inference is around 2 seconds for 128^3 resolution. Crucially, the inference time for the more complicated model with mass effect stays the same as for the Fisher-Kolmogorov tumor model. This emphasizes the key practical contribution of the proposed method in that it allows constant time model personalization for an arbitrary tumor model complexity.

4. Conclusion

We present a learnable brain tumor model personalization methodology. We demonstrate that it is feasible to learn an inverse model in a supervised fashion from a data set of numerical simulations. We show that the choice of output can crucially affect network's performance - predicting time-independent combinations of parameters notably outperforms time-dependent counterparts. Such time-independent parametrization is not limited to the PDEs considered here, and thus the proposed *Learn-Morph-Infer* pipeline can be adapted to other inverse problems in natural science and engineering disciplines. For the brain

tumor growth model, the *Learn-Morph-Infer* pipeline provides real-time performance of the parametric inference. Most importantly, the personalization time is stable across tumor models of different numerical complexity. These performance benefits pave the way for clinical testing of various mathematical tumor descriptions on a large cohort of patients.

5. Acknowledgment

I. Ezhov and S. Shit are supported by the Translational Brain Imaging Training Network under the EU Marie Sklodowska-Curie program (Grant ID: 765148). B. Menze, B. Wiestler, and F. Kofler are supported through the SFB 824, subproject B12, by DFG through TUM International Graduate School of Science and Engineering, GSC 81, and by the Institute for Advanced Study, funded by the German Excellence Initiative. The authors acknowledge funding from DCoMEX (Grant agreement ID: 956201). B. Menze is supported by Helmut Horten Foundation.

References

B. B. Avants, N. Tustison, and G. Song. Advanced normalization tools (ants). *Insight j*, 2(365):1–35, 2009.

M. Dax, S. R. Green, J. Gair, J. H. Macke, A. Buonanno, and B. Schölkopf. Real-time gravitational wave science with neural posterior estimation. *Physical review letters*, 127(24):241103, 2021.

I. Ezhov, T. Mot, S. Shit, J. Lipková, J. C. Paetzold, F. Kofler, F. Navarro, C. Pellegrini, M. Kollovieh, M. Metz, B. Wiestler, and B. Menze. Geometry-aware neural solver for fast bayesian calibration of brain tumor models. 2020.

I. Ezhov et al. Neural parameters estimation for brain tumor growth modeling. In *MICCAI*, pages 787–795. 2019.

E. Geremia, B. H. Menze, M. Prastawa, M.-A. Weber, A. Criminisi, and N. Ayache. Brain tumor cell density estimation from multi-modal mr images based on a synthetic tumor growth model. In *International MICCAI Workshop on Medical Computer Vision*, pages 273–282. Springer, 2012.

H. L. Harpold, E. C. Alvord Jr, and K. R. Swanson. The evolution of mathematical modeling of glioma proliferation and invasion. *Journal of Neuropathology & Experimental Neurology*, 66(1):1–9, 2007.

K. He, X. Zhang, S. Ren, and J. Sun. Deep residual learning for image recognition. In *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2016. doi: 10.1109/CVPR.2016.90.

C. Hogea et al. An image-driven parameter estimation problem for a reaction–diffusion glioma growth model with mass effects. *J Math Biol*, 56(6):793–825, 2008.

D. A. Hormuth, S. L. Eldridge, J. A. Weis, M. I. Miga, and T. E. Yankeelov. Mechanically coupled reaction-diffusion model to predict glioma growth: methodological details. In *Cancer Systems Biology*, pages 225–241. Springer, 2018.

D. A. Hormuth, A. M. Jarrett, X. Feng, and T. E. Yankeelov. Calibrating a predictive model of tumor growth and angiogenesis with quantitative mri. *Annals of biomedical engineering*, 47(7):1539–1551, 2019.

D. A. Hormuth et al. Image-based personalization of computational models for predicting response of high-grade glioma to chemoradiation. *Scientific Reports*, 11(1):1–14, 2021.

P. R. Jackson, J. Juliano, A. Hawkins-Daarud, R. C. Rockne, and K. R. Swanson. Patient-specific mathematical neuro-oncology: Using a simple proliferation and invasion tumor model to inform clinical practice. *BMB*, 77(5): 846–856, mar 2015.

M. Kasim et al. Up to two billion times acceleration of scientific simulations with deep neural architecture search. *arXiv preprint arXiv:2001.08055*, 2020.

B. Kim et al. Deep fluids: A generative network for parameterized fluid simulations. *Comput. Graph. Forum*, 38:59–70, 2018.

D. P. Kingma and J. Ba. Adam: A method for stochastic optimization. In *ICLR 2015*, 2015.

F. Kofler et al. Brats toolkit: Translating brats brain tumor segmentation algorithms into clinical and scientific practice. *Frontiers in Neuroscience*, 14, 2020.

E. Konukoglu et al. Extrapolating glioma invasion margin in brain magnetic resonance images: Suggesting new irradiation margins. *MedIA*, 14(2):111 – 125, 2010a.

E. Konukoglu et al. Image guided personalization of reaction-diffusion type tumor growth models using modified anisotropic eikonal equations. *IEEE Transactions on Medical Imaging*, 29(1):77–95, 2010b.

M. Le et al. Personalized radiotherapy planning based on a computational tumor growth model. *IEEE Transactions on Medical Imaging*, 36(3):815–825, mar 2017.

J. Lipkova et al. Personalized radiotherapy design for glioblastoma: Integrating mathematical tumor models, multimodal scans and bayesian inference. *IEEE Transactions on Medical Imaging*, pages 1–1, 2019.

G. Lorenzo and et al. Quantitative in vivo imaging to enable tumor forecasting and treatment optimization. *arXiv preprint arXiv:2102.12602*, 2021.

I. Loshchilov and F. Hutter. Decoupled weight decay regularization. In *International Conference on Learning Representations*, 2019. URL `https://openreview.net/forum?id=Bkg6RiCqY7`.

J.-M. Lueckmann et al. Flexible statistical inference for mechanistic models of neural dynamics. In *NeurIPS*, pages 1289–1299, 2017.

M. Lê, H. Delingette, J. Kalpathy-Cramer, E. R. Gerstner, T. Batchelor, J. Unkelbach, and N. Ayache. Mri based bayesian personalization of a tumor growth model. *IEEE Transactions on Medical Imaging*, 35(10):2329–2339, 2016. doi: 10.1109/TMI.2016.2561098.

A. Mang, A. Toma, T. A. Schuetz, S. Becker, T. Eckey, C. Mohr, D. Petersen, and T. M. Buzug. Biophysical modeling of brain tumor progression: From unconditionally stable explicit time integration to an inverse problem with parabolic pde constraints for model calibration. *Medical Physics*, 39(7Part1): 4444–4459, 2012.

B. Menze et al. A generative approach for image-based modeling of tumor growth. In *IPMI*, pages 735–747, 2011. ISBN 978-3-642-22092-0.

B. H. Menze et al. The multimodal brain tumor image segmentation benchmark (brats). *IEEE transactions on medical imaging*, 34(10):1993–2024, 2014.

G. Papamakarios and I. Murray. Fast ϵ-free inference of simulation models with bayesian conditional density estimation. In *NeurIPS*, pages 1028–1036. 2016.

V. Patel and L. Hathout. Image-driven modeling of the proliferation and necrosis of glioblastoma multiforme. *Theoretical Biology and Medical Modelling*, 14(1): 1–9, 2017.

S. Pati, V. Sharma, H. Aslam, S. Thakur, H. Akbari, A. Mang, S. Subramanian, G. Biros, C. Davatzikos, and S. Bakas. Tmod-09. glioblastoma biophysical growth estimation using deep learning-based regression. *Neuro-Oncology*, 22 (Supplement_2):ii229–ii229, 2020.

M. Raissi, P. Perdikaris, and G. E. Karniadakis. Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *Journal of Computational Physics*, 378:686–707, 2019.

T. Rohlfing, N. M. Zahr, E. V. Sullivan, and A. Pfefferbaum. The sri24 multichannel atlas of normal adult human brain structure. *Human brain mapping*, 31(5):798–819, 2010.

K. Scheufele, A. Mang, A. Gholami, C. Davatzikos, G. Biros, and M. Mehl. Coupling brain-tumor biophysical models and diffeomorphic image registration. *Computer methods in applied mechanics and engineering*, 347:533–567, 2019.

K. Scheufele, S. Subramanian, and G. Biros. Automatic mri-driven model calibration for advanced brain tumor progression analysis. *arXiv: Medical Physics*, 2020.

V. Sitzmann et al. Implicit neural representations with periodic activation functions. *arXiv preprint arXiv:2006.09661*, 2020.

B. Stevens and T. Colonius. Finitenet: A fully convolutional lstm network architecture for time-dependent partial differential equations. *arXiv preprint arXiv:2002.03014*, 2020.

R. Stupp, M. Brada, M. Van Den Bent, J.-C. Tonn, and G. Pentheroudakis. High-grade glioma: Esmo clinical practice guidelines for diagnosis, treatment and follow-up. *Annals of oncology*, 25:iii93–iii101, 2014.

S. Subramanian, A. Gholami, and G. Biros. Simulation of glioblastoma growth using a 3d multispecies tumor model with mass effect. *Journal of mathematical biology*, 79(3):941–967, 2019.

S. Subramanian, K. Scheufele, N. Himthani, and G. Biros. Multiatlas calibration of biophysical brain tumor growth models with mass effect. *arXiv preprint arXiv:2006.09932*, 2020a.

S. Subramanian, K. Scheufele, M. Mehl, and G. Biros. Where did the tumor start? an inverse solver with sparse localization for tumor growth models. *Inverse Problems*, 36(4):045006, 2020b.

K. R. Swanson, E. C. Alvord Jr, and J. Murray. A quantitative model for differential motility of gliomas in grey and white matter. *Cell proliferation*, 33(5):317–329, 2000.

N. Thuerey, K. Weißenow, L. Prantl, and X. Hu. Deep learning methods for reynolds-averaged navier–stokes simulations of airfoil flows. *AIAA Journal*, 58(1):15–26, 2020.

B. Tunc, D. Hormuth, G. Biros, and T. E. Yankeelov. Modeling of glioma growth with mass effect by longitudinal magnetic resonance imaging. *IEEE Transactions on Biomedical Engineering*, 2021.

B. Tunç, D. A. Hormuth II, G. Biros, and T. E. Yankeelov. Modeling of glioma growth with mass effect by longitudinal magnetic resonance imaging. *IEEE Transactions on Biomedical Engineering*, 68(12):3713–3724, 2021.

T. E. Yankeelov, N. Atuegwu, D. Hormuth, J. A. Weis, S. L. Barnes, M. I. Miga, E. C. Rericha, and V. Quaranta. Clinically relevant modeling of tumor growth and treatment response. *Science translational medicine*, 5(187): 187ps9–187ps9, 2013.

H. Young, R. Baum, U. Cremerius, K. Herholz, O. Hoekstra, A. Lammertsma, J. Pruim, P. Price, et al. Measurement of clinical and subclinical tumour response using [18f]-fluorodeoxyglucose and positron emission tomography: review and 1999 eortc recommendations. *European journal of cancer*, 35(13): 1773–1782, 1999.