

A Generative Approach for Image-Based Modeling of Tumor Growth

Bjoern H. Menze^{1,2}, Koen Van Leemput^{1,3,4}, Antti Honkela⁵, Ender Konukoglu⁶ Marc-André Weber⁷, Nicholas Ayache², and Polina Golland¹

¹ Computer Science and Artificial Intelligence Laboratory,
Massachusetts Institute of Technology, USA

² Asclepios Research Project, INRIA Sophia-Antipolis, France

³ Dept. of Radiology, Massachusetts General Hospital, Harvard Medical School, USA

⁴ Department of Information and Computer Science, Aalto University, Finland

⁵ Helsinki Institute for Information Technology HIIT, University of Helsinki, Finland

⁶ Machine Learning and Perception Group, Microsoft Research, Cambridge, UK

⁷ Department of Diagnostic Radiology, Heidelberg University Hospital, Germany

Abstract. Extensive imaging is routinely used in brain tumor patients to monitor the state of the disease and to evaluate therapeutic options. A large number of multi-modal and multi-temporal image volumes is acquired in standard clinical cases, requiring new approaches for comprehensive integration of information from different image sources and different time points. In this work we propose a joint generative model of tumor growth and of image observation that naturally handles multi-modal and longitudinal data. We use the model for analyzing imaging data in patients with glioma. The tumor growth model is based on a reaction-diffusion framework. Model personalization relies only on a forward model for the growth process and on image likelihood. We take advantage of an adaptive sparse grid approximation for efficient inference via Markov Chain Monte Carlo sampling. The approach can be used for integrating information from different multi-modal imaging protocols and can easily be adapted to other tumor growth models.

1 Introduction

Processes related to tumor growth can be modeled at different scales ranging from signaling at sub-cellular level, via multi-cellular processes determining metabolic properties of the tumor, to the gross bio-mechanical behavior of tumor-affected tissue at a macroscopic scale. Tumor models in medical image analysis rely almost exclusively on information from morphologic images and consequently focus on the macroscopic phenomena of tumor evolution. An important class of macroscopic tumor models is based on the reaction-diffusion equations [1–8]. In this paper, we propose an efficient Bayesian framework for image-based inference in this type of tumor growth models.

Tumor models are used to describe effects a lesion has on surrounding healthy tissue. Modeling this effect helps to improve inter-subject image and atlas registration [8–10] and to improve tissue segmentation in the presence of the lesion [11, 12]. As the tumor location it is not consistent across patients, tumor

shape is often employed for constructing spatial priors. For example, a large number of simulated tumors can be used to learn characteristic tumor-induced deformations [13]. For data of individual patients growth models also help to estimate deformation fields and tissue perturbations [8, 10]. Most of these methods take a static approach to modeling tumors for *removing* the tumor-induced effects from the analysis of the subject’s brain itself.

Other biophysical tumor models describe explicitly the evolution of the lesion. For glioma most such models follow a reaction-diffusion framework [1]. They assume an infiltrative growth of the tumor cells, may consider differences in cell diffusion in white and gray matter [2], and model locally anisotropic migration patterns by integrating information from diffusion tensor images (DTI) [3]. Some models also include the mechanical effect of the lesion on surrounding structures by modeling the interaction between the tumor mass and the invaded tissue [3, 5], or couple the growth dynamics with phenomenological image models for Gadolinium uptake and changes in DTI [6]. The primary application for these image-based dynamical tumor models is in simulating tumor growth, either to analyze macroscopic patterns of disease progression [4], or to generate realistically appearing tumor images to be used for validation of tumor segmentation methods [6].

Many dynamical tumor models could serve as a framework for integrating complex information from different image modalities and longitudinal data sets. But unlike the static whole brain models, the image-based personalization remains challenging for most forward simulators. As a consequence, some studies only provide qualitative measures for analyzing lesions and their progression [14], employ strict assumptions on the relationship between the tumor cell density and image observation [2], or focus on a theoretical treatment of the inverse problem under the assumption of having appropriate observables available at a later stage [5]. The PDE-constrained optimization approach in [5] relates the tumor growth model to landmark-based registration using a reaction-diffusion-advection model. This approach is similar to the whole brain modeling [8, 10, 13], but depends critically on the tumor-tissue interaction model. Alternatively, a traveling wave approximation of the reaction-diffusion model can be used to estimate model parameters from the temporal evolution of the tumor front [7]. Unfortunately, this approach only provides the speed of growth. Furthermore, all of the previous methods focus on point-estimates of the process. This presents significant challenges in presence of noisy observations and uncertainty in evolution models.

In this paper, we develop a novel approach for personalizing tumor models, for integrating multi-modal data, and for quantifying uncertainty. We formulate a generative model that includes a biophysical tumor growth model and statistical observation models for different image modalities. We devise an efficient inference scheme based on the forward model and applicable in high-dimensional observation settings. We demonstrate the resulting method in experiments with both synthetic and clinical patient data. We envision that a joint tumor and




Fig. 1. Joint process and observation model. The physical model defines the state variable \mathbf{u} as a function of model parameters θ_p and initial conditions θ_x and θ_t . The imaging models relate \mathbf{u} to the image observations \mathbf{y}^m and \mathbf{y}^s , representing functional maps and tumor segmentations, respectively. Variables are represented by circles, parameters are shown as squares. We estimate the model parameters θ_p to characterize the state of disease visible from images \mathbf{y} .

image modeling will close the gap between functional image interpretation and disease modeling, and will provide new directions for therapy optimization.

2 Tumor growth model

Our approach includes a physical process model that describes the progression of the disease through the evolution of the tumor cell concentration \mathbf{u} and a probabilistic imaging model that captures the relationship between the latent state \mathbf{u} and the image observations \mathbf{y} . Figure 1 presents the full graphical model described in this section.

2.1 Physical process model

We let $\mathbf{u} = (u_1, \dots, u_I)^T$ be the latent state variables where $u_i \in [0, 1]$ is the tumor cell concentration in voxel i ($1 \leq i \leq I$). We model the temporal evolution of the tumor as an inhomogeneous anisotropic diffusion governed by the Fisher-Kolmogorov equation [1]:

$$\frac{\partial \mathbf{u}}{\partial t} = \nabla_x (\mathbf{D} \nabla_x \mathbf{u}) + \rho \cdot \mathbf{u} (1 - \mathbf{u}), \quad (1)$$

where ∇_x represents the spatial derivative operator. The equation describes the growth of a glioma as an infiltrative process with diffusivity tensor \mathbf{D} and self-limiting growth with proliferation rate ρ . We further assume that $\mathbf{D} = D \cdot \hat{\mathbf{D}}$, where $\hat{\mathbf{D}}$ is an patient-specific diffusion tensor, observed via DTI, and D is a global diffusivity parameter [3]. We apply Neumann border conditions, and assume that tumor cells migrate in the white and gray matter only, with higher diffusivity in the white matter ($D_w \gg D_g$). To personalize the model, we construct estimates of the model parameters $\theta_p = \{D, \rho\}$ individually for every patient.

By integrating Eq. (1) over time – starting from the initial state $\mathbf{u}(t = 0)$ when the first image of the tumor was acquired – we obtain a 4D functional model $U(x, y, z, t)$ that describes the evolution $\mathbf{u}(t)$ at any time point t . We choose a parametric representation for the initial state by placing a seed u^{init} in a voxel located at $\theta_x = \{x, y, z\}$ and by growing the tumor for the time interval θ_t . We form the deterministic tumor evolution model:

$$p(\mathbf{u}(t) | \theta_x, \theta_t, \theta_p) = \delta(U(x, y, z, t; \theta_x, \theta_t, \theta_p) - \mathbf{u}(x, y, z, t)), \quad (2)$$

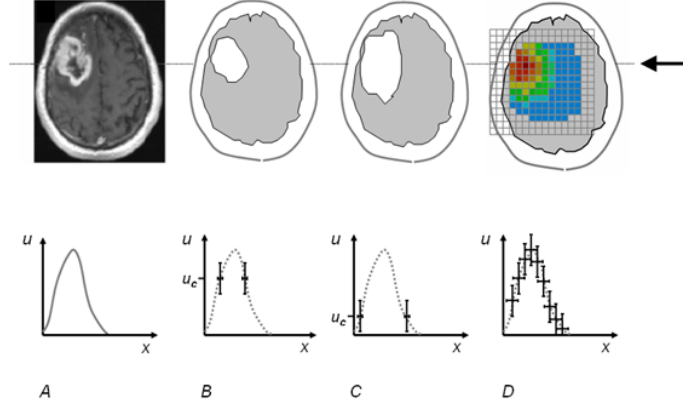


Fig. 2. Example relationships between image information \mathbf{y} (top) and tumor cell density \mathbf{u} (bottom). *A*: actual T₁gad image and the schematic tumor cell distributions along the line indicated by the arrow in the top row. *B*, *C*: exemplary modalities where only the visible tumor segmentation \mathbf{y}^s can be correlated with specific values u_c of the state variable \mathbf{u} . *D*: exemplary modality where all entries of the functional map \mathbf{y}^m can be correlated with the latent state variable \mathbf{u} .

where δ is Dirac's delta indicating that \mathbf{u} is an exact solution of the functional $U(x, y, z, t)$. We choose to use a deterministic model without a stochastic component in Eq. (1) or noise term in Eq. (2) as it represents a stronger constraint on the growth process. We incorporate non-deterministic behavior into our framework via a probabilistic model of the imaging process, as described below.

2.2 Imaging models

Observations $\mathbf{y} = (y_1, \dots, y_T)^T$ are acquired at each voxel i at specific times t , representing partial observations of the process. We evaluate every image volume $\mathbf{y}(t)$ with the corresponding state vector $\mathbf{u}(t)$. For notational simplicity we drop the time index in the imaging models below. As illustrated in Fig. 2, we assume that there are two different types of observations that can be related to \mathbf{u} : continuous *functional maps* and binary *tumor segmentations* which are the most common observation.

Functional maps Functional maps \mathbf{y}^m contain continuous values $y_i^m \in \mathbb{R}$. They represent imaging modalities with functional parameters that can be directly correlated with the latent physiological variable \mathbf{u} using a (possibly non-linear) forward observational operator F . Formally,

$$p(\mathbf{y}^m | \mathbf{u}, \theta_y^m) = \prod_i p(y_i^m | \mathbf{u}, \theta_y^m) = \prod_i \mathcal{N}(y_i^m; \mu_i, \sigma_m^2), \quad (3)$$

where $\mathcal{N}(\cdot; \mu, \sigma^2)$ denotes a normal distribution with mean μ and variance σ^2 . Eq. 3 represents a noisy observation of

$$\mu_i = F(b^c \cdot \mathbf{u})_i. \quad (4)$$




Fig. 3. Parameter space of the tumor model with tumor cell diffusivity D and proliferation rate ρ . Left: Parameterization for high- and low-grade glioma [4, 16]. The speed of growth v can be obtained from times series of images [7]. It provides complementary information to tumor shape. Right: Shapes of equally large tumors for different parameterizations. Isolines of $\mathbf{u} = .05$ and $\mathbf{u} = .6$ are shown in yellow. We use the information from both time series and tumor shape jointly to characterize infiltrative and proliferative behavior of the tumor.

We let $\theta_y^m = \{b^c, \sigma_m^2\}$ be the parameters of this model, where b^c is the coefficient in a linear model relating \mathbf{u} and \mathbf{y}^m in a first order approximation, relevant for example for magnetic resonance spectroscopic images (MRSI). A function $F(\cdot)$ could be, for example, a known nonlinear smoothing and subsampling operator, modeling the point spread function of the imaging process and different spatial resolution of \mathbf{y}^m and \mathbf{u} . Examples for \mathbf{y}^m include metabolic maps from MRSI [15], or apparent diffusion coefficient maps from DTI.

Binary segmentations Tumor segmentations contain discrete labels \mathbf{y}^s , $y_i^s \in \{0, 1\}$. Since tumor outlines are commonly associated with constant tumor infiltration, we model observation \mathbf{y}^s for a given tumor cell concentration \mathbf{u} as a Bernoulli random variable:

$$p(\mathbf{y}^s | \mathbf{u}, \theta_y^s) = \prod_i p(y_i^s | \mathbf{u}, \theta_y^s) = \prod_i \alpha_i^{y_i^s} \cdot (1 - \alpha_i)^{1 - y_i^s}, \quad (5)$$

where α_i is the probability of observing characteristic tumor-induced changes:

$$\alpha_i = .5 + .5 \cdot \text{sign}(u_i - u^c) \left(1 - e^{-\frac{(u_i - u^c)^2}{\sigma_s^2}} \right), \quad (6)$$

which is a double logistic sigmoid function with parameters $\theta_y^s = \{u^c, \sigma_s^2\}$. We essentially assume that the tumor cell infiltration is invisible below the threshold u^c but modifies the image signal in a predictable way after surpassing u^c [2]. Parameter σ_s^2 transforms the hard threshold into a smoother decision, also reflecting the uncertainty in the threshold u^c .

2.3 Joint model

Combining the deterministic tumor growth model in Eq. (2) with the image observation models in Eqs. (3)–(5) and letting $\mathbf{Y} = [\mathbf{y}_1, \dots, \mathbf{y}_k]$ denote the col-

lection of all k image observations acquired at N time points t_n , we obtain

$$\begin{aligned}
p(\mathbf{Y}|\theta_x, \theta_t, \theta_p, \theta_y) &= \prod_{t_n} p(\mathbf{y}(t_n)|\theta_x, \theta_t, \theta_p, \theta_y) \\
&= \prod_{t_n} \int_{\mathbf{u}(t_n)} p(\mathbf{y}(t_n)|\mathbf{u}(t_n), \theta_y) p(\mathbf{u}(t_n)|\theta_x, \theta_t, \theta_p) d\mathbf{u}(t_n) \\
&= \prod_{t_n} \int_{\mathbf{u}(t_n)} \prod_k p(\mathbf{y}_k(t_n)|\mathbf{u}(t_n), \theta_y) p(\mathbf{u}(t_n)|\theta_x, \theta_t, \theta_p) d\mathbf{u}(t_n), \quad (7)
\end{aligned}$$

assuming that all observations are conditionally independent given the latent state \mathbf{u} . We adapt a factored prior $p(\theta_x, \theta_t, \theta_p, \theta_y) = p(\theta_x)p(\theta_t)p(\theta_p)p(\theta_y)$ and choose uniform distributions for all parameters. Similar to [4], we choose the range for the evolution parameters θ_p as illustrated in Fig. 3. We use experimental evidence [2, 15] to set the range of the imaging parameters θ_y and use the life expectancy of a healthy person to set the range for the temporal parameter θ_t . We assume that the tumor started growing from a location within the hyperintense areas in T₂ MRI at $t = 0$ to set a prior on the spatial parameter θ_x . We obtain the joint posterior distribution of the parameters using Bayes' rule:

$$p(\theta_x, \theta_t, \theta_p, \theta_y|\mathbf{Y}) \propto p(\mathbf{Y}|\theta_x, \theta_t, \theta_p, \theta_y)p(\theta_x, \theta_t, \theta_p, \theta_y), \quad (8)$$

from which we compute the posterior distribution of θ_p via marginalization

$$p(\theta_p|\mathbf{Y}) = \int_{\theta_x} \int_{\theta_t} \int_{\theta_y} p(\theta_x, \theta_t, \theta_p, \theta_y|\mathbf{Y}) d\theta_x d\theta_t d\theta_y. \quad (9)$$

$p(\theta_p|\mathbf{Y}) = p(D, \rho|\mathbf{Y})$ is the desired quantity of diagnostic interest in our application.

3 Inference

The posterior distribution of $\theta_p = \{\rho, D\}$ allows us to summarize a large set of image observations \mathbf{y} through diffusivity D and cell doubling rate ρ (Fig. 3). We aim at visualizing $p(\theta_p|\mathbf{Y})$ by drawing samples from it using an efficient Markov Chain Monte Carlo (MCMC) sampling strategy. MCMC constructs a random walk through the parameter space which converges towards the solution. The method avoids intractable integrals in Eq. (7). Unfortunately, the repeated evaluation of Eq. (9) for this walk requires the costly forward integration of the physical model in Eq. (1) for any sample from θ_p . To overcome this challenge, we separate the integration over the model parameters in Eq. (9): we integrate over parameters θ_t and θ_y which can be calculated sequentially at any step when propagating the forward model, and rely on MCMC only to sample from θ_x and θ_p .

Sequential integration of θ_t and θ_y Given the parameters of the tumor growth model θ_p and a starting point θ_x , we propagate the physical model in Eq. (1) based the deterministic model U in Eq. (2). During the forward integration of Eq. (1), the integral

$$p(\theta_p, \theta_x | \mathbf{Y}) = \int_{\theta_y} \int_{\theta_t} p(\theta_x, \theta_t, \theta_p, \theta_y | \mathbf{Y}) d\theta_t d\theta_y \quad (10)$$

can be calculated efficiently in a single sweep, by evaluating the image likelihood $p(\mathbf{y} | \mathbf{u}, \theta_y)$ at any step forward in time, and subsequently integrating over all steps. This procedure represents an integration over θ_t on a regularly spaced grid. Similarly, we integrate over the parameters θ_y^m and θ_y^s on a regularly spaced grid. This step can be performed very quickly once $\mathbf{u}(t_n)$ has been calculated for the different observations $\mathbf{y}(t_n)$. This procedure approximates Eq. (10) by a sum:

$$p(\theta_p, \theta_x | \mathbf{Y}) \approx \sum_{\theta_t^{(g)}} \sum_{\theta_y^{(g)}} p(\theta_x, \theta_t^{(g)}, \theta_p, \theta_y^{(g)} | \mathbf{Y}). \quad (11)$$

Efficient sampling of θ_x and θ_p Even with the proposed sequential integration scheme, evaluating $p(\theta_p, \theta_x | \mathbf{Y})$ repeatedly for different values $\{\theta_p, \theta_x\}$ during MCMC sampling remains prohibitively time consuming. We therefore rely on a representation of $p(\theta_p, \theta_x | \mathbf{Y})$ that can be sampled more efficiently. In particular, we approximate $p(\theta_p, \theta_x | \mathbf{Y})$ through $\tilde{p}(\theta_p, \theta_x | \mathbf{Y})$ using a so-called *sparse grid* basis [17, 18] in the first step. This approximation allows us to sample the space of $p(\theta_p, \theta_x | \mathbf{Y})$ in a *structured* way, with less redundant computations, and systematically increasing the accuracy of the approximation with every sample. More specifically, we interpolate between G pre-computed sampling points:

$$\tilde{p}(\theta_p, \theta_x | \mathbf{Y}) = \sum_{g=1}^G p(\theta_p^{(g)}, \theta_x^{(g)} | \mathbf{Y}) \cdot \Phi_s(\theta_p, \theta_x), \quad (12)$$

relying on a hierarchical linear basis function Φ_G that spans the space of θ_p and θ_x . This approximation is constructed from a small number of evaluations of $p(\theta_p, \theta_x | \mathbf{Y})$, sampled at specific grid locations $\{\theta_p^{(g)}, \theta_x^{(g)}\}$ within the range of $p(\theta_x)$ and $p(\theta_p)$. The interpolated posterior distribution can be evaluated at little computational cost. In the second step, we use the pre-computed approximation $\tilde{p}(\theta_p, \theta_x | \mathbf{Y})$ when evaluating a large number of samples from $p(\theta_p, \theta_x | \mathbf{Y})$ via MCMC. Efficiently sampling from $\tilde{p}(\theta_p, \theta_x | \mathbf{Y})$ also allows us to tune the sampling algorithm more easily in the presence of local minima. In the final step, we construct

$$\tilde{p}(\theta_p | \mathbf{Y}) = \int_{\theta_x} \tilde{p}(\theta_p, \theta_x | \mathbf{Y}) d\theta_x \quad (13)$$

by aggregating values of θ_p from our samples θ_x, θ_p . In order to minimize the number of necessary samples from $p(\theta_p, \theta_x | \mathbf{Y})$, i.e., the number of forward integrations of the physical model, we choose the sparse grid collocation from [19,

20], but with a depth first search strategy for local refinement. To sample efficiently from Eq. (12) we use the Delayed Rejection Adaptive Metropolis (DRAM) variant of the Metropolis-Hastings MCMC sampler [21].

4 Experiments

We evaluate our method on a series of images acquired for monitoring patients with low-grade glioma. The low-grade glioma is typically only visible as hyperintense lesion in T_2 /FLAIR images. Patients are monitored for the occurrence of T_1 gad enhancements indicating the transition to high-grade disease and to be treated immediately. However, clinical evidence suggests that the dynamics of the tumor growth can predict this transition at a much earlier stage [16] (Fig. 3, left). We hypothesize that model parameters θ_p may provide more accurate information about the state of disease. Here we evaluate our approach under this clinical objective on synthetic ground truth and real clinical data. We also compare standard MCMC sampling with the proposed adaptive sparse grid MCMC sampling.

Implementation We implement the tumor growth model in Eq. (1) by employing the preconditioned conjugate gradient descent to solve the state equations [22]. We sequentially evaluate Eq. (10) while integrating Eq. (1). In every forward step we evaluate the image likelihood, integrating on a grid over $u_c(T_2) = .01 \dots .2$ and $u_c(T_1\text{gad}) = .6 \dots .8$ with $\sigma_2^s = .05 \dots 100$ and for metabolic maps for $b^c(\text{Choline}) = 1 \dots .8$ and $b^c(\text{NAA}) = -8 \dots 1$ with $\sigma_m^2 = .5 \dots 2$. We impose general physiological constraints such as a maximum time of growth of 50 years and a maximal sustainable tumor volume of 150 cm^3 .

We perform integration by sampling in a five dimensional space with $\{D, \rho\} = \theta_p$ and $\{x, y, z\} = \theta_x$ and always initialize the sampling in the center of the valid parameter space. We evaluate the samples via MCMC both with and without sparse grid representation. In the first case (*direct MCMC*) we perform a random walk with 2000 samples. We use the last 1500 samples for evaluating the statistic of the parameters. Similar to [21], we use 100 samples to adapt the proposal distribution, and propagate up to two steps into suboptimal directions of the sampling space. We obtain rejection rates of 40%-60% with approximately 3000-5000 integrations of Eq. (10). In the second case (*sparse grid MCMC*) we evaluate Eq. (10) about the same number of times, i.e., for 3000 basis vectors, but at positions in the parameter space defined by an adaptive sparse grid with “maximum norm” basis [18]. We then perform MCMC sampling under the same settings as for the *direct MCMC*, but use the approximated posterior, interpolated from the sparse grid basis, as the function to be evaluated for each sample. The *direct* sampling takes about 6-12 hours for each test case on a standard personal computer. The *sparse grid* sampling takes a similar amount of computing, but can be easily parallelized, resulting in an algorithm that is 8-10 times faster. Subsequent MCMC sampling using the sparse grid interpolation is accomplished within minutes.

Data In the first experiment, we model a synthetic low-grade glioma. We use DTI and tissue segmentations from a healthy segment of a patient data set to evolve a synthetic tumor with $D = 10^{-2}$ and $\rho = 10^{-2.3}$ for $\theta_t = 2000$ days, starting from a tumor seed point u^{init} with 10% infiltration in a 1 mm^3 volume. We assume $D = D_w = 10^3 D_g$ [4, 7]. We model a second observation by propagating the model for another 180 days. For both observations we model the segmentations from T₂/FLAIR images using a threshold of $u_c(\text{T}_2) = .095$. We set $u_c(\text{T}_1\text{gad}) = .695$. This is higher than the maximum value of \mathbf{u} at both time points and our synthetic T₁gad images do not show contrast agent-enhanced regions. We model metabolic maps from MRSI by smoothing \mathbf{u} with a 1 cm wide Gaussian kernel and subsampling the volume on a grid with 1 cm voxels and model two metabolite maps with coefficients $b^c(\text{Cho}) = 1.8$ and $b^c(\text{NAA}) = -4.6$. In the second experiment, we model a developing high-grade tumor. We use the same setting as above, but with $D = 10^{-0.2}$ and $\rho = 10^{-1.3}$. The tumor is evolved for 250 days for the first observation, and another 90 days for the second. Figure 4 (left) shows the second time point. All observations are modeled as above; T₁gad images do not show enhancements in this experiment either, and the tumor still appears as a low-grade glioma.

In addition to the two test cases with ground truth for D and ρ , we evaluate our model on two clinical data sets. The first data set comprises six sets of images acquired approximately every 3 months over 1.5 years. The second data set comprises four sets of images, acquired every 3-6 months in a similar time span. Available are in both cases DTI, T₂/FLAIR and T₁gad images. T₂/FLAIR images show a visible progression of the tumor front, while T₁gad is free of Gadolinium enhancements. The lesion is segmented manually in three 2D slices intersecting with the tumor center in the first case; it is manually segmented in all slices for the second case. To ensure the diffusion tensor to be free of tumor-induced effects, we use the mirrored DTI of the contra-lateral, disease-free hemisphere to evolve the tumor. The second data set was provided by the authors of [7], who reported the speed of tumor growth in the white matter of $v_w = 2\sqrt{D_w\rho} = 2\sqrt{.2 \cdot .008} \text{ mm/d}$ for this patient.

Results Fig. 4 shows the adaptive sparse grid for the parameters of the growth model θ_p (left) and the tumor seed point θ_x (right). Green dots represent sampling points used to construct the approximation of the posterior in Eq. (12). The right image shows the central slice of the synthetic *high grade* data set. Indicated by dots are $\{x, y\}$ positions of the location parameters θ_x . Dot size indicates the number of model evaluations with varying $\{z, D, \rho\}$ for the given $\{x, y\}$. We find most of the adaptively chosen locations to be close to the real location of the tumor (pink cross). This also holds true for the parameter space spanned by diffusivity D and proliferation ρ in the left image. Here the location close to the actual parameter (pink cross) is the one evaluated for the highest number of different seed points. The contour map shows the likelihood evaluated on a regular grid with the true seed point θ_x and is in good agreement with the adaptively sampled locations on the sparse grid. The most likely region is a diagonal line in D - ρ space, similar to areas of constant shape in Fig. 3 (right) and

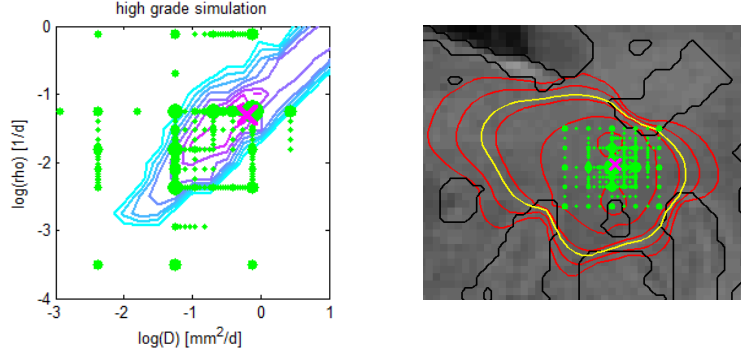


Fig. 4. Adaptive sampling of the parameter space for the synthetic high-grade data set. Left: Sampling points (green) for the model parameters $\theta_p = \{D, \rho\}$. The contour map shows the posterior distribution $p(\theta_p|\mathbf{Y})$ evaluated on a regular grid (cyan–low, pink–high). Right: Sampling points (green) for the $\{x, y\}$ coordinates of the tumor seed θ_x . The Fig. also shows isolines of tumor cell density (red), the predicted extensions of the T2 hyper-intense area (yellow) and tissue boundaries (black). In both figures the size of the green sampling points indicates how often the indicated parameter was evaluated under different combinations. The ground truth is indicated by the pink cross. Most adaptively chosen sampling points are close to the ground truth.

orthogonal to regions of constant speed of growth in Fig. 3 (left). This indicates that information about the shape of the tumor is effectively used to infer model parameters in our proposed approach.

Fig. 5 reports estimates of the model parameters D and ρ for all four test cases and both the direct sampling (blue) and the sparse grid sampling (green). For the first two data sets (A , B), the sparse grid sampling is in good agreement with the ground truth (pink cross). The mean value of the direct sampling (blue/black circle) is relatively close to ground truth and the mean of the sparse grid sampling (green/black circle). We find the sparse grid results to be much more compact with noticeably less variation. This observation is also true for the patient data (C , D). We used the current number of 3000 samples for sparse grid MCMC only to allow for a direct comparison of both approaches. Already an interpolation with a few hundred samples led to results similar to those in Fig. 3, demonstrating the benefits of the sparse sampling approach for estimating parameters of the tumor model.

Results in Fig. 5 can also be interpreted in terms of the diagnostic task. In subfigure C , results suggest a “classical” low-grade tumor, well in agreement with the expected parameters for low-grade patients (Fig. 3, left). In subfigure D , we observe a more infiltrative and rapidly growing tumor. This is in good agreement with the clinical observation of the tumor expanding several centimeters per year and showing a star-shaped pattern, often associated with poor prognosis. The data in D had already been evaluated in [7] and the authors had estimated the speed of the tumor front shown as the red pink diagonal line in D . We observe a good agreement of our results with this earlier estimate, but obtain the additional information that the expansion of the tumor is not due to a fast




Fig. 5. MCMC sampling results in the space spanned by model parameters D and ρ , for the four experiments. Green samples are obtained from the sparse grid interpolation Eq. (12), blue-purple samples come from the direct sampling in Eq. (10). Black circles indicate means of the two distributions. Ground truth for A and B are indicated by the pink cross. In D the previously estimated speed of growth [7] is shown by the pink line. The sparse grid sampling approximation performs better than the direct MCMC (A - B). Estimates correlate well with results from [7], but provide a more accurate characterization of the state of disease (D).

infiltration, but a rapid doubling of tumor cells with a direr prognosis for this particular patient.

5 Conclusions

The proposed tumor modeling approach links the physical process model with the statistical image observation model and enables efficient estimation of model parameters in a Bayesian framework, potentially also to be used with model components at different scales and with different sources of information. Preliminary results are promising. The approach can readily be adapted to incorporate information from other functional imaging modalities, and to model mass effect and tissue translations, or tissue alterations resulting from therapy.

Acknowledgements. This work was supported by the German Academy of Sciences Leopoldina (Fellowship Programme LPDS 2009-10), the Academy of Finland (133611), INRIA CompuTumor, NIH NIBIB NIMIC U54-EB005149, NIH NCCR NAC P41-RR13218, NIH NINDS R01-NS051826, NIH R01-NS052585, NIH R01-EB006758, NIH R01-EB009051, NIH P41-RR014075 and the NSF CAREER Award 0642971.

References

1. Chaplain, M.A.J., Stuart, A.M.: A mathematical model for the diffusion of tumour angiogenesis factor into the surrounding host tissue. *J Math Appl Med Biol* **8** (1991) 191–220
2. Swanson, K.R., Alvord, E.C., Murray, J.D.: A quantitative model for differential motility of gliomas in grey and white matter. *Cell Prolif* **33** (2000) 317–329
3. Clatz, O., Sermesant, M., Bondiau, P.Y., Delingette, H., Warfield, S.K., Malandain, G., Ayache, N.: Realistic simulation of the 3-D growth of brain tumors in MR images coupling diffusion with biomechanical deformation. *IEEE TMI* **24** (2005) 1334–1346
4. Alvord, E.C., Swanson, K.R.: Using mathematical modeling to predict survival of low-grade gliomas. *Ann Neurol* **61** (2007) 496–497

5. Hoge, C., Davatzikos, C., Biros, G.: An image-driven parameter estimation problem for a reaction-diffusion glioma growth model with mass effects. *J Math Biol* **56** (2008) 793–825
6. Prastawa, M., Bullitt, E., Gerig, G.: Simulation of brain tumors in MR images for evaluation of segmentation efficacy. *MedIA* **13** (2009) 297–311
7. Konukoglu, E., Clatz, O., Menze, B.H., Weber, M.A., Stieltjes, B., Mandonnet, E., Delingette, H., Ayache, N.: Image guided personalization of reaction-diffusion type tumor growth models using modified anisotropic Eikonal equations. *IEEE TMI* **29** (2010) 77–95
8. Gooya, A., Biros, G., Davatzikos, C.: Deformable registration of glioma images using EM algorithm and diffusion reaction modeling. *IEEE TMI* **30** (2011) 375–390
9. Kyriacou, S.K., Davatzikos, C., Zinreich, S.J., Bryan, R.N.: Nonlinear elastic registration of brain images with tumor pathology using a biomechanical model. *IEEE TMI* **18** (1999) 580–592
10. Zacharaki, E.I., Hoge, C.S., Shen, D., Biros, G., Davatzikos, C.: Non-diffeomorphic registration of brain tumor images by simulating tissue loss and tumor growth. *Neuroimage* **46** (2009) 762–774
11. Hamamci, A., Unal, G., Kucuk, N., Engin, K.: Cellular automata segmentation of brain tumors on post contrast MR images. In: *Proc MICCAI*. (2010) 137–146
12. Corso, J.J., Sharon, E., Dube, S., El-Saden, S., Sinha, U., Yuille, A.: Efficient multilevel brain tumor segmentation with integrated Bayesian model classification. *IEEE TMI* **9** (2008) 629–40
13. Mohamed, A., Zacharakib, E.I., Shena, D., Davatzikos, C.: Deformable registration of brain tumor images via a statistical model of tumor-induced deformation. *MedIA* **10** (2006) 752–763
14. Cobzas, D., Mosayebi, P., Murtha, A., Jagersand, M.: Tumor invasion margin on the Riemannian space of brain fibers. *Proc MICCAI* **12** (2009) 531–539
15. Ganslandt, O., Stadlbauer, A., Fahlbusch, R., Kamada, K., Buslei, R., Blumcke, I., Moser, E., Nimsky, C.: ¹H-MRSI integrated into image-guided surgery: correlation to standard MR imaging and tumor cell density. *Neurosurg* **56** (2005) 291–298
16. Pallud, J., Mandonnet, E., Duffau, H., Galanaud, D., Taillandier, L., Capelle, L.: Prognostic value of initial magnetic resonance imaging growth rates for World Health Organization grade II gliomas. *Ann Neurol* **60** (2006) 380–383
17. Bungartz, H.J., Griebel, M.: Sparse grids. *Acta Numerica* **13** (2004) 147–269
18. Klimke, A., Wohlmuth, B.: Piecewise multilinear hierarchical sparse grid interpolation in MATLAB. *ACM Trans Math Software* **31** (2005) 1
19. Ma, X., Zabaras, N.: An efficient Bayesian inference approach to inverse problems based on an adaptive sparse grid collocation method. *Inverse Problems* **25** (2009) 035013 (27pp)
20. Zarabas, N.: Solving stochastic inverse problems: A sparse grid collocation approach. In Biegler, L., ed.: *Large-scale inverse problems and quantification of uncertainty*. Wiley, Chichester, UK (2011)
21. Haario, H., Laine, M., Mira, A., Saksman, E.: DRAM: Efficient adaptive MCMC. *Statistics and Computing* **16** (2006) 339–354
22. McCorquodale, P., Colella, P., Johansen, H.: A Cartesian grid embedded boundary method for the heat equation in irregular domains. *J. Comp. Phys.* **173**(2) (2001) 620–635