

Bayesian Personalization of Brain Tumor Growth Model

Matthieu Lê¹, Hervé Delingette¹, Jayashree Kalpathy-Cramer²,
Elizabeth R. Gerstner³, Tracy Batchelor³,
Jan Unkelbach⁴, and Nicholas Ayache¹

¹ Asclepios Project, INRIA Sophia Antipolis, France

² Martinos Center for Biomedical Imaging, Harvard–MIT Division of Health Sciences and Technology, Charlestown, MA, USA

³ Department of Neurology, Massachusetts General Hospital, Boston, Massachusetts, USA

⁴ Department of Radiation Oncology, Massachusetts General Hospital and Harvard Medical School, Boston, MA, USA

Abstract. Recent work on brain tumor growth modeling for glioblastoma using reaction-diffusion equations suggests that the diffusion coefficient and the proliferation rate can be related to clinically relevant information. However, estimating these parameters is difficult due to the lack of identifiability of the parameters, the uncertainty in the tumor segmentations, and the model approximation, which cannot perfectly capture the dynamics of the tumor. Therefore, we propose a method for conducting the Bayesian personalization of the tumor growth model parameters. Our approach estimates the posterior probability of the parameters, and allows the analysis of the parameters correlations and uncertainty. Moreover, this method provides a way to compute the evidence of a model, which is a mathematically sound way of assessing the validity of different model hypotheses. Our approach is based on a highly parallelized implementation of the reaction-diffusion equation, and the Gaussian Process Hamiltonian Monte Carlo (GPHMC), a high acceptance rate Monte Carlo technique. We demonstrate our method on synthetic data, and four glioblastoma patients. This promising approach shows that the infiltration is better captured by the model compared to the speed of growth.

1 Introduction

Glioblastomas (GBM) are one of the most common types of primary brain tumors. In addition to being infiltrative by nature, they are also the most aggressive of brain tumors. Reaction-diffusion equations have been widely used to model the GBM growth, using a diffusion and a logistic proliferation term,

$$\frac{\partial u}{\partial t} = \underbrace{\nabla(D \cdot \nabla u)}_{\text{Diffusion}} + \underbrace{\rho u(1 - u)}_{\text{Logistic Proliferation}} \quad (1)$$

$$D \nabla u \cdot \vec{n} \Big|_{\partial \Omega} = 0 \quad (2)$$




Fig. 1. For the two time points: a) T1Gd with the abnormality segmented in orange. b) T2-FLAIR with the abnormality segmented in red. c) Tumor cell density using the maximum a posteriori parameters with the 80% and 16% threshold outlined in black. d) Comparison between the clinician segmentations and the model segmentations.

Equation (1) describes the spatio-temporal evolution of the tumor cell density u , which infiltrates neighboring tissues with a diffusion coefficient D , and proliferates with a net proliferation rate ρ . Equation (2) enforces Neumann boundary conditions on the brain domain Ω . Equation (1) admits solutions which asymptotically behave like traveling waves with speed $v = 2\sqrt{D\rho}$ and infiltration length $\lambda = \sqrt{D/\rho}$ [1]. The infiltration length is related to the spatial rate of decay of the density u . Preliminary works suggest that D and ρ can help in assessing the response to therapy [2].

The personalization of D and ρ is based on serial acquisitions of Magnetic Resonance Images (MRIs). The tumor causes abnormalities on T1 post Gadolinium (T1Gd) and T2-FLAIR MRIs, which can be monitored over time. To relate the tumor cell density u to those images, it is usually assumed that the visible borders of these abnormalities correspond to a threshold of the tumor cell density [1]. Since the T2-FLAIR abnormality is larger than the T1Gd abnormality, the T2-FLAIR abnormality borders correspond to a lower threshold of tumor cell density (Fig. 1). Therefore, given MRIs of a patient at 2 time points, it is possible to extract the segmentations of the 4 visible abnormalities. Those can be used to personalize a tumor growth model, i.e. to find the set of parameters $\theta = (D, \rho)$ which best fits those segmentations.

This problem was studied by Harpold *et al.* [1] who described a heuristic to relate the volumes of the four available segmentations to the speed $v = 2\sqrt{D\rho}$ and infiltration length $\lambda = \sqrt{D/\rho}$, based on the asymptotic behavior of the solutions of equation (1), hence allowing for the personalization of D and ρ . Konukoglu *et al.* [3] also elaborated a personalization strategy based on BOBYQA, a state-of-the-art derivative-free optimization algorithm. However, this latter method only considers one MRI modality per time point, and fails to simultaneously identify D and ρ . These estimation procedures provide single-solutions which are sensitive to the initialization, and do not provide any understanding about potential correlation between parameters.

To circumvent these drawbacks, we propose a method to estimate the posterior probability of $\theta = (D, \rho)$ from the T1Gd and T2-FLAIR abnormalities at 2 time




Fig. 2. a) Two synthetic time points with the simulated T1Gd and T2-FLAIR abnormalities (orange and red resp.). b) Gaussian process interpolation of E_{pot} using the evaluation of the posterior at well chosen points (in red). c) Kernel density estimate of the posterior using 1000 samples. Results of the BOBYQA optimization with multiple initializations in blue, true parameters in red.

points. We show how this method provides information on the confidence in the estimation, and provides a global understanding of the correlation between the parameters. We further demonstrate how the knowledge of the posterior distribution can be used to compare different model hypotheses. The method is based on the computation of the statistical evidence: the probability of observing the data given the model. The Bayes factor, defined as the ratio of the evidence of two different models, allows one to compare the models and select which better explains the data. This idea is applied to determine which tumor cell density threshold is the most adequate for the T2-FLAIR abnormality.

Naive Monte Carlo would be impractical due to computational overload. Some strategies rely on approximations of the posterior probability with sparse grids [4], approximations of the forward model based on reduced order models [5,6], or on expensive grid methods [7]. We present here a method based on the highly parallelized implementation of the reaction-diffusion equation using the Lattice Boltzmann Method (LBM), and the Gaussian Process Hamiltonian Monte Carlo (GPHMC), a high acceptance rate Monte Carlo technique.

2 Method

2.1 Tumor Growth Model

We consider the model defined in equations (1,2). Following previous work [1], the diffusion tensor is defined as $D = d_w \mathbb{I}$ in the white matter and $D = d_w/10 \mathbb{I}$ in the gray matter, where \mathbb{I} is the 3x3 identity matrix. We further identify the parameter D with d_w . As such, the diffusion is heterogeneous and locally isotropic. This model reproduces the infiltrative nature of the GBM, takes into account anatomical barriers (ventricles, sulci, falx cerebri), and the tumor's preferential progression along white matter tracts such as the corpus callosum. The reaction-diffusion equation is implemented using the Lattice Boltzmann Method (LBM) [8]. The LBM are parallelized such that simulating 30 days of growth, with a time step of 0.1 day, takes approximately 80 seconds on a 60 core machine for a 1mm

isotropic brain MRI. The initialization of the tumor cell density $u(t = t_1, \mathbf{x})$ at the time of the first acquisition is of particular importance, as it impacts the rest of the simulation. In this work, the tumor tail extrapolation algorithm described in [9] is used. The tumor cell density is computed outward (and inward) of the T1Gd abnormality borders as a static approximation of the wave-like solution of equation (1) with parameter θ . Consequently, the T1Gd abnormality falls exactly on the threshold of the tumor cell density at the first time point.

2.2 GPHMC

Given parameter θ , the model is initialized by extrapolating the tumor cell density from the T1Gd segmentation at time t_1 . Its dynamics is then simulated using the LBM until time t_2 . The thresholded tumor cell density is then compared to the segmentations S of the T1Gd at time t_2 and the T2-FLAIR at time t_1 and t_2 (Fig. 1). To cast the problem in a probabilistic framework, we follow Bayes rule: $P(\theta|S) \propto P(S|\theta) P(\theta)$. The likelihood is modeled as $P(S|\theta) \propto \exp(-d(S, u)^2/\sigma^2)$, where the distance $d(S, u)$ is the mean symmetric Hausdorff distance between the border of the segmentations S and the isolines of the simulated tumor cell density u . The noise level is quantified through σ . Samples are drawn from the posterior probability using the Gaussian Process Hamiltonian Monte Carlo [10].

Hamiltonian Monte Carlo (HMC). HMC is a Markov Chain Monte Carlo algorithm which uses a refined proposal density function based on the Hamiltonian dynamics [10]. The problem is augmented with a momentum variable $\mathbf{p} \sim \mathcal{N}(\mathbf{0}, \mathbb{I})$. By randomly sampling \mathbf{p} , we define the current state (θ, \mathbf{p}) . HMC is designed to propose a new state (θ^*, \mathbf{p}^*) with constant energy $H = E_{\text{pot}} + E_{\text{kin}}$, with potential energy $E_{\text{pot}} = -\log(P(S|\theta)P(\theta))$, and kinetic energy $E_{\text{kin}} = 1/2 \|\mathbf{p}\|_2^2$. The new state is the result of the Hamiltonian dynamics: $d\theta_i/dt = \partial H/\partial \mathbf{p}_i$ and $d\mathbf{p}_i/dt = -\partial H/\partial \theta_i$. The new state (θ^*, \mathbf{p}^*) is accepted with probability $\mathcal{A} = \min[1, \exp(-H(\theta^*, \mathbf{p}^*) + H(\theta, \mathbf{p}))]$. The conservation of the energy during the Hamiltonian dynamics, up to the numerical discretization approximation, insures a high acceptance rate \mathcal{A} .

Gaussian Process Hamiltonian Monte Carlo (GPHMC). In the HMC, computing the Hamiltonian dynamics requires a significant amount of model evaluations. To circumvent this difficulty, E_{pot} is approximated with a Gaussian process [10] (Fig. 2). During an initialization phase, E_{pot} is evaluated by running the forward model at random locations. Those evaluation points are used to define the Gaussian process approximating E_{pot} - details in [10]. HMC is then run using the Gaussian process interpolation of E_{pot} to compute the Hamiltonian dynamics. Given that the Gaussian process well captures E_{pot} , the GPHMC benefits from the high acceptance rate of the HMC.

Posterior Density. $P(\theta|S, \mathcal{M})$ is computed using a kernel density estimation on the drawn samples $\{\theta_i\}_{i=1}^N$. A bounce-back boundary condition is applied to ensure the density integrates to 1 on the bounded parameter space.

2.3 Bayes Factor: Extension of the Chib's Method

The evidence of the model \mathcal{M} can be expressed using Chib's method [11]

$$P(S|\mathcal{M}) = \frac{P(S|\boldsymbol{\theta}, \mathcal{M})P(\boldsymbol{\theta}|\mathcal{M})}{P(\boldsymbol{\theta}|S, \mathcal{M})} \propto \exp\left(-\frac{d(S, Y_{\boldsymbol{\theta}})^2}{\sigma^2}\right) \frac{P(\boldsymbol{\theta}|\mathcal{M})}{P(\boldsymbol{\theta}|S, \mathcal{M})} \quad (3)$$

In other words, the probability of observing the segmentation given the model (evidence) is equal to the likelihood times the prior, normalized by the posterior. The Chib's method evaluates this equation on a sample $\boldsymbol{\theta}$ of high posterior probability, like its mode. We propose instead a more robust estimation of the evidence, as an extension of the Chib's method, by averaging it over the samples $\{\boldsymbol{\theta}_i\}_{i=1}^N$ drawn from the posterior distribution,

$$P(S|\mathcal{M}) = \mathbb{E}_{\boldsymbol{\theta}|S, \mathcal{M}} [P(S|\mathcal{M})] \propto \frac{1}{N} \sum_{i=1}^N \exp\left(-\frac{d(S, Y_{\boldsymbol{\theta}_i})^2}{\sigma^2}\right) \frac{P(\boldsymbol{\theta}_i|\mathcal{M})}{P(\boldsymbol{\theta}_i|S, \mathcal{M})} \quad (4)$$

This estimator is less prone to errors when the kernel density estimation of the posterior approximately captures the real density at the modes. In order to compare two models \mathcal{M}_1 and \mathcal{M}_2 , we can compute the Bayes factor $B = (P(S|\mathcal{M}_1)P(\mathcal{M}_1))/(P(S|\mathcal{M}_2)P(\mathcal{M}_2))$. If B is substantially larger than 1, it is evidence that \mathcal{M}_1 is better suited to explain the data than \mathcal{M}_2 . This is a Bayesian alternative to the frequentist hypotheses testing approach.

3 Experiments and Results

We present results on synthetic and real data. The tumor cell density visibility thresholds are assumed to be 80% for the T1Gd abnormality and 16% for the T2-FLAIR abnormality [12]. The parameters are constrained such that $D \in [0.02, 1.5] \text{ mm}^2/\text{days}$, and $\rho \in [0.002, 0.2] \text{ days}^{-1}$ [1]. The prior $P(\boldsymbol{\theta})$ is assumed uniform within this bounded box. After trials, the noise level has been manually tuned to $\sigma = 10 \text{ mm}$. Our results are compared to the direct minimization of $d(S, Y)$ using BOBYQA. BOBYQA is run 9 times with 9 different initializations in the parameter space - around 20 iterations per BOBYQA were observed. Finally, we compare two models using two different tumor cell density visibility thresholds for the T2-FLAIR abnormality frontier: 16% for the first model \mathcal{M}_1 , and 2% for the second model \mathcal{M}_2 , following different studies [1]. The Bayes factor between \mathcal{M}_1 and \mathcal{M}_2 is computed to compare the models.

Synthetic Data. Starting from a manually drawn T1Gd abnormality segmentation, the growth of a tumor is simulated on an atlas of white and gray matter during 30 days with $D = 1.0 \text{ mm}^2.\text{days}^{-1}$ and $\rho = 0.18 \text{ days}^{-1}$. The first and last images are thresholded in order to define the 4 segmentations (Fig. 2 left). 1000 samples are drawn with the GPHMC, using a GP interpolation of E_{pot} with

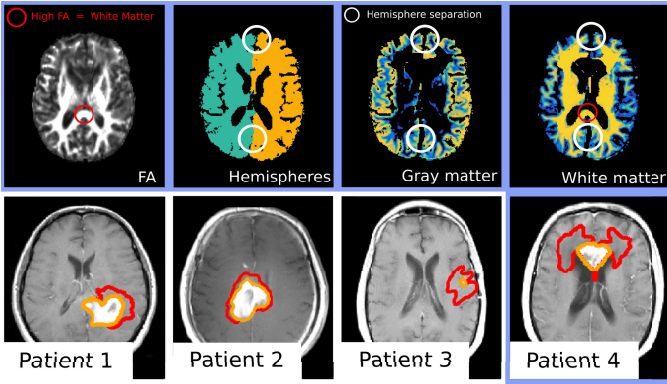


Fig. 3. Top: Data processing using the Fractional Anisotropy (FA) and the hemisphere segmentation to ensure the separation of the hemispheres in the white and gray matter. Bottom: T1Gd of the 4 real cases for the first time point. The T1Gd (resp. T2-FLAIR) abnormality is outlined in orange (resp. red).

35 points (Fig. 2 middle). The acceptance rate is 90%, demonstrating that the Gaussian process successfully captured E_{pot} . The best BOBYQA solutions are in very good agreement with the real parameters (Fig. 2 right). The existence of BOBYQA solutions somewhat distant from the true solution, demonstrates the need for multiple initializations of the algorithm. The true parameter is outlined by high posterior density isocontour, and the posterior reflects little correlation between the parameters (unimodal and somewhat peaked posterior). Finally, with equation (4), we computed a Bayes factor $B = P(S|\mathcal{M}_1)/P(S|\mathcal{M}_2) = 3$, indicating that \mathcal{M}_1 is better suited than \mathcal{M}_2 in the synthetic case.

GBM Cases. The method is applied to the 4 GBM patients shown in Fig. 3 (bottom panel). The T1Gd and T2-FLAIR abnormalities are segmented by a clinician. For each time point, the T2-FLAIR is rigidly registered to the T1Gd. The rigidly registered T2-FLAIR at time t_2 is then non-linearly registered to the T2-FLAIR at time t_1 . All the segmentations are then mapped to the T1Gd image space at time t_1 where the white matter, gray matter and cerebrospinal fluid (CSF) are segmented (Fig. 3 top panel). Close attention is paid to the separation of the hemispheres to prevent tumor cells from invading the contralateral hemisphere. Voxels at the boundary of the hemispheres that show low Fractional Anisotropy value are tagged as CSF to ensure the separation of the hemispheres, while keeping the corpus calosum as white matter (Fig. 3).

The mean acceptance rate for the 4 patients is 74% (Fig. 1) which shows that the Gaussian process successfully interpolate E_{pot} . The BOBYQA solutions, although reasonable, fail to describe the spatial dynamics of the posterior (Fig. 4). This emphasizes the need for a global understanding of the posterior’s behavior. We see that the model can capture the infiltration length of the tumor growth: the posterior is elongated along the line of constant infiltration $\lambda = \sqrt{D/\rho}$, es-




Fig. 4. Posterior estimation for the 4 real cases. The blue dots are the three best BOBYQA solutions. The red dot is the maximum a posteriori of the drawn samples θ_{MAP} . The dashed black line is the line of constant infiltration $\rho = D/\lambda^2$.

Table 1. Presentation of the patients.

	$t_2 - t_1$ (days)	Acc. rate (%)	Bayes factor
Patient 1	105	90	2.8
Patient 2	29	80	1.8
Patient 3	26	51	0.49
Patient 4	29	75	1.2

pecially for patients 1, 2 and 3 (see the black dashed line on Fig. 4). A small infiltration λ corresponds to a T2-FLAIR abnormality close to the T1Gd abnormality (patient 1 Fig. 3) while a larger infiltration corresponds to more distant abnormalities (patient 4 Fig. 3). However, the model less successfully captures the speed of the growth, as the posterior does not present a clear mode on the constant infiltration line. For patient 3, the presence of two modes suggests that the speed of growth might be explained by two sets of parameters. Moreover, the presence of two modes is responsible for the relatively low acceptance rate of 51%. The posterior distribution is more complex and therefore harder to interpolate with the Gaussian process. For patient 4, the posterior is more peaked and suggests that both the infiltration and speed can be explained by the model. The simulation corresponding to the sample θ of lowest potential energy E_{pot} is presented on Fig. 1 (right). We note $\theta_{\text{MAP}} = (D_{\text{MAP}}, \rho_{\text{MAP}})$ the drawn sample corresponding to the lowest potential energy E_{pot} . It is interesting to note that θ_{MAP} can be slightly skewed on the edge of high posterior isocontour (patient 3 or 4 on Fig. 1). This is caused by the very high sensitivity of the output to the parameters' value, as well as the likelihood model which allows for a certain noise level σ . The Bayes factor between \mathcal{M}_1 and \mathcal{M}_2 was computed for the different patients (Fig. 1). We can see that for patients 1, 2, and 4, the Bayes factor leans toward validating the model \mathcal{M}_1 with a mean of 1.9. Patient 3 presents a lower Bayes factor supporting model \mathcal{M}_2 , but this patient presents an odd posterior, which might indicate that the reaction-diffusion model less successfully captures the dynamics of the growth.

4 Conclusion

We presented a Bayesian personalization of the parameters of a tumor growth model on 4 patients. The computation of the posterior provides far more information than the direct optimization, and takes only five times longer. Our method provides a global understanding of the correlation between the parameters, and we showed that the model captures the infiltration way better than the speed of growth of the tumor. Furthermore, it provides a way to compute the Bayes factor and compare the validity of different model hypotheses, which showed that a threshold of 16% was probably better suited for the data. In the future, we want to apply this methodology to larger dimensional problems. Namely, the noise level σ could be included as an unknown, and depend on the image modality. Finally, we believe that such a method could be applied to personalized treatment planning.

Acknowledgments. Part of this work was funded by the European Research Council through the ERC Advanced Grant MedYMA 2011-291080.

References

1. Harpold, H.L., Alvord Jr., E.C., Swanson, K.R.: The evolution of mathematical modeling of glioma proliferation and invasion. *J. of Neurop. & Exp. Neurol.* (2007)
2. Neal, M.L., Trister, A.D., Ahn, S., Baldock, A., et al.: Response classification based on a minimal model of GBM growth is prognostic for clinical outcomes and distinguishes progression from pseudoprogression. *Cancer research* (2013)
3. Konukoglu, E., Clatz, O., Menze, B.H., Stieltjes, B., Weber, M.A., Mandonnet, E., et al.: Image guided personalization of reaction-diffusion type tumor growth models using modified anisotropic eikonal equations. *MedIA* 29(1), 77–95 (2010)
4. Menze, B.H., Van Leemput, K., Honkela, A., Konukoglu, E., Weber, M.-A., Ayache, N., Golland, P.: A generative approach for image-based modeling of tumor growth. In: Székely, G., Hahn, H.K. (eds.) *IPMI 2011. LNCS*, vol. 6801, pp. 735–747. Springer, Heidelberg (2011)
5. Konukoglu, E., Relan, J., Cilingir, U., Menze, B.H., Chinchapatnam, P., Jadidi, A., et al.: Efficient probabilistic model personalization integrating uncertainty on data and parameters: Application to eikonal-diffusion models in cardiac electrophysiology. *Progress in Biophysics and Molecular Biology* 107(1), 134–146 (2011)
6. Neumann, D., et al.: Robust image-based estimation of cardiac tissue parameters and their uncertainty from noisy data. In: Golland, P., Hata, N., Barillot, C., Hornegger, J., Howe, R. (eds.) *MICCAI 2014, Part II. LNCS*, vol. 8674, pp. 9–16. Springer, Heidelberg (2014)
7. Wallman, M., Smith, N.P., Rodriguez, B.: Computational methods to reduce uncertainty in the estimation of cardiac conduction properties from electroanatomical recordings. *MedIA* 18(1), 228–240 (2014)
8. Yoshida, H., Nagaoka, M.: Multiple-Relaxation-Time LBM for the convection and anisotropic diffusion equation. *Journal of Computational Physics* 229(20) (2010)

9. Konukoglu, E., Clatz, O., Bondiau, P.Y., Delingette, H., Ayache, N.: Extrapolating glioma invasion margin in brain magnetic resonance images: Suggesting new irradiation margins. *MedIA* 14(2), 111–125 (2010)
10. Rasmussen, C.E.: Gaussian processes to speed up hybrid monte carlo for expensive bayesian integrals. In: *Bayesian Statistics*, vol. 7 (2003)
11. Chib, S.: Marginal likelihood from the Gibbs output. *Journal of the American Statistical Association* 90(432), 1313–1321 (1995)
12. Swanson, K., Rostomily, R., Alvord, E.: A mathematical modelling tool for predicting survival of individual patients following resection of glioblastoma: a proof of principle. *British Journal of Cancer* 98(1), 113–119 (2008)