

A Classification-based Glioma Diffusion Model using MRI Data

Marianne Morris, Russell Greiner, Jörg Sander, Albert Murtha¹, Mark Schmidt

Department of Computing Science, University of Alberta,
Edmonton, AB, T6G 2E8, Canada

{*marianne, greiner, joerg, schmidt*}@cs.ualberta.ca

¹ Department of Radiation Oncology, Cross Cancer Institute,
11560 - University Ave, Edmonton, AB, T6G 1Z2, Canada

Abstract. Gliomas are diffuse, invasive brain tumors. We propose a 3D classification-based diffusion model, CDM, that predicts how a glioma will grow at a voxel-level, on the basis of features specific to the patient, properties of the tumor, and attributes of that voxel. We use Supervised Learning algorithms to learn this general model, by observing the growth patterns of gliomas from other patients. Our empirical results on clinical data demonstrate that our learned CDM model can, in most cases, predict glioma growth more effectively than two standard models: uniform radial growth across all tissue types, and another that assumes faster diffusion in white matter.

Keywords: machine learning, brain tumors, glioma, diffusion models, prediction

1 Introduction

Primary brain tumors originate from a single glial cell in the nervous system, and grow by invading adjacent cells, often leading to a life-threatening condition. Proper treatment requires knowing both where the tumor mass is, and also where the occult cancer cells have infiltrated in nearby healthy tissue. Some conventional treatments implicitly assume the tumor will grow radially in all directions — *e.g.*, the standard practice in conformal radiotherapy involves irradiating a volume that includes both the observed tumor, and a uniform 2cm margin around this border [7, 8]. Swanson’s model [17] claims the tumor growth rate is 5 times faster in white matter, versus grey matter. Our empirical evidence, however, shows that neither model is particularly accurate.

We present an alternative approach to modeling tumor growth: use data from a set of patients to *learn* the parameters of a diffusion model. In particular, given properties of the patient, tumor and each voxel (based on MRI scans; see Fig. 1[a–g]) at one time, our CDM system predicts the tumor region at a later time (Fig. 1[h]). This model can help define specific treatment boundaries that would replace the uniform, conventional 2cm margin. It can also help find regions where radiologically occult cancer cells concentrate but do not sufficiently enhance on




Fig. 1. Axial slices of brain tumor patient: (a) T1-weighted scan; (b) T1-weighted scan using gadolinium contrast; (c) T2-weighted scan; (d) white matter; (e) grey matter; (f) CSF — cerebrospinal fluid; (g) segmented patient tumor; (h) predicted patient tumor, after adding 30,000 voxels in 3D, overlaid on T1c (light grey represents the true positives, dark grey the false positives and black the false negatives).

the MRI scan. Therefore, the model can help define the appropriate radiation doses to deliver to the relevant regions adjacent to the visible tumor.

Section 2 reviews standard glioma diffusion models, and Section 3 formally defines the diffusion models we are considering. Finally, Section 4 describes our experiments testing CDM, comparing it with two other models: naïve uniform growth and tissue-based diffusion. Additional details are in [1, 2, 13].

2 Related Work

In recent decades, glioma growth modeling has offered important contributions to cancer research, shedding light on tumor growth behaviour and helping improve treatment methods. In this section, we describe two types of tumor modeling: volumetric at the macroscopic level, and models based on white matter invasion.

2.1 Macroscopic and Volumetric Modeling

Mathematical modeling of gliomas at the macroscopic level has represented the traditional framework in predicting glioma diffusion, using growth and proliferation parameters. We review three of these models:

Kansal *et al.* [10] simulate the *gompertzian growth*, which views the tumor as a population of cells and the growth as a dynamic process where proliferating and inactive classes of cells interact. Kansal *et al.* use cellular automata to model the different states of tumor cells, from dividing cells at the periphery, to non-proliferating, and finally to the necrotic state at the centre of the tumor. This model is designed to predict the growth of glioblastoma multiforme (GBM), the most aggressive, grade IV gliomas. The model does not account for various tumor grades, brain anatomy, nor the infiltrating action of cancer cells in tissue near the tumor.

Tabatabai *et al.* [20] simulate asymmetric growth as in real tumors and accommodate the concept of increasing versus decreasing tumor radii (due to treatment effects), but do not account for various clinical factors involved in malignant diffusion. Instead, their model describes tumors as self-limited systems, not incorporating the interactions between healthy and cancer cells at the tumor border and the competition of cells inside the tumor. This is not a realistic representation of clinical cancer diffusion.

Zizzari's model [22] describes the proliferation of GBMs using tensor product splines and differential equations, the solutions of which give the distribution of tumor cells with respect to their spatio-temporal coordinates. Zizzari extends his growth model to introduce a treatment planning tool that incorporates a supervised learning task. However, his growth predictions are based only on geometric issues, and do not consider biological factors nor patient information.

2.2 Glioma Modeling based on White Matter Invasion

The trend in glioma research is to study biological and clinical factors involved in cancer diffusion through healthy tissue. Recent models provide a more promising direction, which can also help provide more effective treatment. In this section, we review models that incorporate the heterogeneity of brain tissue and histology of cancer cells.

Swanson *et al.* [17] develop a model based on the differential motility of glioma cells in white versus grey matter, suggesting that the diffusion coefficient in white matter is 5 times that in grey matter. This model was extended to simulate virtual gliomas [19] and to assess the effectiveness of chemotherapy delivered to different tissue types in the brain [18]. This modeling is different from our CDM system as we do not *a priori* assume the cancer diffusion rates in different tissue types, but rather our system can *learn* glioma diffusion behaviour from clinical data.

Price *et al.* [15] use T2-weighted scans and Diffusion Tensor Imaging (DTI) to determine whether DTI can identify abnormalities on T2 scans. Regions of interest particularly include white matter adjacent to the tumor, and areas of

abnormality on DTI that appeared normal on T2 images. Results demonstrated further glioma invasion of white matter tracts near the observed tumor.

Clatz *et al.* [4] propose a model that simulates the growth of GBM based on an anatomical atlas that includes white fibre diffusion tensor information. The model is initialized with a tumor detected on the MRI scan of a patient, and results are evaluated against the tumor observed six months later. However, model results are reported for only one patient, leaving in question how it performs on a variety of patients, and with various tumor types.

2.3 Discussion

Each of the glioma diffusion models presented above describes the geometrical growth of gliomas as evolving objects. Few of these models use the biological complexity of cancerous tumors, the heterogeneity of the human brain anatomy, or the clinical factors of malignant invasion. Moreover, none of these earlier systems attempts to *learn* general growth patterns from existing data, nor are they capable of predicting growth of various tumor grades (as opposed to methods specifically designed to predict GBM growth only).

The literature does suggest that the following factors should help us predict how the tumor will spread — *i.e.*, whether the tumor is likely to infiltrate to a new voxel:

- *Anatomical features of the brain*: regions that represent pathways versus brain structures that act as a boundary to the spreading action of the malignant cells.
- *Properties of the tumor*: the grade of the tumor (as high-grade gliomas grow much faster than low-grade ones); the location of the tumor within the brain (as the shape of the tumor depends on surrounding anatomical structures).
- *Properties of the voxels* (at the periphery of the tumor where there can be interaction between malignant and normal cells): its tissue type — grey versus white matter; whether it currently contains edema¹.

We incorporate these diffusion factors as learning features into our ‘general’ diffusion model, CDM. The remainder of this paper describes the diffusion models we implemented, presents the experiments, and evaluates the performance of the three models given our dataset of MRI scans.

3 Diffusion Models

In general, a diffusion model (Fig. 2) takes as input an image whose voxels are each labeled with: the current “voxel label”, \mathbf{VL} , which is “1” if that voxel is *currently* a tumor and “0” otherwise (see Fig. 1[g])² as well as general information

¹ Swelling due to accumulation of excess fluid.

² Here, expert radiologists have manually delineated the “enhancing regions” of tumors based on their MRI scans. Note this does not include edema, nor any other labels. We then spatially interpolate each patient image to fill inter-slice gaps and to obtain voxels of size 8mm³.

```

1. Diffusion( VoxelLabel: VL; GeneralInfo: e; int: s )
   % VL[i, j, k]=1 if position <i, j, k> is a tumor
   % Initially VL corresponds to current tumor
   % When algorithm terminates, VL will correspond to tumor containing "s" additional voxels
2.   total_count := 0
3.   Do forever:
4.     Compute  $N := \left\{ \langle x, y, z \rangle \mid \text{VL}[x, y, z] = 0 \ \& \ \begin{pmatrix} \text{VL}[x + 1, y, z] = 1 \ \vee \\ \text{VL}[x - 1, y, z] = 1 \ \vee \\ \vdots \\ \text{VL}[x, y, z - 1] = 1 \end{pmatrix} \right\}$ 
5.     For each location  $v_i \in N$ 
6.       Determine if  $v_i$  becomes a tumor
7.       If so,
8.         Set VL[ $v_i$ ] := 1
9.         total_count++;
10.      If (total_count == s), return

```

Fig. 2. Generic Diffusion Model

$\mathbf{e} = \mathbf{e}_{Patient} \cup \mathbf{e}_{Tumor} \cup \{\mathbf{e}_i\}_i$ about the patient $\mathbf{e}_{Patient}$, the tumor \mathbf{e}_{Tumor} and the individual voxels \mathbf{e}_i (see Section 3.1). The third input is an integer s that tells the diffusion model how many additional voxels to include. See line 1 of Fig. 2. The output is the prediction of the next s additional voxels that will be incorporated into the tumor, represented as a bit-map over the image. For example, if the tumor is currently 1000 voxels and the doctor needs to know where the tumor will be, when it is 20% larger — *i.e.*, when it is 1200 voxels — he would set $s = 200$.

A diffusion model first identifies the set of voxels N just outside the border of the initial tumor; see line 4 of Fig. 2. In the following diagram



(where each **X** cell is currently a tumor), N would consist of the voxels labeled v1 through v5, but not v6 nor v7 (as we are not considering diagonal neighbors). In the 3D case, each voxel will have 6 neighbors.

The diffusion model then iterates through these candidate voxels, $v_i \in N$. If it decides that one has become a tumor, it then updates VL (which implicitly updates the tumor/healthy border) and increments the total number of “transformed voxels”; see lines 5–9 of Fig. 2. After processing all of these neighbors (in parallel), it will then continue transforming the neighbors of this newly enlarged boundary. If a voxel is not transformed on one iteration, it remains eligible to be transformed on the next iteration. When the number of transformed voxels matches the total s , the algorithm terminates, returning the updated VL assignment (Fig. 2, line 10).

The various diffusion models differ only in how they determine if v_i has become tumor — line 6 of Fig. 2. The *uniform growth* model, UG, simply includes every “legal” voxel it finds (where a voxel is legal if it is part of the brain, as opposed to skull, eye, etc.). The *tissue-based* model, GW, assumes the growth rate for white matter is 5 times faster than for grey matter [17], and 10 times faster than other brain tissue. Here, whenever a neighboring voxel v_i is white matter, it is immediately included. If v_i is grey matter (other tissue), its count is incremented by 0.2 (resp., 0.1). GW does not allow diffusion into the skull. This is easy to determine as the \mathbf{e}_i part of the `GeneralInfo` \mathbf{e} specifies the tissue type of each v_i voxel, as computed by SPM [6] (see Fig. 1[d-f]).

3.1 CDM Diffusion Model

Our CDM model is more sophisticated. First, its decision about each voxel depends on a number of features, based on:

- the patient, $\mathbf{e}_{Patient}$:** the age (which may implicitly indicate the tumor grade).
- the tumor, \mathbf{e}_{Tumor} :** volume-area ratio, edema percentage, and volume increase.
- each individual voxel $\{\mathbf{e}_i\}_i$:** various attributes for every voxel v_i — spatial coordinates, distance-area ratio, minimum euclidean distance from the tumor border, whether the voxel is currently in an edema region, white matter, grey matter, or CSF (automatically determined by SPM [6]), and image intensities of T1, T1-contrast and T2 axial scans [3] (obtained both from the patient’s scan and a standard template³ [9] — after normalization and registration using SPM [5]).
- neighborhood of each voxel $\{\mathbf{e}_i\}_i$:** attributes of each of the 6 neighbors of the voxel — whether a neighbor voxel n_j is edema, white matter, grey matter, or CSF, and image intensities from the template’s T2 and T1-contrast.

(The webpage [1] provides more details about each of these features, as well as some explicit examples.)

CDM then uses a probabilistic classifier to compute the probability q_i that one tumor neighbor v_i of a tumor voxel will become tumorous, $q_i = P_{\Theta}(\ell(v_i) = \text{Tumor} \mid \mathbf{e}_{Patient}, \mathbf{e}_{Tumor}, \mathbf{e}_i)$. Some voxels can have more than *one* such tumor-neighbors; *e.g.*, in diagram (1), the voxels $\mathbf{v}1$, $\mathbf{v}2$ and $\mathbf{v}5$ each have a single tumor-neighbor, while $\mathbf{v}3$ and $\mathbf{v}4$ each have 2. Each tumor-neighbor of the voxel v_i has a q_i chance to transform this v_i ; hence if there are k such neighbors, and each acts independently, the probability that v_i will be transformed on this iteration is $p_i = 1 - (1 - q_i)^k$. CDM will then transform this voxel to be a tumor with probability p_i . We then assign it to be a tumor if $p_i > \tau$ using a probability threshold of $\tau = 65\%$.⁴ CDM performs these computations in parallel — hence on the first iteration, even if $\mathbf{v}3$ is transformed, $\mathbf{v}4$ still has only 2 tumor-neighbors

³ Several images of a normal brain of an individual, averaged and registered to the same coordinate system.

⁴ We experimented with several thresholds, and chose this $\tau = 0.65$ value as it provided the best observed accuracy.

(on this iteration). We discuss below how CDM learns the parameters Θ used in $P_{\Theta}(\cdot)$.

4 Experiments

We empirically evaluated the three models, UG, GW and CDM, over a set of 17 patients. For each patient, we had two sets of axial scans R_1 and R_2 taken at different times, each with known tumor regions. Let s_i refer to the size of the tumor in scan R_i . For each patient, we then input that patient’s initial scan (R_1) to each model, and asked it to predict the next $s = s_2 - s_1$ voxels that would be transformed. We then compare the predicted voxels with the truth — *i.e.*, the tumor region of the second scan, R_2 .

To measure the quality of each model, let “ nt ” be a set of tumor cells for the patient that are actually transformed (*i.e.*, this is the “truth”, associated with R_2) and “ pt_{χ} ” be the cells that the χ model predicts will be transformed. We then use the standard measures: “precision” of χ (on this patient) is $\frac{|nt \cap pt_{\chi}|}{|pt_{\chi}|}$ and “recall” is $\frac{|nt \cap pt_{\chi}|}{|nt|}$. In our case, as our diffusion models stop when $|pt_{\chi}| = |nt| = s$, the precision and recall values will be the same⁵ (see tables in [1, 13]). We report results in terms of the “F-measure” = $\frac{2 \times \text{precision} \times \text{recall}}{\text{precision} + \text{recall}}$ [21], where F-measure = precision = recall, for each patient.

While UG and GW are completely specified, CDM must first be trained. We use a “patient level” cross-validation procedure: That is, we trained a learner (*e.g.*, Logistic Regression [12] or SVM [14]) on 16 patients, then tested on the 17th. Each training instance corresponded to a single voxel v_i around the initial tumor in the first scan R_1 , with features $\mathbf{e}_{Patient}$, \mathbf{e}_{Tumor} , and \mathbf{e}_i , and with the label of “1” if this voxel was in the tumor in R_2 , or “0” otherwise. Training voxels represent the set difference between the tumor in R_1 and R_2 for each patient (*i.e.*, the region that a ‘perfect’ diffusion model would consider), in addition to a 2-voxel border around the tumor in R_2 to account for the segmentation error margin at the tumor border. The total number of training voxels was approximately $\frac{1}{2}$ million for the 17 patients. Notice this training is at the voxel level, and is only implicitly based on the diffusion approach (in that this is how we identified the specific set of training voxels).

Results appear in Fig. 3 and in [1, 13]. Below we analyze these results in terms of best, typical, and special cases; describe system performance versus tumor grade; and statistically assess of the three models.

4.1 Feature Selection

Here, we consider finding the best subset of the 75 features described in Section 3.1, called S_0 . We first computed the Information Gain (IG) of each feature,

⁵ In some patients, precision and recall can be slightly different if the algorithm terminates prematurely, *i.e.*, before reaching the target size of the tumor.




Fig. 3. Empirical Results

The F-measure for the three Models with 17-fold “patient-level” CV (Note F-measure = precision = recall, for each patient — see Section 4). Results correspond to the output of a logistic regression classifier, learned with feature set S_1 . The name of each patient identifies their tumor grades — Astrocytoma grade I (A) and grade II (A.GBM) that progressed into GBM, Mixed Oligo-astrocytoma grade II (MOA), Anaplastic astrocytoma grade III (AA), and the most common GBM.

then ranked the features based on their IG scores. We observed that patient-specific tissue features have the lowest IG scores (likely due to SPM’s segmentation errors and the presence of tumors in patients’ scans). We formed two subsets of features based on the IG scores and the feature type (*e.g.*, tumor-specific, tissue-based features, spatial coordinates, etc.). The first subset S_1 contains 28 features only; it excluded all patient-specific tissue features since these have lower IG scores (see [1, 13]), as well as spatial coordinates and template-specific tissue features, to help generalize the learned tumor growth model (*i.e.*, without making any assumptions about the spatial location of the tumor). The second subset S_2 contains 47 features, excluding only CSF features as these are associated with the lowest IG scores, likely due to errors in SPM’s tissue segmentation process. (Note tumors do not grow into CSF regions, *e.g.*, ventricles⁶, but induced tumor pressure can deform them, which allows tumors to appear in a region that *had* been ventricles, etc.)

By excluding tissue-based features from S_1 , we allow the model to perform more accurately for subjects whose tumors have altered the basic brain anatomy — *e.g.*, tumors that have deformed the ventricles, such as patients *A.GBM_4* and *GBM_12* (see Fig. 4). But accuracy slightly decreased for scenarios that rely on specific training information (*i.e.*, voxel locations and tissue information). Since S_2 includes spatial and tissue information, classifiers that used these features

⁶ Cavities in the brain filled with cerebrospinal fluid (CSF).




Fig. 4. Tumor-induced pressure deforms the ventricles in patients *A.GBM_4* and *GBM_12* (two image slices for each patient).

performed almost the same as S_0 . Fig. 3 reports results obtained when training on S_1 feature set only. Results with the other feature sets appear in [1, 13].

4.2 Tumor Growth Patterns Learned from the Data

This section presents a few of the patterns we observed based on a learned model. Given a Logistic Regression model, we compute — from a classifier-labeled held-out set of patient *GBM_7* (22,600 voxels that were normal but became tumor on the next scan) — the conditional probability that a voxel v , has attribute a , given its class label, *i.e.*, $P(a(v) = a_i | class(v))$.

We observe that among the voxels that were labeled by the classifier as ‘tumor’, 70% are edema, 29% have $T2 \geq 0.75$, 42% have $T1 < 0.5$, 41% are grey matter, and 44% white matter (note each probability is computed independently of other attributes), and in voxels labeled as ‘non-tumor’, the probabilities are 10%, 6%, 87%, 17%, and 0.5%, respectively. Additional details are in [1].

Based on the learned model, we then derive the conditional probability that v will become ‘tumor’ given its properties of edema, T2, and grey or white matter: $P(class(v) = 'tumor' | edema(v) = 1, T2(v) \geq 0.75, tissue(v) = g/w) = 99\%$.

Hence, we conclude from the learned model that voxels located in edema regions (bright on T2, dark on T1 scans) and in the grey or white matter (the last being a diffusion pathway for tumor cells) are more likely to become diseased.

4.3 Typical, Best, and Special Case Results

Patients *GBM_1*, *GBM_2*, and *GBM_3* represent typical case results, where CDM performs more accurately than UG and GW by at least a small percentage. In these cases, the tumor tends to grow along the edema as glioma cells have already infiltrated into the peritumoral edema regions. These diffuse occult cells did not enhance at first on T1-contrast images as these cells may exist only in very low concentration. But on the next scan of the patient, enhancing tumors appeared in these regions as glioma cells built up into detectable masses.

Infiltration of glioma cells in edema regions is particularly more obvious on the MRI scans for patient *GBM_7* (Fig. 5), which represents the best case results as CDM models tumor diffusion more accurately than UG and GW, by 20% and 12% respectively (see Fig. 3 and tables in [1, 13]).




Fig. 5. MR T1-contrast images of Patient *GBM_7*, showing lower to higher axial brain slices from left to right, augmented with shades of grey corresponding to results from CDM model: initial tumor volume is colored white, true positives are light grey, false positives are dark grey, and false negatives are black (see color version in [1]).

In typical and best case scenarios, the prediction is based on what the classifier recognizes as ‘tumor’, which are often voxels located in edema regions. Glioma cell infiltration in peritumoral edema may be even more detectable if the truth volume was obtained from a patient scan before that patient underwent a surgical procedure or received radiation treatment.

Patients *GBM_10*, *GBM_12*, and *GBM_13* are examples of special tumor growth cases where tumors do not follow usual diffusion patterns (*e.g.*, the tumor shrinks due to treatment and recurs a few months later in regions near the original mass). In these cases, CDM performed the same as the standard models. The effect of treatment is present in all of our data, but is more prominent in these patients.

4.4 Model Performance versus Tumor Grade

Our dataset consists of four different glioma grades ranging from low-grade astrocytomas to the most invasive GBM. GBMs are the most common among glioma patients, and represent $\frac{2}{3}$ of our data. Because CDM is a general learning model, it is not restricted to predicting a particular tumor grade, but it requires a fair representation of various tumor types in training data. Currently, low-grade tumors are under-represented in our data since they are less common among glioma patients.

Also, CDM’s prediction is based on probabilities assigned by classifiers to the unlabeled voxels. High-probability tumor voxels are likely to be located in peritumoral edema regions (edema features have the highest IG scores), particularly more pronounced in high-grade, larger tumors (*e.g.*, patients *GBM_1*, *GBM_3*, and *GBM_7*). This is because peritumoral edema regions harbour diffuse malignant cells that infiltrated through tissue near the visible tumor. These malignant cells form detectable tumor masses over time.

4.5 Statistical Evaluation of the Three Models

Over the 17 patients, the average leave-one-out recall (\equiv precision) values for the CDM, UG and GW models are 0.598, 0.524 and 0.566 respectively. We ran

a t -test [16] for paired data to determine if these average values are statistically significant from one another, at the 95% confidence interval (*i.e.*, $p < 0.05$).

- Comparing CDM versus UG, the t value is 4.14 meaning the probability of the null hypothesis (*i.e.*, values are not significantly different) is 0.001. In this case, we reject the null hypothesis and conclude that the average recall obtained with CDM and UG are significantly different.
- Comparing CDM versus GW, the t value is 3.61 meaning the probability of the null hypothesis is 0.002, which suggests that the average recall obtained with CDM and GW are significantly different as well.

Given the above t -test results, we conclude that our CDM model is performing more accurately, in general, than either of UG and GW.

4.6 Computational Cost of the Three Models

CDM requires several preprocessing steps of the MRI scan followed by feature extraction (which require approximately one hour). Given a segmented tumor, and a learned classifier (*e.g.*, Logistic Regression), CDM produces its prediction of tumor growth in 10 minutes on average⁷. UG and GW require the same data processing, and produce their predictions in 1 and 10 minutes on average, respectively. Note UG performs the least number of iterations.

5 Contributions and Future Work

Our team has produced a system that can automatically segment tumors based on their MRI images [2]; we are currently using this system to produce tumor volume labels for hundreds of patients, over a wide variety of tumor types and grades. We plan to train our diffusion model on this large dataset. We will also experiment with other learning algorithms, including Conditional and Support Vector Random Fields [11], as these may better account for neighborhood interdependencies between tumor and normal voxels. We will also investigate other attributes, *e.g.*, estimated tumor growth rate, and features from other types of data such as Magnetic Resonance Spectroscopy. We may also incorporate diagonal neighbors in the diffusion algorithm, which may help improve the accuracy, and will also help decrease the number of iterations required to grow the tumor, making the algorithm more efficient.

Contributions: This paper has proposed a classification-based model, CDM, to predict glioma diffusion, which *learns* ‘general’ diffusion patterns from clinical data. (To the best of our knowledge, this is the first such system.) We empirically compare CDM with two other approaches: a naïve uniform growth model (UG) and a tissue-based diffusion model (GW), over pairs of consecutive MRI scans. Our results, on real patient data (as opposed to simulating virtual tumors [19]), show statistically that CDM is more accurate. See [1] for more details.

⁷ This average is computed over our set of 17 patients, characterized by a wide variety of tumor sizes, including a few that require a very large number of additional voxels to grow, and were therefore more computationally costly.

References

1. Brain Tumor Growth Prediction: <http://www.cs.ualberta.ca/~btgp/ai06.html>
2. Brain Tumor Growth Project: <http://www.cs.ualberta.ca/~btgp>
3. Brown M., and Semelka R., MRI Basic Principles and Applications (2003) Wiley, Hoboken, NJ, USA
4. Clatz O., Bondiau P., Delingette H., *et al.*, In Silico Tumor Growth: Application to Glioblastomas. *MICCAI 2004, LNCS 3217* (2004) 337–345
5. Friston K., Ashburner J., Frith C., Poline J., Heather J., and Frackowiak R., Spatial Registration and Normalization of Images Human Brain Mapping **2** (1995) 165–189
6. Friston K., and Ashburner J., Multimodal Image Coregistration and Partitioning — a Unified Framework **6** (1997) NeuroImage 209–217
7. Halperin E., Bentel G., and Heinz E., Radiation Therapy Treatment Planning in Supratentorial Glioblastoma Multiforme, *Int. J. Radiat. Oncol. Biol. Phys.* **17** (1989) 1347–1350
8. Hochberg F., and Pruitt A., Assumptions in the Radiotherapy of Glioblastoma, *Neurology* **30** (1980) 907–911
9. Holmes C., Hoge R., Collins L., *et al.*, Enhancement of MR images using registration for signal averaging. *J Comput Assist Tomogr*, **22**(2) (1998) 324–333
10. Kansal A., Torquato S., Harsh G., *et al.*, Simulated brain tumor growth dynamics using a three-dimensional cellular automaton. *J Theor Biol.* **203** (2000) 367–382
11. Lee C.H., Greiner R., and Schmidt M., Support Vector Random Fields for Spatial Classification. *Proceedings of the 9th European Conference on Principals and Practices of Knowledge Discovery in Data (PKDD 2005)*
12. le Cessie S. and van Houwelingen J., Ridge Estimators in Logistic Regression. *Applied Statistics*, **41**(1) (1992) 191–201
13. Morris M., *Classification-based Glioma Diffusion Modeling*. MSc Thesis, University of Alberta (2005)
14. Platt J., Fast Training of Support Vector Machines using Sequential Minimal Optimization. *Advances in Kernel Methods — Support Vector Learning*, eds., Schoelkopf B., Burges C., and Smola A. (1998) MIT Press
15. Price S., Burnet N., Donovan T., *et al.*, Diffusion tensor imaging of brain tumors at 3T: a potential tool for assessing white matter tract invasion? *Clinical Radiology*, **58** (2003) 455–462
16. Student, The Probable Error of a Mean. *Biometrika*, **6** (1908) 1–25
17. Swanson K., Alvord E., Murray J., A Quantitative Model for Differential Motility of Gliomas in grey and White Matter, *Cell Prolif*, **33** (2000) 317–329
18. Swanson K., Alvord E., Murray J., Quantifying efficacy of chemotherapy of brain tumors with homogeneous and heterogeneous drug delivery. *Acta Biotheor*, **50**(4) (2002) 223–237
19. Swanson K., Alvord E., Murray J., Virtual brain tumors (gliomas) enhance the reality of medical imaging and highlight inadequacies of current therapy, *British Journal of Cancer*, **86** (2002) 14–18
20. Tabatabai M., Williams D. and Bursac Z., Hyperbolastic growth models: theory and application. *Theor Biol Med Model*, **2**(1) (2005), 14
21. Van Rijsbergen C. J., *Information Retrieval* (1979) second edition, London, Butterworths.
22. Zizzari A., *Methods on Tumor Recognition and Planning Target Prediction for the Radiotherapy of Cancer* (2004) PhD Thesis, University of Magdeburg