Exploring trade-offs in treatment planning for brain tumor cases with a probabilistic definition of the clinical target volume.

Gregory Buti[1,2], Nadya Shusharina[2], Ali Ajdari[2], Edmond Sterpin[1,3], Thomas Bortfeld[2]

[1]UCLouvain, Institute of Experimental and Clinical Research, Center of Molecular Imaging, Radiotherapy and Oncology, Avenue Hippocrate 54 - box B1.54.07, 1200 Brussels, Belgium

[2]Massachusetts General Hospital and Harvard Medical School, Department of Radiation Oncology, Division of Radiation Biophysics, 100 Blossom St, Boston, MA 02114, USA

[3]KU Leuven, Department of Oncology, Laboratory of Experimental Radiotherapy, UZ Herestraat 49 - box 7003, 3000 Leuven, Belgium

Version typeset August 19, 2022

E-mail: gregory.buti@uclouvain.be

Abstract

Purpose: This study demonstrates how a novel probabilistic CTV concept—the clinical target distribution (CTD)—can be used to navigate the trade-off between target coverage and organ sparing with a semi-interactive treatment planning approach.

Methods: Two probabilistic treatment planning methods are presented that use tumor probabilities to balance tumor control with organ-at-risk (OAR) sparing. The first method explores OAR dose reduction by systematically discarding $x\%$ of CTD voxels with an unfavorable dose-to-probability ratio from the minimum dose coverage objective. The second method sequentially expands the target volume from the GTV edge, calculating the CTD coverage vs. OAR sparing trade-off after dosing each expansion. Each planning method leads to estimated levels of tumor control under specific statistical models of tumor infiltration: an independent tumor islets model and contiguous circumferential tumor growth model. The methods are illustrated by creating proton therapy treatment plans for two glioblastoma patients with the clinical goal of sparing the hippocampus and brainstem. For probabilistic plan evaluation, the concept of a *dose-expected-volume* histogram is introduced, which plots the dose to the expected tumor volume $\langle v \rangle$ considering tumor probabilities.

Results: Both probabilistic planning approaches generate a library of treatment plans to interactively navigate the planning trade-offs. In the first probabilistic approach, a significant reduction of hippocampus dose could be achieved by excluding merely 1% of CTD voxels without compromising expected tumor control probability (TCP) or CTD coverage: the hippocampus D_2 dose reduces with 9.5 Gy and 5.3 Gy for Patient 1 and 2, while the TCP loss remains below 1%. Moreover, discarding up to 10% of the CTD voxels does not significantly diminish the expected CTD dose, even though evaluation with a binary volume suggests poor CTD coverage. In the second probabilistic approach, the expected CTD $D_{\langle 98 \rangle}$ and TCP depend more strongly on the extent of the high-dose region: the target volume margin can not be reduced by more than 2 mm if one aims at keeping the expected CTD $D_{\langle 98 \rangle}$ loss and TCP loss under 1 Gy and 2%, respectively. Therefore, there is less potential for improved OAR sparing without compromising TCP or expected CTD coverage.

Conclusions: This study proposes and implements treatment planning strategies to explore trade-offs using tumor probabilities.

Keywords— CTV uncertainty, probabilistic planning, proton therapy, glioblastoma

I. Introduction

A probabilistic definition of the clinical target volume (CTV) is increasingly garnering attention in the medical physics community due to its ability to deal with the stochastic nature of microscopic disease.[13,31] Compared to a binary CTV, a probabilistic CTV or clinical target distribution (CTD) represents the likelihood that a sub-volume or voxel is tumorous.[24] The rationale underlying the CTD approach is that it can guide an optimization algorithm or treatment planner in making compromises between conflicting clinical goals such as tumor dosage and sparing healthy tissues. Moreover, the CTD may be advantageous in reducing the inter-user variability of target volume definition.[4,32] The application of the CTD to the robust optimization of lung and prostate cases can be found in Buti et al.[9] and Ferjančič et al.[12], respectively.

In the studies mentioned above, the probabilistic target concept was implemented by varying the importance weights of the target objectives according to the tumor probabilities. By doing so, dosing the higher probability sub-volumes is prioritized over lower probabilities in the presence of

a dose-limiting structure. This CTD implementation is similar to the way Baum et al.[5] treated geometric uncertainties in the optimization, i.e., the coverage probabilities for CTV and OAR voxels were estimated under geometric errors and incorporated explicitly in the planning objective function. Alternatively, the study of Bortfeld et al.[7] uses the tumor probabilities to evaluate a tumor control probability (TCP) function. The TCP expression differed depending on whether the tumor volume was assumed to be composed of independent or dependent voxels. The voxel-dependency assumption relates to the choice of the microscopic tumor propagation model that underlies the CTD definition. Either the tumor is assumed to propagate via independent tumor islets (for independent voxels) or a process of contiguous circumferential growth (for dependent voxels). Under both assumptions, using the TCP as an objective function leads to a non-convex optimization problem. Interestingly, the non-linear dose response of the TCP gives rise to 'sacrificing' behavior; with a limitation on the integral dose, dosing the higher probability sub-volumes with a therapeutic dose level and not dosing the lower probability sub-volumes at all yields the highest expected tumor control.

Unfortunately, the non-convexity renders direct optimization of the TCP function unfeasible for more complex cases. To address this issue, the problem can be reformulated in the limit of high dose levels where the tumor control is reasonable.[2,7] However, such an approach does not translate well to the presence of low doses to potentially tumorous voxels, particularly when dose constraints in nearby critical organs must be respected. This study investigates the use of tumor probabilities in treatment planning, without making the restrictive high dose approximation of other studies. We present two novel planning methods that take advantage of the sacrificing behavior presented in Bortfeld et al.[7] to spare a nearby critical OAR. The goals of this study are twofold: (a) develop semi-interactive planning approaches to explore the treatment compromises, akin to multi-criteria optimization, and (b) show that each planning method leads to expected levels of tumor control under specific statistical assumptions of voxel-dependency. We demonstrate the potential of these methods by creating treatment plans for two glioblastoma multiforme (GBM) cases, exploring the trade-off between CTD dose and sparing of the hippocampus and brainstem with intensity-modulated proton therapy.

II. METHODS

II. Methods

This study proposes two different probabilistic planning methods based on CTD tumor probabilities that can be used to balance target coverage with OAR sparing. Both methods explore the possibility of OAR dose reduction by delivering no or low dose to lower probability CTD voxels, but differ in the way the dose gets distributed across these lower probability voxels. The two assumptions that are investigated are a fully independent assumption and a dependent voxel assumption. Although both assumptions follow from inherently statistical principles, one can interpret them as describing a tumor propagation process occurring via independent tumor islets and contiguous circumferential growth, respectively.[7]

II.A. Clinical target distribution

The CTD is defined by a discrete number of evenly-spaced iso-probability surfaces called « shells » to which a probability is assigned.[24] The shell probability represents the likelihood that there is tumor presence outside that shell, in a patient population. The innermost shell coincides with the GTV contour, while the outermost shell is the surface beyond which we assume no tumor presence. We refer to the volume between two adjacent shells as a « layer ».

For the studied GBM cases, the CTD shells were obtained with an automated 3D expansion algorithm, implemented in the research version of Raystation 10A (RaySearch Laboratories, Stockholm, Sweden).[25] Starting from the GTV, isotropic expansions were generated at a 2 mm increment, taking into account barrier structures. The barrier structures included the ventricles, falx cerebri, tentorium cerebelli, brainstem, and the outer surface of the brain. The surfaces of the GTV expansions yielded the shells, while the voxels in between two expansion surfaces produced the layers. Fig. 1 shows the GBM cases with a set of expansion contours.

After assigning the shell probabilities, the voxel-level probabilities follow directly from the voxel-dependency assumption underlying the CTD definition. Fig. 2 presents a schematic representation for a simplified CTD. The first probabilistic planning method considers the CTD voxels to be mutually independent, that is, the probability that there is tumor in one CTD voxel does not depend on any other voxel. Here, microscopic tumor propagation manifests itself through the occurrence of tumor islets at quasi random locations with a probability given by the CTD, but without any connection or correlation between the islets, , see Fig. 2 (a) Let p be the probability that

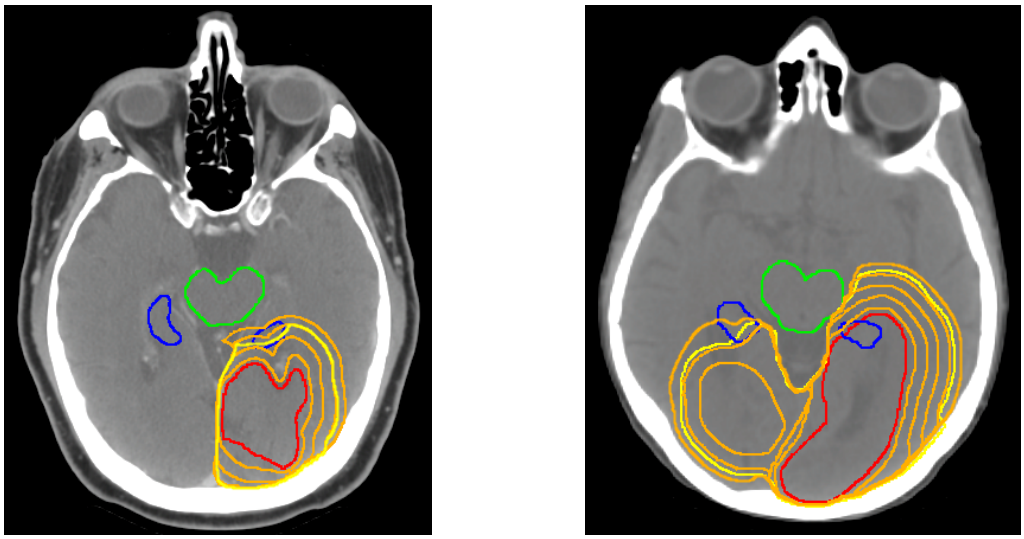

(a) Patient 1. (b) Patient 2.

Figure 1: CT images of two GBM patients with a probabilistic definition of the CTV. Representative iso-probability contours are shown in orange. Other delineated contours indicate the clinically-approved CTV (yellow), GTV (red), hippocampi (blue) and brainstem (green).

there is tumor in a voxel and r the shell probability. Shusharina et al.[24] derived the relationship of p to r under voxel-independence as:

$$p_h = 1 - \left(\frac{1 - r_h}{1 - r_{h+1}}\right)^{\frac{1}{N_h}} \tag{1}$$

with N_h the number of voxels in layer h. Eq. 1 describes the increasing probability of being tumorous as a voxel is part of an inner layer vs. an outer layer. For GBM cases, with layers spaced 2 mm apart, p typically has a value ranging from 10^{-5} (inner layers) to 10^{-6} (outer layers), depending on the shell probability and the number of voxels in the layer.

The second probabilistic planning method considers voxel-dependency of the tumor volume with two assumptions: (1) if a voxel is tumorous, all voxels along the shortest path between it and the GTV are also tumorous. In other words, tumor cells cannot tunnel through layers, and (2) if one voxel in a layer is tumorous, all voxels in the same layer are tumorous. In other words, the tumor cells propagate outward via a tumor front. The above two assumptions characterize tumor growth through a contiguous circumferential front process, see Fig. 2 (b). Appendix A presents the derivation of p as a function of r. We find that the voxel-level probabilities equal the shell probabilities:

$$p_h = r_h. \tag{2}$$

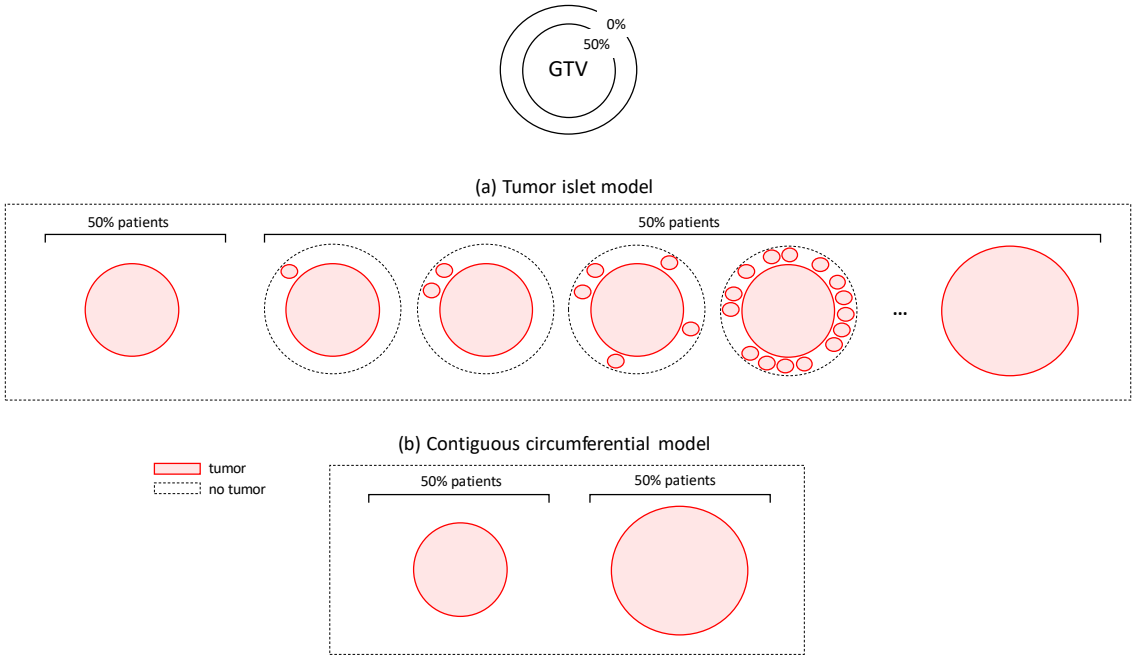

Figure 2: Schematic representations of the underlying tumor model assumptions. Consider a simplified case of a two-layer CTD, defined by two iso-probability shells: an outer shell of 0% and inner shell of 50% (= GTV contour). The probabilities are defined in such a way that 0% of patients have tumor outside of the 0%-shell and 50% of patients have tumor outside the 50%-shell. The possible tumor realizations depend on the assumed microscopic tumor propagation model. Let's consider a population of patients with the probability of tumor described by the above CTD. (a) in the tumor islet model, 50% of patients have tumor realized solely in the GTV, and 50% of patients have tumor in the GTV + tumor in any combination of independent islets from 1 islet to all islets that make up the outer layer. The tumor islets can be voxels or larger tumor sub-volumes, (b) in the contiguous circumferential growth model, only two tumor realizations are possible. 50% of patients have tumor realized in the GTV, while 50% of patients have tumor in the GTV + outer layer.

II.B. Probabilistic treatment planning method assuming an independent tumor islet model

Bortfeld et al.[7] shows that when the dose to an idealized target is limited, not dosing the lowest probability voxels yields the highest expected tumor control. Here, we take advantage of this 'all-or-nothing' dose property to design a method that balances CTD dose with OAR dose using tumor probabilities. For instance, one could decide not to treat 1% of potentially tumorous voxels near a critical OAR if the expected loss of tumor control remains limited.

Deciding the optimal set of CTD voxels to discard, to strike the best balance between CTD coverage and OAR sparing, requires formulating the problem as an instance of mixed-integer programming. However, this method could be computationally prohibitive, given the large size of the underlying problem for typical clinical cases. The optimization problem can be approximated by following a procedure similar to the typical methods for handling (integer) dose-volume histogram (DVH) constraints in treatment planning. DVH constraints are typically implemented by projecting the dose distribution onto the nearest dose distribution that fulfills the constraint at each iteration step.[6,8,11] This concept can be applied to our optimization problem as follows: let K be the number of voxels that cannot be covered with the prescribed dose (d^{presc}) and M the total number of CTD voxels. Based on the 'nearest dose distribution' idea, we select the $M-K$ voxels with the smallest positive ratio of $(d^{\text{presc}}-d)/p$. The idea is that CTD voxels are preferred where (1) the dose is close to the prescription, and (2) the tumor probability is high, given that dosing higher-probability voxels leads to greater tumor control.

The main difficulty lies in the fact that the appropriate number of voxels to sacrifice is not known *a priori*. This planning challenge is investigated by following a multi-criteria optimization approach. The balance between CTD coverage and OAR sparing will be explored interactively by generating a library of treatment plans that aim at dosing different fractions of CTD voxels, e.g., from 100% down to 90% in steps of 1% (as shown in Fig. 3a). For each treatment plan, the ideal dose projection approach is followed by excluding the $100\%-x\%$ CTD voxels with an undesirable dose-to-probability ratio to satisfy a maximum dose objective for an OAR.

II.B.1. Justification using a simplified tumor control probability function

In this section, we show that optimizing the dose to a discrete number of CTD voxels is equivalent to aiming for a pre-determined TCP level, under an independent tumor islet model. The voxel-wise TCP in volume v_i after irradiation with the dose d_i, is typically expressed with a double exponential function:

$$\text{TCP}_i = e^{-\rho v_i e^{-\alpha d_i}}, \tag{3}$$

where we have omitted the quadratic dose effects. ρ is the tumor cell density, v_i is the volume, and α the radio-sensitivity parameter. We consider here a simplified TCP expression by approximating

Eq. 3 with a step function:

$$\mathrm{TCP}_i \simeq \begin{cases} 0, & d_i < d^{\mathrm{ref}} \\ 1, & d_i \geq d^{\mathrm{ref}}, \end{cases} \tag{4}$$

with d^{ref} the dose at 50% tumor control. According to this approximation, the dose has either a curative ($d_i \geq d^{\mathrm{ref}}$) or non-curative ($d_i < d^{\mathrm{ref}}$) effect. For the commonly used values $\alpha \approx 0.4$ Gy^{-1} and $\rho \approx 10^7/\mathrm{cm}^3$, a reasonable estimate of d^{ref} is 20–30 Gy. d^{ref} equal to 24.1 Gy was chosen to match 50% tumor control for a $0.65 \times 0.65 \times 2.5$ mm^3 (CT resolution) voxel. The step function approximation should be reasonably valid if the shape of the voxel control probability curve is sufficiently steep.

As derived in Alber and Thorwarth[2], Shusharina et al.[24] and Bortfeld et al.[7], the TCP expression for a target volume composed of mutually independent voxels with probabilities p_i is:

$$\mathrm{TCP} = \prod_i \left(1 - p_i + p_i \mathrm{TCP}_i\right). \tag{5}$$

Applying the step function approximation gives:

$$\mathrm{TCP} \approx \prod_{i, d_i < d^{\mathrm{ref}}} \left(1 - p_i\right), \tag{6}$$

where the product applies to all voxels outside of the GTV with a dose below d^{ref}. Given that the probabilities are very small (10^{-5}–10^{-6}), this expression can be linearized by ignoring the higher-order terms in p:

$$\mathrm{TCP} \approx 1 - \sum_{i, d_i < d^{\mathrm{ref}}} p_i. \tag{7}$$

Eq. 7 implies that for the TCP to not fall below a pre-defined level, the sum of the probabilities ($\sum p_i$) over the foregone voxels—where the d^{ref} dose level is not reached—must not exceed a maximum threshold. Hence, we can aim for a pre-defined TCP level by optimizing a discrete set of potentially tumorous voxels that should receive a dose larger than d^{ref}. Or inversely, scaling the number of CTD voxels that should receive the prescribed dose varies the expected TCP.

II.C. Probabilistic treatment planning method assuming a contiguous circumferential growth model

The second probabilistic planning method balances CTD coverage with OAR sparing by dosing only the voxels within a certain expansion volume of the GTV. In other words, the voxels outside of the

GTV expansion are potentially not treated to spare a critical OAR. The design of this probabilistic dosing approach follows from the idea that delivering dose to inner layers is more preferential than outer layers, given that their tumor probabilities are higher. Similar to the method presented in Section II.B., interactive plan navigation is possible by varying the extent of CTD coverage. Here, a library of treatment plans is optimized, where each treatment plan aims at delivering the prescribed dose to an increasing number of CTD layers, from the GTV to the volume encompassed by the outermost shell. Therefore, a growing expansion of the GTV gets dosed from one treatment plan to the next. Fig. 3b shows an example of a set of pre-computed dose distributions generated for this probabilistic planning approach.

II.C.1. Justification using a simplified tumor control probability function

We show that dosing the voxels within an expansion volume of the GTV is equivalent to aiming for a pre-determined TCP level assuming a contiguous circumferential growth model. Under contiguous circumferential growth, the probability that a layer is tumorous is equal to the probability that a single voxel in that layer is tumorous (Eq. 15). Taking this into consideration, an equivalent derivation can be followed as presented in Bortfeld et al.[7] but for a 3D CTD composed of a series of layers.[1] The TCP can be expressed as the sum of H product terms (rather than H single voxels as in the 1D case):

$$\begin{aligned}\mathrm{TCP} = \; & (1-p_1)\prod_{i\in\mathrm{GTV}}\mathrm{TCP}_i + (p_1-p_2)\prod_{i\in\mathrm{GTV}}\mathrm{TCP}_i\prod_{i\in\mathcal{L}_1}\mathrm{TCP}_i \\ & + \cdots + p_{H-1}\prod_{i\in\mathrm{GTV}}\mathrm{TCP}_i\prod_{i\in\mathcal{L}_1}\mathrm{TCP}_i\prod_{i\in\mathcal{L}_2}\mathrm{TCP}_i\cdots\prod_{i\in\mathcal{L}_{H-1}}\mathrm{TCP}_i,\end{aligned} \tag{8}$$

where $\mathcal{L}_h$ denotes CTD layer h. After applying the step function approximation, Eq. 8 simplifies if the dose in at least one voxel is limited, $d_i < d^{\mathrm{ref}}$. In the particular case that a foregone voxel is part of layer h, the $H-h$ product terms which evaluate this voxel are zero. The general expression then simplifies to:

$$\begin{aligned}\mathrm{TCP} \approx \; & (1-p_1)+(p_1-p_2)+(p_2-p_3)+\cdots+(p_{h-1}-p_h). \\ \approx \; & 1-p_h.\end{aligned} \tag{9}$$

Eq. 9 implies that for the TCP to not fall below a pre-defined level, the probability of the innermost layer which contains at least one foregone voxel must not exceed a maximum threshold. In

[1] [7], pp. 3, section 2.2 with the chain of voxels replaced by a series of layers.

other words, the treatment planning method presented in Section II.C. aims at generating a set of treatment plans with different TCP values, depending on the extent of the GTV expansion volume that is dosed.

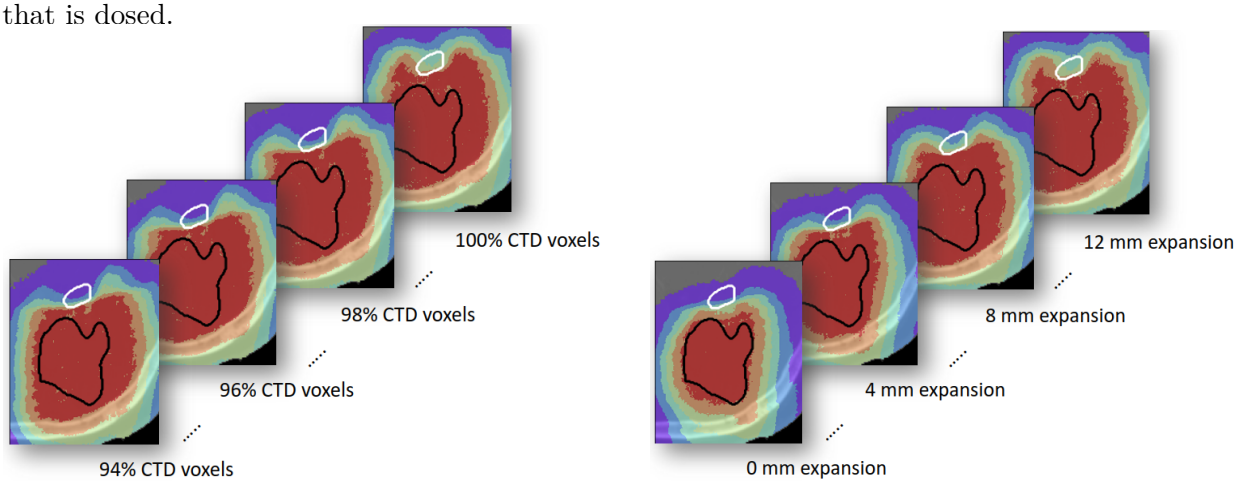

(a) Strategy 1. (b) Strategy 2.

Figure 3: Multi-criteria style navigation of dosimetric trade-offs. Libraries of treatment plans are pre-computed through which the trade-off between CTD coverage and sparing of nearby OARs can be explored. An example is shown for a GBM patient. The dose distributions overlay the CT anatomy with the GTV and OAR depicted in black and white, respectively. Left: treatment plans vary the % of CTD voxels that we aim at dosing with the prescription dose (from 90% to 100%). Right: treatment plans vary the GTV expansion we aim at dosing uniformly (from 0 mm to 12 mm for the example patient).

II.D. Evaluation

The treatment plans are evaluated with both conventional and probabilistic metrics. The conventional metrics do not take into account tumor probabilities and rely solely on evaluating binary volumes with dose-volume histograms (DVH). The probabilistic metrics include information on the tumor probabilities through the TCP and the expected dose coverage to the CTD. The TCPs are computed with Eqs. 5 and 8 for each respective tumor propagation model. To avoid any biases, the double-exponential TCP model (Eq. 3) will be compared to the step function approximation. For the expected dose coverage, the concept of a conventional DVH is extended to a *dose-expected-volume histogram* (DEVH). The DEVH was first suggested in Shusharina et al.[24] as a way to evaluate the dose coverage of probabilistic volumes.[2] A formal definition of the DEVH is presented in Appendix B. The DEVH plots the expected volume $\langle v \rangle$ as a function of the dose for all voxels within the CTD, i.e., the volume of voxels with non-zero tumor probability. The intuition behind

[2]Shusharina et al.[24], pp. 11, section 4.

the DEVH can be explained with the following example: if the probability of tumor in a voxel is twice as large as in another voxel, then it must be valued twice as much in the differential dose-volume histogram (given that the voxel will be tumorous in twice as many patients, in a patient population). Naturally, if the probability of tumor is 0%, then this voxel will not be counted at all. The dose coverage to the expected target volume, e.g., $\langle v \rangle = 98\%$, can be read directly from the graph as $\mathrm{D}_{\langle 98 \rangle}$. Note that for binary volumes, the DEVH reduces to the conventional DVH.

II.E. Treatment planning system

Proton therapy treatment plans were created with an in-house treatment planning system called OpenTPS, coded in Python. OpenTPS is coupled to the Monte Carlo dose engine MCsquare for the dose calculation.[26,27] Beamlets were computed on a $2 \times 2 \times 2.5$ mm^3 dose grid with 5E4 protons. The final dose calculations were performed with the number of protons necessary to reach a sub-2% statistical noise level in CT resolution. All optimizations were performed with the limited-memory Broyden–Fletcher–Goldfarb–Shann (L-BFGS) solver, provided by the Scipy package.[19,22]

The optimization problems are formulated as minimizing the weighted sum of quadratic planning objectives of the form:

$$\begin{cases} f(\boldsymbol{x}) = \frac{1}{N_s} \sum\limits_{i \in s} \max\{0, d^{\min} - d_i(\boldsymbol{x})\}^2 \text{ or,} \\ f(\boldsymbol{x}) = \frac{1}{N_s} \sum\limits_{i \in s} \max\{0, d_i(\boldsymbol{x}) - d^{\max}\}^2, \end{cases} \tag{10}$$

with N_s the number voxels within the structure.

The dose projection idea from Section II.B. was implemented in the optimization as follows: at each iteration step, the CTD voxels are ranked from low to high, based on their $(d^{\text{presc}} - d)/p$ value. The structure composed of the $(100\% - x\%)$ lowest valued voxels is then included as the target volume in the next iteration step of the optimization.

II.F. Patient cases

The probabilistic planning methods are illustrated by creating treatment plans for two GBM patients previously treated at Massachusetts General Hospital, as shown in Fig 1. GBM tumors are characterized by their fast-growing and highly infiltrative nature in the surrounding brain tissue, with a GTV-to-CTV margin of 2 cm typically used in clinical practice.[21] As a consequence, it can

be challenging to create a treatment plan that both delivers the prescribed dose to the CTV *and* avoids critical structures, without compromising the CTV. These planning challenges are investigated using the example of hippocampal avoidance of the GBM treatments.[17,23,33] Both GBM cases were selected to have the brainstem and at least one hippocampus near the GTV (see Fig. 1). The clinical goals were to spare the brainstem and both hippocampi for Patient 1, while for Patient 2, the goal was to spare the brainstem and contralateral hippocampus. The maximum dose levels for the OARs were set as 12 Gy and 54 Gy for the hippocampus and brainstem respectively. The OAR maximum dose levels were included with a lower weight compared to the weight of the minimum dose objective of the target. For all plans, a ring structure was added around the target volume to encourage dose conformity and control the dose fall-off. The planning objectives are summarized in Table 1

Table 1: Planning objectives included in the optimization.

ROI	**type**	**ref. dose**	**weight**
CTD/CTV	min	60	5
CTD/CTV	max	63	5
Hippocampus	max	12	0.1
Brainstem	max	54	0.1
Ring	max	0	0.01

The clinically-approved manually drawn CTVs closely matched the 10.5 mm and 16 mm expansions for Patients 1 and 2, respectively. Following the ICRU 83 report,[1] the surfaces of the 10.5 mm and 16 mm GTV expansions were assigned the 10%-probability iso-surface for each respective patient. The probabilities of other GTV expansions were chosen to linearly decrease from 100% inside the GTV to 10% at the CTV edge. The 0-probability iso-surface, i.e., the outermost shell, corresponds to the 12 mm and 18 mm expansion contours for Patient 1 and 2, respectively.

The proton beam angles were selected by a certified medical physicist and prescription dose was set to 60 Gy.

III. Results

The goal is to highlight the trade-offs present using the concept of sacrificing voxels, i.e. removing CTD voxels from the minimum dose coverage objectives in order to achieve a notable dose reduction to the hippocampus and brainstem.

III.A. Probabilistic treatment planning method assuming an independent tumor islet model

For both patients, a library of treatment plans was generated where each plan aimed at dosing a fraction of the CTD, from 100% to 90% in steps of 1%. The CTD plans were compared with a reference plan that aims at dosing the CTV with the same dose constraints.

Fig. 4 and 5 compare the dose distributions and DVHs for two representative treatment plans where the goal was to dose 100% and 99% of the CTD, to the CTV plan. The main trade-offs are depicted in Fig. 6. Both conventional DVH metrics (CTD D_{98} and hippocampus D_2) and probabilistic metrics (TCP and expected $D_{\langle 98 \rangle}$) were calculated for the entire set of treatment plans.

In the CTD plans, the dose to the hippocampus can be significantly reduced by discarding 1% of the voxels from the dose coverage objectives (D_2 of 49.9 Gy to 40.4 Gy for ipsilateral hippocampus for Patient 1 and D_2 of 44.0 Gy to 30.5 Gy for the contralateral hippocampus for Patient 2). Discarding up to 5% of voxels voxels has the effect of reducing the hippocampus dose further, to 34.8 Gy and 26.5 Gy for Patient 1 and 2, respectively. In the CTV plan, the hippocampus D_2 dose is similar to dosing 100% of the CTD, i.e., 48.4 Gy and 45.3 Gy for Patient 1 and 2, respectively.

For CTD coverage, the CTD D_{98} decreases from 56.7 Gy to 56.3 Gy for Patient 1 and from 57.1 Gy to 56.5 Gy for Patient 2, after underdosing 1% of CTD voxels. In terms of the expected CTD coverage, the $D_{\langle 98 \rangle}$ is approximately 60 Gy if we aim at dosing 100% of the CTD and remains nearly equal to 60 Gy, even as 10% of CTD voxels are discarded during optimization. Therefore, no significant loss of expected CTD coverage is observed when a fraction of the CTD is underdosed to spare the hippocampus and brainstem. Similarly, the TCP has a maximum value of 97.0% for both patients when we aim at dosing 100% of the CTD. By discarding 1% of CTD voxels, the TCP reduces to only 96.5% and 96.2% for Patient 1 and 2; and down to 96.2% and 95.9% after

discarding 5% of CTD voxels. In general, the approximate TCP follows the same behavior as the exact TCP but systematically overestimates its value.

The CTV plan shows slightly lower CTD dose coverage with a D_{98} dose level of 54.8 Gy and 53.4 Gy for Patient 1 and 2, respectively. The CTD $D_{\langle 98 \rangle}$ is almost identical with 58.4 Gy and 59.7 Gy for Patient 1 and 2, respectively. Similarly, the CTV plan has an expected TCP that is comparable to the maximum value obtained with the CTD plans, i.e., 97.1% and 96.8% for Patient 1 and 2, respectively.

Fig. 7 shows the DVHs in three outer layers with the goal to investigate in which layers the optimizer prefers to discard voxels. Two treatment plans are compared for each patient: (a) aiming at dosing 100% of the CTD, and (b) aiming at dosing 5% of the CTD, that is, 5% of the CTD voxels are discarded. The selected layers are spaced 2 mm apart and have a thickness of 2 mm. As shown, the difference between the DVHs are much larger for inner layers compared to outer layers, indicating that the optimizer underdoses mostly the layers with lower probability voxels.

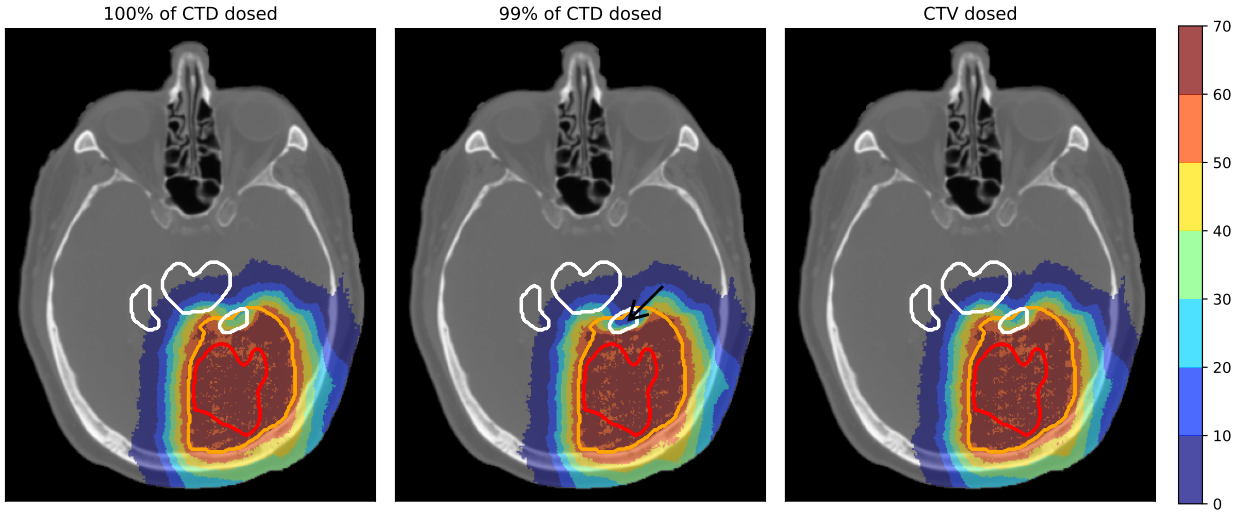

(a) Patient 1.

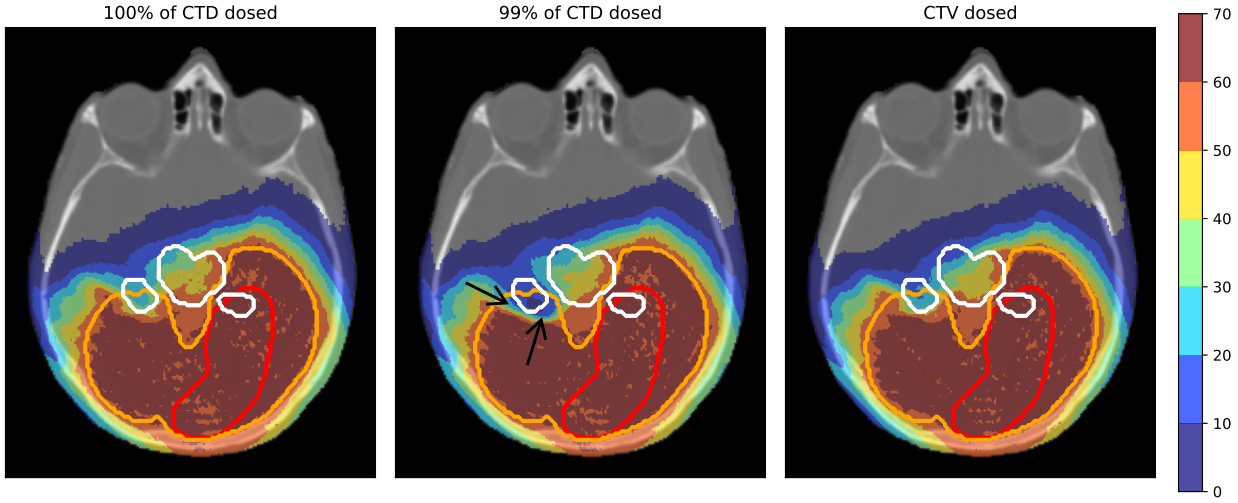

(b) Patient 2.

Figure 4: Dose distributions for treatment plans optimized under the independent voxel assumption: (a) Patient 1, with the clinical goal to spare the ipsilateral hippocampus and brainstem; (b) Patient 2, with the clinical goal to spare the contralateral hippocampus and brainstem. For each patient, three plans are shown where the goal was to dose 100% of the CTD (left), 99% of the CTD (middle), and the CTV (right). The CTD volume is shown in orange. The GTV and OARs are displayed in red and white, respectively. The most significant dose differences are indicated by the black arrows.

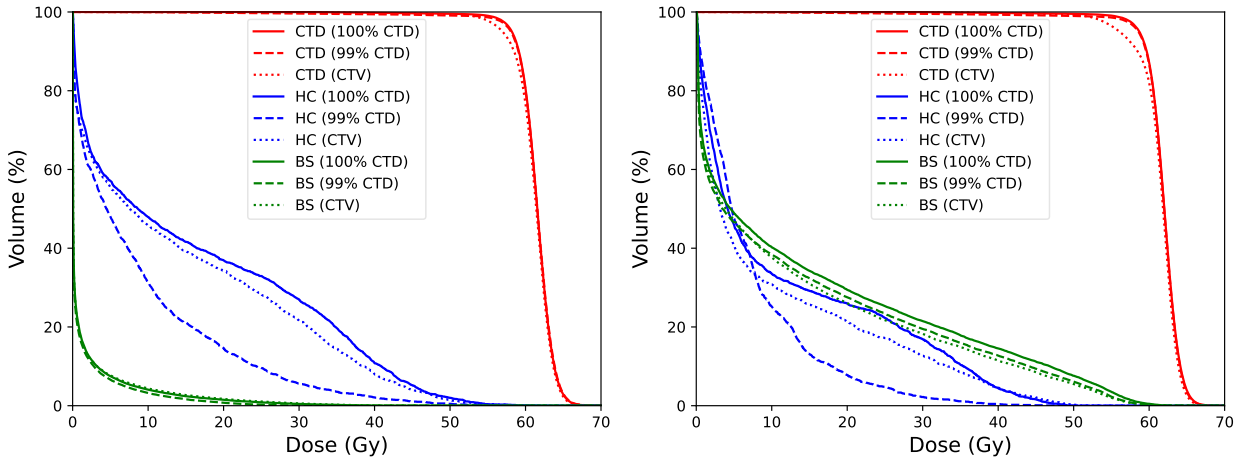

(a) Patient 1. (b) Patient 2.

Figure 5: Dose-volume histograms of the CTD, hippocampus (HC), and brainstem (BS) for three treatment plans as indicated in the parenthesis: dosing 100% of the CTD (solid), dosing 99% of the CTD (dashed), and dosing the CTV (dotted). The DVH of the CTD is computed for the binary volume within the 0-probability iso-surface.

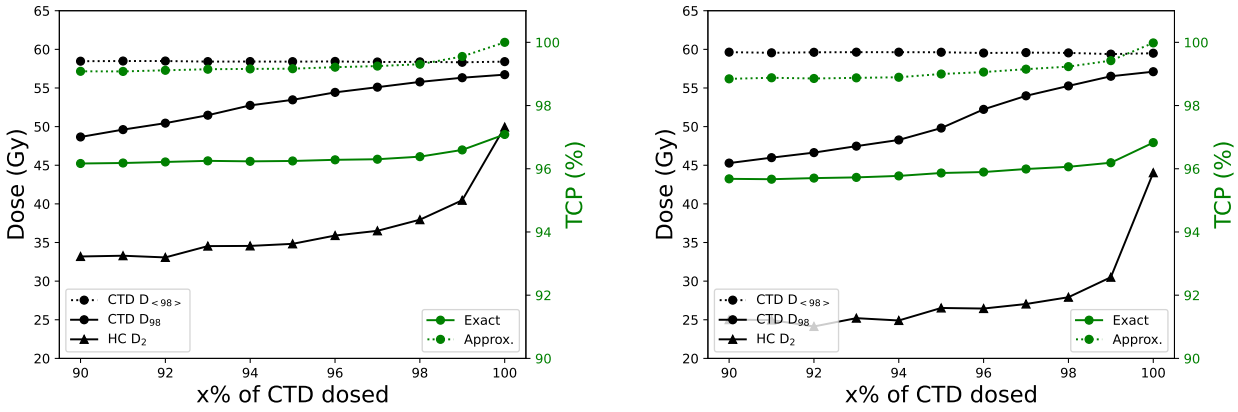

(a) Patient 1. (b) Patient 2.

Figure 6: Conventional and probabilistic evaluation metrics calculated for a set of treatment plans that aim at dosing a fraction of the CTD, from 100% to 90%. Left axis (black): CTD, expected CTD and hippocampus (HC) dose. Right axis (green): exact and approximate TCP in solid and dotted lines, respectively. CTD D_{98} is evaluated with a binary volume while CTD $D_{\langle 98 \rangle}$ represents the expected CTD dose by taking into account tumor probabilities.

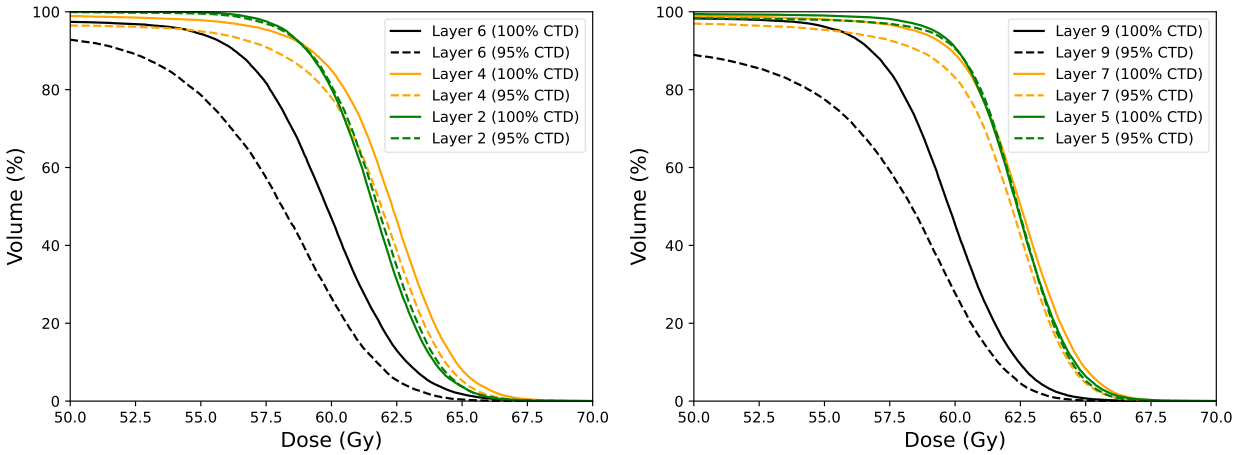

(a) Patient 1. (b) Patient 2.

Figure 7: Dose-volume histograms of three layers of voxels. The treatment plan that aims at dosing 100% of the CTD (solid) is compared to dosing 95% of the CTD (dashed); plans indicated in parenthesis. Left: Layers 6, 4 and 2 correspond to layers in between GTV expansions 12 mm - 10 mm, 8 mm - 6 mm, and 4 mm - 2 mm, respectively. Right: Layers 9, 7, and 5 correspond to layers in between GTV expansions 18 mm - 16 mm, 14 mm - 12 mm, and 8 mm - 6 mm.

III.B. Probabilistic treatment planning method assuming a contiguous circumferential growth model

As before, a library of treatment plans is presented for each patient. Here, the trade-off between CTD coverage and OAR dose is investigated by dosing the voxels within a GTV expansion volume, from no expansion to the maximum GTV expansion bounded by the outermost shell of the CTD, in steps of 2 mm. As a reference, the CTV plan is equal to the treatment plans that aim at dosing 10 mm and 16 mm, for Patient 1 and 2, respectively.

Figs. 8 and 9 show dose distributions and DVHs for two representative treatment plans: a 8 mm GTV expansion and 12 mm GTV expansion for Patient 1 and a 14 mm GTV expansion and 18 mm GTV expansion for Patient 2. Fig. 10 presents conventional DVH metrics (CTD D_{98} and hippocampus D_2) and probabilistic metrics (CTD $D_{\langle 98 \rangle}$ and TCP) as a function of the GTV expansion that should receive the prescribed dose.

The ipsilateral hippocampus dose reduces from its maximum D_2 value of 49.0 Gy to consecutively 48.4, 47.0 and 41.9 when the GTV expansion reduces with 2 mm, 4 mm and 6 mm for

Patient 1. For Patient 2, no D_2 dose reduction is observed for the contralateral hippocampus after reducing the GTV expansion with the first 4 mm. Reducing the GTV expansion with more that 6 mm reduces the hippocampus D_2 dose with only 3.3 Gy. In the CTV plan, the hippocampus D_2 dose is 48.4 Gy and 45.3 Gy for Patient 1 and 2, respectively.

For the same set of treatment plans, the CTD D_{98} reduces from 56.7 Gy to 54.8 Gy for Patient 1 and from 57.0 Gy to 53.4 Gy for Patient 2, after contracting the target volume with 2 mm from the maximum (= CTV plan). Dosing an 8 mm GTV expansion yields a CTD D_{98} of 47.4 Gy for Patient 1 while dosing a 14 mm GTV expansion yields a CTD D_{98} of 47.3 Gy for Patient 2. The expected dose coverage follows the same trend but its decline is less steep: starting from the maximum GTV expansion, the CTD $D_{\langle 98 \rangle}$ is 57.8 Gy and reduces to 57.5 Gy and 54.7 Gy by dosing a 10 mm (= CTV plan) and 8 mm GTV expansion for Patient 1. For Patient 2, the maximum CTD $D_{\langle 98 \rangle}$ is 58.7 Gy and reduces to 57.9 Gy and 54.7 Gy by dosing a 16 mm (= CTV plan) and 14 mm GTV expansion. For the expected TCP, a maximum TCP of 71.4% and 56.9% is achieved for Patient 1 and 2, respectively. Discarding a 2 mm outer layer has the effect of reducing the TCP to 69.6% and 55.4% for each respective patient. Again, the approximate TCP follows the same trend as the exact TCP but systematically overestimates its value.

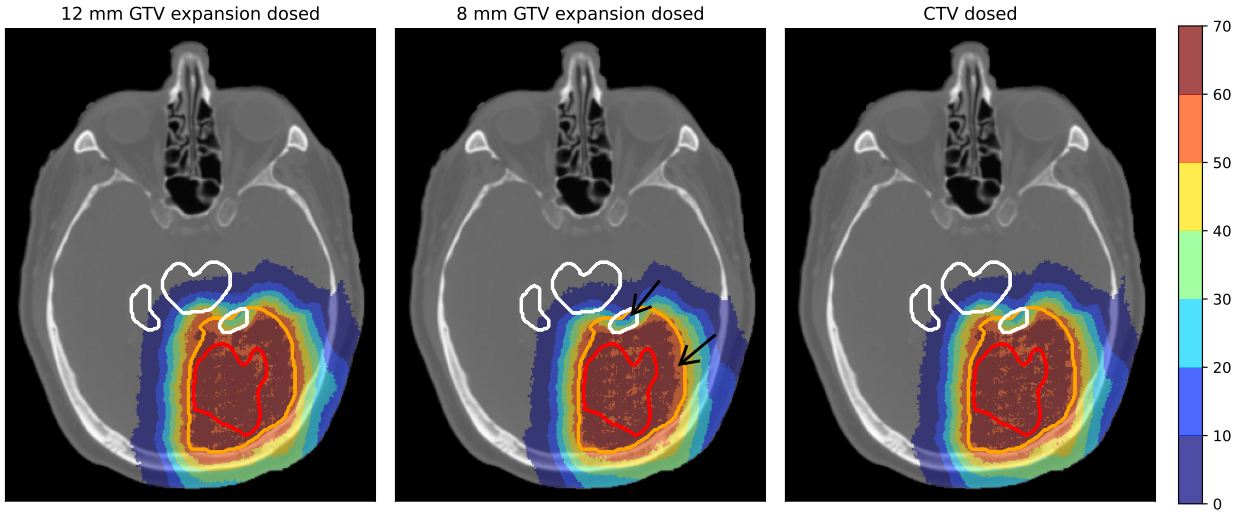

(a) Patient 1.

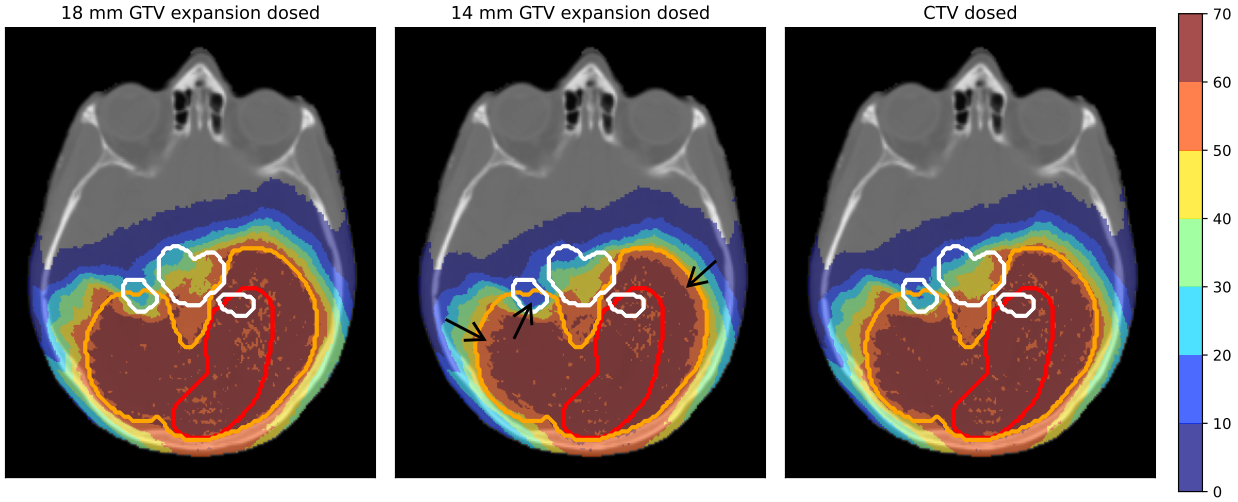

(b) Patient 2.

Figure 8: Dose distributions for treatment plans optimized under the dependent voxel assumption: (a) Patient 1, with the clinical goal to spare the ipsilateral hippocampus; (b) Patient 2, with the clinical goal to spare the contralateral hippocampus. For each patient, two plans are shown, with a 6 mm GTV expansion and 10.5 mm GTV expansion (Patient 1) and a 6 mm GTV expansion and 16 mm GTV expansion (Patient 2). The CTD volume is shown in orange. The GTV and OARs are displayed in red and white, respectively. The dose differences are indicated by the black arrows.

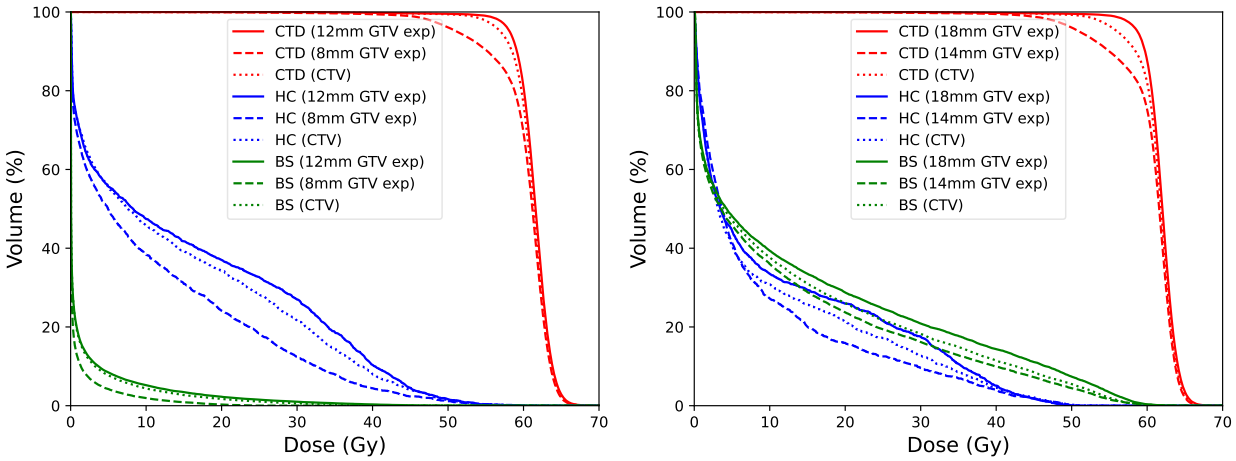

(a) Patient 1. (b) Patient 2.

Figure 9: Dose-volume histograms of the CTD, hippocampus (HC), and brainstem (BS) for three treatment plans as indicated in the parenthesis, a 12 mm GTV expansion (solid) and 8 mm GTV expansion (dashed) for Patient 1 and an 18 mm GTV expansion (solid) and 14 mm GTV expansion (dashed) for Patient 2, compared to CTV plan (dotted) The DVH of the CTD is computed for the binary volume within the 0-probability iso-surface.

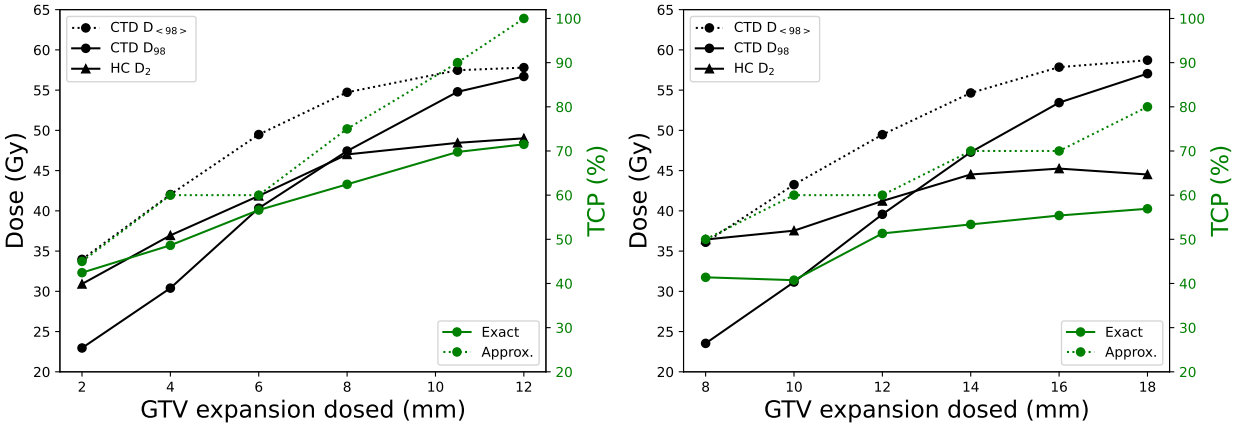

(a) Patient 1. (b) Patient 2.

Figure 10: Conventional and probabilistic evaluation metrics calculated for a set of treatment plans that aim at dosing an increasing expansion of the GTV, in steps of 2 mm. Left axis (black): CTD, expected CTD and hippocampus (HC) dose. Right axis (green): exact and approximate TCP in solid and dotted lines, respectively. CTD $D_{\langle 98 \rangle}$ is evaluated with a binary volume while CTD $D_{\langle 98 \rangle}$ represents the expected CTD dose taking into account tumor probabilities.

IV. Discussion

This study extends the application of the clinical target distribution (CTD) to clinical cases with dose-limiting OARs. We propose two probabilistic planning methods that apply the concept of sacrificing voxels, where CTD voxels are either dosed with the prescription or not dosed at all, to explore the possibility of sparing a nearby OAR.

Our probabilistic planning methods differ from clinical practice where conflicts between dose constraints and CTV coverage are often solved by manually editing the CTV, to mitigate the overlap with a nearby critical OAR. As a result, the treatment plan may suggest acceptable target coverage even though tumor control is limited, given that the CTV was compromised. Yet, manual edits are often necessary as binary target volumes do not discriminate between the relative importance of the voxels, potentially resulting in inadequate trade-offs. In this study, a framework is proposed through which a treatment planner or physician can visualize dose compromises in near real-time. Specifically, a quantitative assessment of the risk involved with underdosing the target to spare a nearby OAR is provided. One can find similar dose compromises in clinical practice, but without reporting of the expected loss of tumor control.

The expected dose coverage was reported with dose-expected-volume histograms (DEVH). The DEVH attempts to address the limitations of shell-based metrics in order to evaluate treatment plans. Limitations include (a) shell DVHs neglect any dose information outside of the shell, even though the probability of tumor may be finite, and (b) the shells do not discriminate between the relative importance of the evaluated voxels. The DEVH, in contrast, takes into account tumor probabilities within the entire volume and summarizes the dose information into a single curve. DEVH takes into account that the target volume is a stochastic quantity, with the probability of the volume realization defined by the CTD. However, the tumor volume realizations do not need to be sampled explicitly and can be incorporated directly in the differential DVH. This DVH formulation stands in contrast to other probabilistic DVHs (see for example dose-coverage histogram in Gordon et al.[14]) that need explicit scenario simulations, evaluated with a fixed region-of-interest.

Based on the probabilistic metrics (TCP and DEVH), a significant difference in behavior is observed between evaluating a treatment plan under the independent tumor islet model or the contiguous circumferential growth model. For the independent tumor islet model, the studied patient cases indicate that the expected CTD dose does not deteriorate from underdosing a small

number of CTD voxels, even though evaluation with a binary volume suggests poor CTD coverage. This effect can be attributed to the fact that for clinical cases (with a large number of voxels), the CTD voxel probabilities are low under voxel independence, 10^{-5}–10^{-6}. Similarly, the estimated TCP loss does not fall below 1% from the maximum, owing to the voxels' low probabilities and the small number of voxels necessary to spare the hippocampus. It must be noticed that the reported TCP values depend on the set of chosen TCP parameters (α, β, and ρ). The planning method, as presented in this study, was calibrated to have a voxel TCP near 100% at the prescription dose level. The results would become more sensitive to dose differences if the model parameters were chosen to have a starting point near the steep region of the TCP curve. However, the general behaviour of the models are not expected to be different with a change of TCP parameters if the TCP dose-response curve remains sufficiently steep.

For the contiguous circumferential growth model, the expected CTD dose and TCP depend more strongly on the number of CTD layers that are dosed. This means that there is less potential for improved OAR sparing (by reducing the extent of the high-dose region), without compromising significantly on tumor control. The TCP could be improved by escalating the dose within the CTD or setting a higher priority to the minimum CTD dose objectives. However, the main aim of this study was to show by which principle we allow voxels of the CTD to be underdosed to achieve a notable dose reduction to an OAR. Notice that this approach fundamentally differs from dose painting studies. In dose painting, the prescription is typically differentiated based on a varying tumor cell density or tumor response.[15,16,29] The underlying assumption is that one knows with absolute certainty which voxels are tumorous, including the ones in the microscopic tumor extension. Introducing a variation in the tumor cell density can lead to fundamentally different dose distributions than when the probability of tumor presence is considered.

We have not addressed the question of which modeling assumption (independent tumor islet or contiguous circumferential growth) should be pursued in practice. Bortfeld et al.[7] demonstrates that the difference between dependent and independent solutions becomes negligible for high dose levels. In contrast, we observe big differences in behavior between both models of tumor propagation, as the high dose limit is no longer valid in the presence of low dose to CTD voxels. The dependence of the underlying assumptions on the optimal probabilistic planning method emphasizes the importance of choosing a realistic assumption. Such a decision primarily depends on the growth pattern of microscopic disease in the surrounding tissues. Tumor growth characteristics can be obtained from histopathological measurements and recurrence analyses after treatment. The

data published on this topic suggests that solid tumor types preferentially spread through microscopic tumor islets[3,10,20] while reaction-diffusion based tumor growth models predict tumor fronts for diffuse tumors such as GBM.[18,28,30] Unfortunately, the availability of data necessary to validate these models remains limited. A practical alternative could be to have a physician or computer algorithm (e.g. trained by diagnostic images) inform the CTD definition and estimate the probability of tumor being present.

The studied GBM cases illustrate that the optimal clinical compromise depends on patient-specific factors such as the proximity of the OAR to the target. Therefore, the results should not be viewed as recommendations, given that they are not generalizable to the broader patient population, nor is this the study's intention. Instead, the GBM cases were selected to exemplify how planning trade-offs can be explored with tumor probability information. Nevertheless, the principle of identifying the sub-volume of the CTD to treat, motivated by non-linear dose-effect relation of TCP models, can be applied to other tumor types and is not limited to GBMs.

V. Conclusion

To the best of our knowledge, this is the first study that proposes and implements probabilistic treatment planning strategies to navigate the trade-off between tumor control and organ sparing using a clinical target distribution (CTD). The potential for OAR sparing is explored interactively by varying the extent to which the CTD is dosed. We show how proposed planning method can be justified under specific statistical models of tumor propagation—an independent tumor islet model or a contiguous circumferential growth model—and evaluate its impact on the estimated level of tumor control.

Acknowledgements

The authors thank Lena Nenoff (Massachusetts General Hospital, Harvard Medical School) for the beam angle selection of the treatment plans. Gregory Buti is supported by the Télévie Grant from the Belgian 'Fonds National pour la Recherche Scientifique' F.R.S-FNRS under grant number 7650220F. Research reported in this publication was supported by the National Cancer Institute of the United States under grant number R01CA266275, and by the Therapy Imaging Program

(TIP) funded by the Federal Share of program income earned by Massachusetts General Hospital on C06CA059267, Proton Therapy Research and Treatment Center. The content is solely the authors' responsibility and does not necessarily represent the official views of the National Institutes of Health.

A Tumor probabilities under the assumption of contiguous circumferential tumor growth

Let $s_h = 1 - r_h$ denote the probability that there is no tumor outside shell h and q_h the probability that a voxel in layer h is not tumorous, $q_h = 1 - p_h$. The probability that there is no tumor outside shell h must be equal to the probability of no tumor in layer h *AND* no tumor outside shell $h+1$:

$$s_h = P((\text{no tumor in } \mathcal{L}_h) \; AND \; (\text{no tumor outside shell } h+1)). \tag{11}$$

In the contiguous circumferential growth model, these events are dependent, therefore,

$$s_h = P(\text{no tumor in } \mathcal{L}_h) \cdot P(\text{no tumor outside shell } h+1 \mid \text{no tumor in } \mathcal{L}_h). \tag{12}$$

Under contiguous growth, the process of tunneling of tumor cells along the layers is not allowed therefore:

$$P(\text{no tumor outside shell } h+1 \mid \text{no tumor in } \mathcal{L}_h) = 1. \tag{13}$$

And for each voxel $i \in \mathcal{L}_h$, the following relationship must hold:

$$\begin{aligned} P(\text{no tumor in } \mathcal{L}_h) = \; & P((\text{no tumor in voxel } i) \; AND \; (\text{no tumor in } \mathcal{L}_h \setminus \text{voxel } i)) \\ = \; & P(\text{no tumor in voxel } i) \\ & \cdot P(\text{no tumor in } \mathcal{L}_h \setminus \text{voxel } i \mid \text{no tumor in voxel } i). \end{aligned} \tag{14}$$

By definition, $P(\text{no tumor in voxel } i) = q_h$. Moreover, under the circumferential growth assumption, all voxels in the same layer are correlated, i.e., $P(\text{no tumor in } \mathcal{L}_h \setminus \text{voxel } i \mid \text{no tumor in voxel } i) = 1$. Therefore, Eq. 14 reduces to:

$$P(\text{no tumor in } \mathcal{L}_h) = q_h. \tag{15}$$

Inserting this in Eq. 12, we find that:

$$s_h = q_h, \tag{16}$$

or,

$$p_h = r_h. \tag{17}$$

B Dose-Expected-Volume Histogram, DEVH

Let κ be a binning of the dose axis, such that $\kappa : \mathbb{R}_0^+ \to \mathbb{N}_0$, $D \to \kappa(D)$ maps the dose value D to the bin number k. We assume the binning is such that $\kappa(0) = 0$ and $\kappa(D_{\max}) = K$.

The normal way to define the dose-volume histogram in a volume (such as the tumor target volume) with N voxels is through

$$h_k = \sum_{i=1}^{N} v_i \delta_{k,\kappa(d_i)}, \tag{18}$$

where v_i is the volume of voxel i, and $\delta_{k,k'}$ is the Kronecker delta symbol. The typical integral form of the dose-volume histogram, DVH, is then obtained by summing the bins from k to K:

$$\mathrm{DVH}_k = \sum_{j=k}^{K} h_j. \tag{19}$$

For probabilistic volumes, we assume that there are M possible realizations of tumor (combinations of voxels that are tumorous) with distinct probabilities. Here, we refer to the realizations of tumor as 'scenarios', where a voxel i can be tumorous in some scenarios and not tumorous in other scenarios. The overall expected-volume histogram is obtained by summing over all M scenarios as follows:

$$\langle h_k \rangle = \frac{1}{M} \sum_{m=1}^{M} \sum_{i=1}^{N} v_i^m \delta_{k,\kappa(d_i)}. \tag{20}$$

The index m stands for the scenario and M is the total number of scenarios. The volume v_i^m equals v_i if and only if voxel i is tumorous in scenario m, otherwise $v_i^m = 0$.

A more compact way of writing the expected-volume histogram is as follows:

$$\langle h_k \rangle = \sum_{i=1}^{N} p_i v_i \delta_{k,\kappa(d_i)}. \tag{21}$$

The integral form of the expected-volume histogram is again obtained by integrating bins from k to K, as above, i.e.:

$$\mathrm{DEVH}_k = \sum_{j=k}^{K} \langle h_j \rangle. \tag{22}$$

For the CTD definition presented in this study, the CTD probabilities (p_i) are assigned *a priori*. Therefore, Eq. 22 can be evaluated directly without the need to perform explicit scenario simulations. Moreover, the DEVH expression remains invariant under the dependent/independent voxel assumption (which gets realized solely by the difference in p_i).

References

[1] The international commission on radiation units and measurements. *Journal of the ICRU*, 10(1):NP.2–NP, April 2010. doi: 10.1093/jicru/ndq001. URL https://doi.org/10.1093/jicru/ndq001.

[2] Markus Alber and Daniela Thorwarth. Multi-modality functional image guided dose escalation in the presence of uncertainties. *Radiotherapy and Oncology*, 111(3):354–359, June 2014. doi: 10.1016/j.radonc.2014.04.016. URL https://doi.org/10.1016/j.radonc.2014.04.016.

[3] Rudi Apolle, Maximilian Rehm, Thomas Bortfeld, Michael Baumann, and Esther G.C. Troost. The clinical target volume in lung, head-and-neck, and esophageal cancer: Lessons from pathological measurement and recurrence analysis. *Clinical and Translational Radiation Oncology*, 3:1–8, April 2017. doi: 10.1016/j.ctro.2017.01.006. URL https://doi.org/10.1016/j.ctro.2017.01.006.

[4] Rudi Apolle, Steffen Appold, Henk P Bijl, Pierre Blanchard, Johan Bussink, Corinne Faivre-Finn, Jonathan Khalifa, Anne Laprie, Yolande Lievens, Indira Madani, et al. Inter-observer variability in target delineation increases during adaptive treatment of head-and-neck and lung cancer. *Acta Oncologica*, 58(10):1378–1385, 2019.

[5] Christoph Baum, Markus Alber, Mattias Birkner, and Fridtjof Nüsslin. Robust treatment planning for intensity modulated radiotherapy of prostate cancer based on coverage probabilities. *Radiotherapy and Oncology*, 78(1):27–35, January 2006. doi: 10.1016/j.radonc.2005.09.005. URL https://doi.org/10.1016/j.radonc.2005.09.005.

[6] Thomas Bortfeld. Clinically relevant intensity modulation opitimization using physical criteria. In *XII International Conference on the Use of Computers in Radiation Therapy, 1997*. Medical Physics Publishing, 1997.

[7] Thomas Bortfeld, Nadya Shusharina, and David Craft. Probabilistic definition of the clinical target volume—implications for tumor control probability modeling and optimization. *Physics in Medicine & Biology*, 66(1):01NT01, January 2021. doi: 10.1088/1361-6560/abcad8. URL https://doi.org/10.1088/1361-6560/abcad8.

[8] A.L. Boyer and J. Unkelbach. Intensity-modulated radiation therapy planning. In *Comprehensive Biomedical Physics*, pages 431–470. Elsevier, 2014. doi: 10.1016/b978-0-444-53632-7.00914-x. URL https://doi.org/10.1016/b978-0-444-53632-7.00914-x.

[9] Gregory Buti, Kevin Souris, Ana Maria Barragán Montero, John Aldo Lee, and Edmond Sterpin. Introducing a probabilistic definition of the target in a robust treatment planning framework. *Physics in Medicine & Biology*, 66(15):155008, July 2021. doi: 10.1088/1361-6560/ac1265. URL https://doi.org/10.1088/1361-6560/ac1265.

[10] Sorcha Campbell, Ian Poon, Dan Markel, Dan Vena, Kevin Higgins, Dan Enepekides, Simon Rapheal, John Wong, Ghassan Allo, Eric Morgen, Nader Khaoum, Ben Smith, Judith Balogh, Robert MacKenzie, Jean Davidson, Dan Wang, and Martin Yaffe. Evaluation of microscopic disease in oral tongue cancer using whole-mount histopathologic techniques: Implications for the management of head-and-neck cancers. *International Journal of Radiation Oncology∗Biology∗Physics*, 82(2):574–581, February 2012. doi: 10.1016/j.ijrobp.2010.09.038. URL https://doi.org/10.1016/j.ijrobp.2010.09.038.

[11] Yair Censor, Thomas Bortfeld, Benjamin Martin, and Alexei Trofimov. A unified approach for inversion problems in intensity-modulated radiation therapy. *Physics in Medicine and Biology*, 51(10):2353–2365, April 2006. doi: 10.1088/0031-9155/51/10/001. URL https://doi.org/10.1088/0031-9155/51/10/001.

[12] Peter Ferjančič, Uulke A van der Heide, Cynthia Ménard, and Robert Jeraj. Probabilistic target definition and planning in patients with prostate cancer. *Physics in Medicine & Biology*, 66 (21):215011, October 2021. doi: 10.1088/1361-6560/ac2f8a. URL https://doi.org/10.1088/1361-6560/ac2f8a.

[13] Claudio Fiorino, Robert Jeraj, Catharine H. Clark, Cristina Garibaldi, Dietmar Georg, Ludvig Muren, Wouter van Elmpt, Thomas Bortfeld, and Nuria Jornet. Grand challenges for medical physics in radiation oncology. *Radiotherapy and Oncology*, 153:7–14, December 2020. doi: 10.1016/j.radonc.2020.10.001. URL https://doi.org/10.1016/j.radonc.2020.10.001.

[14] J. J. Gordon, N. Sayah, E. Weiss, and J. V. Siebers. Coverage optimized planning: Probabilistic treatment planning based on dose coverage histogram criteria. *Medical Physics*, 37(2):550–563, January 2010. doi: 10.1118/1.3273063. URL https://doi.org/10.1118/1.3273063.

[15] Eric Grönlund, Silvia Johansson, Anders Montelius, and Anders Ahnesjö. Dose painting by numbers based on retrospectively determined recurrence probabilities. *Radiotherapy and Oncology*, 122(2):236–241, February 2017. doi: 10.1016/j.radonc.2016.09.007. URL https://doi.org/10.1016/j.radonc.2016.09.007.

[16] Jian-Yue Jin, Feng-Ming Kong, Dezhi Liu, Lei Ren, Haisen Li, Hualiang Zhong, Benjamin Movsas, and Indrin J. Chetty. A TCP model incorporating setup uncertainty and tumor cell density variation in microscopic extension to guide treatment planning. *Medical Physics*, 38 (1):439–448, December 2010. doi: 10.1118/1.3531543. URL https://doi.org/10.1118/1.3531543.

[17] Tomas Kazda, Radim Jancalek, Petr Pospisil, Ondrej Sevela, Tomas Prochazka, Miroslav Vrzal, Petr Burkon, Marek Slavik, Ludmila Hynkova, Pavel Slampa, and Nadia N Laack. Why and how to spare the hippocampus during brain radiotherapy: the developing role of hippocampal avoidance in cranial radiotherapy. *Radiation Oncology*, 9(1), June 2014. doi: 10.1186/1748-717x-9-139. URL https://doi.org/10.1186/1748-717x-9-139.

[18] E. Konukoglu, O. Clatz, B.H. Menze, B. Stieltjes, M.-A. Weber, E. Mandonnet, H. Delingette, and N. Ayache. Image guided personalization of reaction-diffusion type tumor growth models using modified anisotropic eikonal equations. *IEEE Transactions on Medical Imaging*, 29(1): 77–95, January 2010. doi: 10.1109/tmi.2009.2026413. URL https://doi.org/10.1109/tmi.2009.2026413.

[19] Dong C. Liu and Jorge Nocedal. On the limited memory BFGS method for large scale optimization. *Mathematical Programming*, 45(1-3):503–528, August 1989. doi: 10.1007/bf01589116. URL https://doi.org/10.1007/bf01589116.

[20] Xue Meng, Xindong Sun, Dianbin Mu, Ligang Xing, Li Ma, Baijiang Zhang, Shuqiang Zhao, Guoren Yang, Feng-Ming (Spring) Kong, and Jinming Yu. Noninvasive evaluation of microscopic tumor extensions using standardized uptake value and metabolic tumor volume in non-small-cell lung cancer. *International Journal of Radiation Oncology∗Biology∗Physics*, 82 (2):960–966, February 2012. doi: 10.1016/j.ijrobp.2010.10.064. URL https://doi.org/10.1016/j.ijrobp.2010.10.064.

[21] Maximilian Niyazi, Michael Brada, Anthony J. Chalmers, Stephanie E. Combs, Sara C. Erridge, Alba Fiorentino, Anca L. Grosu, Frank J. Lagerwaard, Giuseppe Minniti, René-Olivier Mir-

imanoff, Umberto Ricardi, Susan C. Short, Damien C. Weber, and Claus Belka. ESTRO-ACROP guideline "target delineation of glioblastomas". *Radiotherapy and Oncology*, 118(1): 35–42, January 2016. doi: 10.1016/j.radonc.2015.12.003. URL https://doi.org/10.1016/j.radonc.2015.12.003.

[22] Jorge Nocedal. Updating quasi-newton matrices with limited storage. *Mathematics of Computation*, 35(151):773–782, 1980. doi: 10.1090/s0025-5718-1980-0572855-7. URL https://doi.org/10.1090/s0025-5718-1980-0572855-7.

[23] Silvia Scoccianti, Beatrice Detti, Davide Gadda, Daniela Greto, Ilaria Furfaro, Fiammetta Meacci, Gabriele Simontacchi, Lucia Di Brina, Pierluigi Bonomo, Irene Giacomelli, Icro Meattini, Monica Mangoni, Sabrina Cappelli, Sara Cassani, Cinzia Talamonti, Lorenzo Bordi, and Lorenzo Livi. Organs at risk in the brain and their dose-constraints in adults and in children: A radiation oncologist's guide for delineation in everyday practice. *Radiotherapy and Oncology*, 114(2):230–238, February 2015. doi: 10.1016/j.radonc.2015.01.016. URL https://doi.org/10.1016/j.radonc.2015.01.016.

[24] Nadya Shusharina, David Craft, Yen-Lin Chen, Helen Shih, and Thomas Bortfeld. The clinical target distribution: a probabilistic alternative to the clinical target volume. *Physics in Medicine & Biology*, 63(15):155001, July 2018. doi: 10.1088/1361-6560/aacfb4. URL https://doi.org/10.1088/1361-6560/aacfb4.

[25] Nadya Shusharina, Jonas Söderberg, David Edmunds, Fredrik Löfman, Helen Shih, and Thomas Bortfeld. Automated delineation of the clinical target volume using anatomically constrained 3d expansion of the gross tumor volume. *Radiotherapy and Oncology*, 146:37–43, May 2020. doi: 10.1016/j.radonc.2020.01.028. URL https://doi.org/10.1016/j.radonc.2020.01.028.

[26] Kevin Souris. Mcsquare - http://www.openmcsquare.org/ accessed june 2019.

[27] Kevin Souris, John Aldo Lee, and Edmond Sterpin. Fast multipurpose monte carlo simulation for proton therapy using multi- and many-core CPU architectures. *Medical Physics*, 43(4):1700–1712, March 2016. doi: 10.1118/1.4943377. URL https://doi.org/10.1118/1.4943377.

[28] Kristin R. Swanson, Carly Bridge, J.D. Murray, and Ellsworth C. Alvord. Virtual and real brain tumors: using mathematical modeling to quantify glioma growth and invasion. *Journal*

of the Neurological Sciences, 216(1):1–10, December 2003. doi: 10.1016/j.jns.2003.06.001. URL https://doi.org/10.1016/j.jns.2003.06.001.

[29] Jan Unkelbach, Bjoern H Menze, Ender Konukoglu, Florian Dittmann, Nicholas Ayache, and Helen A Shih. Radiotherapy planning for glioblastoma based on a tumor growth model: implications for spatial dose redistribution. *Physics in Medicine and Biology*, 59(3):771–789, January 2014. doi: 10.1088/0031-9155/59/3/771. URL https://doi.org/10.1088/0031-9155/59/3/771.

[30] Jan Unkelbach, Bjoern H Menze, Ender Konukoglu, Florian Dittmann, Matthieu Le, Nicholas Ayache, and Helen A Shih. Radiotherapy planning for glioblastoma based on a tumor growth model: improving target volume delineation. *Physics in Medicine and Biology*, 59(3):747–770, January 2014. doi: 10.1088/0031-9155/59/3/747. URL https://doi.org/10.1088/0031-9155/59/3/747.

[31] Jan Unkelbach, Thomas Bortfeld, Carlos E. Cardenas, Vincent Gregoire, Wille Hager, Ben Heijmen, Robert Jeraj, Stine S. Korreman, Roman Ludwig, Bertrand Pouymayou, Nadya Shusharina, Jonas Söderberg, Iuliana Toma-Dasu, Esther G.C. Troost, and Eliana Vasquez Osorio. The role of computational methods for automating and improving clinical target volume definition. *Radiotherapy and Oncology*, 153:15–25, December 2020. doi: 10.1016/j.radonc.2020.10.002. URL https://doi.org/10.1016/j.radonc.2020.10.002.

[32] Shalini K. Vinod, Michael G. Jameson, Myo Min, and Lois C. Holloway. Uncertainties in volume delineation in radiation oncology: A systematic review and recommendations for future studies. *Radiotherapy and Oncology*, 121(2):169–179, November 2016. doi: 10.1016/j.radonc.2016.09.009. URL https://doi.org/10.1016/j.radonc.2016.09.009.

[33] Chan Woo Wee, Kyung Su Kim, Chae-Yong Kim, Jung Ho Han, Yu Jung Kim, and In Ah Kim. Feasibility of hippocampus-sparing VMAT for newly diagnosed glioblastoma treated by chemoradiation: pattern of failure analysis. *Radiation Oncology*, 15(1), May 2020. doi: 10.1186/s13014-020-01552-0. URL https://doi.org/10.1186/s13014-020-01552-0.