Robust PET-guided intensity-modulated radiation therapy

H. Li[a)]
*Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road,
Toronto, Ontario M5S 3G8, Canada*

J. P. Bissonnette and T. Purdie
*Radiation Oncology, Princess Margaret Cancer Center, 610 University Avenue,
Toronto, Ontario M5G 2M9, Canada and Techna Institute for the Advancement of Technology for Health,
124 - 100 College Street, Toronto, Ontario M5G 1P5, Canada*

T. C. Y. Chan
*Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road,
Toronto, Ontario M5S 3G8, Canada and Techna Institute for the Advancement of Technology for Health,
124 - 100 College Street, Toronto, Ontario M5G 1P5, Canada*

Purpose: Functional image guided intensity-modulated radiation therapy has the potential to improve cancer treatment quality by basing treatment parameters such as heterogeneous dose distributions information derived from imaging. However, such heterogeneous dose distributions are subject to imaging uncertainty. In this paper, the authors develop a robust optimization model to design plans that are desensitized to imaging uncertainty.

Methods: Starting from the pretreatment fluorodeoxyglucose-positron emission tomography scans, the authors use the raw voxel standard uptake values (SUVs) as input into a series of intermediate functions to transform the SUV into a desired dose. The calculated desired doses were used as an input into a robust optimization model to generate beamlet intensities. For each voxel, the authors assume that the true SUV cannot be observed but instead resides in an interval centered on the nominal (i.e., observed) SUV. Then the authors evaluated the nominal and robust solutions through a simulation study. The simulation considered the effect of the true SUV being different from the nominal SUV on the quality of the treatment plan. Treatment plans were compared on the metrics of objective function value and tumor control probability (TCP).

Results: Computational results demonstrate the potential for improvements in tumor control probability and deviation from the desired dose distribution compared to a nonrobust model while maintaining acceptable tissue dose.

Conclusions: Robust optimization can help design treatment plans that are more stable in the presence of image value uncertainties. © *2015 American Association of Physicists in Medicine.*
[http://dx.doi.org/10.1118/1.4926845]

Key words: robust, optimization, PET, biological, IMRT

1. INTRODUCTION

Current clinical IMRT treatment planning is largely based on computed tomography (CT) imaging, which provides geometric information about the tumor and surrounding organs-at-risk (OARs). Treatment planning using CT imaging implicitly assumes that the tumor is biologically homogeneous and aims to deliver a uniform dose[1] as this results in the optimal tumor control probability (TCP).[2,3]

Functional imaging such as positron emission tomography (PET) can provide insight into the heterogeneity of the tumor.[4] This heterogeneity of the tumor requires a nonuniform dose distribution in order to maximize TCP.[5,3] Nonuniform dose distributions are also being explored in a large multicenter clinical trial (RTOG1106). How to best utilize PET information to determine a heterogeneous dose distribution remains an open question. A number of researchers have proposed a tracer-independent linear relationship to transform the PET signal from a biological image to a desired heterogeneous dose distribution,[1,6–9] but it has also been suggested that a linear transformation may lead to unnecessary overestimations of the desired dose.[10] Yang and Xing[11] calculated heterogeneous dose distributions required to maximize TCP given voxel-specific radiobiological parameters under the linear–quadratic (LQ) model. South *et al.*[12] developed a theoretical framework to derive heterogeneous dose distributions based on functional imaging. A follow-up study by South *et al.*[13] applied this theoretical framework to PET imaging using fluorodeoxyglucose (FDG) as a tracer to ultimately describe a heterogeneous distribution reflecting radiosensitivity. We follow a similar framework as South *et al.*[13] in this study to translate the PET image map to a desired heterogeneous dose distribution.

Boellaard[14] presents a comprehensive list of uncertainties affecting PET images, including the patient's blood glucose level, patient motion, inflammation, uptake period, scan acquisition parameters, image reconstruction parameters, region

of interest, and blood glucose level correction. These uncertainties, in turn, affect our understanding of the true biological activity. Thus, any treatment planning paradigm that uses PET imaging to derive a heterogeneous dose distribution must consider the inherent uncertainty in the PET signal.

A few studies have considered incorporating uncertainty into a biologically based treatment planning process. Kåver *et al.*[15] compared the use of stochastic optimization and margins on the biological parameters to maximize the expected probability of uncomplicated treatment. Witte *et al.*[16] addressed geometric uncertainties when using heterogeneous dose distributions in a probabilistic optimization model, where random and systematic (positional) errors were approximated using isotropic Gaussian kernels.

Robust optimization methods form another class of methods that aim to mitigate the effects of uncertainty in treatment planning. Such methods have been applied to nonbiological IMRT treatment planning for interfraction uncertainty (e.g., setup error)[17,18] and intrafraction uncertainty (e.g., tumor motion).[19–26] Robust optimization methods have also been developed for intensity-modulated proton therapy.[27–30]

In this paper, we develop the first robust optimization approach for PET-based treatment planning. We assume that the PET signal represents tumor heterogeneity, thus requiring a voxel-specific heterogeneous dose distribution. However, uncertainty in the PET signal can affect the overall treatment quality; if the true signal is higher than measured, we may be underdosing the tumor and if it is lower than measured, we may miss out on an opportunity for improved sparing. Our approach accounts for uncertainty in the PET signal using a cardinality-based robust optimization model.[31] Our model is based on dose and does not optimize TCP directly. TCP is measured after the optimization models are solved. We apply our framework to a clinical lung case with PET information from Princess Margaret Cancer Centre in Toronto, Canada. We do not explicitly consider the effect of motion blurring on the dose-influence matrix, but the effects of motion on the PET signal can be implicitly captured in our model.

2. METHODS AND MATERIALS

We assume that a voxel-specific standardized uptake value (SUV) obtained from a FDG-PET image is used to generate a heterogeneous desired dose. We aim to achieve this distribution using both nonrobust and robust optimization methods. The two methods are then compared through a simulation study considering changes in SUV.

FDG uptake is hypothesized to be affected by a number of biological processes including radioresistance,[32] proliferation,[33] cell density,[34] and hypoxia.[35–37] Hotspots of FDG have also been found to be associated with areas of recurrence, which affects overall survival.[32,38] The optimization model we present is tracer-independent and only requires a voxel-specific desired dose distribution. The relationship we assume between FDG and desired dose is given in the Appendix. Note that as long as the values of the tracer observed in the image can be converted to a desired dose distribution, it can be used in our framework.

2.A. Model of SUV uncertainty

We assume that the true SUV is not directly observed due to uncertainty in the image.[14] We define θ_i to be the nominal (i.e., observed) SUV in voxel i of the PET image. The true (i.e., unobserved) SUV for voxel i, $\tilde{\theta}_i$ resides in an interval $\mathcal{U}_i = [\theta_i - \hat{\theta}_i, \theta_i + \hat{\theta}_i]$, where $\hat{\theta}_i$ is the maximum absolute deviation from the nominal SUV for voxel i.

The values θ_i and $\hat{\theta}_i$ will be used to determine a dose distribution for treatment planning. Since the true SUV is unknown, the dose that should have been given may be different from the nominal dose distribution, which is derived from the observed SUV. Thus, if we design a treatment using the nominal heterogeneous dose distribution, we may underdose or overdose certain voxels with respect to the dose that should have been given.

2.B. Mathematical formulation

We develop a robust optimization framework for treatment planning in the presence of PET signal uncertainty based on the "budget of protection" model described in Ref. 31. Our goal is to demonstrate a proof of concept for the value of robustness in the presence of PET signal uncertainty. As such, we use a simple penalty-based linear model that approximates a treatment planning formulation but omits many of the more sophisticated features of clinical formulations for the sake of simplicity.

Let i, j, and k index voxels, beamlets, and structures, respectively. Let $\mathcal{I}_k$ be the set of all voxels in structure k, O be the index set of all OARs, $\mathcal{I}_\mathcal{T}$ be the set of all tumor voxels, $\mathcal{I}_O$ be the set of all OAR voxels, and $\mathcal{J}$ be the set of all beamlets. Let y_i and z measure the underdose to voxel i and maximum underdose to the clinical target volume (CTV), respectively. Let w_j be the intensity of beamlet j. Let D_{ij} be the influence matrix describing the dose from unit intensity of beamlet j to voxel i. Let $f(\theta_i)$ be the desired dose to voxel i, given SUV θ_i. Let U_i and L_i be the upper and lower bounds on the dose to voxel i, respectively. Let μ_k be an upper bound on the mean dose to OAR k. We will refer to Formulation (1) as the nominal (no uncertainty) model. A solution to this formulation is called the *nominal solution* or *nominal plan*,

$$\underset{w,y,z}{\text{minimize}} \quad \lambda^- \sum_{i \in \mathcal{I}_\mathcal{T}} y_i + \lambda z + \lambda^O \sum_{i \in \mathcal{I}_O} \sum_{j \in \mathcal{J}} D_{ij} w_j \tag{1a}$$

$$\text{subject to} \quad -y_i \leq \sum_{j \in \mathcal{J}} D_{ij} w_j - f(\theta_i), \quad \forall i \in \mathcal{I}_\mathcal{T}, \tag{1b}$$

$$z \geq y_i, \quad \forall i \in \mathcal{I}_\mathcal{T}, \tag{1c}$$

$$L_i \leq \sum_{j \in \mathcal{J}} D_{ij} w_j \leq U_i, \quad \forall i \in \mathcal{I}_O \cup \mathcal{I}_\mathcal{T}, \tag{1d}$$

$$\frac{1}{|\mathcal{I}_k|} \sum_{i \in \mathcal{I}_k} \sum_{j \in \mathcal{J}} D_{ij} \leq \mu_k, \quad \forall k \in O, \tag{1e}$$

$$y_i, w_j \geq 0, \quad \forall i \in \mathcal{I}_\mathcal{T}, j \in \mathcal{J}. \tag{1f}$$

Objective (1a) minimizes a weighted combination of the total absolute deviation (underdose) from the voxel-specific desired dose for all target voxels and the total dose to the

OARs. Parameters $\lambda^-, \lambda, \lambda^O$ are the weights for penalizing underdose, maximum underdose, and OAR dose, respectively. Constraint (1b) models the underdose computation. Constraint (1c) models the maximum underdose computation. Constraint (1d) bounds the upper and lower doses to every OAR voxel. Constraint (1e) limits the mean dose for each OAR. Since Formulation (1) assumes no uncertainty in the underlying radiobiological parameters, a treatment that is able to deliver at least the desired dose to all CTV voxels will generate a voxel TCP of 0.999 99, according to our TCP model in the Appendix.

The robust formulation accounts for SUV uncertainty and builds on Formulation (1). First, let $f(\tilde{\theta}_i)$ be the desired dose for voxel i associated with the uncertain SUV $\tilde{\theta}_i$, which we use in place of $f(\theta_i)$ in constraint (1b). Second, to facilitate the comparison of our robust model with the nominal model, we require the robust model to deliver a mean dose to the tumor that is the same as what is delivered in the nominal model. Thus, after solving the nominal model, we calculate the mean tumor dose, which we define as Ω, and include a constraint which holds the mean CTV dose delivered by the robust solution equal to Ω. This ensures an unbiased comparison between different models because a simple boost dose to all voxels in the tumor will result in an increase in TCP. By holding the mean dose constant, we know that the improvement in TCP is a result of the redistribution of dose rather than a boost dose. Of course, a boost dose can be included to further improve TCP if possible. These two changes to the nominal model are implemented by eliminating constraint (1b) in Formulation (1) and adding the following constraints:

$$-y_i \le \sum_{j \in \mathcal{J}} D_{ij} w_j - f(\tilde{\theta}_i), \quad \forall \tilde{\theta}_i \in \mathcal{U}_i, i \in \mathcal{I}_{\mathcal{T}}, \tag{2a}$$

$$\frac{1}{|\mathcal{I}_{\mathcal{T}}|} \sum_{i \in \mathcal{I}_{\mathcal{T}}} \sum_{j \in \mathcal{J}} D_{ij} w_j = \Omega. \tag{2b}$$

Since $\mathcal{U}_i$ is an interval, a formulation including constraint (2a) is not directly solvable as a linear program. However, we can easily reformulate it into a tractable linear program.

Our model assumes that each voxel's SUV may deviate from its nominal value; however, the total number of voxels that will deviate (and conspire to produce a worst-case effect) is bounded. We introduce a parameter $\gamma_i \in [0,1]$ for each constraint involving $\mathcal{U}_i$ to model the extent of the SUV change of voxel i. For example, $\gamma_i = 0$ models the belief that voxel i will remain at its nominal SUV θ_i and therefore only require the corresponding nominal desired dose, while $\gamma_i = 1$ models the belief that voxel i will realize its worst-case (largest) SUV $\theta_i + \hat{\theta}_i$ and therefore require a higher dose to achieve the same TCP. Formulation (3) shows the linear robust formulation. A solution to this formulation is called the *robust solution* or *robust plan*,

$$\begin{aligned} \underset{w,y,z}{\text{minimize}} \quad & \lambda^- \sum_{i \in \mathcal{I}_{\mathcal{T}}} y_i + \lambda z + \lambda^O \sum_{i \in \mathcal{I}_O} \sum_{j \in \mathcal{J}} D_{ij} w_j \\ \text{subject to} \quad & -\sum_{j \in \mathcal{J}} D_{ij} w_j + f(\theta_i) + \gamma_i \big(f(\theta_i + \hat{\theta}_i) - f(\theta_i)\big) \le y_i, \\ & \forall i \in \mathcal{I}_{\mathcal{T}}, \quad z \ge y_i, \quad \forall i \in \mathcal{I}_{\mathcal{T}}, \end{aligned}$$

$$L_i \le \sum_{j \in \mathcal{J}} D_{ij} w_j \le U_i, \quad \forall i \in \mathcal{I}_O \cup \mathcal{I}_{\mathcal{T}},$$

$$\frac{1}{|\mathcal{I}_k|} \sum_{i \in \mathcal{I}_k} \sum_{j \in \mathcal{J}} D_{ij} \le \mu_k, \quad \forall k \in O,$$

$$\frac{1}{|\mathcal{I}_{\mathcal{T}}|} \sum_{i \in \mathcal{I}_{\mathcal{T}}} \sum_{j \in \mathcal{J}} D_{ij} w_j = \Omega,$$

$$y_i, w_j \ge 0, \quad \forall i \in \mathcal{I}_{\mathcal{T}}, j \in \mathcal{J}. \tag{3}$$

If we allow γ_i to be decision variables that can be optimized simultaneously with the beamlet intensities w_j, the optimization engine will chose values for γ_i in an optimistic (i.e., not worst-case) manner. Instead, to ensure that Formulation (3) is protecting against a worst-case realization of the SUV, we formulate auxiliary optimization problem (4) that identifies the worst-case combination of voxels to change SUV and chooses the γ_i values accordingly within an overall budget of Γ. Model (4) only considers the worst-case underdose as that is more important in terms of TCP robustness,

$$\begin{aligned} \underset{\gamma}{\text{maximize}} \quad & \sum_{i \in \mathcal{I}_{\mathcal{T}}} \gamma_i \big(f(\theta_i + \hat{\theta}_i) - f(\theta_i)\big) \\ \text{subject to} \quad & \sum_{i \in \mathcal{I}_{\mathcal{T}}} \gamma_i \le \Gamma, \\ & 0 \le \gamma_i \le 1, \quad \forall i \in \mathcal{I}_{\mathcal{T}}. \end{aligned} \tag{4}$$

The value Γ can be interpreted as the maximum number of voxels that we expect to change their SUV to their worst-case values. Therefore, Γ need not be larger than $\Gamma_{\max} := |\mathcal{I}_{\mathcal{T}}|$ (the number of tumor voxels). Choosing $\Gamma = 0$ forces all $\gamma_i = 0$, and robust model (3) reduces to nominal model (1). Choosing $\Gamma = |\mathcal{I}_{\mathcal{T}}|$ results in the most conservative model, where we assume all tumor voxels can change their SUV to their worst-case values. Given a particular Γ, auxiliary problem (4) is solved first to find the optimal γ_i^* values, which are then used as input to Formulation (3). Note that Formulation (4) is an instance of the continuous knapsack problem, for which it is known that the greedy solution is optimal. That is, if the values of $f(\theta_i + \hat{\theta}_i) - f(\theta_i)$ are ranked in descending order, the optimal solution is $\gamma_i^* = 1$ for $i = 1, \ldots, \lfloor \Gamma \rfloor$ and $\gamma_{\lceil \Gamma \rceil}^* = \Gamma - \lfloor \Gamma \rfloor$ (in the case Γ is not integer). Hence, the effort required to solve Formulation (4) is minimal. Overall, the parameter Γ provides flexibility and allows a treatment planner to adjust his or her level of conservatism when designing the robust treatment.

2.C. Patient data and optimization parameters

The exhale phase of a combined 4D FDG-PET/CT dataset was exported into the computational environment for radiotherapy research (CERR), in order to generate the dose influence matrices necessary for optimization.[39] The 4D data alleviate blur due to respiratory motion. The SUV of each voxel from the PET image was converted into a voxel-specific desired dose distribution as described in the Appendix.

The penalty parameters were set as follows: $\lambda^- = 1000$, $\lambda = 100\,000$, $\lambda^O = 1$. Mean and maximum dose constraints were obtained from the QUANTEC series of papers[40] and clinical protocols at Princess Margaret Cancer Center (see Table II).

For the target, we set an upper limit of $1.4\times$ the desired dose on each voxel.

For the robust formulation, we set $\Omega = 83.25$ Gy, which is the mean tumor dose of the solution to the nominal formulation. We employed a coplanar equispaced 7-beam configuration for the two optimization models. We assumed a beamlet size of 1×1 cm. The models were solved using CPLEX 12 on a remote computing cluster using a node with a 2.27 Ghz Intel Xeon 10 core processor and 252 GB of RAM. All models solved in under one min.

The original CT image had a resolution of $0.0977\times0.0977 \times 0.2$ cm. The original PET scan had a resolution of $0.3906 \times 0.3906 \times 0.3270$ cm. The SUV of the PET voxels was mapped to their closest corresponding voxel in the CT image. Then the CT scan was downsampled to obtain a final planning resolution of $0.3906\times0.3906\times0.2$ cm. This resulted in 7,372 voxels in the CTV and 207 938 voxels in the OARs. Since the desired dose distribution is generated to obtain a voxel TCP of 0.999 99, a plan that is able to deliver at least the desired dose to all CTV voxels will obtain an overall CTV TCP of 0.93 ($0.999\,99^{7372}$).

2.D. Simulating SUV changes to evaluate optimization results

Once we determine an optimal fluence map from the nominal and robust formulations, we evaluate the dosimetric consequences of the realization of different SUVs. We present two approaches to simulate SUV changes, depending on whether the voxels are uncorrelated or correlated.

2.D.1. Uncorrelated voxels

The first simulation assumes that randomly chosen voxels will realize their worst-case SUV (the maximum value $\theta_i + \hat{\theta}_i$ in the interval). This will result in underdosing if we only provided the nominal dose to the voxels. Boellaard[14] provides the typical ranges and maximum effect of many factors that can affect a PET signal, and although they may all affect the image, some of their interactions may cancel out. In the absence of more detailed data about the statistical distribution of each factor's effect, we set $\hat{\theta}_i$ to $0.6\theta_i$ as a conservative estimate of the worst-case effect. We simulated 2000 realizations for this scenario, each representing a random set of voxels that realize a change to their worst-case SUVs.

2.D.2. Correlated voxels

The second method to simulate SUV changes assumes that voxels in close proximity will be correlated and thus their SUV deviations will move together. We consider the CTV to be composed of seven nested shells each of 4 mm thickness. We assume the outermost shell will have the highest uncertainty due to its proximity to neighboring nontarget structures. The uncertainty decreases as we move from the outermost shell to the innermost shell. Table I outlines the size of the uncertainty half-interval (i.e., $\hat{\theta}_i$) for each shell. We only considered voxels potentially increasing in SUV. The simulation chooses all

TABLE I. Size of the uncertainty half-interval for each shells in the correlated voxels simulation. Shell 1 is the outermost shell and shell 7 is the innermost shell.

	Shell						
	1	2	3	4	5	6	7
Half-interval size	$0.3\theta_i$	$0.25\theta_i$	$0.2\theta_i$	$0.15\theta_i$	$0.1\theta_i$	$0.05\theta_i$	$0.05\theta_i$

voxels to realize a positive change from their nominal SUV, thus representing a correlated effect. We investigate the effects if all voxels realize a positive change to a fraction of their maximum SUV within the specified range. Specifically, all voxels will realize a SUV of $\theta_i + \delta\hat{\theta}_i$ for various $\delta \in [0,1]$. Note that this experiment represents any source of uncertainty that causes the voxel SUVs to move in the same direction.

3. RESULTS

3.A. Uncorrelated voxels

First, we consider the impact of uncertainty and protection level Γ on the objective function values (i.e., a weighted measure of deviation from the desired dose distribution) and TCP. Recall that $\Gamma = 0$ is equivalent to the nominal plan, while positive Γ values correspond to robust solutions. Figure 1 depict histograms of the objective function and TCP values, where the relative frequencies are generated from simulating random sets of voxels to experience a change from their nominal SUV to their worst-case value.

Figure 1 shows that the nominal plan generates worse objective function values and has more variability in these values, compared to the robust plans. A higher objective function value means the delivered dose distribution deviates more from the desired dose distribution. Increased spread in the histogram means that the performance of the nominal plan is much more variable in the presence of uncertainty. Note for $\Gamma > 0.4\Gamma_{max}$, the resulting histogram is essentially the same as the one for $\Gamma = 0.4\Gamma_{max}$. The results for TCP calculated on the CTV are similar. The robust plans generally have higher TCP values and less variability in TCP especially under the worst-case uncertainty scenarios. Again, negligible gains in TCP are realized when Γ is increased above $0.4\Gamma_{max}$.

Table II provides more granular results for the performance of the nominal and robust plans with respect to OAR dose. V_x is the fractional volume of the ROI that exceeds a dose of x Gy. Max cm$^3 \geq x$ Gy is the maximum volume in cubic centimeters of the ROI that is allowed to exceed x Gy. The first row of the table shows the DVH criteria currently used at The Princess Margaret Cancer Center. The remaining rows show the DVH values corresponding to the robust plans for varying levels of Γ. We were able to design treatment plans that satisfied the clinical limits using both the nominal and robust models, although as expected the robust plans had slightly higher (but still acceptable) lung dose.

Naturally, there is a trade-off between TCP and OAR sparing. In Fig. 2, the x-axis shows the mean dose to the lung,
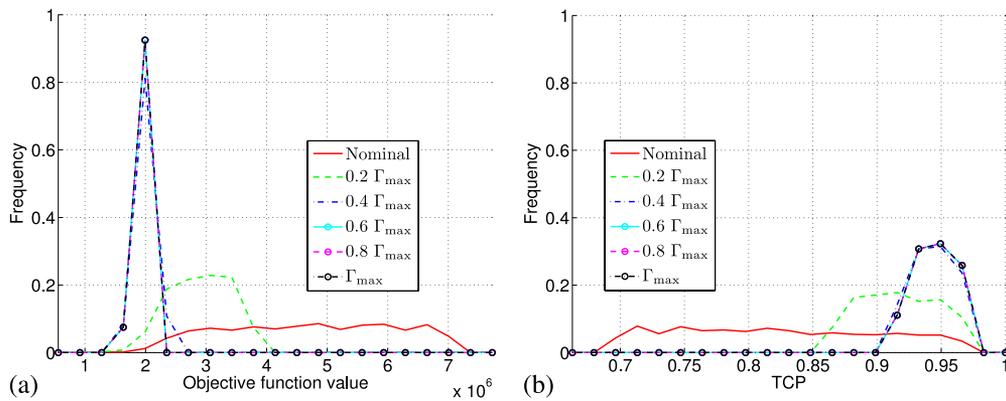
FIG. 1. Comparison of realized objective function and TCP values between nominal and robust plans under worst-case SUV uncertainty. (a) Objective function values and (b) TCP values.

while the *y*-axis shows the mean TCP value over all simulated realizations. The dose to the lung only changes with the value of Γ and not with each simulated realization of SUV (since the simulated SUV changes only occur in the target). As previously observed, there is an increase in minimum TCP and decrease in variance as we increase the value of Γ. However, the mean lung dose increases as well.

3.B. Correlated voxels

Figure 3 shows the results of the simulation for when all the voxels realize a SUV of $\theta_i + \delta\hat{\theta}_i$ for $\delta = 0.1, 0.2, \ldots, 1.0$. We observe that the nominal solution performance worsens with a much steeper slope in both objective function and TCP values, while the robust solutions are much less sensitive to the amount of SUV variation. Similar to the uncorrelated case, there is a threshold above which increasing Γ does not further improve robustness: approximately $0.6\Gamma_{max}$ for the objective function value and $0.2\Gamma_{max}$ for TCP.

4. DISCUSSION

It is important to note that the robust model maintains the same level of mean tumor dose as the nominal model. Simply increasing the beamlet intensities from the nominal solution

by a uniform factor will lead to an improvement in TCP since dose to the tumor will increase. However, this action will also violate some of the maximum dose limits on the OARs. On the other hand, our robust solution is able to obtain an improvement in TCP through the intelligent redistribution of dose, rather than a naïve increase of it. In practice, redistribution of the dose should be done only if a desired minimum dose can be maintained in the tumor. Ideally, a boost dose would be given as well to maximize the benefit. We demonstrate the use of a redistribution here simply to have a fair comparison between the robust and nominal solutions. Even with this redistribution, the minimum dose in the CTV was at least 97% of the minimum desired dose (corresponding to SUV of 0). Overall, the redistribution was fairly modest—taking a small amount of dose from many colder voxels to give to fewer hot voxels—but still yielded a nontrivial improvement in TCP. It is important to maintain a certain minimum dose level to the target, especially when dose may be redistributed away from voxels that may actually be "hotter" than they appear on the image. Also, although all OAR dose/QUANTEC limits are met for all robust solutions, a more difficult case may require the relaxation of some constraints or a penalized objective on normal tissue dose.

Figure 2 shows that while increasing Γ (the conservatism of the robust solution) improves TCP, the marginal returns are

TABLE II. Computed DVH metrics over varying Γ values. The italicized row indicates the clinical limits.

	Lung				Esophagus		Spinal canal	Heart	
	Mean (Gy)	V_5	V_{10}	V_{20}	V_{40}	Max (Gy)	Max (Gy)	V_{40}	cc $\geq$ 75 Gy
Clinical req.	*18*	*0.5*	*0.4*	*0.3*	*0.6*	*65*	*50*	*0.6*	*2*
$0\Gamma_{max}$	7.62	0.21	0.16	0.11	0.31	65.00	50.00	0.01	0.00
$0.1\Gamma_{max}$	7.78	0.21	0.16	0.12	0.32	65.00	50.00	0.01	0.00
$0.2\Gamma_{max}$	8.32	0.25	0.18	0.12	0.32	65.00	50.00	0.01	0.00
$0.3\Gamma_{max}$	8.55	0.26	0.19	0.13	0.37	65.00	50.00	0.01	0.00
$0.4\Gamma_{max}$	8.84	0.28	0.21	0.13	0.33	65.00	50.00	0.01	0.00
$0.5\Gamma_{max}$	8.90	0.28	0.21	0.14	0.34	65.00	50.00	0.01	0.00
$0.6\Gamma_{max}$	8.90	0.28	0.21	0.14	0.34	65.00	50.00	0.01	00.00
$0.7\Gamma_{max}$	8.90	0.28	0.21	0.14	0.34	65.00	50.00	0.01	0.00
$0.8\Gamma_{max}$	8.90	0.28	0.21	0.14	0.34	65.00	50.00	0.01	0.00
$0.9\Gamma_{max}$	8.90	0.28	0.21	0.14	0.34	65.00	50.00	0.01	0.00
Γ_{max}	8.90	0.28	0.21	0.14	0.34	65.00	50.00	0.01	0.00
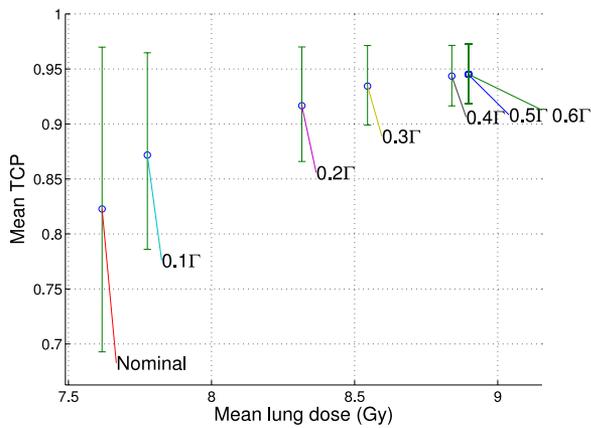
Fig. 2. The trade-off between mean TCP and mean lung dose. The circle is the mean TCP value, and the upper and lower bars represent the maximum and minimum values of the realized TCP, respectively.

diminishing. Thus, we get the most increase in TCP relative to the increase in OAR dose from a small Γ value. Recall that histograms of the objective function value and TCP did not change for values of Γ greater than $0.4\Gamma_{max}$ in the uncorrelated voxel scenarios. One does not need to formulate a very conservative treatment in order to derive most of the benefit of robustness. In other words, choosing a modest value for the protection parameter Γ may provide most of the benefit of protecting against every possible voxel changing its SUV, possibly to worst-case values. This is a promising result because it means that we may be able to achieve improved robustness without giving up too much in terms of healthy tissue dose. This type of robustness result is common in the literature, where accounting for a modest amount of uncertainty is sufficient to change the structure of the robust solution so that other uncertain scenarios are now also protected against (e.g., Ref. 20).

There is another reason why accounting for a modest amount of robustness seems to provide most of the protection of a fully conservative solution, specific to the models in this paper. The reason is because the solution to the auxiliary problem essentially ranks the voxels in decreasing order of worst-case deviation from the nominal SUV and selects voxels to protect in that order. Therefore, larger values of Γ will include the protection of voxels that deviate less and less from

their nominal SUV. Another byproduct of the current approach is that we are essentially protecting against the "hottest" voxels. Because of our assumption in the uncorrelated case that the size of the uncertainty half-interval is an increasing function of nominal SUV, the rank-ordering induced by the auxiliary problem can be viewed as not only a rank-ordering of the voxels with the largest potential deviation from their nominal SUV but also of the voxels with large observed SUV. Thus, our results can also speak to the potential of a much simpler intervention that does not depend on the particular formulation and SUV-to-dose conversion function: boosting the dose to the hottest voxels, at the expense of the coolest ones, while maintaining mean target dose constant to ensure a fair comparison. Subsequent boosting of the entire dose distribution to ensure a consistent minimum target dose may highlight further gains achievable by the robust model, but at the expense of increased OAR dose.

It is important to emphasize that the Γ thresholds observed in this paper may not be generalizable to every patient. The purpose of this paper is to develop a framework that is generalizable and not to determine a general value of Γ. The value of such an optimization framework is to provide the decision maker with flexibility to tailor a robust solution for the particular patient at hand. Note that an appropriate value of Γ will depend on the size of the uncertainty set. Future research could examine using this framework to identify certain "classes" of patients with similar characteristics that can utilize a similar value of Γ.

Many robust optimization approaches exist in the IMRT and IMPT literatures that are conceptually similar to our approach in that they construct a model for the underlying uncertain phenomenon and incorporate this model into the optimization. All of these approaches have the common goal of desensitizing the resulting solution (i.e., fluence map) to the uncertainty. One distinguishing feature of the approach in this paper is our use of Γ to model a budget of uncertainty in an auxiliary problem, whose optimal solution is provided as input into the primary model. Though such a robust model has been applied in other domains, this is the first application of such a model in radiation therapy optimization.

Since the CTV adds a margin around the gross tumor volume (GTV), it includes part of the non-PET avid region due to
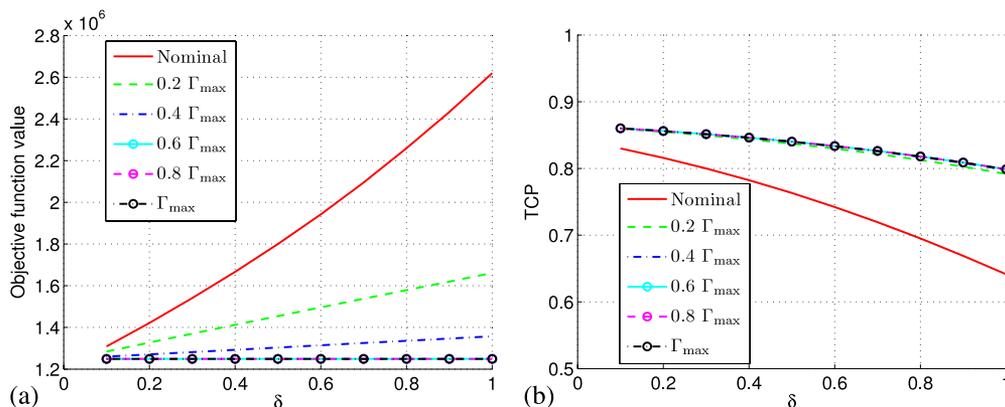
Fig. 3. Fraction of worst-case uncertainty under simulation of the correlated voxels case: (a) objective function value and (b) TCP.

the presumed low clonogen density expected in microscopic disease. As a result, some CTV voxels have a very low nominal SUV that even if doubled would have no effect on the desired dose according to the conversion functions in the Appendix. The robust model prioritizes voxels with a high SUV, so the low SUV voxels are only taken into account at higher Γ values. Thus, we see very little change in the TCP histograms at higher Γ values. If the conversion function from SUV to desired dose was smooth, instead of having the thresholding seen in the Appendix, we would see more of a change in the objective function value and TCP histograms at higher values of Γ. Nevertheless, the changes would still be marginally decreasing with increasing Γ.

In this study we did not consider the accumulation of the delivered dose in the presence of respiratory motion or positional errors. However, in principle, our robust model could be modified to account for these errors given the availability of 4D PET/CT, which helps alleviate motion blur. Our framework addresses uncertainty in general and can accept any noise scenario as input into the model. The specific SUV-to-dose conversion function that we use, which is outlined in the Appendix, was developed based on several studies in the literature. There is much active research ongoing in this area and future work should consider the robustness of our optimization approach to different calibration functions, as new experimental evidence is uncovered. Overall, we believe that our robust approach and model of uncertainty can support future work in robust functional image guided treatment planning optimization.

5. CONCLUSION

In this paper, we present the first robust optimization model for functional image guided IMRT, aimed at mitigating the effects of uncertainty in the PET signal. In our model, we assume voxels in the CTV have a nominal SUV from which they can deviate. Our model uses a simplified treatment planning formulation that gives the treatment planner an adjustable parameter that can be used to adjust the conservatism of the resulting solution. Our framework can accommodate any biological tracer once their values are converted to dose. Our experimental results, based on an assumed nonlinear SUV-to-dose conversion function, show that robust plans can exhibit not only improved objective function and TCP values but also reduced variance of these values compared to the nominal plan. Aligned with results observed in other robust optimization studies, our results suggest that setting the protection parameter to a modest value may provide most of the benefit of robustness without the higher healthy tissue dose characteristic of more conservative solutions.

APPENDIX: THEORETICAL RADIOBIOLOGICAL FRAMEWORK

The following theoretical framework, based on several models in the literature,[13] generates a heterogeneous desired dose distribution from PET SUVs and allows us to evaluate our model's performance. This is not meant to represent a clinically validated function to map PET values to dose. Of course, should such a relationship be discovered in the future, our framework would be readily able to accept the "true" conversion function. Note that the optimization model only requires the dose distribution and does not depend on any biological function.

To generate the desired dose for each voxel, we start with an observed SUV, convert it to a normalized biological feature, generate the α, β parameters of the LQ model of cell kill, and finally compute the voxel-specific desired dose needed to achieve a certain TCP.

1. Radiobiological modeling

We calculate the tumor control probability of voxel i, TCP_i, according to the LQ model.[2,11,13,41] TCP_i is given by

$$\text{TCP}_i = \exp\left[-\rho_i V_i \exp\left(-\alpha_i d_i - \beta_i \frac{d_i^2}{r} + \frac{\log 2}{T_p}(T - T_k)\right)\right]. \tag{A1}$$

Parameters α_i and β_i are the linear and quadratic coefficients of the LQ model for voxel i, respectively. Parameter T_p is potential cell doubling time in days, T is the overall treatment time, T_k is the number of days until repopulation begins, ρ_i is the number of clonogens per cm^3 in voxel i, V_i is the volume of voxel i, r is the number of fractions, and d_i is the total dose to a voxel i. The TCP for the CTV is calculated as the product of TCP_i for each voxel i in the CTV.

2. Conversion from SUV to voxel-specific desired dose

While the true function to convert FDG SUV to a desired dose is unknown, we will assume a particular functional form for the purposes of model evaluation. The optimization is independent of the dose function; it only assumes that a dose can be calculated given the SUV. The dose distribution will be derived to maximize the TCP under the LQ model.

We will assume that the normalized biological feature in this case is oxygenation. Following South *et al.*,[13] we assume that the PET signal is negatively correlated to oxygenation. First, we convert SUV θ_i, to oxygenation $p(\theta_i)$, according to the following equation:

$$p(\theta_i) = \min\left\{\max\left\{\phi_1 \exp\left(\frac{\phi_2}{\theta_i}\right), p_{\min}\right\}, p_{\max}\right\}. \tag{A2}$$

Parameters $p_{\min}$ and $p_{\max}$ were chosen to be 1 and 100, respectively. The function $p(\theta_i)$ is shown in Fig. 4(a). We chose
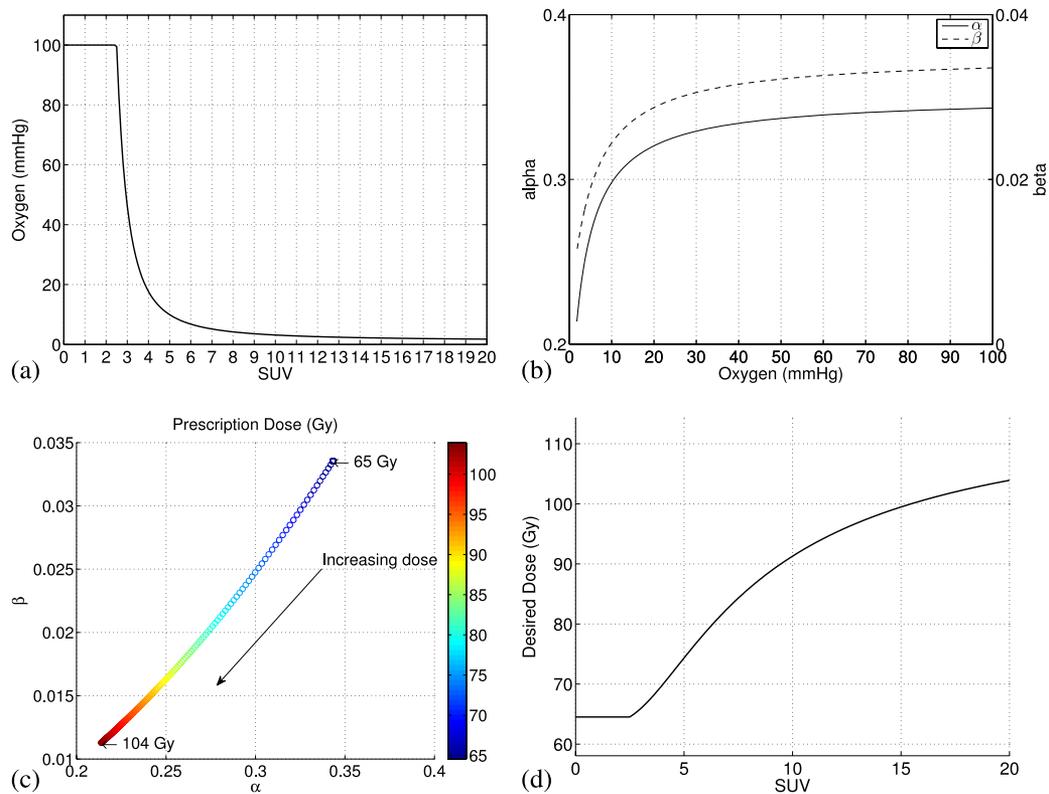
FIG. 4. Conversion functions: (a) SUV to oxygen; (b) oxygen to radiosensitivity parameters; (c) linear–quadratic model parameters to desired dose; and (d) SUV to desired dose.

parameters ϕ_1 and ϕ_2 to fit the overall relationship between SUV and dose so that at a SUV of 2.5, the dose would be approximately 67 Gy and at a SUV of 20, the dose would be approximately 105 Gy [see Fig. 4(d)]. The values chosen for ϕ_1, ϕ_2 were 1 and 14, respectively. A SUV of 2.5 was used because it is a commonly used threshold for tumor delineation.[42] We recognize that a simple threshold for delineation has a number of issues regarding variability. While normalized SUV (N-SUV) is an approach that attempts to reduce potential variations in the measurement, such as varying blood glucose levels,[43] precautions have been taken preemptively to avoid effects in our imaging protocol (i.e., patient fasting, controlled blood glucose levels). It should also be noted that a number of SUV thresholds have been validated with surgical samples.[44,45]

For each voxel i, we convert partial oxygen pressure into the factors, A_i and B_i, which modify the intrinsic radiosensitivity parameters, α_i and β_i, respectively,[46,47] according to the following equations:

$$A_i = \left(\frac{1}{\mathrm{OER}_{\alpha_{max}}} \right) \left[\frac{(p_i \mathrm{OER}_{\alpha_{max}} + K_m)}{p_i + K_m} \right], \quad \text{(A3)}$$

$$B_i = \left(\frac{1}{\mathrm{OER}_{\beta_{max}}} \right)^2 \left[\frac{(p_i \mathrm{OER}_{\beta_{max}} + K_m)}{p_i + K_m} \right]^2. \quad \text{(A4)}$$

Parameters $\mathrm{OER}_{\alpha_{max}}$ and $\mathrm{OER}_{\beta_{max}}$ are the maximum oxygen enhancement ratios (OERs) for α and β, respectively.[10] Parameter K_m is the partial oxygen pressure at which half-maximum sensitization is reached and p_i is the partial oxygen pressure for voxel i. To derive voxel-specific radiosensitivity

parameters α_i and β_i, the intrinsic parameters, α and β, are then multiplied by their respective A_i and B_i: $\alpha_i = A_i \alpha$ and $\beta_i = B_i \beta$. Figure 4(b) shows the relationship between oxygenation and the LQ parameters.

The modified radiosensitivity parameters determine a voxel-specific dose that results in a specified TCP, according to the method described by Yang and Xing.[11] Figure 4(c) shows the conversion from the LQ parameters to a voxel-specific dose that results in a voxel TCP of 0.999 99. Figure 4(d) summarizes all the steps into a single plot, showing the relationship between the input SUV and output dose.

For this study, the intrinsic radiobiological parameters for the lung tumor were set as follows:[13,48,49] $\rho = 10^7$ clonogens/cm^3, $\alpha = 0.35$ Gy^{-1}, $\beta = 0.035$, $\frac{\alpha}{\beta} = 10$, $T_p = 3$ days (equivalent to a cell proliferation rate of 0.1386 days^{-1}), $\mathrm{OER}_{\alpha_{max}} = 2.5$, $\mathrm{OER}_{\beta_{max}} = 3$, $K_m = 3.28$, $T_k = 28$ days. Partial oxygen pressure for a voxel was assumed to be bounded within the range of $p \in [1,100]$ mmHg.[13] The treatment plan parameters were set as follows: $T = 44$ days, $r = 32$ fractions.

a)Author to whom correspondence should be addressed. Electronic mail: heyse.li@mail.utoronto.ca

[1]B. Vanderstraeten, W. De Gersem, W. Duthoy, W. De Neve, and H. Thierens, "Implementation of biologically conformal radiation therapy (BCRT) in an algorithmic segmentation-based inverse planning approach," Phys. Med. Biol. **51**, N277–N286 (2006).
[2]S. Webb and A. E. Nahum, "A model for calculating tumour control probability in radiotherapy including the effects of inhomogeneous distributions of dose and clonogenic cell density," Phys. Med. Biol. **38**, 653–666 (1999).
[3]D. Levin-Plotnik and R. J. Hamilton, "Optimization of tumour control probability for heterogeneous tumours in fractionated radiotherapy treatment protocols," Phys. Med. Biol. **49**, 407–424 (2004).

[4]C. C. Ling, J. Humm, S. Larson, H. Amols, Z. Fuks, S. Leibel, and J. A. Koutcher, "Towards multidimensional radiotherapy (MD-CRT): Biological imaging and biological conformality," Int. J. Radiat. Oncol., Biol., Phys. **47**, 551–560 (2000).

[5]A. Brahme and A. K. Argren, "Optimal dose distribution for eradication of heterogeneous tumors," Acta Oncol. **26**, 377–385 (1987).

[6]L. Xing, C. Cotrutz, S. Hunjan, A. L. Boyer, E. Adalsteinsson, and D. Spielman, "Inverse planning for functional image-guided intensity-modulated radiation therapy," Phys. Med. Biol. **47**, 3567–3578 (2002).

[7]M. Alber, F. Paulsen, S. M. Eschmann, and H. J. Machulla, "On biologically conformal boost dose optimization," Phys. Med. Biol. **48**, N31–N35 (2003).

[8]S. K. Das *et al.*, "Feasibility of optimizing the dose distribution in lung tumors using fluorine-18-fluorodeoxyglucose positron emission tomography and single photon emission computed tomography guided dose prescriptions," Med. Phys. **31**, 1452–1461 (2004).

[9]B. Vanderstraeten, W. Duthoy, W. De Gersem, W. De Neve, and H. Thierens, "[18F] fluoro-deoxy-glucose positron emission tomography ([18F] FDG-PET) voxel intensity-based intensity-modulated radiation therapy (IMRT) for head and neck cancer," Radiother. Oncol. **79**, 249–258 (2006).

[10]I. Toma-Dasu, J. Uhrdin, A. Daşu, and A. Brahme, "Therapy optimization based on non-linear uptake of PET tracers versus linear dose painting," in *World Congress on Medical Physics and Biomedical Engineering, Munich, Germany, September 7-12, 2009, Munich, Germany* (Springer, Berlin Heidelberg, 2009), pp. 221–224.

[11]Y. Yang and L. Xing, "Towards biologically conformal radiation therapy (BCRT): Selective IMRT dose escalation under the guidance of spatial biology distribution," Med. Phys. **32**, 1473–1484 (2005).

[12]C. P. South, M. Partridge, and P. M. Evans, "A theoretical framework for prescribing radiotherapy dose distributions using patient-specific biological information," Med. Phys. **35**, 4599–4611 (2008).

[13]C. P. South, P. M. Evans, and M. Partridge, "Dose prescription complexity versus tumor control probability in biologically conformal radiotherapy," Med. Phys. **36**, 4379–4388 (2009).

[14]R. Boellaard, "Standards for PET image acquisition and quantitative data analysis," J. Nucl. Med. **50**, 11S–20S (2009).

[15]G. Kåver, B. K. Lind, J. Löf, A. Liander, and A. Brahme, "Stochastic optimization of intensity modulated radiotherapy to account for uncertainties in patient sensitivity," Phys. Med. Biol. **44**, 2955–2969 (1999).

[16]M. Witte, G. Shakirin, A. Houweling, H. Peulen, and M. van Herk, "Dealing with geometric uncertainties in dose painting by numbers: Introducing the ΔVH," Radiother. Oncol. **100**, 402–406 (2011).

[17]M. Chu, Y. Zinchenko, S. G. Henderson, and M. B. Sharpe, "Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty," Phys. Med. Biol. **50**, 5463–5477 (2005).

[18]A. Olafsson and S. J. Wright, "Efficient schemes for robust IMRT treatment planning," Phys. Med. Biol. **51**, 5621–5642 (2006).

[19]T. C. Y. Chan, T. Bortfeld, and J. N. Tsitsiklis, "A robust approach to IMRT optimization," Phys. Med. Biol. **51**, 2567–2583 (2006).

[20]T. Bortfeld, T. C. Y. Chan, A. Trofimov, and J. N. Tsitsiklis, "Robust management of motion uncertainty in intensity-modulated radiation therapy," Oper. Res. **56**, 1461–1473 (2008).

[21]T. C. Y. Chan and V. V. Mišić, "Adaptive and robust radiation therapy optimization for lung cancer," Eur. J. Oper. Res. **231**, 745–756 (2013).

[22]C. McCann, T. G. Purdie, H. Rehbinder, A. Lundin, A. J. Hope, A. Bezjak, and J.-P. Bissonnette, "Sci—Thurs PM: Planning—06: A simple, robust IMRT optimization method for lung cancer, accounting for tissue heterogeneity and intra-fraction lung tumour motion," Med. Phys. **36**(9), 4319 (2009).

[23]T. C. Y. Chan, H. Mahmoudzadeh, and T. G. Purdie, "A robust-CVaR optimization approach with application to breast cancer therapy," Eur. J. Oper. Res. **238**, 876–885 (2014).

[24]P. A. Mar and T. C. Y. Chan, "Adaptive and robust radiation therapy in the presence of drift," Phys. Med. Biol. **60**, 3599–3615 (2015).

[25]H. Mahmoudzadeh, J. Lee, T. C. Y. Chan, and T. G. Purdie, "Robust optimization methods for cardiac sparing in tangential breast IMRT," Med. Phys. **42**, 2212–2222 (2015).

[26]V. V. Mišić and T. C. Y. Chan, "The perils of adapting to dose errors in radiation therapy," PLoS One **10**, e0125335 (2015).

[27]J. Unkelbach, T. C. Y. Chan, and T. Bortfeld, "Accounting for range uncertainties in the optimization of intensity modulated proton therapy," Phys. Med. Biol. **52**, 2755–2773 (2007).

[28]A. Fredriksson, A. Forsgren, and B. Hårdemark, "Minimax optimization for handling range and setup uncertainties in proton therapy," Med. Phys. **38**, 1672–1684 (2011).

[29]W. Liu, X. Zhang, Y. Li, and R. Mohan, "Robust optimization of intensity modulated proton therapy," Med. Phys. **39**, 1079–1091 (2012).

[30]W. Chen, J. Unkelbach, A. Trofimov, T. Madden, H. Kooy, T. Bortfeld, and D. Craft, "Including robustness in multi-criteria optimization for intensity-modulated proton therapy," Phys. Med. Biol. **57**, 591–608 (2012).

[31]D. Bertsimas and M. Sim, "The price of robustness," Oper. Res. **52**, 35–53 (2004).

[32]H. J. W. L. Aerts *et al.*, "Identification of residual metabolic-active areas within individual NSCLC tumours using a pre-radiotherapy 18 Fluorodeoxyglucose-PET-CT scan," Radiother. Oncol. **91**, 386–392 (2009).

[33]H. Vesselle, R. A. Schmidt, J. M. Pugsley, M. Li, S. G. Kohlmyer, E. Vallières, and D. E. Wood, "Lung cancer proliferation correlates with [F-18] fluorodeoxyglucose uptake by positron emission tomography," Clin. Cancer Res. **6**, 3837–3844 (2000).

[34]C. Dooms, A. van Baardwijk, E. Verbeken, R. J. van Suylen, S. Stroobants, D. De Ruysscher, and J. Vansteenkiste, "Association between 18F-fluoro-2-deoxy-D-glucose uptake values and tumor vitality: Prognostic value of positron emission tomography in early-stage non-small cell lung cancer," J. Thorac. Oncol. **4**, 822–828 (2009).

[35]R. Airley, J. Loncaster, S. Davidson, M. Bromley, S. Roberts, A. Patterson, R. Hunter, I. Stratford, and C. West, "Glucose transporter glut-1 expression correlates with tumor hypoxia and predicts metastasis-free survival in advanced carcinoma of the cervix," Clin. Cancer Res. **7**, 928–934 (2001).

[36]A. van Baardwijk *et al.*, "The maximum uptake of 18F-deoxyglucose on positron emission tomography scan correlates with survival, hypoxia inducible factor-1α and GLUT-1 in non-small cell lung cancer," Eur. J. Cancer **43**, 1392–1398 (2007).

[37]U. G. A. Sattler and W. Mueller-Klieser, "The anti-oxidant capacity of tumour glycolysis," Int. J. Radiat. Biol. **85**, 963–971 (2009).

[38]A. Abramyuk, S. Tokalov, K. Zöphel, A. Koch, K. S. Lazanyi, C. Gillham, T. Herrmann, and N. Abolmaali, "Is pre-therapeutical FDG-PET/CT capable to detect high risk tumor subvolumes responsible for local failure in non-small cell lung cancer?," Radiother. Oncol. **91**, 399–404 (2009).

[39]J. O. Deasy, A. I. Blanco, and V. H. Clark, "CERR: A computational environment for radiotherapy research," Med. Phys. **30**, 979–985 (2003).

[40]L. B. Marks, E. D. Yorke, A. Jackson, R. K. Ten Haken, L. S. Constine, A. Eisbruch, S. M. Bentzen, J. Nam, and J. O. Deasy, "Use of normal tissue complication probability models in the clinic," Int. J. Radiat. Oncol., Biol., Phys. **76**, S10–S19 (2010).

[41]J. F. Fowler, "The linear-quadratic formula and progress in fractionated radiotherapy," Br. J. Radiol. **62**, 679–694 (1989).

[42]D. Hellwig, T. P. Graeter, D. Ukena, A. Groeschel, G. W. Sybrecht, H.-J. Schaefers, and C.-M. Kirsch, "18F-FDG PET for mediastinal staging of lung cancer: Which SUV threshold makes sense?," J. Nucl. Med. **48**, 1761–1766 (2007).

[43]F.-M. S. Kong, K. A. Frey, L. E. Quint, R. K. Ten Haken, J. A. Hayman, M. Kessler, I. J. Chetty, D. Normolle, A. Eisbruch, and T. S. Lawrence, "A pilot study of [18F] fluorodeoxyglucose positron emission tomography scans during and after radiation-based therapy in patients with nonsmall-cell lung cancer," J. Clin. Oncol. **25**, 3116–3123 (2007).

[44]A. van Baardwijk *et al.*, "PET-CT–based auto-contouring in nonsmall-cell lung cancer correlates with pathology and reduces interobserver variability in the delineation of the primary tumor and involved nodal volumes," Int. J. Radiat. Oncol., Biol., Phys. **68**, 771–778 (2007).

[45]K. Wu *et al.*, "PET CT thresholds for radiotherapy target definition in nonsmall-cell lung cancer: How close are we to the pathologic findings?," Int. J. Radiat. Oncol., Biol., Phys. **77**, 699–706 (2010).

[46]T. Alper, "The modification of damage caused by primary ionization of biological targets," Radiat. Res. **5**, 573–586 (1956).

[47]P. Howard-Flanders and T. Alper, "The sensitivity of microorganisms to irradiation under controlled gas conditions," Radiat. Res. **7**, 518–540 (1957).

[48]M. Mehta, R. Scrimger, R. Mackie, B. Paliwal, R. Chappell, and J. Fowler, "A new approach to dose escalation in nonsmall-cell lung cancer," Int. J. Radiat. Oncol., Biol., Phys. **49**, 23–33 (2001).

[49]S. Y. El Sharouni, H. B. Kal, and J. J. Battermann, "Tumour control probability of stage III inoperable non-small cell lung tumours after sequential chemo-radiotherapy," Anticancer Res. **25**, 4655–4661 (2005).