OPTIMIZING THE SCENARIO POSITIONS FOR ROBUST RADIATION THERAPY TREATMENT PLANNING

Albin FREDRIKSSON*, Anders FORSGREN†, and Björn HÅRDEMARK‡

Technical Report TRITA-MAT-2012-OS5
Department of Mathematics
Royal Institute of Technology
October 10, 2012

Abstract

Complex external beam radiation therapy treatment plans are susceptible to errors. Robust optimization methods that take uncertainties into account during the optimization can alleviate the effects of the errors. For many robust methods, the possible error realizations are discretized into scenarios. The methods then aim at achieving high quality plans in all scenarios simultaneously. The choice of scenarios has high impact on the quality of the plans: when the optimization is aimed at high target coverage in scenarios that are incompatible with healthy tissue sparing in other scenarios, the plan quality in all scenarios may suffer. It can then be beneficial to reduce the range of errors the scenarios represent. To accomplish this, we propose a method that determines how large errors that can be accounted for in robust optimization of intensity-modulated photon and proton therapy. It does so by optimizing the scenario positions along predefined directions simultaneously with the direct machine parameter optimization, under constraints on target coverage and healthy tissue sparing. For both modalities, scenario position optimization was applied to a prostate case and a lung case. The optimization reduced the ranges of errors that were accounted for compared to the a priori ranges. The determined scenario positions were used in robust optimization with fixed scenarios. This resulted in plans that satisfied a substantially larger number of constraints in the determined scenarios and moreover a larger number of constraints in the a priori scenarios than the plans optimized with a priori scenario positions.

1. Introduction

In external beam radiation therapy, misalignment of the patient relative to the beams can lead to large differences between the planned and the delivered dose distributions. The conventional approach to account for errors is to apply margins during treatment planning [18]. Planning is then performed towards delivery of high doses to an enlarged target volume. For complex cases, conventional margins do not

*Optimization and Systems Theory, Department of Mathematics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden; and RaySearch Laboratories, Sveavägen 25, SE-111 34 Stockholm, Sweden. E-mail: albfre@kth.se and albin.fredriksson@raysearchlabs.com.

†Optimization and Systems Theory, Department of Mathematics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden

‡RaySearch Laboratories, Sveavägen 25, SE-111 34 Stockholm, Sweden.

always provide the intended robustness against uncertainties [12]. Methods that utilize information about the uncertainties in the optimization and optimize measures such as the expected objective value or the worst case objective value appear to lead to more robust plans in general [5, 17]. Unfortunately, these methods too have their shortcomings: Expected value optimization applied to quadratic penalty functions provides no robustness guarantees against systematic errors [4]. Moreover, the conservativeness of worst case optimization may lead to unnecessarily compromised plan quality for some error realizations, since when there is no low penalty solution for a specific error realization, the method provides little or no incentive to further improve the penalty for other realizations. The expected value and worst case optimization methods also intensify the conflicts between target coverage and sparing of organs at risk (OARs), since they do not only accommodate conflicts between goals in the case of no errors, but also between goals across different error realizations. When the optimization minimizes a weighted sum of functions quantifying how well the treatment goals are fulfilled, the multiple conflicts yield the trade-off between target coverage and OAR sparing more dependent on the weights than conventionally, but also make the impact of the weights less intuitive. This renders the robust treatment planning more difficult than conventional planning.

When target coverage in one scenario is incompatible with OAR sparing in another scenario, it might be beneficial to retract some of the scenarios to positions in less conflict. In this paper, a method is proposed that determines how large errors that can be accounted for in robust optimizations. The method optimizes the positions of the scenarios under the constraint that all treatment goals be satisfied. When the magnitude of the errors that can be accounted for has been determined, the error positions found are used in worst case optimizations. In a worst case optimization, the objective is minimized in the worst of a number of predefined scenarios. The presented method is applied to a prostate case and a lung case, both subject to systematic setup errors and treated with intensity-modulated radiation therapy (IMRT), i.e., photon therapy, and intensity-modulated proton therapy (IMPT).

Optimization of the expected value of functions penalizing physical quantities such as the deviation of dose has been performed previously for IMRT by, e.g., Unkelbach and Oelfke [17] with respect to random and systematic setup errors and in IMPT by Unkelbach et al. [16] with respect to systematic range and setup errors. Pflugfelder et al. [14] also accounted for systematic range and setup errors in IMPT, but applied the penalty functions to a worst case dose distribution, where each voxel considered independently received its worst dose from a number of different scenarios. In a previous paper, we used worst case (or "minimax") optimization, in which the penalty of the objective function applied to the dose distribution of the worst scenario is minimized, for patient cases treated with IMPT subject to systematic range and setup errors [5]. Such optimization is also the basis for the robust method used in the present paper. Chen et al. [3] used minimax optimization for IMPT accounting for systematic range and setup errors in a linear programming multicriteria optimization framework, which enabled scaling between nominal and robust solutions to the optimization problem. Generally, the methods that incorporate information about the uncertainties into the optimization enable similar or

more robust target coverage and better sparing of healthy tissues than conventional methods that use margins to account for uncertainties. However, a comparison of IMPT treatment planning by expected value minimization, worst case penalty minimization, and minimization of the conditional value-at-risk showed that expected value minimization, while requiring a less complex optimization formulation, may result in worse target coverage than the more conservative methods [4].

The methods for robustness cited above optimize uncertainty measures with the ranges of errors accounted for specified prior to the optimization. The method proposed in the present paper differs from those methods in that it changes the ranges of errors that are accounted for during the optimization. Changing the size of the region within which to be robust has similarities to some previous methods. Gordon and Siebers [8] updated the sizes of the planning margins for IMRT plans iteratively until a coverage probability was met. Gordon et al. [7] proposed a method based on dose-coverage histograms, which show coverage probabilities as functions of dose for the regions of interest (ROIs), similar to how dose-volume histograms show the volumes as functions of dose. Using dose-coverage histograms, they introduced optimization functions aimed towards achieving specified target coverage and OAR sparing probabilities. Moore et al. [13] developed a probabilistic treatment planning method to account for systematic setup errors, which considered multiple setup error scenarios and tried to achieve some specified target coverage probability. When the method was unable to reach the specified probability, the level was iteratively updated. The method proposed in the present paper differs from these coverage probability-based methods in that it optimizes on the coverage probability directly. It can thereby reach the highest coverage probability up to the desired level for which the treatment goals can be satisfied.

2. Method

A scenario position optimization problem is formulated that is aimed at identifying the scenarios that are compatible with the treatment goals. In this problem, the scenario positions are optimized under the constraint that all goals must be satisfied. The optimal solution provides the compatible scenario positions, which can then be included in a standard robust optimization with fixed scenario positions. Optimizing the scenarios can be thought of as isocenter optimization for multiple scenarios simultaneously.

2.1. Uncertainties and scenarios

Systematic setup errors are considered. These are modeled as translations of the beam isocenters. During the optimization, the possible error realizations are discretized into a number of scenarios and dose is computed for each of the included scenarios.

The main goal of robust radiation therapy treatment plan optimization is to satisfy all treatment goals in all scenarios. If the goals are formulated as optimization constraints this becomes a feasibility problem. Sometimes it is not possible to satisfy all constraints simultaneously in a given scenario. This is however not a reason to

drop that scenario entirely, for it might still be worthwhile to try to satisfy some of the constraints in that scenario: Consider, as an example, the geometry depicted in Figure 1 that has two realizable scenarios. Assume that the ROIs move in a static dose distribution. Then there cannot exist a dose distribution that simultaneously achieves full OAR sparing in the first scenario and target coverage in the second scenario. Therefore, one might decide to drop the aim for OAR sparing in the first scenario. It would however be detrimental if one simultaneously dropped the aim for target coverage in the first scenario, since this aim is not in conflict with any other goal. In order to include the possibility of target coverage in a scenario in which the OAR sparing is in conflict with some other goal, it is necessary to consider each structure individually in each scenario.

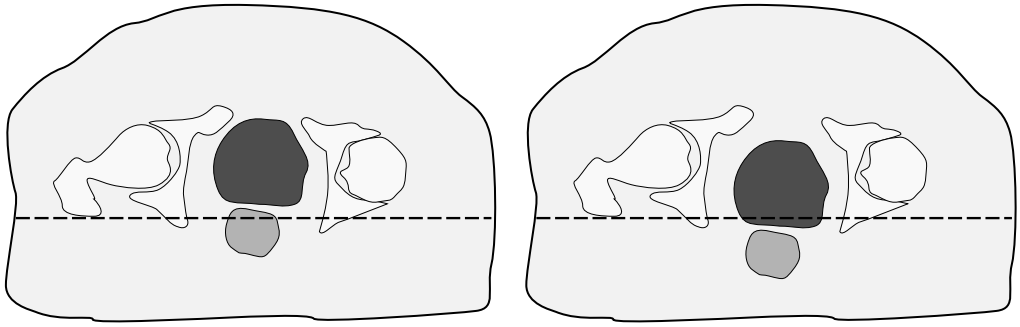

(a) Scenario 1 (b) Scenario 2

Figure 1. An example case with two scenarios. The target (darkgray) and the OAR (lightgray) are indicated. OAR sparing in scenario 1 is incompatible with target coverage in scenario 2, but target coverage in scenario 1 is not in conflict with any goal, so it is still suitable to aim for target coverage in scenario 1.

2.2. Mathematical formulation

In order to determine the scenario positions where all goals can be simultaneously satisfied, a scenario position optimization problem is formulated. The random variable picking the realization of the systematic error is denoted by S, which is thus a random variable vector in $\mathbb{R}^3$. The aim of the scenario position optimization is to maximize the probability that S falls within the volume in which the criteria for all ROIs enumerated by the set $\mathcal{R}$ are simultaneously satisfied. This volume is modeled by a discrete set of scenarios.

Consider scenarios corresponding to shifts in the n directions $p_i \in \mathbb{R}^3$ for index $i = 1, \ldots, n$, and, for each ROI $r \in \mathcal{R}$ and each direction $i = 1, \ldots, n$, let the scalar factor $\alpha_{r,i} \in [0, 1]$ determine the position of the scenario for ROI r along the direction p_i. Let $\alpha_{r,0} = 0$ for all $r \in \mathcal{R}$ and let p_0 be arbitrary. Note that the set $\mathcal{R}$ could enumerate criteria instead of ROIs, so that different scenario positions were considered for different criteria of a given ROI. The volume for which the criteria of ROI $r \in \mathcal{R}$ are satisfied is parameterized by $\alpha_r \in \mathbb{R}^{n+1}$ and $p \in \mathbb{R}^{3\times(n+1)}$, and is denoted by $C(\alpha_r, p) \subseteq \mathbb{R}^3$. The volume in which the criteria of all ROIs are satisfied

then becomes $\bigcap_{r\in\mathcal{R}} C(\alpha_r, p)$, so the objective function to be maximized becomes

$$\mathbb{P}\left(S \in \bigcap_{r\in\mathcal{R}} C(\alpha_r, p)\right),$$

where $\mathbb{P}(A)$ denotes the probability of the event A occurring.

Besides the main goal to satisfy the criteria of all ROIs simultaneously, it is—as discussed in Section 2.1 above—beneficial to satisfy the criteria of each ROI also in scenarios in which the criteria of other ROIs cannot be simultaneously satisfied. Therefore, terms $\mathbb{P}(S \in C(\alpha_r, p))$ for $r \in \mathcal{R}$ encouraging enlarging the volume covered for each ROI r individually are weighted into the objective by a factor $0 < \rho \ll 1$.

Denote by $d(x, q)$ the dose distribution as a function of the optimization variables x from the set $\mathcal{X}$ of feasible optimization variables and of the displacement $q \in \mathbb{R}^3$. Then, the scenario position optimization problem can then be formulated as

$$\begin{aligned} \underset{\alpha,x}{\text{maximize}} \quad & \mathbb{P}\left(S \in \bigcap_{r\in\mathcal{R}} C(\alpha_r, p)\right) + \rho \sum_{r\in\mathcal{R}} \mathbb{P}(S \in C(\alpha_r, p)) \\ \text{subject to} \quad & f_r(d(x, \alpha_{r,i} p_i)) \leq 0, \qquad r \in \mathcal{R}, \quad i = 0, \ldots, n, \\ & \alpha_{r,i} \in [0, 1], \qquad\qquad\quad r \in \mathcal{R}, \quad i = 1, \ldots, n, \\ & x \in \mathcal{X}, \end{aligned} \tag{2.1}$$

where f_r is a penalty function representing ROI $r \in \mathcal{R}$ such that f_r evaluates to zero when the criteria of the ROI are satisfied. The constraint on f_r for $i = 0$ enforces the criteria in the nominal scenario, since $\alpha_{r,0} p_0 = (0\ 0\ 0)^T$. Because the dose is nonconvex in the scenario positions in general (and, for IMRT, also in machine parameters such as the multi-leaf collimator leaf positions), this is a nonconvex optimization problem. Common nonlinear programming algorithms aim for locally optimal solutions. Due to the nonconvexity, the scenario positions resulting after such an optimization cannot be guaranteed to be globally optimal. While the problem is nonconvex, it is likely that the partial derivatives of the constraints with respect to the scenario positions are positive for conformal dose distributions, since the target coverage and OAR sparing often deteriorates as the distance of a scenario from the nominal position is increased.

In the above formulation, it is assumed that if a constraint is satisfied in the nominal scenario and for $\alpha_{r,i} p_i$ for $i = 1, \ldots, n$, it is satisfied in the full volume $C(\alpha_r, p)$. If this is not the case, additional constraints for other positions in $C(\alpha_r, p)$ ought to be included in the optimization.

The optimization criteria of the scenario position optimization problem should represent the worst clinically acceptable criteria, i.e., the criteria such that there is little meaning in treating the patient unless they are satisfied. These criteria are not intended to fine tune the trade-offs between conflicting goals, and should be contrasted to the conventional optimization goals constituting the objective function, which reflect the goals that are desirable, but not necessary, to satsify.

2.3. Probability computation

The setup errors are assumed to be normally distributed with zero mean and covariance matrix $\Sigma = \sigma^2 I$. If there were only a single direction along which to move (i.e., $n = 2$, $p_1, p_2 \in \mathbb{R}$, and $p_1 = -\beta p_2$ for some $\beta > 0$), the probability density function would be $(\sigma\sqrt{2\pi})^{-1} e^{-t^2/(2\sigma^2)}$, so the probability covered by p and α_r for a ROI $r \in \mathcal{R}$ would be

$$\mathbb{P}(S \in [-\alpha_{r,1}|p_1|, \alpha_{r,2}|p_2|]) = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\alpha_{r,1}|p_1|}^{\alpha_{r,2}|p_2|} e^{-t^2/(2\sigma^2)}\, dt = \frac{1}{2} \sum_{i=1}^{2} \operatorname{erf}\left(\frac{\alpha_{r,i}|p_i|}{\sigma\sqrt{2}}\right),$$

where the error function $\operatorname{erf}(x)$ is defined by

$$\operatorname{erf}(x) = (2/\sqrt{\pi}) \int_0^x e^{-t^2}\, dt.$$

For simplicity, the probability covered by the scenarios of ROI $r \in \mathcal{R}$ when there are multiple directions is assumed to be the sum of such single direction probabilities, viz.,

$$\mathbb{P}(S \in C(\alpha_r, p)) = \frac{1}{n} \sum_{i=1}^{n} \operatorname{erf}\left(\frac{\alpha_{r,i}\|p_i\|_2}{\sigma\sqrt{2}}\right),$$

where the length of p_i is measured by the norm instead of the absolute value, since p_i is now in $\mathbb{R}^3$. This assumption says that the different directions contribute independently to the coverage probability. The probability covered by the scenarios of all ROIs in $\mathcal{R}$ is then given by

$$\mathbb{P}\left(S \in \bigcap_{r\in\mathcal{R}} C(\alpha_r, p)\right) = \frac{1}{n} \sum_{i=1}^{n} \min_{r\in\mathcal{R}} \operatorname{erf}\left(\frac{\alpha_{r,i}\|p_i\|_2}{\sigma\sqrt{2}}\right).$$

The function $\operatorname{erf}(x)$ is concave in $x \in \mathbb{R}_+$ and the min operator preserves concavity. Hence, since the objective function is maximized, the objective of problem (2.1) under this probability computation assumption does not introduce additional non-convexity to the optimization problem.

If it can be assumed that only one ROI is compromised in each scenario direction $i = 1, \ldots, n$ while each other ROI $r \in \mathcal{R}$ keeps its corresponding value of $\alpha_{r,i}$ at its upper bound throughout the optimization (assuming that it is initialized at its upper bound), the solutions resulting from using the objectives

$$\frac{1}{n} \sum_{i=1}^{n} \min_{r\in\mathcal{R}} \operatorname{erf}\left(\frac{\alpha_{r,i}\|p_i\|_2}{\sigma\sqrt{2}}\right) \qquad \text{and} \qquad \sum_{r\in\mathcal{R}} \frac{1}{n} \sum_{i=1}^{n} \operatorname{erf}\left(\frac{\alpha_{r,i}\|p_i\|_2}{\sigma\sqrt{2}}\right)$$

are the same. This is the case for the example geometry in Figure 1, since the anterior scenario does not compromise the target coverage whereas the posterior scenario does not compromise the OAR sparing. That this is the case is assumed subsequently, and therefore only the latter term is used as objective in the scenario position optimizations.

2.4. Optimizing margins

The formulation (2.1) that optimizes the scenario positions is intended to be as general as possible, and is not restricted to a specific modality or patient geometry. The generality comes at increased computational cost compared to conventional problems. Margins are commonly used to account for uncertainties in photon therapy treatments [10, 15]. For cases when margins can be assumed to be sufficient, the displacements $\alpha_{r,i}p_i$ for $i = 1, \ldots, n$ and a given ROI $r \in \mathcal{R}$ could be used to parameterize the margin of the ROI, so that the margins would be optimized instead of the scenario positions. Since this approach would not require multiple constraints for each structure nor dose computations in multiple scenarios, it would potentially be faster than the scenario position optimization.

2.5. Computational study

Scenario position optimization was implemented in the RayStation 2.8 treatment planning system (RaySearch Laboratories, Stockholm, Sweden). The optimization in RayStation is performed by a sequential quadratic programming algorithm using quasi-Newton updates of an approximation of the Hessian of the Lagrangian. A similar method is described by Gill et al. [6]. The scenario position optimization was applied to a prostate case and a lung case. Both IMRT and IMPT plans were optimized for the two cases. For the prostate case, a five-field IMRT treatment with equispaced beams beginning at 0°, and a two-field IMPT treatment with beams at 90° and 270° were optimized. For the lung case, a seven-field IMRT treatment with equispaced beams beginning at 0°, and a two-field IMPT treatment with beams at 35° and 250° were optimized. For both cases, the dose grid resolutions were $2.5 \times 2.5 \times 2.5$ mm^3 and the standard deviations of the systematic setup errors were $\sigma = 5$ mm. It was assumed that it was desired to account for errors of at most 2σ, i.e., shifts of up to 1 cm. Transversal slices of the patients are shown in Figure 2.

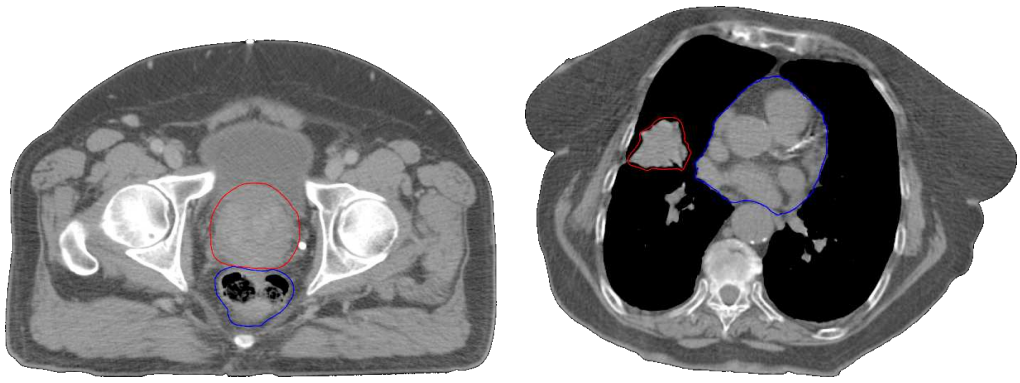

(a) Prostate case (b) Lung case

Figure 2. Transversal slice of the patient cases. The targets (red) and the main OARs (blue) are indicated as contours.

When the scenario positions had been determined by the scenario position op-

timization, plans were optimized with worst case optimization [5] with the scenario positions fixed at the locations determined by the previous optimization. These plans were compared to plans optimized with the a priori scenario positions of distance 2σ from the nominal scenario. Since the different ROIs can have different scenario positions after the scenario position optimization, the scenarios must be separated per ROI also in the standard robust optimization. Given the scenario directions p and the scenario positions α_r for ROI $r \in \mathcal{R}$, the standard robust optimization problem is thus formulated as

$$\begin{aligned} \underset{x\in\mathcal{X}}{\text{minimize}} \quad & f(d(x)) + \sum_{r\in\mathcal{R}} \max_{i=0,\dots,n} f_r(d(x, \alpha_{r,i}p_i)) \\ \text{subject to} \quad & g(d(x)) \leq 0, \\ & g_r(d(x, \alpha_{r,i}p_i)) \leq 0, \qquad\qquad r \in \mathcal{R}, \quad i = 0, \dots, n, \end{aligned} \tag{2.2}$$

where $d(x)$ denotes the nominal dose distribution $d(x, (0\ 0\ 0)^T)$, f is the weighted sum of nominal objective constituents, f_r is the weighted sum of robust objective constituents for ROI $r \in \mathcal{R}$, g represents the nominal constraints, and g_r represents the robust constraints of ROI $r \in \mathcal{R}$. The optimization functions constituting the objectives and constraints used in this paper are defined mathematically in Appendix A.

During the optimizations, scenario doses for both IMRT and IMPT were calculated using the nominal mapping from fluence to dose, but with the fluence maps shifted (and bilinearly interpolated) according to the displacement of the scenarios. An error along a beam direction was assumed to affect the resulting beam dose only according to the inverse-square law, so that when the patient moves away from the treatment unit, the beam dose is scaled downwards. Derivatives with respect to the scenario positions were approximated by finite differences. Thus, the number of dose calculations increased by one per iteration for each scenario but the nominal.

The scenario position optimization was initialized from plans reached by optimization with fixed scenario positions for seven iterations. For IMRT, these seven iterations were fluence map optimization iterations, and the resulting plan was segmented before the scenario positions were optimized in combination with direct step-and-shoot optimization [9].

The scenario position optimization for IMRT was performed using a fast dose calculation algorithm based on singular value decomposition of pencil beam kernels [2]. For the optimizations with fixed scenario positions, the same algorithm was primarily used, but at intermediate iterations accurate dose was computed by a collapsed cone dose calculation algorithm [1] and the subsequent optimization was performed on the dose of the fast algorithm incremental from the accurate dose.

For IMPT, the dose was computed using the pencil beam dose calculation algorithm of RayStation, which takes heterogeneities into account, also within the cross-section of each spot. The line spacing and the energy layer separation (in water equivalent media) were both set to 5 mm, but to improve upon the approximate scenario dose calculation, auxiliary spots were computed for 2.5 mm line spacing. The additional spots were not included as variables in the optimization. In Unkel-

bach et al. [16], auxiliary spots were used in a similar way, but nearest neighbor interpolation was applied to the shifted spots.

In the robustness evaluation, the beam isocenters were shifted before dose was computed. There was thus a difference in the scenario dose calculation during the optimization and that in the evaluation, since the shifted fluence used during the optimization results in slight inaccuracies due to the different divergence.

3. Results

3.1. Prostate case

The prostate case was simplified to resemble to the example shown in Figure 1. Thus, only setup shifts in the anterior and posterior directions were considered.

3.1.1. Optimization problem

The goals for the target and the rectum of the prostate case were assumed to require robustness. Other goals were included for the nominal scenario only. Five scenarios were included for dose calculation: the nominal scenario, and posterior and anterior shifts for the target and for the rectum, so $n = 2$. The scenario position vectors were set to $p_1 = (0\ 1\ 0)^T$ cm and $p_2 = (0\ -1\ 0)^T$ cm and the optimization problem was formulated similar to (2.1). Its minimum robust requirements, as well as its nominal requirements, are presented in Table 1. The objective was to maximize the probability covered by the posterior and anterior shifts for the target and for the rectum.

Table 1. Robust and nominal constraints for the prostate case.

Robust constraints			Nominal constraints		
Structure	Function	Dose level [Gy]	Structure	Function	Dose level [Gy]
Target	Min dose	70	Bladder	Max 20 % DVH	70
Target	Min 98 % DVH	74	L. femoral head	Max dose	40
Rectum	Max 45 % DVH	40	R. femoral head	Max dose	40
Rectum	Max 20 % DVH	60	External	Max dose	82
Rectum	Max 5 % DVH	78			

3.1.2. Optimized scenarios

The posterior and anterior scenarios for the prostate and the rectum were included as variables in the IMRT and IMPT optimizations. The optimizations did not result in changes to the posterior isocenter shift for the target or the anterior isocenter shift for the rectum, but kept these at their maximum positions of 1 cm. The two other shifts were modified for both modalities. Figure 3 displays the progress of the modified scenario positions and the maximum constraint violation during the optimization.

The optimizations first retract the scenario positions rapidly to improve upon the feasibility. Plans are considered feasible if the maximum constraint violation is less

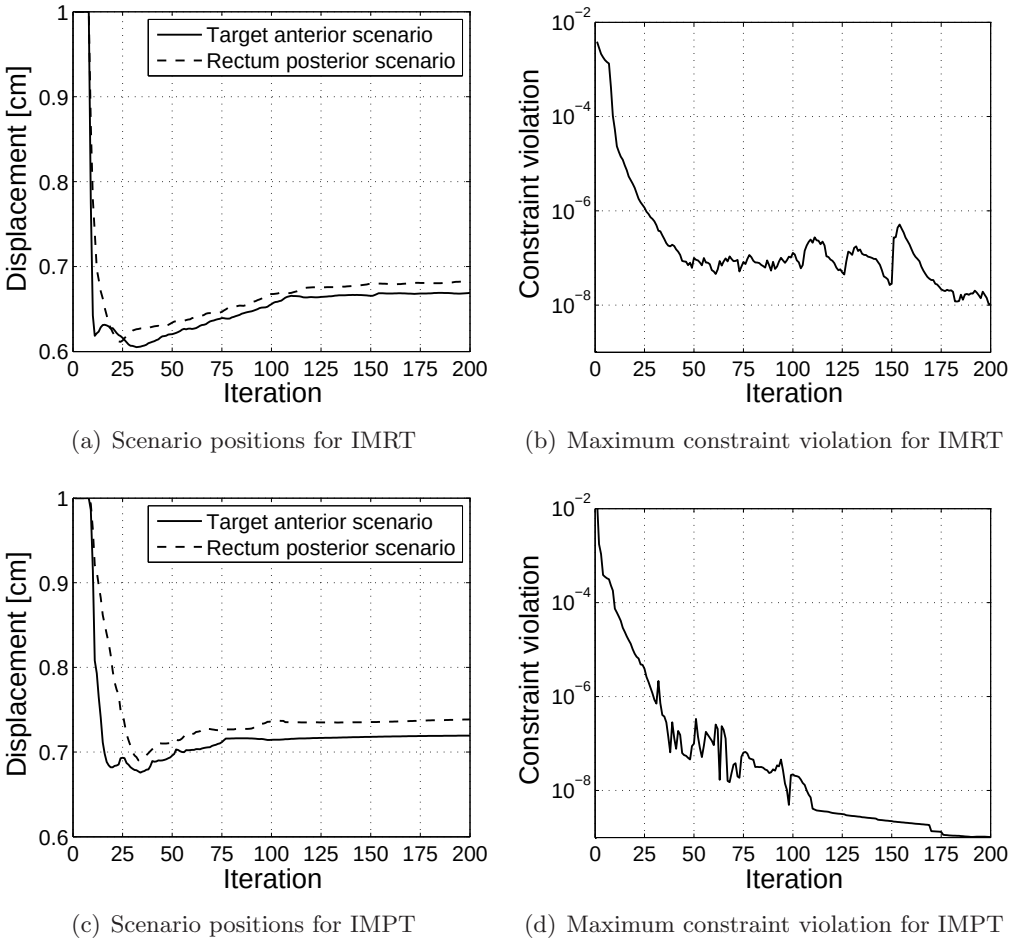

(a) Scenario positions for IMRT

(b) Maximum constraint violation for IMRT

(c) Scenario positions for IMPT

(d) Maximum constraint violation for IMPT

Figure 3. Progress of the scenario positions and maximum constraint violation during the prostate case optimizations.

than 10^{-6}. The IMRT as well as the IMPT plans become feasible after iteration 25. At the same time, the scenario positions are being slowly pushed outwards, which improves upon the objective value. For IMRT, the scenario position optimization resulted in the positions 0.67 cm and 0.68 cm for respectively the target anterior and the rectum posterior scenarios. For IMPT, it resulted in the positions 0.72 cm and 0.74 cm.

3.1.3. Feasible scenario positions

In order to determine what could best be expected from a scenario position optimization, optimizations with fixed scenario positions were performed for an enumeration of possible scenario positions for the two ROIs. The posterior isocenter shift for the target and the anterior isocenter shift for the rectum were always set to 1 cm, while the anterior shift for the target and the posterior shift for the rectum were considered for each point in a 8×8 point regular discretization of the $[0.5, 0.85] \times [0.5, 0.85]$ cm^2 box. For each point in the grid, an optimization was performed for 100 iterations. The combinations of scenarios that resulted in a feasible solution after standard robust optimization (with fixed scenario positions during the optimization) are shown in Figure 4. The nonlinear constraints were considered satisfied if they evaluated to less than 10^{-6}. Bilinear interpolation was used to approximate the cutoff when it occurred between points.

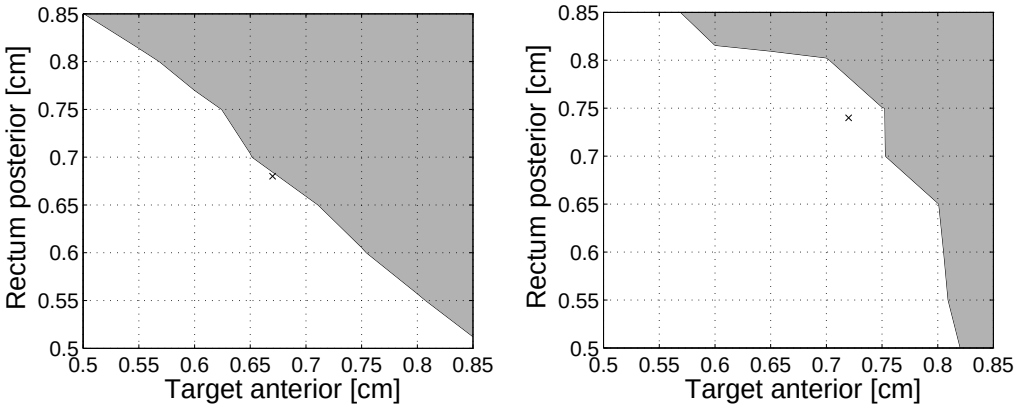

(a) IMRT (b) IMPT

Figure 4. Feasible region (white) for the prostate case after optimization for 100 iterations. The crosses (x) denote the positions found by the scenario position optimization.

For IMRT, the optimized scenario positions came very close to what was deemed feasible for optimization with fixed scenario positions. For IMPT, the difference was a few tenth of a millimeter.

3.1.4. Robust plans with selected scenarios

The optimized scenario positions were used as fixed positions in robust optimizations with standard robust goals according to the formulation (2.2). The constraints from Section 3.1.1 were kept, but mean dose objectives of unity weight were introduced to reduce the doses to all healthy ROIs.

The standard robust optimization can be warm started from the solution to the scenario position optimization problem. Here, however, the standard robust optimization was started from scratch to make the plan more comparable to the reference plan using the a priori scenario positions (1 cm in the posterior as well as the anterior direction for each ROI requiring robustness), which was also optimized. Since with the a priori scenarios the constraints could not be satisfied, the target and rectum goals were relaxed into objective constituents with weights 100 for that optimization. DVHs in different scenarios for these plans are shown in Figure 5. The DVHs are based on doses computed with shifted beam isocenters, so there is a slight difference between these doses and the ones used during the optimization, which were computed with the nominal isocenters but with shifted fluences. For IMRT as well as IMPT, the plans with optimized scenarios neglect the 1 cm shift for the target but are in return able to perform better in the other scenarios than the optimization with a priori scenarios. The same is true of the -1 cm shift for the OAR.

To evaluate how well the different methods satisfied the robust constraints, each robust constraint of Table 1 (with volume level relaxed by 0.5 %) was evaluated in each of the optimized scenario positions as well as in the a priori scenario positions. For each of these two scenario groups and each method, the number of satisfied constraints was summed over the scenarios. The results are shown in Table 2. The IMRT and IMPT plans optimized with the scenarios determined by the scenario position optimization satisfied respectively 5 and 3 constraints more in the optimized scenarios than the plans optimized with a priori constraints. They also satisfied respectively 3 and 1 constraints more in the a priori scenarios.

Table 2. The number of satisfied constraints for the prostate case over the three optimized scenarios and over the three a priori scenarios. Five constraints and three scenarios make the maximum 15.

	No. of satisfied constraints	
Method	Optimized scenarios	A priori scenarios
IMRT with optimized scenarios	13	11
IMRT with a priori scenarios	8	8
IMPT with optimized scenarios	12	10
IMPT with a priori scenarios	9	9

3.2. Lung case

For the lung case, not only the anterior and posterior, but also the left, right, superior, and inferior scenarios were included in the scenario position optimization.

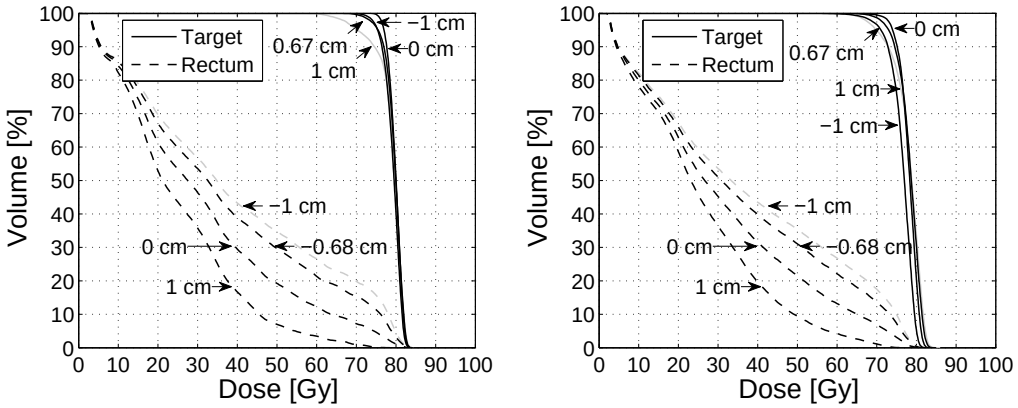

(a) IMRT plan using optimized scenario positions

(b) IMRT plan using a priori scenario positions

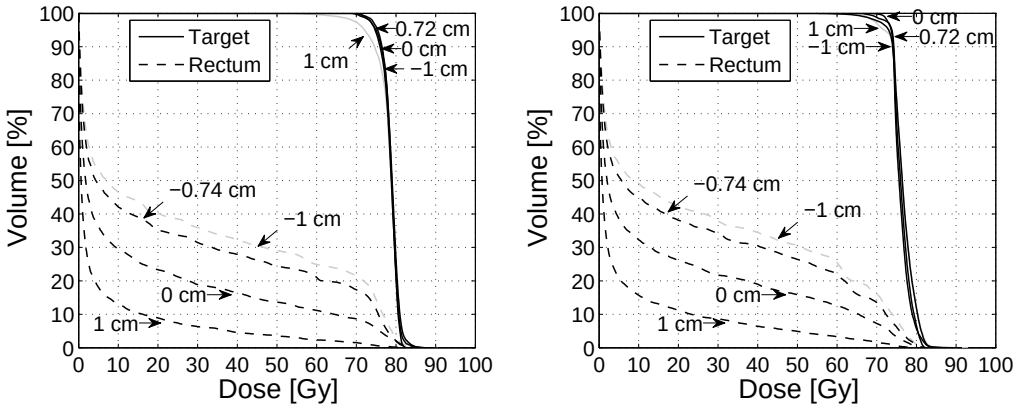

(c) IMPT plan using optimized scenario positions

(d) IMPT plan using a priori scenario positions

Figure 5. Prostate case DVHs for the plans using optimized scenario positions and the plans using a priori scenario positions. The shifts in the anterior direction are annotated. Black curves correspond to the optimized scenario positions and gray curves correspond to the a priori scenario positions (± 1 cm in the anterior direction).

3.2.1. Optimization problem

The target and the heart were considered as the only goals of the lung case requiring robustness. There were 13 scenarios included for dose calculation: the nominal scenario, and left, right, posterior, anterior, superior, and inferior scenarios for the target and for the heart, so $n = 6$. The scenario vectors were set to $p_1 = (1\ 0\ 0)^T$ cm, $p_2 = (-1\ 0\ 0)^T$ cm, $p_3 = (0\ 1\ 0)^T$ cm, $p_4 = (0\ -1\ 0)^T$ cm, $p_5 = (0\ 0\ 1)^T$ cm, $p_6 = (0\ 0\ -1)^T$ cm and the optimization problem was formulated similar to (2.1). Its minimum robust requirements, as well as its nominal requirements, are presented in Table 3. The objective was to maximize the probability covered by the scenarios for the target and for the heart.

Table 3. Robust and nominal constraints for the lung case.

Robust constraints			Nominal constraints		
Structure	Function	Dose level [Gy]	Structure	Function	Dose level [Gy]
Target	Min dose	68	Lung	Max 37 % DVH	20
Target	Min 98 % DVH	70	External	Max dose	77
Target	Max dose	77			
Heart	Max 1 % DVH	40			

3.2.2. Optimized scenarios

The left, right, posterior, anterior, superior, and inferior scenarios for the target and the heart were included as variables in the IMRT and IMPT optimizations. The progress of the scenario positions and the maximum constraint violation during the optimizations are shown in Figure 6.

As for the prostate case, the optimizations first retract the scenario positions rapidly to improve upon the feasibility. For both IMRT and IMPT, the maximum constraint violation drops below 10^{-6} before iteration 75. At the same time, the scenario positions are being slowly pushed outwards, which improves upon the objective function value. The resulting scenario positions are shown in Table 4.

Table 4. Optimized scenario positions for the lung case.

	IMRT		IMPT	
Direction	Target scenario [cm]	Heart scenario [cm]	Target scenario [cm]	Heart scenario [cm]
Left	1.0	0.35	0.94	0.65
Right	0.5	1.0	0.81	1.0
Posterior	1.0	1.0	0.87	1.0
Anterior	1.0	1.0	0.73	1.0
Superior	1.0	1.0	0.70	1.0
Inferior	0.92	1.0	1.0	0.75

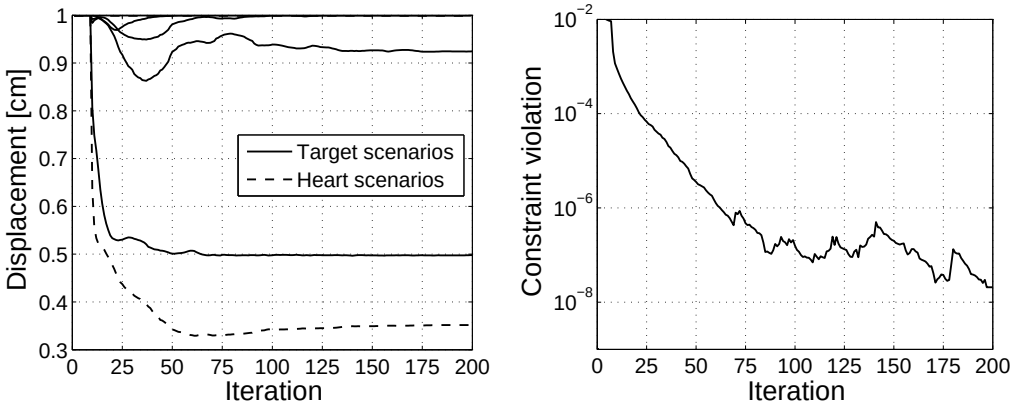

(a) Scenario positions for IMRT

(b) Maximum constraint violation for IMRT

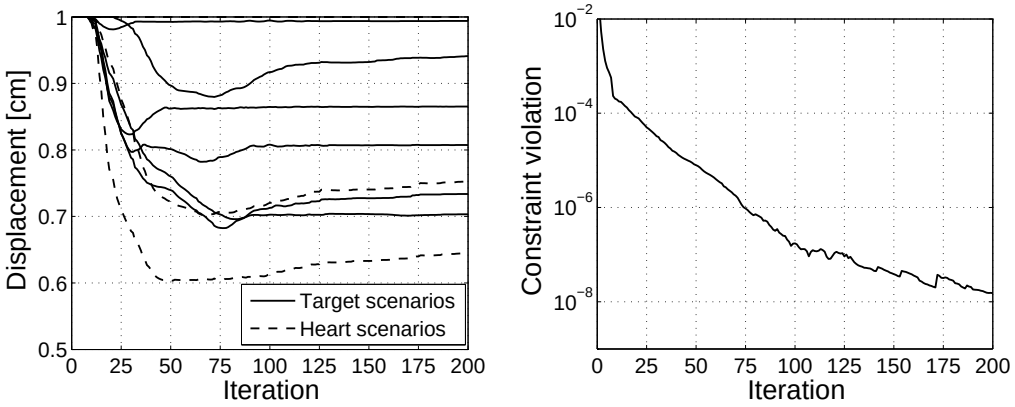

(c) Scenario positions for IMPT

(d) Maximum constraint violation for IMPT

Figure 6. Progress of the scenario positions and maximum constraint violation during the lung case optimizations.

3.2.3. Feasible scenario positions

Due to its multidimensionality, the feasible region of the scenario positions for the lung case optimization problem cannot be as easily determined as that of the prostate case problem. To determine whether the scenario position optimizations resulted in unnecessarily retracted scenario positions, plans with fixed scenario positions were optimized. First, the scenario positions as determined by the scenario position optimizations, shown in Table 4, were used. Second, these scenario positions were pushed 0.1 mm outwards, but to a maximum of 1 cm, and these extended positions were used. In the first case, the IMRT as well as the IMPT optimization resulted in feasible plans, with maximum constraint violation less than 10^{-6}. With the scenarios pushed outwards, neither the IMRT nor the IMPT optimization resulted in a feasible solution. This indicates that the scenario positions found could not be much improved upon. It does not, however, prove global optimality, since it is possible that some of the scenario positions could be pushed outwards, or that the positions could be redistributed in a way that would improve upon the objective value while admitting a feasible solution.

3.2.4. Robust plans with selected scenarios

The optimized scenario positions were used as fixed positions in robust optimizations with standard robust goals according to the formulation (2.2). The constraints from Section 3.2.1 were kept, but mean dose objectives of unity weight were introduced to reduce the doses to all healthy ROIs.

As for the prostate case, the standard robust optimization for the lung case was started from scratch to make the plan more comparable to the reference plan using the a priori scenario positions (1 cm in the positive and negative axis directions), which was also optimized. Since with the a priori scenarios the constraints could not be satisfied, the target and heart goals were relaxed into objective constituents with weights 100 for that optimization. DVHs in the different scenarios for these plans are shown in Figure 7. The DVHs are based on doses computed with shifted beam isocenters, so there is a slight difference between these doses and the ones used during the optimization, which were computed with the nominal isocenters but with shifted fluences. They show that the heterogeneous density of the lung affects the IMPT plan more than the IMRT plan. The 1 cm right shift for IMRT and the 1 cm superior shift for IMPT resulted in the worst target coverage for the plans with optimized scenario positions. While the plans with optimized scenario positions neglect some shifts, they are in return able to improve upon the target coverage and OAR sparing in other shifts compared to the optimizations with a priori scenarios.

The number of scenarios in which the robust constraints for the lung case (relaxed by 0.5 %) are satisfied are shown in Table 5, for the scenario positions determined by the optimization as well as for the a priori scenario positions. Since the min 98 % DVH constraint was not satisfied in any scenario for any method other than IMPT with optimized scenario positions at the 70 Gy level, the dose level was reduced to 69.5 Gy. The plans with optimized scenario positions satisfy the constraints in

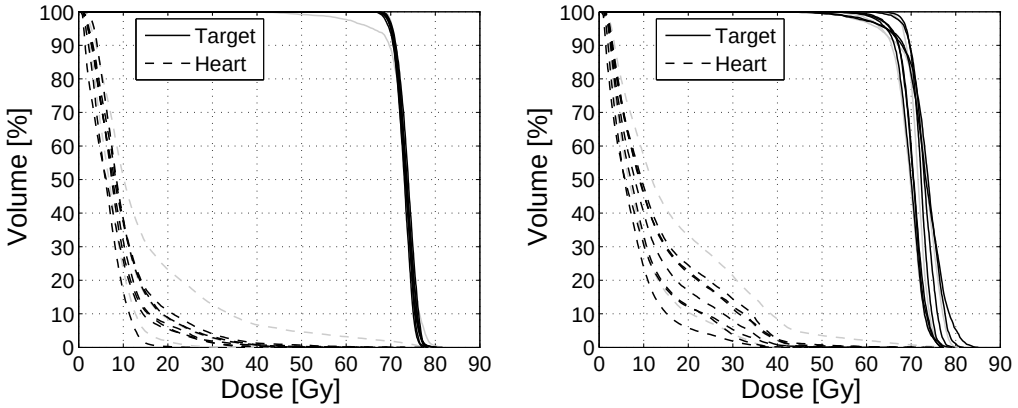

(a) IMRT plan using optimized scenario positions

(b) IMRT plan using a priori scenario positions

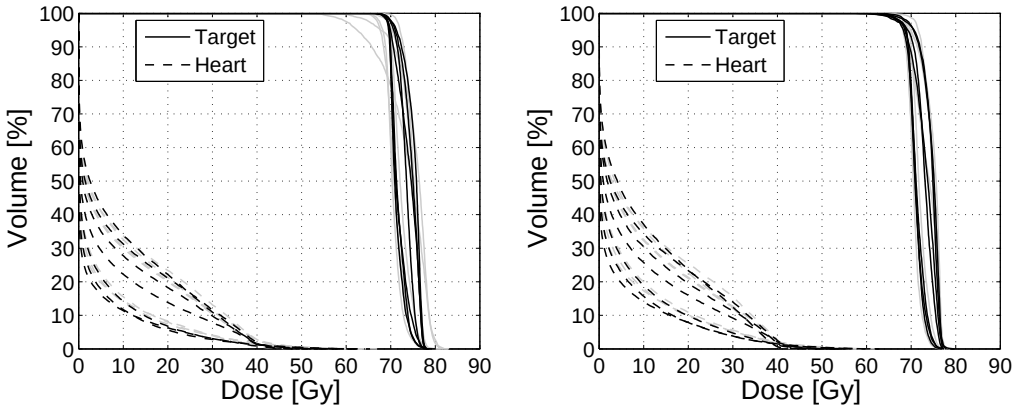

(c) IMPT plan using optimized scenario positions

(d) IMPT plan using a priori scenario positions

Figure 7. Lung case DVHs for the plans using optimized scenario positions and the plans using a priori scenario positions. Black curves correspond to the optimized scenario positions and gray curves correspond to the a priori scenario positions (± 1 cm in all axis directions).

a larger number of scenarios than the plans with a priori scenario positions for all constraints but the target max dose constraint.

To evaluate how well the different methods satisfied the robust constraints, each robust constraint of Table 3 (with volume level relaxed by 0.5 %) was evaluated in each of the optimized scenario positions as well as in the a priori scenario positions. Since no method satisfied the target max DVH constraint in any scenario, its dose level was relaxed from 70 Gy to 69.5 Gy in the evaluation. For each of these two scenario groups and each method, the number of satisfied constraints was summed over the scenarios. The results are shown in Table 5. The IMRT and IMPT plans optimized with the scenarios determined by the scenario position optimization satisfied respectively 13 and 10 constraints more in the optimized scenarios than the plans optimized with a priori constraints. They also satisfied respectively 12 and 2 constraints more in the a priori scenarios.

Table 5. The number of satisfied constraints for the lung case over the seven optimized scenarios and over the seven a priori scenarios. Four constraints and seven scenarios make the maximum 28.

	No. of satisfied constraints	
Method	Optimized scenarios	A priori scenarios
IMRT with optimized scenarios	21	19
IMRT with a priori scenarios	8	7
IMPT with optimized scenarios	20	11
IMPT with a priori scenarios	10	9

4. Discussion

For the simplified goals of the prostate case, the proposed method resulted in the intuitively correct solution: The posterior isocenter shift for the target and the anterior isocenter shift for the rectum did not move from their maximum positions, as could be expected since these scenarios were not in conflict with other scenarios. The other two shifts were moved to become compatible. By neglecting the 1 cm anterior shift for the target and the 1 cm posterior shift for the rectum, the scenario position optimization enabled better solutions with respect to all other shifts than the solutions of the robust optimization with a priori, incompatible, scenario positions.

For the lung case, the intuitively correct solution would move the target right scenario and the heart left scenario. The scenario position optimization did, but other scenarios were moved as well. As for the prostate case, neglecting the worst scenario positions in some directions enabled better solutions with respect to the other positions—even to positions close to the worst ones. That multiple scenarios were retracted shows that it is hard to determine before the optimization which scenarios to retract. That they were retracted differently shows that it is hard to determine how much they should be retracted.

The differences between the robust plans with optimized scenarios and the robust

plans with a priori scenarios were larger for IMRT than for IMPT. The IMRT plans with optimized scenarios satisfied a larger number of constraints than the plans with a priori scenarios, also when evaluated in the a priori scenarios. This shows that for IMRT, one sometimes gets more than one asks for, which was also the rationale behind the iterative updates of margins performed by Gordon and Siebers [8]. For IMPT, the difference between the plans with optimized scenarios and those with a priori scenarios when they were evaluated in the a priori scenarios were smaller, but the plans with optimized scenarios still satisfied a larger number of constraints. A reason for these results could be that IMPT has more degrees of freedom than does IMRT, so that it has more possibilities to yield solutions that closely comply with what is requested.

For the prostate case, the enumeration of the scenario positions showed that the positions determined by the scenario position optimization could not be much improved upon while maintaining feasibility. For the lung case, enumerating the positions was deemed too computationally demanding, but when the scenario positions were simultaneously extended by 0.1 mm, no feasible solution was found. This shows that although the scenario position optimization problem is nonconvex, it resulted in solutions close to what could be best achieved for the considered cases.

The assumption that the first term of the objective in (2.1) could be neglected was true for most studied cases, but for the lung case subject to protons, the left scenarios for the target and the OAR were strictly below their upper bounds. Since the target scenario was very close to its upper bound, it is plausible that the difference would be small between the solutions with both scenarios included and with only the minimally shifted scenario considered. In a static dose distribution, the left target scenario would not affect the other scenarios. That it was still retracted could be due to that IMPT dose distributions in heterogeneous media deform as a consequence of setup errors, and that the target left scenario thereby was in conflict with other scenarios. Alternatively, it could be due to that the scenario was retracted initially and that the optimization was unable to push it outwards fast enough. Since the optimization problem is nonconvex, there is no guarantee that the global optimum is found by the method. However, the solution was only 0.6 mm from its maximum value.

In this paper, only setup errors were considered. For IMPT, range uncertainty is another influential error source [11]. If interpolation between the energy layers in the spot grid can be used to approximate the effects of range errors, doing so would provide a parameterization of range uncertainty scenarios that could be used in a scenario position optimization like the one used for setup errors in this paper.

It is important to be aware of the fact that even if it is impossible to be spare an OAR robustly in a given direction, it might still be beneficial to aim for robust target coverage in that direction. This was the case for the two studied cases for both applied modalities. The proposed method achieves this by optimizing individual scenario positions for different structures.

5. Conclusion

A method was proposed that determines which scenario positions to incorporate in robust optimization. The method maximizes the probability that the realized error will be accounted for by optimizing the scenario positions under the constraints that some minimum requirements must be satisfied in all of the scenarios. More specifically, each scenario is given a direction, and the optimization algorithm is allowed to adjust how far out along the given direction each scenario should be located: the farther, the better objective value, but also the harder to satisfy the constraints. The method thus provides a way of determining what can reasonably be asked for.

For a prostate case and a lung case, the optimized scenario positions enabled more robust plans than optimization with a priori scenario positions. When the scenario positions were pushed outwards, the optimization algorithm was unable to find feasible solutions. This shows that, for the cases studied, the scenario position optimization did not lead to overly retracted scenario positions.

Acknowledgement

The research was supported by the Swedish Research Council (VR).

A. Optimization functions

Given the optimization variables $x \in \mathcal{X}$, let $D(v; x)$ parameterize the DVH of some considered ROI as a function of the volume $v \in (0, 1]$. A max DVH optimization function with dose level $\hat{d}$ and volume parameter $\hat{v}$ is given by

$$\int_{\hat{v}}^{1} \left(D(v; x) - \hat{d} \right)_{+}^{2} dv.$$

Min DVH functions are defined analogously, but with the signs of $D(v; x)$ and $\hat{d}$ reversed and with the integration taken over $(0, \hat{v}]$. Max and min dose functions are derived from the corresponding DVH functions with $\hat{v}$ set to respectively 0 and 1. A mean dose function is given by

$$\left(\int_{0}^{1} D(v; x)\, dv \right)^{2}.$$

References

[1] A. Ahnesjö. Collapsed cone convolution of radiant energy for photon dose calculation in heterogeneous media. *Medical Physics*, 16(4):577–592, 1989.

[2] T. Bortfeld, W. Schlegel, and B. Rhein. Decomposition of pencil beam kernels for fast dose calculations in three-dimensional treatment planning. *Medical Physics*, 20(2):311–318, 1993.

[3] 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. *Physics in Medicine and Biology*, 57(3):591–608, 2012.

[4] A. Fredriksson. A characterization of robust radiation therapy treatment planning methods—from expected value to worst case optimization. *Medical Physics*, 39(8):5169–5181, 2012.

[5] A. Fredriksson, A. Forsgren, and B. Hårdemark. Minimax optimization for handling range and setup uncertainties in proton therapy. *Medical Physics*, 38(3):1672–1684, 2011.

[6] P. E. Gill, W. Murray, and M. A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47(1):99–131, 2005.

[7] 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, 2010.

[8] J. J. Gordon and J. V. Siebers. Coverage-based treatment planning: optimizing the IMRT PTV to meet a CTV coverage criterion. *Medical Physics*, 36(3):961–973, 2009.

[9] B. Hårdemark, A. Liander, H. Rehbinder, and J. Löf. Direct machine parameter optimization with RayMachine® in Pinnacle^3®. White paper, RaySearch Laboratories, Stockholm, Sweden, 2003.

[10] International Commission on Radiation Units and Measurements. Report 50: Prescribing, recording, and reporting photon beam therapy. ICRU, 1993.

[11] A. J. Lomax. Intensity modulated proton therapy and its sensitivity to treatment uncertainties 1: the potential effects of calculational uncertainties. *Physics in Medicine and Biology*, 53(4):1027–1042, 2008.

[12] A. J. Lomax. Intensity modulated proton therapy and its sensitivity to treatment uncertainties 2: the potential effects of inter-fraction and inter-field motions. *Physics in Medicine and Biology*, 53(4):1043–1056, 2008.

[13] J. A. Moore, J. J. Gordon, M. S. Anscher, and J. V. Siebers. Comparisons of treatment optimization directly incorporating random patient setup uncertainty with a margin-based approach. *Medical Physics*, 36(9):3880–3890, 2009.

[14] D. Pflugfelder, J. J. Wilkens, and U. Oelfke. Worst case optimization: a method to account for uncertainties in the optimization of intensity modulated proton therapy. *Physics in Medicine and Biology*, 53(6):1689–1700, 2008.

[15] J. C. Stroom, H. C. de Boer, H. Huizenga, and A. G. Visser. Inclusion of geometrical uncertainties in radiotherapy treatment planning by means of coverage probability. *International Journal of Radiation Oncology*Biology*Physics*, 43(4):905–919, 1999.

[16] J. Unkelbach, T. Bortfeld, B. C. Martin, and M. Soukup. Reducing the sensitivity of IMPT treatment plans to setup errors and range uncertainties via probabilistic treatment planning. *Medical Physics*, 36(1):149–163, 2009.

[17] J. Unkelbach and U. Oelfke. Inclusion of organ movements in IMRT treatment planning via inverse planning based on probability distributions. *Physics in Medicine and Biology*, 49(17):4005–4029, 2004.

[18] M. van Herk. Errors and margins in radiotherapy. *Seminars in radiation oncology*, 14(1):56–64, 2004.