Maximizing the probability of satisfying the clinical goals in radiation therapy treatment planning under setup uncertainty

Albin Fredriksson[a)]
RaySearch Laboratories, Sveavägen 44, Stockholm SE-111 34, Sweden

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

Björn Hårdemark
RaySearch Laboratories, Sveavägen 44, Stockholm SE-111 34, Sweden

(Received 19 September 2014; revised 28 April 2015; accepted for publication 18 May 2015; published 10 June 2015)

Purpose: This paper introduces a method that maximizes the probability of satisfying the clinical goals in intensity-modulated radiation therapy treatments subject to setup uncertainty.
Methods: The authors perform robust optimization in which the clinical goals are constrained to be satisfied whenever the setup error falls within an uncertainty set. The shape of the uncertainty set is included as a variable in the optimization. The goal of the optimization is to modify the shape of the uncertainty set in order to maximize the probability that the setup error will fall within the modified set. Because the constraints enforce the clinical goals to be satisfied under all setup errors within the uncertainty set, this is equivalent to maximizing the probability of satisfying the clinical goals. This type of robust optimization is studied with respect to photon and proton therapy applied to a prostate case and compared to robust optimization using an *a priori* defined uncertainty set.
Results: Slight reductions of the uncertainty sets resulted in plans that satisfied a larger number of clinical goals than optimization with respect to *a priori* defined uncertainty sets, both within the reduced uncertainty sets and within the *a priori*, nonreduced, uncertainty sets. For the prostate case, the plans taking reduced uncertainty sets into account satisfied 1.4 (photons) and 1.5 (protons) times as many clinical goals over the scenarios as the method taking *a priori* uncertainty sets into account.
Conclusions: Reducing the uncertainty sets enabled the optimization to find better solutions with respect to the errors within the reduced as well as the nonreduced uncertainty sets and thereby achieve higher probability of satisfying the clinical goals. This shows that asking for a little less in the optimization sometimes leads to better overall plan quality. [http://dx.doi.org/10.1118/1.4921998]

Key words: IMRT, IMPT, optimization, robustness, uncertainty

1. INTRODUCTION

Misalignment of the patient relative to the beams can lead to large differences between the planned and the delivered dose distributions in external beam radiation therapy. The conventional approach to take such errors into account is to apply margins during treatment planning. Planning is then performed to yield high dose to an enlarged target volume, called the planning target volume (PTV), and low doses to enlarged organ at risk (OAR) volumes.

For cases of heterogeneous density, and especially cases subject to ion treatment, conventional margins do not always provide the intended robustness against uncertainties.[1] Methods that explicitly compute dose distributions under a number of scenarios (e.g., patient setup shifts) and optimize all scenario doses simultaneously appear to lead to more robust plans in general.[2–5] Many of the methods intended to create robust plans are based on robust optimization and aim to minimize the objective function value under the worst scenario. If the plan quality under the worst scenario is unacceptable, the optimization algorithm has no incentive to make the quality under other scenarios acceptable, which can lead to unnecessarily poor plan quality under these scenarios.[6] This calls for a method that determines which scenarios to consider in the robust optimization in order to achieve acceptable plan quality under as many scenarios as possible. In the present paper, we devise such a method.

The goal of the method that we propose is to maximize the probability that the setup error will fall within an uncertainty set, subject to the constraint that the clinical goals are satisfied for all setup errors within this set. Because it cannot be known *a priori* whether it is possible to satisfy the clinical goals within a given uncertainty set, the method includes the shape of the uncertainty set as a variable in the optimization. The optimization thus simultaneously determines the treatment plan and the set of errors against which the treatment plan should be robust.

The rationale for modifying the size of the uncertainty set is the same as the rationale for modifying margins. Margins are generally specified in accordance with the magnitudes of measured errors in order to ensure a high probability of

meeting the planning goals.[7] However, in cases for which the conflicting goals between PTV coverage and OAR sparing cannot be simultaneously fulfilled, the margins are sometimes retracted in the directions that correspond to the problematic conflicts.[8] Hence, prostate tumors are often given less margin toward the rectum than in the other directions, not because the prostate is less likely to move posteriorly, but because rectum sparing is considered more important than robustness against such errors. A plan with such a margin is not robust against the target moving posteriorly but on the other hand satisfies the clinical goal of rectum sparing.

Changing the magnitudes of the errors against which the plan should be robust has similarities to some previous methods: Gordon and Siebers[9] iteratively updated the sizes of the PTVs and reoptimized the treatment plans until a specified probability of target coverage was met. Gordon *et al.*[10] and Moore *et al.*[11] considered multiple setup error scenarios and tried to meet the clinical goals for a specified fraction of these. Bohoslavsky *et al.*[12] optimized the expected value of the objective conditioned on that one of the P best outcomes will occur. Rather than trying to reach a predetermined probability or optimizing for a specific quantile, our method optimizes the probability of satisfying the clinical goals directly. It can thereby reach the highest probability for which the goals can be fulfilled. The methods of Yang *et al.*[13] and Sobotta *et al.*[14] consider a number of predetermined error scenarios and optimize over these to maximize the probability of satisfying the clinical goals. The method proposed in the present paper differs from these methods in that it does not use predetermined error scenarios. Instead, it considers uncertainty sets of errors, the dimensions of which can be modified. This makes it feasible to consider a smaller number of error scenarios, which reduces the planning time. Moreover, the ability to continuously modify the dimensions of the uncertainty sets enables higher precision than when the set of considered scenarios is predetermined.

To demonstrate the method that we introduce, we use it to perform retrospective planning on a prostate case using both photon-mediated intensity-modulated radiation therapy (IMRT) and intensity-modulated proton therapy (IMPT). The resulting plans are compared to plans optimized to be robust against all errors within *a priori* selected uncertainty sets.

2. PRELIMINARIES

Before formulating the problem in which the shape of the considered uncertainty set is optimized, we give a standard robust optimization formulation that aims to robustly satisfy the clinical goals of a radiation therapy treatment without modifying the uncertainty set. We also describe how the probability of satisfying the clinical goals can be computed on the basis of an uncertainty set.

2.A. Robust optimization with respect to clinical goals

Robust optimization with respect to clinical goals can be formulated as an optimization problem with constraints enforcing the clinical goals of all region of interests (ROIs) to be satisfied whenever the setup error falls within a given uncertainty set.

To formulate the robust optimization problem mathematically, we denote by $d(x;s)$ the dose distribution as a function of the variables x (machine parameters or spot weights) and the setup shift $s \in \mathbb{R}^3$. For each ROI r from the set $\mathcal{R}$ enumerating the ROIs, we assume that the function f_r of dose evaluates to 0 or less whenever the clinical goals of ROI r are satisfied. The uncertainty set is the set of setup errors under which the clinical goals should be satisfied. It is denoted by U and is a subset of $\mathbb{R}^3$. This enables us to formulate the optimization problem for robustly satisfying the clinical goals as

$$\begin{aligned} &\underset{x\in\mathcal{X}}{\text{minimize}} && g(x) \\ &\text{subject to} && f_r(d(x;s)) \le 0, \quad r \in \mathcal{R}, \quad s \in U, \end{aligned} \tag{2.1}$$

where the objective g is selected to reflect some goal that is secondary to the clinical goals, such as minimizing the dose outside the target, and $\mathcal{X}$ is the set of feasible variables.

2.B. Probability of satisfying the clinical goals

A solution to problem (2.1) satisfies the clinical goals if the setup error falls within the uncertainty set U. The probability that the clinical goals are satisfied is therefore given by the probability that the setup error falls within U. Here, we specify how this probability can be calculated.

We consider radiation therapy treatment subject to systematic setup uncertainty. The coordinate system is defined such that the patient is fixed and that the beam isocenters move as a result of the setup error. The setup error is a random variable vector S that takes on values s in $\mathbb{R}^3$ according to some assumed probability distribution. The value s determines the shift of the beam isocenters.

A solution to problem (2.1) ensures that whenever the setup error S takes on a value s that lies in the uncertainty set U, then $f_r(d(x;s))$ will evaluate to 0 or less, implying that the clinical goals of ROI r are satisfied. The probability that the clinical goals will be satisfied is thus given by

$$\begin{aligned} &\mathbb{P}(\text{the clinical goals of all ROIs are satisfied}) \\ &\quad = \mathbb{P}(f_r(d(x;S)) \le 0 \text{ for all } r \in \mathcal{R}) \\ &\quad = \mathbb{P}(S \in U) \\ &\quad = \int_U \text{pdf}(s)\, ds, \end{aligned} \tag{2.2}$$

where $\mathbb{P}(A)$ denotes the probability of the event A occurring and $\text{pdf}(s)$ is the probability density function of the random variable S.

3. METHOD

When there is no treatment plan such that the clinical goals are satisfied under all setup errors within U, formulation (2.1) often yields unacceptable plans. Here, we present a method that avoids this problem by modifying the size of the uncertainty set U until a plan that satisfies the clinical goals under all setup errors within the modified uncertainty set can be found.

The set U is modified with the goal of maximizing probability equation (2.2).

We begin by introducing an idealized method for optimizing the uncertainty set. A computationally tractable approximation of the idealized method is given in Appendix A.

3.A. Uncertainty set optimization

To enable optimization of the uncertainty set, we introduce a parameterization of U: For each ROI r in $\mathcal{R}$, we consider an individual uncertainty set $U(\alpha_r)$, in which the goals of ROI r should be satisfied. The set $U(\alpha_r)$ is a function of the non-negative vector α_r, each component of which specifies the size of $U(\alpha_r)$ in a positive or negative axis direction. Because we consider errors in 3D, α_r is an element of $\mathbb{R}^6_+$. In nonaxis directions, $U(\alpha_r)$ is interpolated, see Appendix B for a mathematical definition. Figure 1 shows a 2D illustration of an uncertainty set $U(\alpha_r)$. The set U, in which the clinical goals of all ROIs are satisfied, is then given by $U = \cap_{r\in\mathcal{R}} U(\alpha_r)$.

Our goal is to find plans that maximize the probability of satisfying the clinical goals, which is the same as maximizing Eq. (2.2). We include the parameters α_r for all ROIs as variables in the optimization and combine objective equation (2.2) with the constraints of Eq. (2.1) to achieve an idealized uncertainty set optimization problem according to

$$\begin{aligned} \underset{\alpha,x}{\text{maximize}} \quad & \mathbb{P}(S \in U(\alpha_r) \text{ for all } r \in \mathcal{R}) \\ \text{subject to} \quad & f_r(d(x;s)) \le 0, \quad r \in \mathcal{R}, \quad s \in U(\alpha_r), \\ & \alpha_{r,i} \ge 0, \qquad\quad r \in \mathcal{R}, \quad i = 1,\ldots,6, \\ & x \in \mathcal{X}. \end{aligned} \tag{3.1}$$

The objective function measures the probability that the setup error S falls within $U = \cap_{r\in\mathcal{R}} U(\alpha_r)$ and the constraints ensure that for each ROI r, the clinical goals of the ROI are satisfied when the setup error falls within the region $U(\alpha_r)$. Thus, the clinical goals of all ROIs are satisfied when the setup error falls within U. Note that g from Eq. (2.1) is not considered in Eq. (3.1). Once Eq. (3.1) is solved, g can be considered in a second optimization with α fixed.

Problem (3.1) can be seen as roughly maximizing $\alpha_{r,i}$ for all r in $\mathcal{R}$ and $i = 1,\ldots,6$. However, while the constraints for all s in $U(\alpha_r)$ in Eq. (3.1) do not take the probability of the different scenarios into account, the objective function does. Hence, because large setup errors are generally less probable than small ones, it will be worth more in terms of improved objective value to increase $\alpha_{r,i}$ when it is close to zero than when it is large. This means that it is worth more to extend $U(\alpha_r)$ so that it includes regions close to the origin than regions distant from the origin.

In order to achieve a computationally tractable problem, we discretize Eq. (3.1). Specifically, the uncertainty set $U(\alpha_r)$ is discretized into a set of scenarios $\alpha_{r,i}p_i$ for $i = 1,\ldots,6$, where p_i are the positive and negative axes unit directions. The discretization is further explicated in Appendix A.

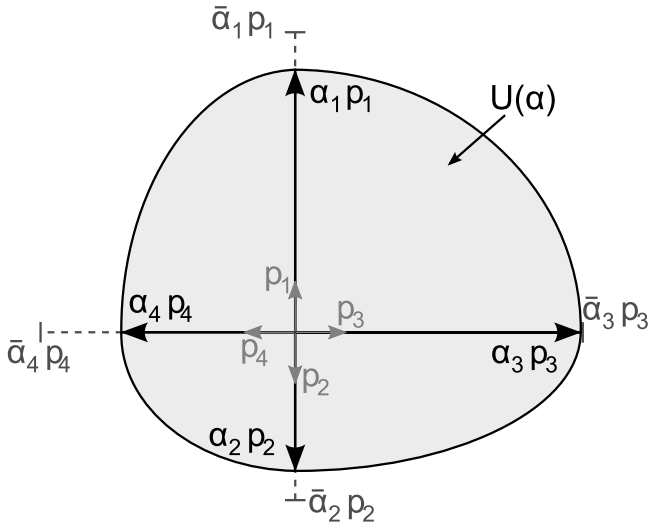

FIG. 1. 2D illustration of the uncertainty set $U(\alpha)$. The clinical goals are enforced whenever the setup error S falls within $U(\alpha)$. The non-negative scalar α_i determines the size of the uncertainty set $U(\alpha)$ in the axis unit direction p_i for $i = 1,\ldots,4$. In other directions, $U(\alpha)$ interpolates elliptically between these points. The scalar $\bar{\alpha}_i$ corresponds to the *a priori* value of α_i.

3.B. Computational study

3.B.1. Patient cases

Practical uncertainty set optimization according to Eq. (A1) was applied to a 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. The dose grid resolution was $2.5\times2.5\times2.5$ mm^3. The setup errors were assumed to follow a truncated normal distribution with standard deviation 0.5 cm, bounded to have length at most 1 cm. A transversal slice of the patient is shown in Fig. 2.

3.B.2. Optimization

Uncertainty set optimization according to Eq. (A1) was implemented in a research version of the RayStation 2.8 treatment planning system (RaySearch Laboratories, Stockholm, Sweden). The optimization in RayStation is performed by a sequential quadratic programming algorithm. A similar method is described by Gill *et al.*[15] The uncertainty set optimizations were started from plans optimized for seven iterations with fixed scenario positions. For IMRT, these iterations were

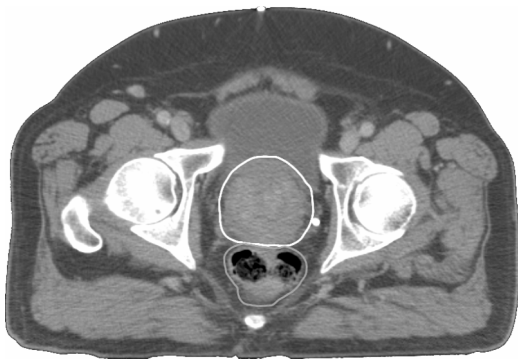

FIG. 2. Transversal slice of the patient case. Contours indicate the target (white) and the rectum (gray).

fluence map optimization iterations. The resulting plans were converted to segments before the scenario positions were optimized in combination with direct step-and-shoot optimization. The optimization functions constituting the objectives and constraints used in the present paper are the standard quadratically penalizing functions defined mathematically by, e.g., Oelfke and Bortfeld.[16]

When problem (A1) has been solved, α determines the largest errors that can be taken into account while the clinical goals are satisfied. When the uncertainty set optimization had determined α, we performed an additional optimization with α fixed to the determined value. The resulting plans were compared to plans optimized with α fixed to the *a priori* value corresponding to setup shifts of 1 cm from the nominal scenario. The optimization problem with fixed values of α is formulated as

$$\begin{aligned} \underset{x\in X}{\text{minimize}} \quad & \sum_{r\in\mathcal{O}} g_r(d(x;0)) \\ \text{subject to} \quad & f_r(d(x;\alpha_{r,i}p_i)) \le 0, \quad r\in\mathcal{R}, \quad i=0,\ldots,n, \\ & f_r(d(x;0)) \le 0, \quad r\in\mathcal{N}, \end{aligned} \tag{3.2}$$

where f_r is the same as in Eq. (A1), $\mathcal{N}$ enumerates the ROIs for which the clinical goals are enforced in the nominal scenario only, $\mathcal{O}$ enumerates the OARs, and g_r penalizes the mean dose of ROI r. This problem thus minimizes the mean dose to the OARs under the nominal scenario while requiring some clinical goals to be satisfied under all scenarios and other clinical goals to be satisfied under the nominal scenario.

3.B.3. Scenario dose calculation

During the optimizations, the scenario doses were calculated using the nominal mapping from fluence to dose, but with the fluence maps shifted (and bilinearly interpolated) according to the displacements of the scenarios.

In the robustness evaluation, the scenario doses were calculated with the patient positions shifted. The scenario dose calculation used during the optimization thus differed from that used in the evaluation. The discrepancy implies that constraints that are satisfied with respect to the doses used during optimization are not necessarily satisfied with respect to the evaluation doses.

Doses were computed by the dose calculation algorithms of RayStation. For IMRT, optimization was performed using RayStation's fast dose engine and followed by computation of final doses using its accurate dose engine. For IMPT, all doses were computed by its pencil beam algorithm. The line spacing and the energy layer separation (in water equivalent media) were both set to 5 mm, but to improve upon the approximate dose calculation with shifted fluences used during the optimization, auxiliary spots were computed for 2.5 mm line spacing, cf. Unkelbach *et al.*[5] The weights of the auxiliary spots were not included as variables in the optimization.

4. RESULTS

We applied the uncertainty set optimization method to a prostate case. In our proof-of-concept, we considered simplified uncertainty sets consisting of error scenarios along one axis only. To actually maximize the probability of satisfying the clinical goals, additional directions and scenarios should typically be included, see Appendix A for more details. We began by solving Eq. (A1) to optimize the uncertainty set and then solved Eq. (3.2) using the optimized uncertainty set and the *a priori* uncertainty set.

The computations were performed under Windows 7 on a 64-bit desktop computer with 24 GB of RAM and an Intel Xeon W3580 processor with four 3.33 GHz cores. Running 200 iterations of robust optimization with respect to clinical goals on form (3.2) for the prostate case took 10 (photons) and 140 (protons) min, while the increased number of dose computations required for approximating the derivatives with respect to α by finite differences leads to optimization times of 20 (photons) and 230 (protons) min for 200 iterations of uncertainty set optimization according to Eq. (A1).

4.A. Uncertainty set optimization

The prostate case was simplified to only include setup shifts in the anterior and posterior directions. The clinical goals for the target and the rectum of the prostate case were assumed to require robustness. Other goals were included in the nominal scenario only. The optimization problem was formulated according to Eq. (A1), using the clinical goals presented in Table I as constraints. The discretization included, in addition to the nominal scenario, two scenarios each for the target and the rectum (anterior and posterior isocenter shifts), each of which was variable in the optimization.

The optimizations changed neither the posterior isocenter shift for the target nor the anterior isocenter shift for the rectum, but kept these at their initial (maximum) positions of 1 cm. Figure 3 displays the progress of the two other scenario positions. The optimizations first retract the target anterior and rectum posterior isocenter shifts rapidly in order to satisfy the

TABLE I. Robust and nominal constraints representing the clinical goals 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			

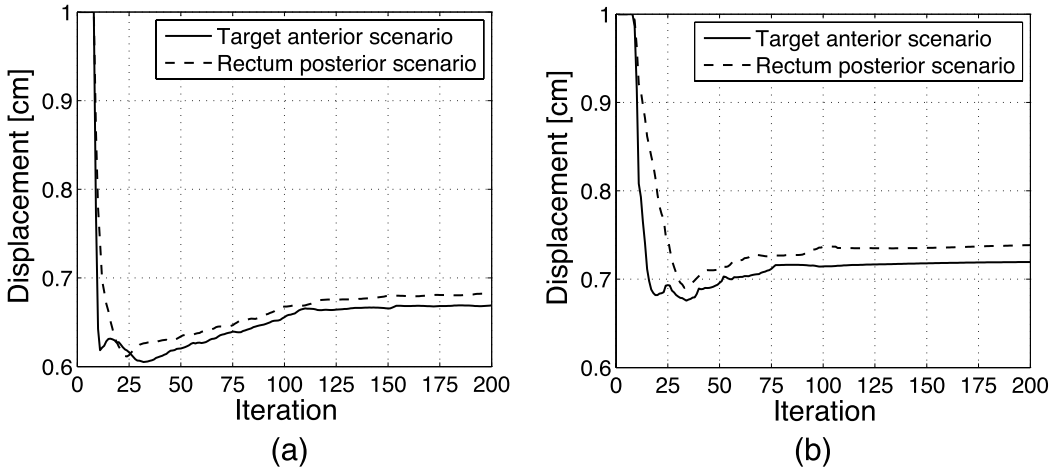

FIG. 3. Displacement of the scenario positions as function of the optimization iteration. The "target anterior scenario" is the scenario in which only the target is considered and the isocenter is shifted in the anterior direction, and the "rectum posterior scenario" the scenario where only the rectum is considered and the isocenter is shifted in the posterior direction. (a) Scenario positions for IMRT and (b) scenario positions for IMPT.

constraints. Shortly after iteration 25, the isocenter shifts are being pushed outward, which improves the objective value. For IMRT, the uncertainty set optimization resulted in the positions 0.67 and 0.68 cm for, respectively, the target anterior and the rectum posterior isocenter shift scenarios. For IMPT, it resulted in the positions 0.72 and 0.74 cm.

4.B. Robust plans with optimized scenarios

The optimized scenario positions were used as fixed positions in standard robust optimizations according to formulation (3.2). This problem was solved also with the scenario positions fixed at the *a priori* locations (1 cm in the posterior and anterior directions). Since, with the *a priori* scenarios, the constraints could not be satisfied, the target and rectum goals were relaxed into a robust objective constituent $\max_{i=0,\ldots,6}\sum_{r\in\mathcal{R}} f_r(d(x;\bar{\alpha}_{r,i}p_i))$ with weight 100 for that optimization. The resulting DVHs are shown in Fig. 4. For both modalities, the plans with optimized scenarios neglect the 1 cm anterior isocenter shift for the target but in return achieve better target coverage under the other scenarios than the plans optimized with *a priori* scenarios.

For each of the methods using optimized and *a priori* uncertainty sets, each robust constraint of Table I (with volume level tolerance 0.5% and dose level tolerance 0.5 Gy) was evaluated under the optimized and the *a priori* scenarios. The number of satisfied robust constraints over these scenarios is shown in Table II.

5. DISCUSSION

Robust optimization aims for plans that are robust over all error scenarios within some uncertainty set. When the constraints cannot be satisfied under some scenarios, the plan quality under all scenarios within the set may suffer. After a slight reduction of the uncertainty set, it is sometimes possible to achieve better plan quality with respect to the scenarios within the smaller set. The results of this paper show that such reduction can moreover enable better plan quality with respect to the majority of the scenarios within the larger uncertainty set.

For the simplified goals of the prostate case, uncertainty set optimization resulted in the intuitively correct solution, which furthermore coincides with a known practice for reducing margins:[8] The posterior isocenter shift scenario for the target and the anterior isocenter shift scenario for the rectum did not move from their maximum positions, as could be expected because these scenarios were not in conflict with other scenarios. The other two scenarios moved to become compatible. By retracting these scenarios, the uncertainty set optimization enabled better solutions with respect to all other scenarios and thereby achieved higher probability of satisfying the clinical goals than optimization with the scenario positions fixed at the *a priori* locations. Asking for a little less in the optimization sometimes leads to better overall plan quality.

Uncertainty set optimization bears similarities to worst case scenario optimization[17] in the sense that constraints enforce fulfillment of the clinical goals under all considered scenarios. However, the objective to be maximized is the probability that the setup error falls within a variable uncertainty set. The optimization thus takes into account the probability distribution and the fact that the probability of errors diminishes with the magnitude of the errors. This is reminiscent of probabilistic planning concepts,[18,5] in which the expected value of the objective function is optimized and the fact that the probability of errors diminishes with the size of the errors is directly reflected by that the objective contributions of such error scenarios are weighted by the small probabilities. A notable difference is that in probabilistic planning, all scenarios are considered in the objective, but weighted by the probability of the scenario occurring, whereas in uncertainty set optimization, scenarios not contained in the uncertainty set are neglected while the clinical goals are required to be fully satisfied for scenarios that are in the uncertainty set.

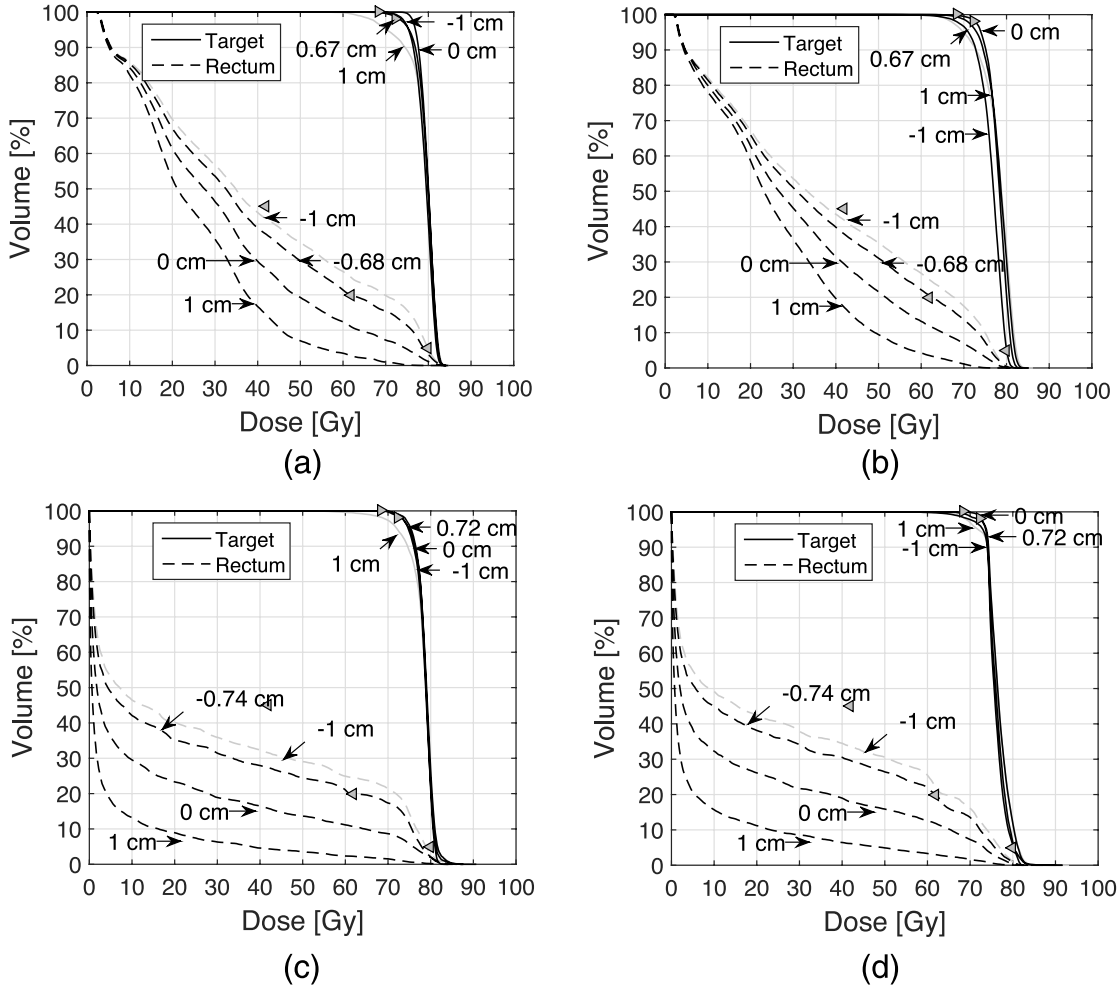

FIG. 4. DVHs for the plans optimized with respect to optimized scenario positions and the plans optimized with respect to *a priori* scenario positions. The setup shifts in the anterior direction are annotated. Black curves correspond to the optimized scenario positions and gray curves correspond to the *a priori* scenario positions (setup shifts of ±1 cm in the anterior direction). Gray triangles represent the robust constraints of Table I. (a) IMRT plan optimized with respect to optimized scenario positions, (b) IMRT plan optimized with respect to *a priori* scenario positions, (c) IMPT plan optimized with respect to optimized scenario positions, and (d) IMPT plan optimized with respect to *a priori* scenario positions.

A limitation of our approach is that we only take systematic setup errors into account. Random errors could be incorporated by dose or fluence blurring (see, e.g., Refs. 11 and 12), but this requires the biological assumption that overdosing an OAR in one fraction can be compensated by a lower dose in another fraction, which we are inclined not to make.

TABLE II. The number of satisfied robust constraints for the prostate case under the three optimized scenarios and under the three *a priori* scenarios of each ROI. Five constraints and three scenarios make the maximum 15 for each group of scenarios.

	No. of satisfied constraints in evaluation		
Method	Optimized scenarios	*A priori* scenarios	Total
IMRT with respect to optimized scenarios	14	11	25
IMRT with respect to *a priori* scenarios	9	9	18
IMPT with respect to optimized scenarios	15	12	27
IMPT with respect to *a priori* scenarios	9	9	18

6. CONCLUSION

A method was introduced that maximizes the probability of satisfying the clinical goals in the presence of setup uncertainty. The method was applied to a prostate case and the resulting plans were compared to plans optimized with respect to errors within *a priori* defined uncertainty sets. The plans of the proposed method fulfilled 1.4 (photons) and 1.5 (protons) times as many clinical goals over the scenarios as the method using *a priori* defined uncertainty sets.

APPENDIX A: PRACTICAL UNCERTAINTY SET OPTIMIZATION

Because $U(\alpha_r)$, as defined in Appendix B, contains infinitely many points (unless $\alpha_r = 0$), formulation (3.1) has

infinitely many constraints, which moreover change with α. Thus, Eq. (3.1) cannot be easily solved. We therefore approximate Eq. (3.1) by discretizing the uncertainty sets $U(\alpha_r)$ into scenarios.

The scenario discretization points are selected as the points $\alpha_{r,i}p_i$, as illustrated in Fig. 1, where p_i for $i = 1,\ldots,6$ are the positive and negative axes unit directions. The point $\alpha_{r,i}p_i$ is referred to as the "scenario position" of scenario i for ROI r. The maximal error to be taken into account in the direction p_i for ROI r is denoted by $\bar{\alpha}_{r,i}p_i$, where $\bar{\alpha}_{r,i}$ is scalar. Hence, $\alpha_{r,i}$ must be less than $\bar{\alpha}_{r,i}$. If unrestricted maximization of the probability is desired, then $\bar{\alpha}_{r,i}$ can be set infinity. The nominal scenario is also included: we let $\alpha_{r,0}p_0$ be fixed and correspond to the nominal scenario by setting $\alpha_{r,0} = 0$ and $\bar{\alpha}_{r,0} = 0$. The optimization problem approximating Eq. (3.1) can now be formulated as the uncertainty set optimization problem

$$\begin{aligned} \underset{\alpha,x}{\text{maximize}} \quad & \mathbb{P}(S \in U(\alpha_r) \text{ for all } r \in \mathcal{R}) \\ \text{subject to} \quad & f_r(d(x;\alpha_{r,i}p_i)) \le 0, \quad r \in \mathcal{R}, \quad i = 0,\ldots,6, \\ & 0 \le \alpha_{r,i} \le \bar{\alpha}_{r,i}, \quad r \in \mathcal{R}, \quad i = 0,\ldots,6, \\ & f_r(d(x;0)) \le 0, \quad r \in \mathcal{N}, \\ & x \in \mathcal{X}, \end{aligned} \tag{A1}$$

where nominal constraints, requiring the clinical goals to be satisfied under the nominal scenario only, have been added for the ROIs enumerated by the set $\mathcal{N}$. Computational details of the objective function are given in Appendix C.

Formulation (A1) approximates Eq. (3.1) by considering only setup error positions along the main axes. The approximation thus includes the assumption that if the goals for ROI r in $\mathcal{R}$ are satisfied in the nominal scenario and for a few points on the boundary of $U(\alpha_r)$ (along the main axes), they will be satisfied in the full uncertainty set $U(\alpha_r)$. Such approximations were empirically found to be adequate in previous studies.[19,20,17] However, when the magnitudes of the errors are large, the approximation will not hold, and solving Eq. (A1) will not correspond to maximizing the probability of satisfying the clinical goals. In such cases, additional constraints for other points in $U(\alpha_r)$ should then be included.

When it is not possible to satisfy the clinical goals of all ROIs simultaneously, it may still be beneficial to try to satisfy the goals for as many ROIs as possible. To this end, an individual term $\mathbb{P}(S \in U(\alpha_r))$ for each ROI r in $\mathcal{R}$, weighted by a small factor, can be added to the objective of Eq. (A1). The factor should be small enough for these terms not to affect the objective of Eq. (A1) substantially.

Note that the dose is generally nonconvex in the setup shift, so problem (A1) is a nonconvex optimization problem. Therefore, the uncertainty sets resulting after optimization cannot be guaranteed to be globally optimal.

The computational cost of scenario-based robust optimization problems is dominated by the cost of scenario dose calculation. In order to solve Eq. (A1), gradients with respect to α must be calculated. These can be approximated by finite differences, which leads to at most a doubling of the number of dose calculations compared to robust optimization with fixed scenarios according to Eq. (3.2).

APPENDIX B: INTERPOLATION OF UNCERTAINTY SET

The vector $\alpha_r \in \mathbb{R}^6_+$ specifies the size of $U(\alpha_r)$ in the positive and negative axes directions. In nonaxis directions, $U(\alpha_r)$ is ellipsoidally interpolated. Formally, $U(\alpha_r)$ is defined by

$$\begin{aligned} U(\alpha_r) = \{(x,y,z) \in \mathbb{R}^3 : {} & x^2/\alpha^2_{r,1+\chi(x<0)} + y^2/\alpha^2_{r,3+\chi(y<0)} \\ & + z^2/\alpha^2_{r,5+\chi(z<0)} \le 1\}, \end{aligned}$$

where $\chi(a)$ is the indicator function that takes the value 1 when the expression a is true and 0 otherwise. A 2D illustration of a set $U(\alpha_r)$ according to this definition is given in Fig. 1.

APPENDIX C: OBJECTIVE FUNCTION

The objective function of Eq. (A1) is the probability that the setup error falls within U, the intersection of $U(\alpha_r)$ for all ROIs r in $\mathcal{R}$. The setup errors are assumed to be normally distributed with zero mean and covariance matrix $\Sigma = \sigma^2 I$. The probability that the setup error S will fall within the set U is therefore

$$\mathbb{P}(S \in U) = \frac{1}{\sigma^3\sqrt{(2\pi)^3}} \int_U e^{-\|s\|^2/2\sigma^2}\, ds. \tag{C1}$$

If there is only a single axis along which to move ($n = 2$, $p_1 = -1, p_2 = 1$), the probability density function becomes $(\sigma\sqrt{2\pi})^{-1}e^{-t^2/2\sigma^2}$. The set U is then given by

$$U = [-\min_{r\in\mathcal{R}} \alpha_{r,1}, \min_{r\in\mathcal{R}} \alpha_{r,2}],$$

so the probability that the goals of all ROIs in $\mathcal{R}$ are satisfied is

$$\begin{aligned} & \mathbb{P}(S \in [-\min_{r\in\mathcal{R}} \alpha_{r,1}, \min_{r\in\mathcal{R}} \alpha_{r,2}]) \\ & \quad = \frac{1}{\sigma\sqrt{2\pi}} \int_{-\min_{r\in\mathcal{R}}\alpha_{r,1}}^{\min_{r\in\mathcal{R}}\alpha_{r,2}} e^{-t^2/2\sigma^2}\, dt. \end{aligned}$$

[a]Electronic mail: albin.fredriksson@raysearchlabs.com

[1]A. Lomax, "Intensity modulated proton therapy and its sensitivity to treatment uncertainties 2: The potential effects of inter-fraction and inter-field motions," Phys. Med. Biol. **53**(4), 1043–1056 (2008).
[2]T. Chan, T. Bortfeld, and J. Tsitsiklis, "A robust approach to IMRT optimization," Phys. Med. Biol. **51**(10), 2567–2583 (2006).
[3]A. Fredriksson, "A characterization of robust radiation therapy treatment planning methods—From expected value to worst case optimization," Med. Phys. **39**(8), 5169–5181 (2012).
[4]D. Pflugfelder, J. Wilkens, and U. Oelfke, "Worst case optimization: A method to account for uncertainties in the optimization of intensity modulated proton therapy," Phys. Med. Biol. **53**(6), 1689–1700 (2008).
[5]J. Unkelbach, T. Bortfeld, B. Martin, and M. Soukup, "Reducing the sensitivity of IMPT treatment plans to setup errors and range uncertainties via probabilistic treatment planning," Med. Phys. **36**(1), 149–163 (2009).
[6]A. Fredriksson and R. Bokrantz, "A critical evaluation of worst case optimization methods for robust intensity-modulated proton therapy planning," Med. Phys. **41**(8), 081701 (11pp.) (2014).
[7]M. van Herk, "Errors and margins in radiotherapy," Semin. Radiat. Oncol. **14**(1), 52–64 (2004).
[8]M. Zhang, V. Moiseenko, and M. Liu, "PTV margin for dose escalated radiation therapy of prostate cancer with daily on-line realignment using

internal fiducial markers: Monte Carlo approach and dose population histogram (DPH) analysis," J. Appl. Clin. Med. Phys. **7**(2), 38–49 (2006).
[9]J. Gordon and J. Siebers, "Coverage-based treatment planning: Optimizing the IMRT PTV to meet a CTV coverage criterion," Med. Phys. **36**(3), 961–973 (2009).
[10]J. Gordon, N. Sayah, E. Weiss, and J. Siebers, "Coverage optimized planning: Probabilistic treatment planning based on dose coverage histogram criteria," Med. Phys. **37**(2), 550–563 (2010).
[11]J. Moore, J. Gordon, M. Anscher, and J. Siebers, "Comparisons of treatment optimization directly incorporating random patient setup uncertainty with a margin-based approach," Med. Phys. **36**(9), 3880–3890 (2009).
[12]R. Bohoslavsky, M. Witte, T. Janssen, and M. Van Herk, "Probabilistic objective functions for margin-less IMRT planning," Phys. Med. Biol. **58**(11), 3563–3580 (2013).
[13]J. Yang, G. Mageras, S. Spirou, A. Jackson, E. Yorke, C. Ling, and C.-S. Chui, "A new method of incorporating systematic uncertainties in intensity-modulated radiotherapy optimization," Med. Phys. **32**(8), 2567–2579 (2005).
[14]B. Sobotta, M. Söhn, and M. Alber, "Robust optimization based upon statistical theory," Med. Phys. **37**(8), 4019–4028 (2010).
[15]P. Gill, W. Murray, and M. Saunders, "SNOPT: An SQP algorithm for large-scale constrained optimization," SIAM Rev. **47**(1), 99–131 (2005).
[16]U. Oelfke and T. Bortfeld, "Inverse planning for photon and proton beams," Med. Dosim. **26**(2), 113–124 (2001).
[17]A. Fredriksson, A. Forsgren, and B. Hårdemark, "Minimax optimization for handling range and setup uncertainties in proton therapy," Med. Phys. **38**(3), 1672–1684 (2011).
[18]M. Bangert, P. Hennig, and U. Oelfke, "Analytical probabilistic modeling for radiation therapy treatment planning," Phys. Med. Biol. **58**(16), 5401–5419 (2013).
[19]M. Casiraghi, F. Albertini, and A. Lomax, "Advantages and limitations of the 'worst case scenario' approach in IMPT treatment planning," Phys. Med. Biol. **58**(5), 1323–1339 (2013).
[20]T. Chan, J. Tsitsiklis, and T. Bortfeld, "Optimal margin and edge-enhanced intensity maps in the presence of motion and uncertainty," Phys. Med. Biol. **55**(2), 515–533 (2010).