IOP Publishing | Institute of Physics and Engineering in Medicine Physics in Medicine & Biology

Phys. Med. Biol. **62** (2017) 1342–1357 doi:10.1088/1361-6560/aa524d

Scenario-based radiation therapy margins for patient setup, organ motion, and particle range uncertainty

Rasmus Bokrantz and Albin Fredriksson

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

E-mail: rasmus.bokrantz@raysearchlabs.com

Received 31 August 2016, revised 30 November 2016
Accepted for publication 7 December 2016
Published 23 January 2017

Abstract

This work extends and validates the scenario-based generalization of margins presented in Fredriksson and Bokrantz (2016 *Phys. Med. Biol.* **61** 2067–82). Scenario-based margins are, in their original form, a method for robust planning under setup uncertainty where the sum of a plan evaluation criterion over a set of scenarios is optimized. The voxelwise penalties in the summands are weighted by a distribution of coefficients defined such that the method is mathematically equivalent to the use of conventional geometric margins if the scenario doses are calculated using the static dose cloud approximation. The purpose of this work is to extend scenario-based margins to general types of geometric uncertainty and to validate their use on clinical cases. Specifically, we outline how to incorporate density heterogeneity in the calculation of coefficients and demonstrate the extended method's ability to safeguard against setup errors, organ motion, and range shifts (and combinations thereof). For a water phantom with a high-density slab partly covering the target, the extended form of scenario-based margins method led to improved target coverage robustness compared to the original method. At most minor differences in robustness were, however, observed between the extended and original method for a prostate and two lung patients, all treated with intensity-modulated proton therapy, yielding evidence that the calculation of weighting coefficients is generally insensitive to tissue heterogeneities. The scenario-based margins were, furthermore, verified to provide a comparable level of robustness to expected value and worst case optimization while circumventing some known shortcomings of these methods.

Keywords: IMPT, margins, robustness, uncertainties, optimization

(Some figures may appear in colour only in the online journal)

1. Introduction

Robust radiation therapy optimization methods have been investigated over the past decade (Unkelbach and Oelfke 2004, Chu *et al* 2005, Chan *et al* 2006, Ólafsson and Wright 2006, Pflugfelder *et al* 2008, Unkelbach *et al* 2009, Fredriksson *et al* 2011, Bohoslavsky *et al* 2013). In order to hedge against the effects of errors, these methods consider a number of error scenarios and aim for good plan quality in all scenarios simultaneously. Scenario-based methods have been demonstrated capable of providing a high degree of target coverage robustness with respect to systematic errors, even under difficult circumstances such as lung cancers treated with multiple individually nonuniform proton fields (Fredriksson *et al* 2011). The question of how to trade target coverage robustness against organ-at-risk (OAR) sparing has received less attention; rather, the tradeoff mechanism has been the haphazard consequence of the selected robust planning method (Fredriksson and Bokrantz 2014).

Geometric margins are the most widely used means of achieving robustness. When geometric margins are used, the tradeoff between target coverage robustness and OAR sparing is made locally, on a voxel-by-voxel basis, with all voxels within a region of interest (ROI) weighted equally (or possibly scaled by the fraction of the voxel that lies within the ROI). This tradeoff mechanism has several benefits: it is simple and well-understood by clinicians, and it promotes steep dose gradients at the edges of the treated volume (unlike if, for instance, coverage of the peripheral parts of targets have a lesser priority, which makes the dose fall-off more gradual). However, there exists a number of situations when a margin expansion is inadequate:

- A margin often fails to provide target coverage robustness against errors that perturb not only the edges but also the central parts of the high-dose region, such as systematic errors for multifield intensity-modulated proton therapy (IMPT) (Albertini *et al* 2011).
- A margin that extends into the air cannot provide target coverage robustness for superficial targets subject to motion, such as breast tumors treated during free breathing.
- Margins are not well-defined for inhomogeneous dose prescriptions, which occur in, e.g. dose-painting guided by functional imaging (Ling *et al* 2000) and adaptive replanning with respect to an accumulated dose (Stewart *et al* 2010).

In the present and a previous paper (Fredriksson and Bokrantz 2016), we devise a robust planning method— *scenario-based margins*—that considers multiple scenarios and can provide target coverage robustness even for difficult cases, but retains the tradeoff mechanism of conventional geometric margins. Moreover, when the static dose cloud approximation (Unkelbach *et al* 2009), which is implicitly assumed with the use of geometric margins, is exact, our scenario-based method is equivalent to geometric margins.

Our previous study was limited to stylized phantom cases in water subject to setup shifts that were perfectly aligned with the dose grid. In the present work, we show that scenario-based margins can be applied to a broad range of geometric uncertainties and that the method does not depend on nonstandard assumptions. Our contributions relative to our prior paper are:

- We provide a further generalized definition of scenario-based margins that is entirely independent of the static dose cloud approximation.
- We provide computationally practical approximation procedures of the further generalized definition with respect to setup uncertainty in heterogeneous media, organ motion uncertainty, and uncertainty in range of particle beams.
- We validate the quality of the proposed approximation procedures by comparative evaluation against worst case (Fredriksson *et al* 2011) and expected value optimization (Unkelbach *et al* 2009) on full-scale patient datasets in a clinical treatment planning system.

2. Methods

We formulate scenario-based margins independently of the static dose cloud approximation. We first give a formulation that applies to general errors. The general formulation is primarily meant for theoretical insights as it is too time and memory consuming to be practicable for realistic instances. Computationally practical realizations are then provided for some important types of geometric uncertainty, namely, setup uncertainty in homogeneous and heterogeneous media, range uncertainty for particle beams, and organ motion uncertainty. The formulations are, in order to keep notation light, outlined with respect to a single clinical target volume (CTV) and a plan evaluation criterion for this structure that is additively separable over the voxels. The formalism, however, applies to multiple structures, including OARs, as well as to other types of criteria (Fredriksson and Bokrantz 2016, section 2.5).

2.1. Notation

The CTV voxels are indexed by the set $\mathcal{C}$ and the voxelwise component of the plan evaluation criterion is given by the function ϕ. The full criterion is thus $\sum_{i\in\mathcal{C}}\phi(d_i)$, where d_i is the dose to voxel i. Uncertainty is represented by a finite set $\mathcal{S}$ of scenarios. The dose to voxel i from beam b under scenario s is denoted by $d_i^{(b)}(s)$ while $d_i(s) = \sum_{b\in\mathcal{B}} d_i^{(b)}(s)$ is the total dose to voxel i under scenario s from the set $\mathcal{B}$ of all beams. A beam-, scenario-, and voxel-specific weight $p_{i,s}^{(b)}$ is selected in a way that compensates for the number of times over the scenarios in $\mathcal{S}$ each voxel i shifted according to scenario s for beam b is contained within the CTV. We refer to the determination of $p_{i,s}^{(b)}$ as a calculation of *overlap volumes*. The notation $|\mathcal{X}|$ denotes cardinality when $\mathcal{X}$ is a finite set and volume when $\mathcal{X}$ is a continuous subset of $\mathbb{R}^3$. The dose-at-volume of a given ROI (the minimum dose such that the isodose volume contains $x\%$ of the ROI) is denoted by D_x.

2.2. The scenario-based formulation of margins

In our previous paper (Fredriksson and Bokrantz 2016, section 2.4), we derived the following formulation of a scenario-based CTV criterion taking beam-specific effects into account:

$$\sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{C}}\phi(d_i(s))\sum_{b\in\mathcal{B}}\frac{d_i^{(b)}(s)}{d_i(s)}p_{i,s}^{(b)}, \tag{1}$$

where $d_i^{(b)}(s)/d_i(s)$ is taken as $1/|\mathcal{B}|$ whenever $d_i(s) = 0$.

The main result in our previous paper is that formulation (1) is equivalent to a nominal criterion on a conventional planning target volume (PTV) if the static dose cloud approximation is exact, i.e. all shifts occur laterally relative to the beams, the material composition of the irradiated volume is homogenous, and the PTV is selected as the union of shifted CTVs. The advantage of (1) over a conventional PTV criterion is that when the static dose cloud approximation is incorrect, the former can still be used to achieve robust plans, using accurately computed scenario doses $d(s)$.

Although we used accurate scenario doses in our previous paper (Fredriksson and Bokrantz 2016), we still assumed the static dose cloud approximation in the calculation of overlap volumes. As a result, the weights $p_{i,s}^{(b)}$ may incorrectly reflect how many times the shifted voxels are within the CTV in reality. In the following sections, we give a general formulation that can be used to calculate overlap volumes for arbitrary uncertainties. We also propose more

computationally practical formulations that can be used to take into account setup uncertainty in heterogeneous media as well as range and organ motion uncertainty.

2.3. General geometric uncertainty

We provide a definition of overlap volumes that is entirely independent of the static dose cloud approximation. For brevity, we assume that scenario doses are calculated by superposition of pencil beam kernels, e.g. kernels per bixel or per scanning spot. This assumption permits us to take into account that the effect of an error varies within a single beam without modification of (1), because the summation over the set of beams $\mathcal{B}$ directly generalizes to a summation over the dose kernels.

The static dose cloud assumption can be lifted if the weights $p_{i,s}^{(b)}$ are calculated with respect to a physically correct description of how the dose kernels deform between scenarios, according to

$$p_{i,s}^{(b)} = 1/\left(\sum_{s' \in \mathcal{S}} \sum_{i' \in \mathcal{C}(s')} A_{i,i'}^{(b)}(s, s')\right). \tag{2}$$

The coefficients $A_{i,i'}^{(b)}(s, s')$ are elements of a $|\mathcal{C}| \times |\mathcal{C}|$ deformation matrix, each describing the fraction of the dose from kernel b to voxel i under scenario s that is deposited inside voxel i' under scenario s'. The set of CTV voxels $\mathcal{C}$ is, for further generality, allowed to vary with the scenario s'. This scenario-dependency requires a substitution of $\mathcal{C}(s)$ for $\mathcal{C}$ in (1). The weights defined by (2) satisfy $1/|\mathcal{S}| \leqslant p_{i,s}^{(b)} \leqslant 1$ because $\sum_{i' \in \mathcal{C}(s')} A_{i,i'}^{(b)}(s, s') \leqslant A_{i,i}^{(b)}(s, s) = 1$. The weights can, furthermore, be verified to coincide with the original definition of overlap volumes (Fredriksson and Bokrantz 2016, equation (2.4)) if the static dose cloud approximation is valid (the matrices $A^{(b)}(s, s')$ become binary-valued under this circumstance).

The calculation and storage of deformation matrices $A^{(b)}(s, s')$ for each dose kernel b and pair of scenarios (s, s') is in general too computationally expensive in practice. In the following sections, we outline computational efficient approximations of the deformations for some commonly encountered types of uncertainty.

2.4. Setup uncertainty in homogenous media

In our previous paper (Fredriksson and Bokrantz 2016, section 2.4), the generalized margins were defined with respect to a set $\mathcal{S}$ of systematic setup shifts. The lateral component of each shift relative to each beam (the component that is perpendicular to the beam central axis) was assumed to be an integer multiple of the voxel side length measured along the lateral axis. Axial components were disregarded because a beam dose is affected by a shift in parallel with the beam central axis mainly by scaling according to the inverse-square law and not by a displacement of the dose. This led to a definition of $p_{i,s}^{(b)}$ in (1) according to

$$p_{i,s}^{(b)} = 1/\left|\{s' \in \mathcal{S} : i + j^{(b)}(s) - j^{(b)}(s') \in \mathcal{C}\}\right|, \tag{3}$$

where $j^{(b)}(s)$ is the offset of the dose grid indices that corresponds to the lateral component of the shift s relative to beam b. The reciprocal of $p_{i,s}^{(b)}$ quantifies the number of scenarios s' under which the shifted voxel $i + j^{(b)}(s)$ overlaps with the shifted CTV $\{i + j^{(b)}(s') : i \in \mathcal{C}\}$.

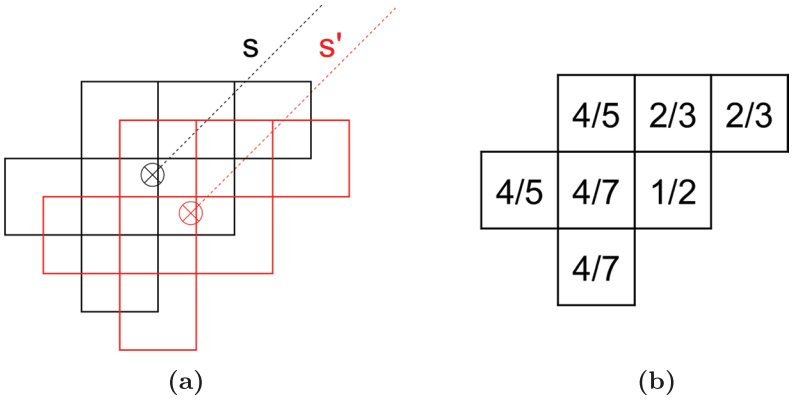

Figure 1. Overlap volume calculation for a 7-voxel CTV subject to a systematic setup shift. The weight $p_{i,s}^{(b)}$ in scenario s for the upper left voxel is $1/(1+1/4) = 4/5$ because of perfect overlap with itself in scenario s and an overlap of one-fourth with itself in scenario s'. (a) Voxelized CTV in the scenarios s and s'. The beam central axis in each scenario is indicated by a dashed line and the isocenter is indicated by '⊗'. (b) Distribution of weights $p_{i,s}^{(b)}$ for each voxel i in the scenario s.

The following notation is used to extend the overlap volume calculation to arbitrary setup errors (leading to fractional voxel overlaps): $u_{\perp}^{(b)}(s)$ in $\mathbb{R}^3$ is the lateral component relative to beam b of the setup shift under scenario s, $\mathcal{V}_i$ is the subset of $\mathbb{R}^3$ that corresponds to voxel i, and $\mathcal{V}_i^{(b)}(s) = \{v + u_{\perp}^{(b)}(s) : v \in \mathcal{V}_i\}$ is this volume shifted with $u_{\perp}^{(b)}(s)$. The overlap volume calculation takes the form:

$$p_{i,s}^{(b)} = \left|\mathcal{V}_i^{(b)}(s)\right| \Big/ \left(\sum_{s' \in \mathcal{S}} \sum_{i' \in \mathcal{C}} \left|\mathcal{V}_i^{(b)}(s) \cap \mathcal{V}_{i'}^{(b)}(s')\right| \right). \tag{4}$$

This calculation amounts to resampling a binary image between two regular grids if the dose grid is regular (the voxels are rectangular cuboids), see figure 1. The weights defined by (4) can be verified to satisfy $1/|\mathcal{S}| \leqslant p_{i,s}^{(b)} \leqslant 1$.

2.5. Setup uncertainty in heterogeneous media

We take tissue heterogeneity in the overlap volume calculation into account by scaling the length of the setup shifts based on the composition of tissues in the entrance channel of the beam. This scaling is the same type of approximation as the one-dimensional pathlength scaling that many pencil beam dose algorithms use to translate dose deposition kernels calculated in water to heterogeneous media (Mohan *et al* 1986, Petti 1992). The heterogeneity-corrected weights $p_{i,s}^{(b)}$ are calculated according to (4), but with the voxel volumes described in water equivalent coordinates where the physical depth—the distance to the patient surface along a ray that emanates from the radiation source—is replaced with the radiological depth—the depth in water that gives the same level of attenuation. The transformation from physical to water equivalent coordinates is given by integration of the relative linear attenuation over the physical depth. The linear attenuation relative to water can be calculated based on mass

density (Siddon 1985), electron density (Seco and Evans 2006), or particle stopping power (Titt *et al* 2015). Figure 2 shows how the transformation to water equivalent coordinates deforms the dose grid. The calculation of heterogeneity-corrected overlap volumes therefore requires resampling between irregular grids.

2.6. Particle range uncertainty

Systematic range errors for particle beams can be modeled by uniform scalings of the density distribution of the irradiated volume (Lomax *et al* 2001). Such a density scaling changes all radiological depths by a corresponding factor, and therefore stretches or contracts the CTV along the beam direction in the water equivalent space, as illustrated in figure 3. The overlap volume calculation for density errors is analogous to (4), but with the calculations performed in water equivalent coordinates. The axial shift $u_{\parallel}^{(b)}(s)$ relative to beam b due to a density scaling of γ is a vector of length $(1-\gamma)z$ directed away from the radiation source for a point of calculation at a radiological depth of z.

2.7. Organ motion uncertainty

Robust planning methods can take organ motion into account by inclusion of scenarios where the dose is calculated on a different image set than the nominal one (McShan *et al* 2006, Mahmoudzadeh *et al* 2015). We let the dose $d(s)$ for an organ motion scenario s be the total dose evaluated on the image (or phase of a time-resolved image) that is specified by s. The use of distinct images as scenarios requires that all ROIs that are considered in the optimization to be delineated on all images.

The calculation of overlap volumes for organ motion scenarios follows an expression of the form of (4), but with the calculation performed with respect to water equivalent coordinates and the set of CTV voxels $\mathcal{C}$ being scenario-dependent. This scenario-dependency requires a substitution of $\mathcal{C}(s)$ for $\mathcal{C}$ in (1). Calculation of an overlap volume for two organ motion scenarios is illustrated in figure 4. Note that optimization with respect to multiple CTV contours on a single image set is a special case of the organ motion scenarios considered here.

2.8. Unified tractable model

The scenario-dependent criterion that implements a CTV-to-PTV margin with respect to a set $\mathcal{S}$ of scenarios that each define a shift of the beam isocenters, a density scaling, and an image with an associated structure set takes the following form:

$$\sum_{s\in\mathcal{S}}\sum_{i\in\mathcal{C}(s)}\phi(d_i(s))\sum_{b\in\mathcal{B}}\frac{d_i^{(b)}(s)}{d_i(s)}p_{i,s}^{(b)}.$$

The overlap volumes $p_{i,s}^{(b)}$ are here defined according to

$$p_{i,s}^{(b)}=\left|\hat{\mathcal{V}}_i^{(b)}(s)\right|\Big/\left(\sum_{s'\in\mathcal{S}}\sum_{k\in\mathcal{C}(s')}\left|\hat{\mathcal{V}}_i^{(b)}(s)\cap\hat{\mathcal{V}}_k^{(b)}(s')\right|\right),$$

where the shifted voxel volumes $\hat{\mathcal{V}}_i^{(b)}(s)$ are defined according to

$$\hat{\mathcal{V}}_i^{(b)}(s)=\{v+u_{\perp}^{(b)}(s)+u_{\parallel}^{(b)}(v;s):v\in\hat{\mathcal{V}}_i\},$$

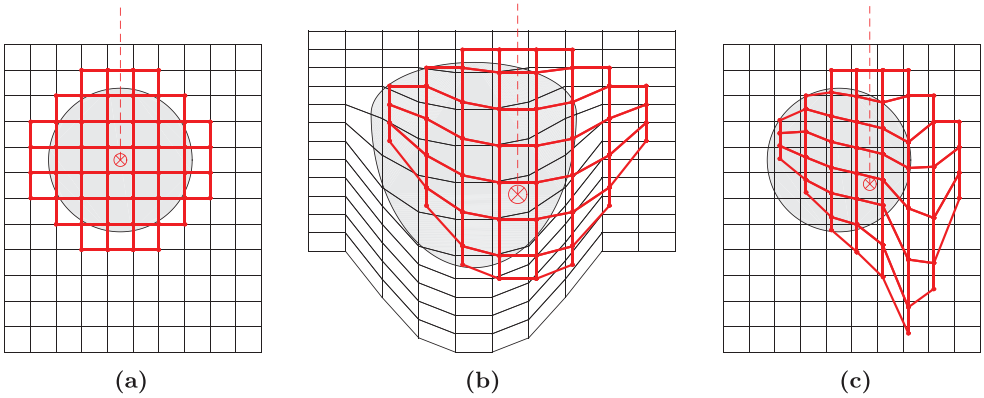

Figure 2. A rigid shift in the water equivalent space amounts to a deformation along isoradiological-depth surfaces in the physical space. The voxelized CTV is indicated by thick solid lines, the beam central axis is indicated by a dashed line, and the isocenter is indicated by '⊗'. The density of the shaded region is twice that of the surrounding material. (a) Nominal CTV in the physical space. (b) CTV in the water equivalent space perturbed by a setup shift. (c) Perturbed CTV in the physical space.

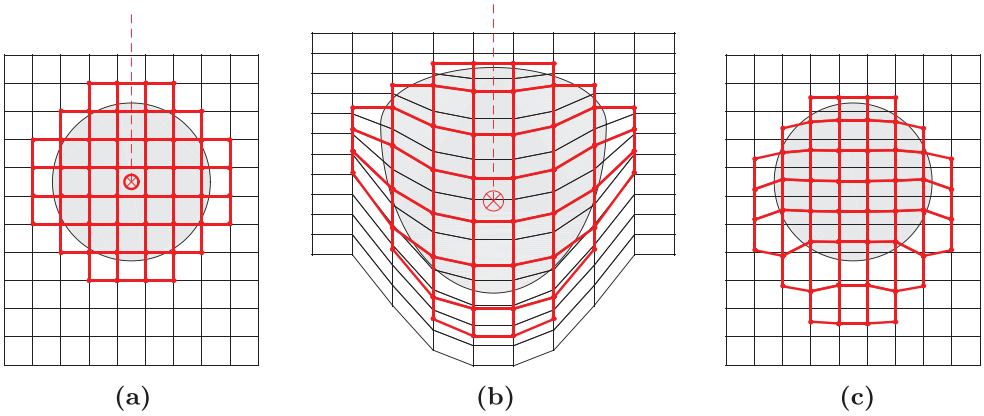

Figure 3. A range shift amounts to a scaling in the water equivalent space and a deformation in the physical space. The voxelized CTV is indicated by thick solid lines, the beam central axis is indicated by a dashed line, and the isocenter is indicated by '⊗'. The density of the shaded region is twice that of the surrounding material. (a) Nominal CTV in the physical space. (b) CTV in the water equivalent space perturbed by a range overshoot. (c) Perturbed CTV in the physical space.

and $\hat{\mathcal{V}}_i^{(b)}$ is the volume encompassed by voxel i expressed in water equivalent coordinates relative to beam b.

2.9. Computational study

We implemented the unified tractable model of scenario-based margins in RayStation v4.5 (RaySearch Laboratories, Stockholm, Sweden) and compared the method to worst case and expected value optimization. The two benchmark methods were defined precisely as in our previous study (Fredriksson and Bokrantz 2016): the worst case objective function value was

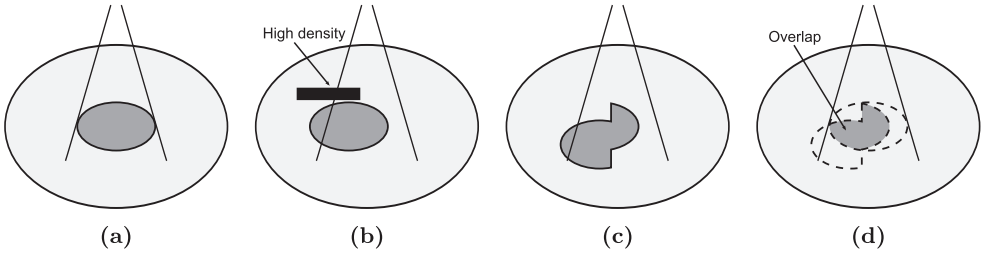

Figure 4. Calculation of CTV (dark gray) overlap with respect to scenarios comprised of independent images. (a) Scenario s. (b) Scenario s'. (c) CTV in scenario s' mapped to scenario s. (d) CTV overlap.

calculated using a smooth power mean approximation in order to obtain continuous derivatives and the expected objective function value was calculated with respect to normally distributed probabilities. The organ motion scenarios were, however, assumed to follow a uniform distribution as there was no well-defined nominal case. Results were also generated for scenario-based margins without heterogeneity correction.

2.9.1. Patient cases and optimization formulations. The evaluation was conducted with respect to treatment planning for IMPT. Treatment plans were optimized with respect to criteria formulated as quadratic penalties on the deviation from the desired dose values. A criterion for an ROI with voxels enumerated by the set $\mathcal{R}$ was formulated as

$$\sum_{i \in \mathcal{R}} \Delta_i (d_i - d^{\text{ref}})^2,$$

where Δ_i is the fraction of this ROI that is contained inside voxel i and d^{ref} is the prescription for the CTV and zero otherwise. The following treatment cases and optimization formulations were investigated (with gantry angles in IEC coordinates):

- The half-slab phantom depicted in figure 5 treated with a single field at 0°:
  - **Objective:** Minimize the value of a criterion for the CTV
  - **Constraint:** Limit the value of a criterion for the external volume within a specified bound
  - **Uncertainty:** Setup and range uncertainty
- A single-image lung case treated with two fields at 0° and 90°:
  - **Objective:** Minimize the value of a criterion for the CTV
  - **Constraint:** Limit the value of a criterion for the external volume within a specified bound
  - **Uncertainty:** Setup uncertainty
- A prostate case treated with two fields at 90° and 270°:
  - **Objective:** Minimize a criterion for each of the CTV, rectum, and external volume, weighted at ratio of 50:10:1 relative to each other
  - **Uncertainty:** Setup uncertainty

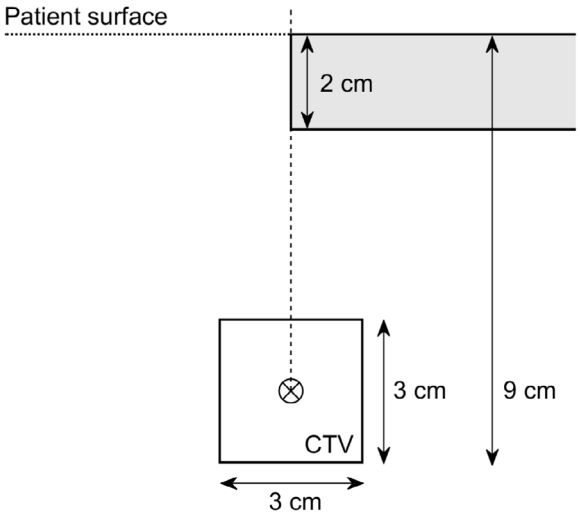

Figure 5. Transveral cut of the treatment setup for the water phantom. The mass density inside the patient volume was 2 g cm^{-3} inside the shaded slab and 1 g cm^{-3} elsewhere.

- A four-dimensional computed tomography (4D-CT) lung case treated with two fields at 180° and 270°:

  – **Objective:** Minimize the value of a criterion for the CTV
  – **Constraint:** Limit the value of a criterion for the external volume to less than or equal to a preselected bound
  – **Uncertainty:** Organ motion and range uncertainty

All beam axes were coplanar with the transversal plane and all calculations were performed in a $3 \times 3 \times 3$ mm^3 dose grid. The bound in the constrained formulations were selected such that coverage and homogeneity for the CTV had to be sacrificed at least partially. We focus on the situation when the CTV dose is limited by the acceptable dose to healthy tissues as the differences between the robust planning methods otherwise vanish (Fredriksson and Bokrantz 2014).

2.9.2. Scenarios. The experiments were designed in order to highlight differences between scenario-based margins and the benchmark methods. Relatively large errors: up to 1 cm isotropic setup error and ±5% range error were therefore considered during optimization. Setup errors were discretized into 27 scenarios: the nominal scenario, one 1 cm shift per positive and negative unit direction, per pairwise combination of unit directions, and per three-way combination of unit directions. Range errors were discretized into three scenarios: the nominal scenario, 5% undershoot, and 5% overshoot. The 4D-CT lung case had ten scenarios that represented different phases of the breathing cycle. Different sources of uncertainty were treated independently of each other, so that the total number of considered scenarios was the product of the number of scenarios of each type.

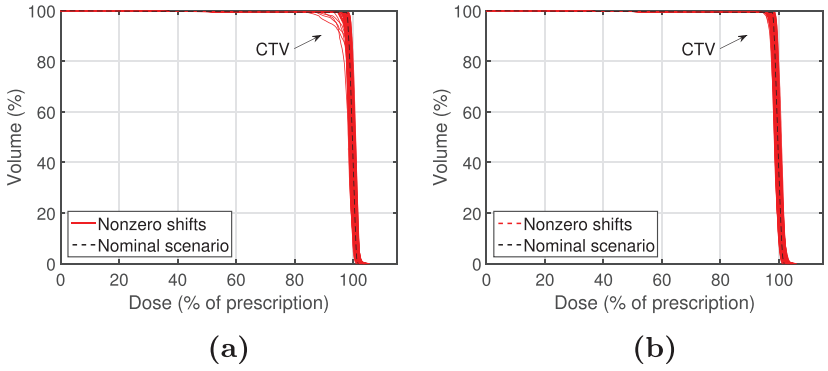

Figure 6. DVHs for the water phantom in all scenarios. The nominal scenario is indicated by a dashed black line and scenarios with a nonzero range or setup shift are indicated by solid red lines. (a) Uncorrected. (b) Heterogeneity-corrected.

3. Results

The plans generated by the different robust planning methods were, in addition to their dose distributions and dose-volume histograms (DVHs), assessed in terms of coverage, quantified as D_{95} (higher values are better), and homogeneity, quantified as $D_5 - D_{95}$ (lower values are better), of the dose to the CTV.

3.1. Water phantom subject to range and setup uncertainty

The water phantom case—where a high-density slab partly covered the target—constitutes a positive control experiment that demonstrates that there exists situations where the heterogeneity correction is important. The DVHs per scenario in figure 6 show that the heterogeneity-corrected margins led to a uniform and high level of coverage and homogeneity over all scenarios whereas the uncorrected form of margins sacrificed these properties in a small subset of the scenarios. Figure 7 shows that the loss of coverage and homogeneity occurred in about 5% of the scenarios and that this sacrifice led to slightly improved dose in the other, less difficult, scenarios. Figure 8 shows that the more consistent coverage for the heterogeneity-corrected plan was due to a more pronounced high-dose 'tongue' distally of the target compared to the uncorrected plan. The dose deposited inside this region does not contribute to coverage in the nominal scenario but is necessary in order to avoid a dose degradation under rightward setup shifts.

3.2. Lung subject to setup uncertainty

It is of paramount importance to take uncertainties into account for IMPT of lung cancers due to the high density difference between the thoracic cavity and the chest wall and the strong density-dependence of the Bragg peak positions. Our results for the single-image lung case show that expected value optimization can lead to unnecessarily poor target coverage robustness. The DVHs in figure 9 show that scenario-based margins and worst case optimization created plans that were largely insensitive to setup errors, whereas expected value optimization led to a plan with nominally better but more variable coverage and homogeneity. Figure 10 shows that scenario-based margins led to marginally better average levels of coverage and homogeneity than worst case optimization, but marginally worse such levels in the worst case. The differences between heterogeneity-corrected and uncorrected scenario-based margins were very minor, but generally in favor of the heterogeneity-corrected method.

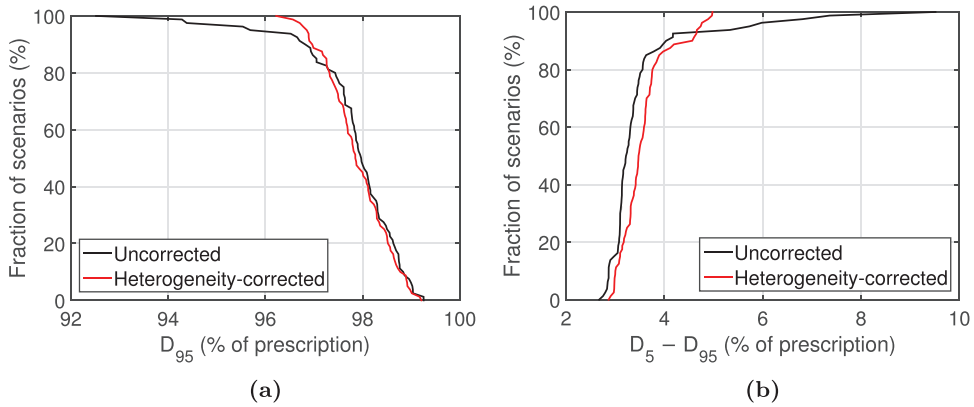

Figure 7. Dose to the CTV for the water phantom. (a) Fraction of scenarios where the level of coverage is greater than or equal to the indicated value. (b) Fraction of scenarios where the level of homogeneity is less than or equal to the indicated value. (a) CTV coverage. (b) CTV homogeneity.

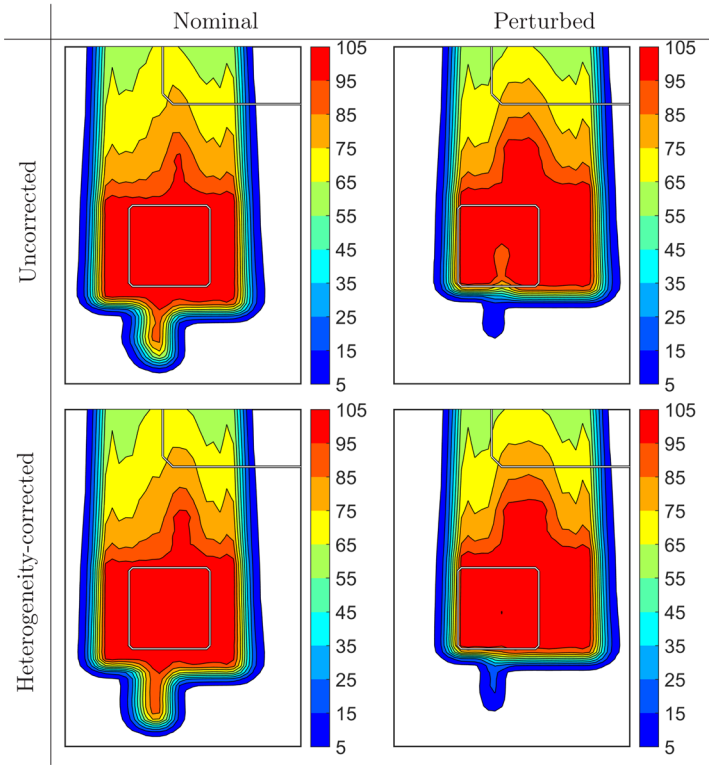

Figure 8. Dose distributions for the water phantom in the nominal scenario and a perturbation scenario (5% range undershoot and a 1 cm rightward setup shift). The doses are in percent of the prescription and are depicted as transversal cuts through the isocenter. The solid white lines indicate the contours of the CTV and high-density slab.

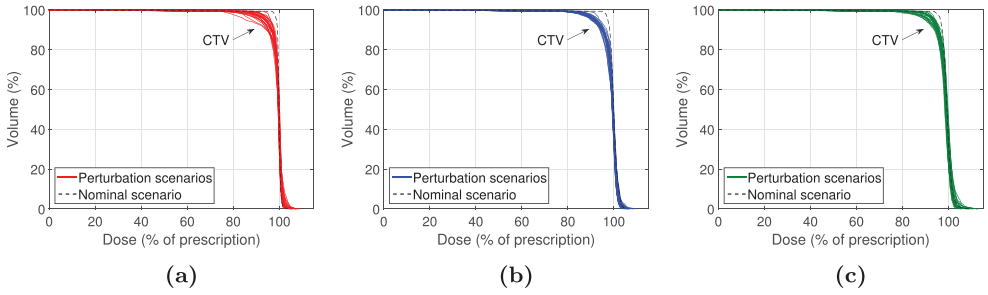

Figure 9. DVHs for the single-image lung case in all scenarios. The nominal scenario is indicated by a dashed black line and scenarios with a nonzero setup shift are indicated by solid lines in color. Results are not depicted for uncorrected scenario-based margins due to close similarity with the heterogeneity-corrected method. (a) Expected value optimization. (b) Worst case optimization. (c) Heterogeneity-corrected scenario-based margins.

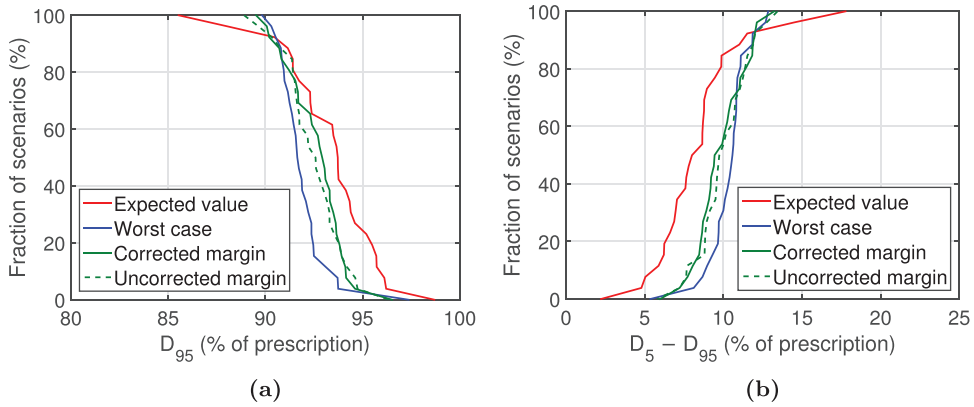

Figure 10. Dose to the CTV for the single-image lung case. (a) Fraction of scenarios where the level of coverage is greater than or equal to the indicated value. (b) Fraction of scenarios where the level of homogeneity is less than or equal to the indicated value. (a) CTV coverage. (b) CTV homogeneity.

3.3. Prostate subject to setup uncertainty

We have previously shown that worst case optimization can lead to unnecessarily poor target coverage in 'easy' scenarios when there are large conflicts between targets and OARs (Fredriksson and Bokrantz 2014, 2016). The results for the prostate patient with a highly weighted rectum sparing criterion replicate this result and demonstrate that the scenario-based margins do not suffer from the same shortcoming. The DVHs in figure 11 show that scenario-based margins as well as expected value optimization gave a high level of coverage and homogeneity except under the anterior shifts that were incompatible with rectal sparing. Worst case optimization, in contrast, sacrificed coverage and homogeneity not only under the anterior shifts but also in the scenarios with an inferior or superior shift. Figure 12 shows that scenario-based margins gave better coverage and homogeneity on average and in the worst case compared to expected value optimization. Expected value optimization led to better such levels in the best case and also spared the rectum to a greater extent. The differences between heterogeneity-corrected and uncorrected scenario-based margins were very minor—even smaller

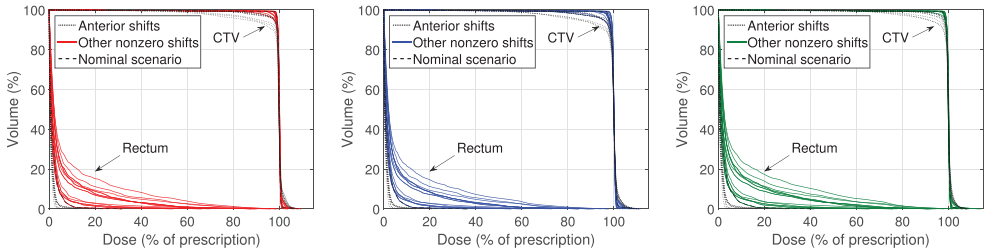

Figure 11. DVHs for the prostate patient in all scenarios. The nominal scenario is indicated by a dashed black line, scenarios with a nonzero anterior shift are indicated by dotted black lines, and the other scenarios are indicated by solid lines in color. Results are not depicted for uncorrected scenario-based margins due to close similarity with the heterogeneity-corrected method. (a) Expected value optimization. (b) Worst case optimization. (c) Heterogeneity-corrected scenario-based margins.

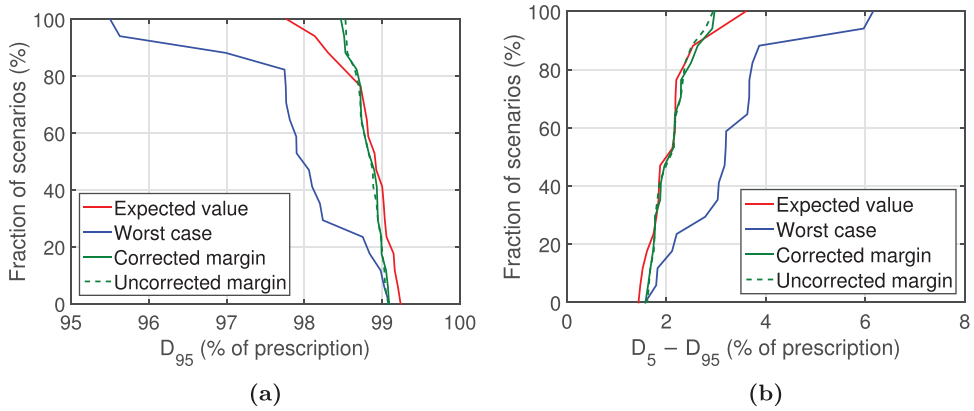

Figure 12. Dose to the CTV for the prostate patient. (a) Fraction of scenarios where the level of coverage is greater than or equal to the indicated value. (b) Fraction of scenarios where the level of homogeneity is less than or equal to the indicated value. (a) CTV coverage. (b) CTV homogeneity.

than for the single-image lung case. This lack of differences constitutes a form of negative control: no major differences between heterogeneity-corrected and uncorrected scenario-based margins were expected because the pelvic region is of relatively homogenous density.

3.4. Lung subject organ motion and range uncertainty

Consideration of setup uncertainty alone is often not sufficient to achieve robustness for lung as the respiratory motion deforms the patient anatomy and can result in highly degraded dose distribution. Our results for the 4D-CT lung case demonstrate scenario-based margins' ability to safeguard against intrafraction motion, as an alternative to a conventional internal target volume. The DVHs in figure 13 and the graphs of coverage and homogeneity in figure 14 are qualitatively similar to those for the single-image lung case: scenario-based margins and worst case optimization created plans that were largely insensitive to the organ motion and range errors whereas expected value optimization created a plan with nominally better but more

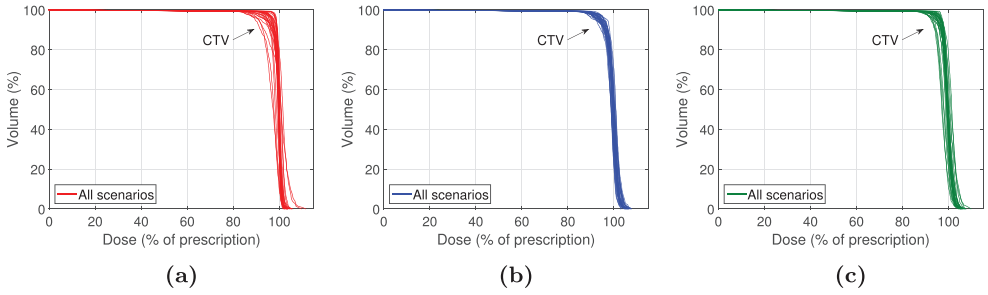

Figure 13. DVHs for the 4D-CT lung case in all scenarios. Results are not depicted for uncorrected scenario-based margins due to close similarity with the heterogeneity-corrected method. (a) Expected value optimization. (b) Worst case optimization. (c) Heterogeinety-corrected scenario-based margins.

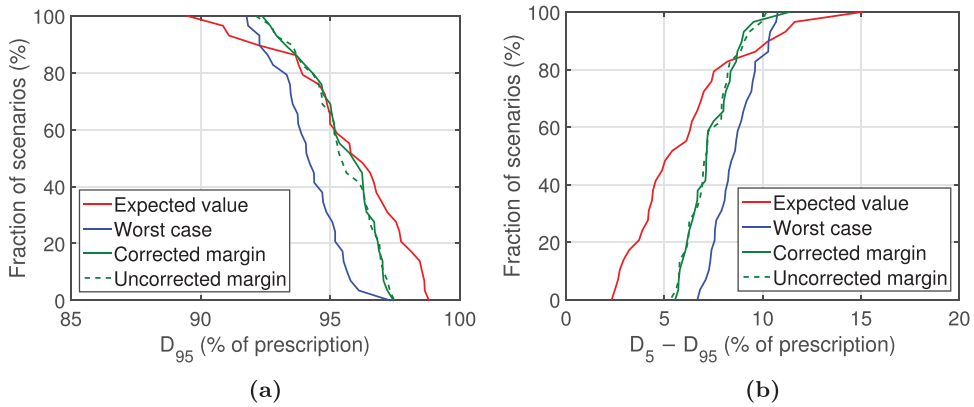

Figure 14. Dose to the CTV for the 4D-CT lung case. (a) Fraction of scenarios where the level of coverage is greater than or equal to the indicated value. (b) Fraction of scenarios where the level of homogeneity is less than or equal to the indicated value. (a) CTV coverage. (b) CTV homogeneity.

variable—and in the worst case worse—coverage and homogeneity. The differences between heterogeneity-corrected and uncorrected scenario-based margins were again very minor.

4. Discussion

The goal of robust radiation therapy planning is to, with high probability, provide a high level of coverage and homogeneity relative to the desired (generally uniform) target dose without placing unnecessary dose within healthy tissue. Several of the scenario-based optimization models in the literature fulfill this goal if there is at most a moderate conflict between the targets and OARs. Scenario-based methods, however, exhibit interesting differences if the dose to targets or OAR sparing must be compromised in at least a subset of the scenarios (Fredriksson and Bokrantz 2014). In the present study, alongside our previous study (Fredriksson and Bokrantz 2016), we found that scenario-based margins avoid the following disadvantages of expected value and worst case optimization under such circumstances:

- Expected value optimization often creates an unnecessarily gentle dose fall-off at the edge of the treated volume, as observed in Fredriksson (2012). A target dose considerably below the prescription contributes little to tumor control and, therefore, poses an unnecessary risk for toxicity. Scenario-based margins creates a steep dose gradient at the edge of the treated volume, similar to geometric margins and worst case optimization.
- Failure to satisfy the clinical goals in one scenario can cause worst case optimization to unnecessarily neglect goal fulfillment also in other, 'easier', scenarios. This was exemplified in Fredriksson and Bokrantz (2014) where the high-dose region was retracted in the inferior and superior directions for a prostate cancer patient when a retracement on the posterior side of the prostate would be sufficient to spare the rectum. Scenario-based margins are not affected by such conflicts that are 'dosimetrically decoupled', similar to geometric margins and expected value optimization.

Expected value optimization's shortcoming is an effect of that the method accepts relatively poor plan quality in an unlikely scenario in order to gain in a more probable scenario. In contrast, scenario-based margins maintain a consistent quality over all scenarios because all voxels within a structure are weighted equally (the weights for a voxel always sum to one over the scenarios). Our computational results, in particular for the two lung patients, verifies this property experimentally. The results for the prostate patient in this work replicates the result from Fredriksson and Bokrantz (2014) that worst case optimization can neglect 'easy' scenarios, and verifies that scenario-based margins do not exhibit this disadvantage. The formulation of scenario-based margins is, moreover, additively separable with respect to multiple objectives. This property makes the method directly compatible with multicriteria treatment planning, in contrast to worst case optimization (Bokrantz and Fredriksson 2014).

In this work, we have extended scenario-based margins to heterogeneous media using a density correction analogous to the pathlength scaling for dose calculation. Such a correction disregards the effect of lateral dose transport and also contains other approximations. More sophisticated techniques for heterogeneity correction (Szymanowski and Oelfke 2002, Jones and Das 2005) are therefore natural extensions of the work presented here. The very minor differences between the corrected and uncorrected forms of scenario-based methods for the clinical patient cases, however, indicate that such refinements might be of limited practical value for our purposes. Quite the opposite, our results indicate that the overlap volume calculation is generally insensitive to density heterogeneity and that therefore use of an approximate heterogeneity correction is justified.

5. Conclusions

We extended scenario-based margins to take density heterogeneities into account and demonstrated the extended method with respect to systematic patient setup, organ motion, and particle range uncertainty. For a half-slab phantom, heterogeneity-corrected scenario-based margins resulted in more robust target coverage than the uncorrected ones. However, for a prostate case and even two lung cases, all treated with IMPT, the differences between heterogeneity-corrected and uncorrected margins were small or minuscule. The more easily implemented uncorrected margins might therefore be sufficiently accurate in practice.

References

Albertini F, Hug E and Lomax A 2011 Is it necessary to plan with safety margins for actively scanned proton therapy? *Phys. Med. Biol.* **56** 4399–413

Bohoslavsky R, Witte M, Janssen T and van Herk M 2013 Probabilistic objective functions for margin-less IMRT planning *Phys. Med. Biol.* **58** 3563–80

Bokrantz R and Fredriksson A 2014 Necessary and sufficient conditions for Pareto efficiency in robust multiobjective optimization (arXiv:1308.4616)

Chan T, Bortfeld T and Tsitsiklis J 2006 A robust approach to IMRT optimization *Phys. Med. Biol.* **51** 2567–83

Chu M, Zinchenko Y, Henderson S and Sharpe M 2005 Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty *Phys. Med. Biol.* **50** 5463–77

Fredriksson A 2012 A characterization of robust radiation therapy treatment planning methods—from expected value to worst case optimization *Med. Phys.* **39** 5169–81

Fredriksson A and Bokrantz R 2014 A critical evaluation of worst case optimization methods for robust intensity-modulated proton therapy planning *Med. Phys.* **41** 081701

Fredriksson A and Bokrantz R 2016 The scenario-based generalization of radiation therapy margins *Phys. Med. Biol.* **61** 2067–82

Fredriksson A, Forsgren A and Hårdemark B 2011 Minimax optimization for handling range and setup uncertainties in proton therapy *Med. Phys.* **38** 1672–84

Jones A and Das I 2005 Comparison of inhomogeneity correction algorithms in small photon fields *Med. Phys.* **32** 766–76

Ling C, Humm J, Larson S, Amols H, Fuks Z, Leibel S and Koutcher J 2000 Towards multidimensional radiotherapy (MD-CRT): biological imaging and biological conformality *Int. J. Radiat. Oncol. Biol. Phys.* **47** 551–60

Lomax A *et al* 2001 Intensity modulated proton therapy: a clinical example *Med. Phys.* **28** 317–24

Mahmoudzadeh H, Lee J, Chan T and Purdie T 2015 Robust optimization methods for cardiac sparing in tangential breast IMRT *Med. Phys.* **42** 2212–22

McShan D, Kessler M, Vineberg K and Fraass B 2006 Inverse plan optimization accounting for random geometric uncertainties with a multiple instance geometry approximation (MIGA) *Med. Phys.* **33** 1510–21

Mohan R, Chui C and Lidofsky L 1986 Differential pencil beam dose computation model for photons *Med. Phys.* **13** 64–73

Ólafsson A and Wright S 2006 Efficient schemes for robust IMRT treatment planning *Phys. Med. Biol.* **51** 5621–42

Petti P 1992 Differential-pencil-beam dose calculations for charged particles *Med. Phys.* **19** 137–49

Pflugfelder D, Wilkens J and Oelfke U 2008 Worst case optimization: a method to account for uncertainties in the optimization of intensity modulated proton therapy *Phys. Med. Biol.* **53** 1689–700

Seco J and Evans P 2006 Assessing the effect of electron density in photon dose calculations *Med. Phys.* **33** 540–52

Siddon R 1985 Calculation of the radiological depth *Med. Phys.* **12** 84–7

Stewart J *et al* 2010 Automated weekly replanning for intensity-modulated radiotherapy of cervix cancer *Int. J. Radiat. Oncol. Biol. Phys.* **78** 350–8

Szymanowski H and Oelfke U 2002 Two-dimensional pencil beam scaling: an improved proton dose algorithm for heterogeneous media *Phys. Med. Biol.* **47** 3313

Titt U, Sell M, Unkelbach J, Bangert M, Mirkovic D, Oelfke U and Mohan R 2015 Degradation of proton depth dose distributions attributable to microstructures in lung-equivalent material *Med. Phys.* **42** 6425–32

Unkelbach J and Oelfke U 2004 Inclusion of organ movements in IMRT treatment planning via inverse planning based on probability distributions *Phys. Med. Biol.* **49** 4005–29

Unkelbach J, Bortfeld T, Martin B and Soukup M 2009 Reducing the sensitivity of IMPT treatment plans to setup errors and range uncertainties via probabilistic treatment planning *Med. Phys.* **36** 149–63