Efficient schemes for robust IMRT treatment planning

Arinbjörn Ólafsson¹ and Stephen J Wright²

¹ Industrial Engineering Department, University of Wisconsin, 1513 University Avenue, Madison, WI 53706, USA

² Computer Sciences Department, University of Wisconsin, 1210 W Dayton Street, Madison, WI 53706, USA

Received 20 January 2006, in final form 6 September 2006

Published 16 October 2006

Online at stacks.iop.org/PMB/51/5621

Abstract

We use robust optimization techniques to formulate an IMRT treatment planning problem in which the dose matrices are uncertain, due to both dose calculation errors and interfraction positional uncertainty of tumour and organs. When the uncertainty is taken into account, the original linear programming formulation becomes a second-order cone program. We describe a novel and efficient approach for solving this problem, and present results to compare the performance of our scheme with more conventional formulations that assume perfect knowledge of the dose matrix.

(Some figures in this article are in colour only in the electronic version)

1. Introduction

Intensity modulated radiation therapy (IMRT) allows tailored doses of radiation to be delivered to a cancer patient via external beams, with the aim of destroying the tumour while sparing surrounding critical structures and normal tissue. After dividing the aperture through which the radiation is delivered into small rectangular regions, thus dividing each beam into so-called beamlets, treatment planning procedures find the set of beamlet weights that most closely matches the dose requirements prescribed by an oncologist. Optimization techniques have proved to be useful for determining optimal beamlet weights. The increasing power and efficiency of these methods has allowed greater refinement and tuning of treatment plans.

A critical role in treatment planning is played by the *dose matrix*, in which each entry describes the amount of radiation deposited by a unit weight (intensity) from a particular beamlet into a particular voxel in the treatment region. Given the properties of the treatment region and the beamlet source and orientation, the dose matrix is calculated by techniques such as convolution/superposition and Monte Carlo. These are numerical/statistical techniques that give approximate results, and hence are one source of uncertainty in the problem. A second source of uncertainty is change in the positions of the internal tissues of the patient between treatment fractions (*interfractional* motion), relative to the beams. These changes can be due to

motion of the tumour and organs or to variations in the patient position on the treatment device between fractions. We focus on dose calculation and positional uncertainties in this paper. We do not deal with other sources of uncertainty, such as imaging errors, physician error, intrafraction patient motion, registration errors, organ shrinkage and possible microscopic extensions of the tumour, though our formulation techniques may be useful in handling some of these additional sources of uncertainty as well.

Two dose calculation methods in wide use today are the convolution/superposition method (Lu *et al* 2005) and the Monte Carlo method (Ma *et al* 1999, Jeraj and Keall 1999, Sempau *et al* 2000). A general discussion of existing dose calculation methods can be found in the report of the Intensity Modulated Radiation Therapy Collaborative Working Group (2001). As described by Kowalok (2004) (chapter 2), the dose in each voxel calculated from a Monte Carlo simulation is normally distributed. Its variance can be approximated by a sample variance arising from the generation procedure.

Most of the existing literature on uncertainty in treatment planning focuses on organ movement and setup errors (Leong 1987, Beckham *et al* 2002, Stapleton *et al* 2005, Baum *et al* 2006, Chu *et al* 2005). In practice, organ movement is often dealt with by the conservative strategy of defining a margin around the expected tumour location and prescribing a high dose for this expanded region. (See our discussion of planning target volume (PTV) and clinical target volume (CTV) in subsection 2.1.) Setup errors can be dealt with in a similar manner. Alternative approaches are described in Leong (1987), Beckham *et al* (2002), and Stapleton *et al* (2005), where the uncertainty in setup is incorporated into the dose matrix calculations. Recent research (Chu *et al* 2005, Unkelbach and Oelfke 2004) has also modelled organ movement as uncertainties in the dose matrix. These authors have considered multi-fraction treatment plans and assumed some knowledge of the distribution of organ positions. van Herk (2004) discusses methods to estimate the probability distribution for organ movement. (We discuss modelling of organ movement further in subsection 2.3.) Other recent work (Chan *et al* 2006) has considered uncertainty in *intrafractional* motion, usually as a result of the patient breathing during delivery.

The focus of this paper is on modelling and efficient solution of a formulation in which the uncertainties from dose calculation and organ/tumour positions are accounted for explicitly and simultaneously. By assuming the probability distributions for both types of uncertainties are known, we can set up and solve a robust optimization problem. Our goal is to find solutions that are referred to in the optimization literature (Ben-Tal and Nemirovski 2000) as *almost reliable* solutions, in which the constraints are guaranteed to be satisfied to a certain level of confidence, given the uncertainty in the data. This methodology represents a significant advance over techniques in which the uncertainty is ignored, as the beamlet weights calculated from such plans may result in unacceptable underdosing of tumours and overdosing of critical regions when the actual realization of the dose matrix differs appreciably from the value used in devising the treatment plan.

We also present an algorithmic advance. Like Chu *et al* (2005), we solve a second-order cone programming (SOCP) formulation, but rather than using standard SOCP software (which is based on interior-point algorithms) we use a sequential linear programming (SLP) approach that is well suited to these problems and apparently more efficient.

The remainder of the paper is structured as follows. In section 2, we outline the underlying linear programming (LP) formulation of the treatment planning problem used in this study and show how to obtain a robust formulation from it. Section 3 describes the sequential linear programming approach for solving the robust formulation. Computational results for a nasopharyngeal case are shown in section 4, while section 5 discusses conclusions and possible future research directions.

2. Linear programming formulations of IMRT

We discuss issues of problem formulation in this section. In subsection 2.1 we introduce some background material and notation. In subsection 2.2 the underlying LP is introduced and in subsection 2.3 we show how the errors can be taken into account by formulating an SOCP.

2.1. Background and notation

For planning purposes, the treatment region in the patient's body is divided into regions known as voxels, indexed as $i = 1, 2, \dots, m$. Each of the apertures through which radiation can be delivered from various angles is divided into rectangular beamlets, which we index as $j = 1, 2, \dots, n$. The dose of radiation delivered to voxel i by a unit weight (intensity) of beamlet j is denoted by A_{ij} ; these quantities can be assembled into a *dose matrix* A of dimension $m \times n$. After assembling the voxel doses x_i into a vector x and the beamlet weights w_j into a vector w , we have the following linear relationship:

$$x = Aw. \quad (2.1)$$

In most approaches to IMRT treatment planning, the beamlet weight vector w is the main variable in the optimization formulation. Clinical data sets used in the computational experiments of Hou *et al* (2003), Llacer *et al* (2003), Alber and Reemtsen (2004) and Zhang *et al* (2004) typically involve 1000 to 5000 beamlets and up to 100 000 voxels.

In treatment planning, the set of voxels $\{1, 2, \dots, m\}$ is partitioned into one or more target volumes $\mathcal{T}$, one or more critical structures $\mathcal{C}$, and a region of normal tissue $\mathcal{N}$. The target volume is often subdivided into a clinical tumour volume (CTV), which is the region most likely to contain the tumour, and a planning tumour volume (PTV), which is a larger volume surrounding the CTV. The PTV includes the region into which the tumour may move during treatment, or into which extensions of the tumour may have spread at a sub-detectable level, so it is usually designated to receive some specified level of radiation, comparable to the CTV. The *critical voxels*, denoted by $\mathcal{C}$, are those that are part of a sensitive structure (such as the spinal cord or a vital organ) that we particularly wish to avoid irradiating. *Normal voxels*, denoted by $\mathcal{N}$, are those in the patient's body that fall into neither category. Ideally, normal voxels should receive little radiation, but it is less important to avoid dose to these voxels than to the critical voxels.

Generally, a dosimetrist tries to devise a plan in which the doses to target voxels are close to a specified value and/or within a given range, while constraints of different types are applied to discourage excessive dosage to critical and normal regions.

2.2. Nominal linear programming formulation

In the following linear programming (LP) formulation, we introduce some additional notation. The vectors $x_{\mathcal{T}}^L$ and $x_{\mathcal{T}}^U$ denote lower and upper bounds on the doses to the target voxels, while $d_{\mathcal{T}}$ represents the prescribed dose to these voxels. When the treatment plan is divided into N fractions, the relevant quantities for each fraction are $x_{\mathcal{T}}^L/N$, $x_{\mathcal{T}}^U/N$, and $d_{\mathcal{T}}/N$, respectively. The per-fraction dose to target voxels in excess of $d_{\mathcal{T}}/N$ is denoted by the vector v , while t represents the dose under this level. The vector $d_{\mathcal{C}}$ represents a set of threshold values for doses to critical voxels, such that a penalty is imposed for the per-fraction dose $x_{\mathcal{C}}$ in excess of $d_{\mathcal{C}}/N$ to these voxels. Penalty vectors for doses to normal voxels, over-threshold doses to critical voxels, and overdose and underdose to target voxels, are denoted by $c_{\mathcal{N}}$, $c_{\mathcal{C}}$, c_v , and c_t , respectively. The dose matrix A is divided into three row submatrices $A_{\mathcal{T}} \in \mathbb{R}^{n_{\mathcal{T}} \times n}$, $A_{\mathcal{N}} \in \mathbb{R}^{n_{\mathcal{N}} \times n}$ and $A_{\mathcal{C}} \in \mathbb{R}^{n_{\mathcal{C}} \times n}$, corresponding to the treatment volumes $\mathcal{T}$, $\mathcal{N}$ and $\mathcal{C}$, respectively.

The linear programming problem for calculating the optimal weight vector for a single fraction can be written as follows:

$$\min_{w, v, t, x_\varepsilon} c_N^T A_N w + c_v^T v + c_t^T t + c_\varepsilon^T x_\varepsilon \quad (2.2a)$$

$$\text{subject to} \quad A_T w \leq \frac{x_T^U}{N} \quad (2.2b)$$

$$A_T w \geq \frac{x_T^L}{N} \quad (2.2c)$$

$$A_C w - x_\varepsilon \leq \frac{d_\varepsilon}{N} \quad (2.2d)$$

$$v - t = A_T w - \frac{d_T}{N} \quad (2.2e)$$

$$w, x_\varepsilon, v, t \geq 0. \quad (2.2f)$$

The formulation is quite flexible, as the bounds and penalties can vary from voxel to voxel and, in particular, the penalty for overdosing a voxel can be different from the penalty for underdose. The threshold d_ε and penalty c_ε for excess doses to the critical voxels can be manipulated to ensure that dose–volume (DV) constraints are satisfied.

The first published usage of LPs in radiation treatment planning was the work by Klepper (1966) and Bahr *et al* (1968). Later works involving LP include Hodes (1974), Morrill *et al* (1990), Langer *et al* (1990) and Rosen *et al* (1991). The linear models vary in their definition of both objectives and constraints. For more recent work, see the surveys by Shepard *et al* (1999) and Reemtsen and Alber (2004) and the references therein.

We refer to the formulation above as a ‘nominal’ LP formulation as it assumes that the dose matrix A is known with certainty, so that the constraints in the model can correspondingly be enforced with certainty. We now describe a formulation that accounts for the fact that A is actually uncertain, and uses our knowledge of this uncertainty to enforce the constraints in a probabilistic fashion.

2.3. Robust linear programming formulation

We consider a model that incorporates two types of uncertainty in the dose matrix.

- Uncertainties due to Monte Carlo dose matrix computations. To model this uncertainty, we treat the dose to each voxel from each beamlet as a random variable, normally distributed with a known mean and variance. Moreover, we assume that these random variables are independent from voxel to voxel.
- Uncertainties due to organ motion and patient positioning errors. These can be modelled by defining a number of scenarios for the positions of the organs relative to the beam and assigning a probability to each scenario. For instance, one scenario could be the ‘nominal’ position identified by imaging, while other scenarios could be assigned by shifting the organs by various amounts in various directions. Each scenario is defined by a different dose matrix. Unless the number of scenarios is very large, there are fewer ‘degrees of freedom’ in this model than in the dose-calculation model. This uncertainty model captures the correlation between doses to adjacent voxels that occur when organ motion and patient positioning errors take place. In general, higher uncertainty will occur in voxels at the boundary between low-dose and high-dose regions, as such voxels will move into high-dose regions in some scenarios and into low-dose regions in other scenarios.

In accounting for the second type of error, we assume as in Chu *et al* (2005) that the dose is delivered in multiple fractions, with the scenario that occurs at each fraction being independent of the scenarios that occur on other fractions. We also assume that the total dose to each voxel is distributed normally, and thus can be completely characterized by its mean and variance. This assumption is justified (Chu *et al* 2005) for the case in which only organ position uncertainty is present by invoking the central limit theorem, assuming independence of the scenario delivered at each fraction and using the fact that the number of fractions N is fairly large (typically 40). We have verified empirically by a number of simulations that the dose distribution appears to remain normal when *both* organ position and dose calculation errors are present.

Our model accounts for these two types of errors by assuming that we have dose matrices $A^1, A^2, \dots, A^L$ associated with the L organ-position scenarios, with associated probabilities $p_1, p_2, \dots, p_L$. We assume in addition that each A^k is a (matrix) random variable rather than a fixed matrix, reflecting the dose calculation uncertainty. Specifically, we assume that each component A_{ij}^k is a normally distributed random variable with expected value $\tilde{A}_{ij}^k$ and variance $(\sigma_{ij}^k)^2$, and that these random variables are independent of each other over $i = 1, 2, \dots, m$ and $j = 1, 2, \dots, n$.

The dose is delivered in N fractions, which we denote by $A^{(q)}$, $q = 1, 2, \dots, N$. Each $A^{(q)}$ takes on the value of one of the random matrices $A^1, A^2, \dots, A^L$, with probability $p_1, p_2, \dots, p_L$, respectively. In practice, the dose matrix A^k for each scenario is calculated once, for each $k = 1, 2, \dots, L$, at the start of the treatment process, and this matrix is used on each fraction on which scenario k occurs. For the expectation of each $A^{(q)}$ we have

$$E(A^{(q)}) = \sum_{k=1}^L p_k \tilde{A}^k =: \hat{A}, \quad q = 1, 2, \dots, N. \quad (2.3)$$

For the expected total dose to voxel i from a given weight vector w , we have

$$\tilde{x}_i = E(x_i) = E\left(\sum_{q=1}^N A_{i\cdot}^{(q)} w\right) = N \sum_{k=1}^L p_k \tilde{A}_{i\cdot}^k w = N \hat{A}_{i\cdot} w, \quad (2.4)$$

where $\hat{A}$ is defined in (2.3). Derivation of the approximate variance of the total dose x_i is shown in appendix A. It has the form

$$\text{Var}(x_i) \approx \|T_i w\|_2^2, \quad (2.5)$$

where the matrix T_i is defined as

$$T_i = \begin{bmatrix} \Sigma_i \\ R_i \tilde{A}_i \end{bmatrix},$$

where

$$\tilde{A}_i = \begin{bmatrix} \tilde{A}_{i\cdot}^1 \\ \tilde{A}_{i\cdot}^2 \\ \vdots \\ \tilde{A}_{i\cdot}^L \end{bmatrix}, \quad (2.6a)$$

$$\Sigma_i = N \text{diag}(\tilde{\sigma}_{i1}, \tilde{\sigma}_{i2}, \dots, \tilde{\sigma}_{in}) \quad \text{with} \quad \tilde{\sigma}_{ij}^2 = \sum_{k=1}^L p_k^2 (\sigma_{ij}^k)^2, \quad (2.6b)$$

$$R_i = N^{1/2} P^{1/2} [I - e p^T], \quad (2.6c)$$

with $p = (p_1, p_2, \dots, p_L)^T$, $e = (1, 1, \dots, 1)^T$ and $P = \text{diag}(p_1, p_2, \dots, p_L)$.

Assuming (as we noted earlier) that the total dose to each voxel is normally distributed, we can write our robust reformulation of the constraint $A_i w \leq b_i/N$ in the nominal LP (2.2) is as follows:

$$N \hat{A}_i w + \Phi^{-1}(1 - \delta) \|T_i w\|_2 \leq b_i,$$

where δ is an upper bound on probability that the constraint will be violated and Φ is the standard normal cumulative distribution function. Introducing the notation $z_{1-\delta} := \Phi^{-1}(1 - \delta)$ and dividing by N , we rewrite this inequality as follows:

$$\hat{A}_i w + \frac{z_{1-\delta}}{N} \|T_i w\|_2 \leq \frac{b_i}{N}. \quad (2.7)$$

Using the notation above we can now write the robust formulation of the LP in (2.2) as follows:

$$\min_{w, x_\varepsilon, v, t} c_N^T A_N w + c_v^T v + c_t^T t + c_\varepsilon^T x_\varepsilon \quad (2.8a)$$

$$\text{subject to } c_U^i(w) \leq 0, \quad i \in \mathcal{T} \quad (2.8b)$$

$$c_L^i(w) \leq 0, \quad i \in \mathcal{T} \quad (2.8c)$$

$$c_C^i(w) - (x_\varepsilon)_i \leq 0, \quad i \in \mathcal{C} \quad (2.8d)$$

$$\hat{A}_T w - v + t - \frac{d_T}{N} = 0, \quad (2.8e)$$

$$w, x_\varepsilon, v, t \geq 0, \quad (2.8f)$$

where

$$c_L^i(w) = - \sum_{j=1}^n \hat{A}_{ij} w_j + \frac{z_{1-\delta}}{N} \|T_i w\|_2 + \frac{(x_T^L)_i}{N}, \quad (2.9a)$$

$$c_U^i(w) = \sum_{j=1}^n \hat{A}_{ij} w_j + \frac{z_{1-\delta}}{N} \|T_i w\|_2 - \frac{(x_T^U)_i}{N}, \quad (2.9b)$$

$$c_C^i(w) = \sum_{j=1}^n \hat{A}_{ij} w_j + \frac{z_{1-\delta}}{N} \|T_i w\|_2 - \frac{(d_\varepsilon)_i}{N}. \quad (2.9c)$$

Because of the presence of the terms $\|T_i w\|_2$, the problem (2.8) is a second-order cone program (SOCP). This is a type of convex optimization problem that has received a good deal of attention in recent years, as it is a useful way to express robust optimization problems and it can be solved efficiently by interior-point methods.

The constraints (2.8b) and (2.8c) ensure that for any given voxel $i \in \mathcal{T}$, the probability of violation of the constraint in question, given that the treatment plan was obtained via some set of realizations of the random dose matrices over the L scenarios, is at most δ . Note that if δ is chosen to be very close to zero (so that $z_{1-\delta}$ is large), the constraints (2.8c) and (2.8b) might represent a significant tightening over the nominal case, possibly leading to an infeasible formulation. That is, there might be no weight vector w that satisfies all these constraints to the desired level of certainty.

For each critical voxel $i \in \mathcal{C}$, constraint (2.8d) ensures that with probability at least $1 - \delta$, the dose to voxel i is at most $(d_\varepsilon)_i/N + (x_\varepsilon)_i$. We continue to use expected dose in the

constraint (2.8e) which gauges the discrepancy between the delivered dose and the prescribed dose vector d_T . It is not crucial to take account of the uncertainty in handling this aspect of the model, which is less critical than avoiding hot and cold spots in the target and overdose to the critical structures.

3. Sequential linear programming algorithm for robust LP

Although good algorithms and software exist for solving second-order cone programs such as (2.8), they tend to require more computational resources than LPs of similar size and structure. Chu *et al* (2005) reported run times of 3.3 h for a formulation similar to (2.8), while the corresponding nominal LP required about 17 min.

However, the particular problem (2.8) has a property that makes it unnecessary to use a general SOCP code. For data sets of interest, the terms $T_i w$ are nonzero for all $i \in \mathcal{T} \cup \mathcal{C}$. Since the constraint functions c_U^i , c_L^i and c_C^i are nonsmooth only at points for which $T_i w = 0$, we can assume smoothness and apply a general algorithm for smooth nonlinear programming. Another property that ensures the success of this approach is that when the terms $\|T_i w\|_2$ are omitted, the nominal LP that remains has a finite solution. We now describe a *sequential linear programming* (SLP) approach for solving (2.8) which is simple and is, we believe, applicable to other robust LP formulations that have the two properties mentioned above.

We start by writing the nonsmooth penalty formulation of (2.8), in which all constraints except the nonnegativity $w \geq 0$ are replaced by nonsmooth penalty terms in the objective:

$$\min_{w \geq 0} \tau(w) \equiv c_N^T A_N w + c_v^T \left(\hat{A}_T w - \frac{d_T}{N} \right)_+ + c_t^T \left(\hat{A}_T w - \frac{d_T}{N} \right)_- + c_\varepsilon^T c_C(w)_+ + \nu (\|c_L(w)_+\|_1 + \|c_U(w)_+\|_1), \quad (3.1)$$

where $(x)_+ = \max(0, x)$ and $(x)_- = \max(0, -x)$. The value of the penalty parameter ν is chosen large enough to ensure that the constraints (2.8b) and (2.8c) are satisfied when the problem (2.8) is feasible. (If the underlying (2.8) is infeasible, then the minimization of the penalty function (3.1) for large ν will yield the w that comes closest to satisfying the constraints in an ℓ_1 sense.)

To ease description of the methodology used to solve this penalty formulation, let us note that this problem has the form

$$\min_{w \geq 0} p(h(w)),$$

where the vector function $h(w)$ contains linear and SOCP-type terms (so is smooth for w of interest) while $p(\cdot)$ is a polyhedral convex function of the general form

$$p(h) = \max_{j=1,2,\dots,J} \phi_j^T h + \beta_j, \quad (3.2)$$

where $\phi_j, \beta_j, j = 1, 2, \dots, J$, are constant vectors and scalars, respectively. Note that $p(h)$ is piecewise linear as a function of h and that $p(h(w))$ is piecewise smooth as a function of w .

In the SLP approach, we obtain a step Δw from the current solution estimate w by solving the following subproblem, for a given positive scalar Δ :

$$\min_{\Delta w} p(h(w) + \nabla h(w)^T \Delta w) \quad \text{subject to} \quad w + \Delta w \geq 0, \quad \|\Delta w\|_\infty \leq \Delta. \quad (3.3)$$

The vector function h is replaced by a linear approximation around the current estimate w , which we can expect to be reasonably close to $p(h(w + \Delta w))$ provided that Δw is not too large. The last constraint is a trust region constraint, which bounds how far we can move away from the current iterate w at this iteration. Essentially, the bound Δ defines the region

on which we ‘trust’ the approximate objective $p(h(w) + \nabla h(w)^T \Delta w)$ to be close to the true objective $p(h(w + \Delta w))$. By making use of the definition (3.2) of $p(h)$, the subproblem (3.3) can be formulated as a linear program. When applied to (3.1), and after conversion to an LP, the subproblem (3.3) has the following form.

$$\min_{\Delta w, v, t, x_\varepsilon, \rho_L, \rho_U} m(\Delta w) \equiv c_N^T A_N(w + \Delta w) + c_v^T v + c_t^T t + c_\varepsilon^T x_\varepsilon + (ve)^T (\rho_L + \rho_U) \text{ s.t.} \quad (3.4a)$$

$$-(\rho_L)_i + \sum_{j=1}^n \left(-\hat{A}_{ij} + \frac{z_{1-\delta}}{\sqrt{N}} \frac{(T_i w)^T (T_i)_{\cdot j}}{\|T_i w\|_2} \right) \Delta w_j + c_L^i(w) \leq 0, \quad i \in \mathcal{T} \quad (3.4b)$$

$$-(\rho_U)_i + \sum_{j=1}^n \left(\hat{A}_{ij} + \frac{z_{1-\delta}}{\sqrt{N}} \frac{(T_i w)^T (T_i)_{\cdot j}}{\|T_i w\|_2} \right) \Delta w_j + c_U^i(w) \leq 0, \quad i \in \mathcal{T} \quad (3.4c)$$

$$\hat{A}_i \Delta w + \frac{z_{1-\delta}}{\sqrt{N}} \frac{(T_i w)^T (T_i \Delta w)}{\|T_i\|_2} - (x_\varepsilon)_i + c_c^i(w) \leq 0, \quad i \in \mathcal{C} \quad (3.4d)$$

$$\hat{A}_T(w + \Delta w) - v + t - \frac{d_T}{N} = 0, \quad (3.4e)$$

$$\Delta w + w \geq 0, \quad (3.4f)$$

$$\|\Delta w\|_\infty \leq \Delta, \quad (3.4g)$$

$$\rho_L, \rho_U, v, t, x_\varepsilon \geq 0, \quad (3.4h)$$

where $e = (1, 1, \dots, 1)^T$ as before.

The complete algorithm is shown as algorithm 1. At each iteration, we compute a quantity s for purposes of a termination test. This quantity consists of a predicted relative decrease in the objective over the current step $((\tau(w^k) - m(w^k + \Delta w^k))/\tau(w^k))$ scaled by the current trust-region radius Δ_k , if the trust-region radius is smaller than 1. We note that the termination test is not critical. Later iterations of algorithm 1 are relatively inexpensive, due to warm-start information being available from earlier iterations, so the time required by the method does not increase dramatically if the algorithm is allowed to run for some extra iterations.

Set $k \leftarrow 0$; choose initial trust region parameter Δ_0 and stopping parameter η ;

Solve for Δw in (3.4) by omitting terms including w ; Set $w^0 \leftarrow \Delta w$;

While $k = 0$ or $(k < 50$ and $s > \eta)$ **do**

Solve the subproblem (3.4) with $w = w^k$ to obtain Δw^k ;

Set $\tilde{w} \leftarrow w^k + \Delta w^k$;

if $\tau(\tilde{w}) < \tau(w^k)$ **then**

Set $w^{k+1} \leftarrow \tilde{w}$;

$\Delta_{k+1} \leftarrow 1.5\Delta_k$;

else

$\Delta_{k+1} = 0.5\|\Delta w^k\|_\infty$;

end

Update $s = \frac{\tau(w^k) - m(\Delta w^k)}{\tau(w^k) \min(1, \Delta_k)}$;

Set $k \leftarrow k + 1$;

end

Algorithm 1 : Sequential linear programming.

4. Computational results

We report computational results obtained on a medium-scale clinical case, a nasopharyngeal tumour, also used by (Wu 2002). The data set is described in subsection 4.1, along with the software, computer and algorithmic parameter values used in the experiments. We then consider two cases. First, in subsection 4.2 we show results in which a PTV is used both in the nominal LP planning and the robust formulation. Second, in subsection 4.3 we show results for which the PTV is omitted from the robust formulation, in the expectation that the more explicit use of positional uncertainty in this formulation will produce less collateral radiation to surrounding tissues than the more conventional use of PTV as a device to handle such uncertainty.

4.1. Problem setup

In the nasopharyngeal data set, treatment is carried out in two axial planes through the patient. The 24 000 voxels are divided into five regions: CTV and PTV; parotids and spinal cord (the critical structures); and the normal region. There are a total of 1989 beamlets, consisting of 39 beamlets from each of 51 angles.

For modelling purposes we used GAMS (Brooke *et al* 1998) with the linear programming solver CPLEX version 9.0. Following our earlier work on efficient LP formulations of treatment planning problems (Ólafsson and Wright 2006), we solve the primal formulation (2.2) of the nominal LP problem using dual simplex with steepest-edge pricing. Experiments were performed on a computer with a 3.0 GHz Intel® Pentium® four-processor PC running CentOS 4. Cache size was 2 MB and total memory was 2 GB. We compare results obtained with the nominal planning procedure of subsection 2.2 and the robust planning methodology of subsection 2.3.

The costs, bounds and thresholds in the nominal problem (2.2) were set as follows. For the costs, we have $c_N = 3e$ and $c_\varepsilon = 10e$. The PTV and CTV components of c_v are set to 10 and 20, respectively, while the corresponding settings for c_t are 15 and 15. For the prescribed dose vector d_T , the CTV components were set to 70 while the PTV components were set to 60. (The unit of dose is Gray (Gy).) Lower bound components x_T^L were set to 55 for PTV and 66 for CTV, while upper bound components are 74 for both PTV and CTV. The critical structure threshold d_ε is set to 5 Gy for the spinal cord and 25 Gy for the parotids. The number of fractions used is $N = 40$.

In algorithm 1 the trust region parameter Δ is initially set to 30 and the stopping parameter η to 10^{-2} . In the subproblem (3.4), we define $\delta = 0.05$, so that $z_{1-\delta} = 1.645$. The penalty parameter is $\nu = 1000$.

We took the output of the Monte Carlo dose calculation engine to be the ‘nominal’ dose matrix A_{ij} , and we use this matrix to formulate the nominal problem (2.2). To formulate the robust problem (2.8), we used A_{ij} as the expected value matrix $\tilde{A}^1$ for the first scenario. Four more scenarios $\tilde{A}^2, \dots, \tilde{A}^5$ were obtained by shifting the target and critical regions by 3.6 mm in the left, right, anterior and posterior directions, as shown in table 1. (The shift happens to be twice the length of the side of a voxel, but non-integer multiples and non-axial directions of shift can be used, without complications.) The probability of the expected position is chosen to be approximately twice the probability of each of the shifted scenarios. The variances σ_{ij}^k were defined as follows:

$$\sigma_{ij}^k = \begin{cases} (0.0002\tilde{A}_{ij}^k)^{1/2}, & \text{for all } i, j, k \text{ such that } \tilde{A}_{ij}^k \neq 0, \\ 0 & \text{otherwise.} \end{cases} \quad (4.1)$$

Figure 1. Total dose to each voxel in a slice using the expected position and the beamlet weights obtained from solving (a) the robust formulation and (b) the nominal LP. The colour coding represents the dose value in Gy, delivered over the total treatment period. Regions are marked in their nominal positions.

Table 1. The probability distribution for the given shift in considered directions.

Scenario	Direction	Shift (mm)	Probability
1	None	0.0	0.32
2	Left	3.6	0.17
3	Right	3.6	0.17
4	Anterior	3.6	0.17
5	Posterior	3.6	0.17

This choice appears to be consistent with the uncertainties actually present in a Monte Carlo calculation; see the discussion in Kowalok (2004) (chapter 2).

Having obtained an optimal weight vector w from either the nominal or robust planning procedures, we evaluate its efficacy by simulating a dose distribution obtained with w over a 40-fraction treatment period. The actual dose matrices A^k , $k = 1, 2, \dots, L$ for each scenario are obtained by choosing A_{ij}^k from a normal distribution with mean $\tilde{A}_{ij}^k$ and standard deviation σ_{ij}^k . We then choose the dose matrix $A^{(q)}$ applied at each fraction $q = 1, 2, \dots, N$ by selecting one of $A^1, A^2, \dots, A^L$ at random according to the probabilities in table 1.

4.2. Using PTV in both nominal and robust plans

In our data set, the specified PTV is quite extensive, suggesting that it was defined clinically to account not just for possible motion of the CTV but also widespread extensions of the tumour. In our first experiment, we retained the PTV in the robust plan, shifting it along with the CTV and critical regions to obtain each scenario in table 1. We would expect the robust plan to have the advantage that the PTV and CTV voxels would be more likely to receive doses within the specified intervals, regardless of which scenarios are realized during the course of treatment. This greater coverage of the target regions may come at the price of increasing dose slightly to the normal and critical regions.

Figure 1 shows the intensity plots obtained by applying the expected dose matrix $\hat{A}$ (see (2.3)) to the optimal weight vectors w obtained from the robust formulation and the

Figure 2. Difference between the dose distributions obtained by using the expected position and the beamlet weights from solving the robust formulation and the nominal LP, that is, the difference between figures 1(a) and (b). The colour coding represents the difference in dose values, measured in Gy, with darker regions being those for which the dose from the robust plan is higher.

nominal formulation. Here and in all similar plots in this section, we show dose to the first of the two axial planes of the treatment region; the dose pattern to the second plane is similar. In these plots, the regions are marked in their unshifted positions. The CTV is the central region that receives high dose; the PTV is the large region surrounding the PTV that receives a slightly lower dose. Parotids are marked at the left and right of the figure, and the spinal cord is below the PTV. The difference in intensities obtained from the two formulations is shown in figure 2.

The differences in dose between robust and nominal solutions are minor, except that the robust solution allocates higher dosage along the upper and lower boundaries of the PTV. This pattern is to be expected, as the robust solution is covering the possibility that the PTV will be shifted up or down during some fractions, and is therefore delivering higher dose to these areas to avoid underdosing too much in these situations. Other authors have noted a similar pattern of increased doses around the target boundary for other robust formulations; see, for example, Chan *et al* (2006) (figures 9 and 10).

Figure 3(a) shows the DVH plot for the dose distributions plotted in figure 1, with the robust dose distribution shown as solid lines and the nominal distribution as dotted lines. This figure suggests that the robust solution respects the lower bounds on doses to CTV and PTV, avoiding the cold spots (due mainly to target motion) that are evident in the nominal solution. We see also that dose to critical structures and normal tissue is slightly elevated in the robust solution, as a consequence of the generally higher dose being delivered to the target regions.

Figure 3(b) contains DVH plots for 25 simulations of the treatment plan, for both nominal and robust weight vectors. (The simulation procedure was described at the end of subsection 4.1.) This plot has the same general features as the expected dose plot figure 3(a), with the robust solution observing the lower bounds to target regions more scrupulously and delivering slightly higher doses to other regions. Table 2 analyses some aspects of these 25 simulations more closely. It shows minimum, maximum and mean values of certain quantities associated with the CTV and PTV, namely, the minimum dose delivered to any voxel in these regions, the maximum dose to any voxel, and the percentage of voxels that violate the lower bounds for these two regions. We note in particular that for the robust solution, the

Figure 3. (a) DVH for the expected position of organs and expected dose matrix. The solid line is the robust case and the dotted line is the nominal solution. (b) 25 cases of simulated dose for robust (solid) and nominal (dotted) plans.

Table 2. Maximum, minimum and mean values of various quantities associated with PTV and CTV, obtained over 25 simulated cases, for both robust and nominal solutions.

		Robust solution			Nominal solution		
		Min	Max	Mean	Min	Max	Mean
CTV	Minimum dose	63.82	66.24	65.16	62.69	64.22	63.39
	Maximum dose	74.08	76.02	74.97	73.88	76.02	74.93
	Violation of lower	0.00%	1.02%	0.41%	3.05%	5.54%	3.98%
PTV	Minimum dose	48.58	55.18	52.27	41.28	47.62	44.20
	Maximum dose	71.01	74.89	72.82	69.28	71.60	70.23
	Violation of lower	0.00%	0.64%	0.21%	2.64%	4.97%	3.84%

average percentage of violated lower bounds is small for the CTV and PTV (0.41% and 0.21%, respectively), while the corresponding violations for the nominal solution are much higher (nearly 4% for both regions). The worst-case violations follow a similar pattern. Note too that the minimum doses to both regions are significantly higher for the robust solution (as indicated by the plots of figure 3), while the maximum doses to the PTV are slightly higher as well.

Figure 4 shows the variance on the doses to each voxel in the target and critical regions (the regions for which variance quantities were used in the robust formulation (3.4)). The voxels are ordered by total variance (defined by (2.5)) for the optimal robust weight vector w , with contributions from the organ position uncertainty and dose calculation uncertainty shown separately. The maximum total variance to any voxel is about 17.6, while the maximum contribution from the organ position component is about 16.8 and the maximum contribution from dose calculation uncertainty is about 1.55. Figure 5 shows how these two contributions to the variance are distributed over the target and critical regions. As both figure 4 and figure 5(b) indicate, the dose calculation error is spread relatively uniformly across the target regions, in accordance with the generally high doses that these regions receive. There is also some spread of the dose calculation variance into the boundary areas of the parotids, again a function of the significant dose received by these regions. For the organ motion contribution, both figures 4 and 5(a) indicate much less uniformity. As expected, this contribution is concentrated around

Figure 4. The variance for each tumour and critical voxel. The three lines are the total variance (blue), organ movement (green) and dose calculation (red) components.

Figure 5. Contribution to variance for target and critical voxels from (a) organ position uncertainty and (b) dose calculation uncertainty.

the boundaries of the target and critical regions, where the voxels are target voxels in some scenarios (receiving high dose) and critical or normal in other scenarios (receiving low dose). The peak contribution from organ motion dominates the peak dose calculation contribution, but on the large majority of voxels, dose calculation has equal or greater significance.

Figure 6 shows probabilities of violating the given nominal constraints to the target region on each voxel, for the optimal robust weight vector. Violation probabilities for lower bounds tend to be greatest near the boundaries of the CTV and PTV, as figure 6(a) shows. The constraint $c_L^i(w) \leq 0$ (see (2.9)) ensures that the probability of violation of voxel i is no greater than $\delta = 0.05$, but as we see in table 3 (discussed below), these constraints are not satisfied at some voxels, because there exist no weight vectors that satisfy all (robust) lower and upper bound constraints simultaneously. The maximum violation probability at any voxel is about 0.46, but there is only one other voxel with a violation probability above 0.06. Figure 6(b) shows that expected upper bound violations are clustered mostly in the middle of

Figure 6. Probability of violating (a) lower and (b) upper bound constraints, for each voxel in the CTV and PTV.

Table 3. Computational results for algorithm 1.

Iter(k)	Δ	$m(\Delta w^k)$	$\tau(w^k)$	s	$\ c_L(w^k)_+\ _\infty$	$\ c_U(w^k)_+\ _\infty$
0	30.0	15 643.91	27 868.37	0.439	0.17	0.06
1	45.0	15 572.45	17 319.63	0.101	0.12	0.02
2	11.8	15 645.08	17 319.63	0.097	0.12	0.02
3	17.7	15 694.85	16 662.56	0.058	0.13	0.01
4	26.6	15 597.25	16 635.47	0.062	0.17	0.01
5	11.5	15 636.40	16 635.47	0.060	0.17	0.01
6	5.8	15 718.65	16 635.47	0.055	0.17	0.01
7	8.7	15 752.10	16 031.75	0.017	0.13	0.01
8	13.0	15 729.47	15 975.64	0.015	0.15	0.01
9	6.5	15 741.01	15 975.64	0.015	0.15	0.01
10	9.7	15 735.25	15 949.58	0.013	0.15	0.01
11	4.9	15 744.97	15 949.58	0.013	0.15	0.01
12	7.3	15 739.79	15 923.91	0.012	0.14	0.00
13	3.6	15 751.47	15 923.91	0.011	0.14	0.00
14	5.5	15 755.75	15 854.55	0.006	0.16	0.00
15			15 854.55		0.16	0.00

the CTV. The maximum violation probability is about 0.06, reflecting the observation from table 3 that the computed robust solution violates the upper bound constraints $c_U^i(w) \leq 0$ for $\delta = 0.05$ only slightly, when at all.

Computational results for algorithm 1, which is used to obtain the robust formulation results of this section, are shown in table 3. The second column shows the trust region bound Δ , which increases and decreases according to whether the step produced a decrease in the penalty function τ or not. The third column shows the objective value $m(\Delta w)$ obtained from solving each LP subproblem (3.4). The fourth column shows the penalty function value τ , defined in (3.1), while the fifth column shows the value at each iteration of the stopping criterion s . The last two columns show the maximum violation of the constraints (2.8c) and (2.8b) at the current point.

The algorithm stopped after 15 iterations and required 27.3 min. (We do not include in this figure the approximately 4.5 min to perform preliminary setup and storage of the data, an

Figure 7. Total dose to each voxel in a slice using the expected position and the beamlet weights obtained from solving (a) the robust formulation without PTV and (b) the nominal LP using a collar-shaped PTV. Colours represent the dose value in Gy. Regions are marked in their nominal positions.

operation that needs to be performed in both the robust and nominal cases.) The total number of LPs solved is 16, since one LP must be solved to obtain the initial point w^0 . The first few iterations generally take three to six times as long as later iterations, as warm-start information is carried over from one LP subproblem to the next. By comparison, the nominal problem required slightly under 3.5 min to solve, not including the 4.5 min setup time.

4.3. Using multiple scenarios to replace PTV

We performed a second experiment in which we discarded the PTV supplied with the clinical data set, and defined a new one that consisted of a ‘collar’ of four voxels in width (approximately 7.2 mm) around the specified CTV, to account for positional uncertainty of the CTV. This PTV was used in the nominal formulation, but our robust formulation discarded the PTV altogether, instead making explicit use of our knowledge of positional uncertainty to irradiate the targets more precisely. The expectation is that the more precise targeting of the CTV, in contrast to the more conventional, conservative strategy of a collaring PTV, should reduce dose to critical structures and normal tissue.

The nominal plan was obtained by solving this case with the parameter settings of section 4.1, except that the lower bound on the (newly defined) PTV was increased from 55 to 60 Gy and the prescribed dose to PTV voxels was increased from 60 to 65 Gy. We increased the value of c_ε from $10e$ to $15e$ to more firmly discourage dose to the critical structures. Finally, to obtain a more even intensity spread in the normal region, c_N is increased to $8e$ for the normal voxels in the area surrounding the CTV, PTV, and parotids (rows 35 to 68). In the robust planning procedure, we discarded the PTV altogether, and shifted the CTV and critical structures by the amounts described in table 1 to obtain the different scenarios. For the nominal formulation, we also increased the value of c_v for the PTV from 10 to 20, to avoid very high doses in this region. To decrease the probability that any CTV voxels violate their bounds, we used the more conservative value $\delta = 0.02$, thus $z_{1-\delta} = 2.054$.

Figure 7 shows the intensity plots for the robust and nominal plans, calculated with the expected dose matrix $\hat{A}$, while figure 8 shows the difference between these intensities. It is apparent that while the nominal solution covers the PTV adequately, the robust solution

Figure 8. Difference between the dose distributions in figure 7. Shadings represent the dose value in Gy (positive values indicate that the dose from the robust formulation is higher).

Figure 9. DVH for the expected position of regions. The solid line is the robust case and the dotted line is the nominal solution. (b) 25 cases of simulated dose for robust (solid) and nominal (dotted) plans.

covers the CTV with some smearing outside this region in a two-voxel band, to account for the positional uncertainty. The difference plot shows one significant part of the normal region (immediately above the CTV) that receives a slightly higher dose in the robust plan. However there are sizeable regions (above left and above right of the CTV, and in the PTV) in which the robust plan has significantly less dose. The highest individual-voxel differences between the two plans are about -62.9 (a voxel for which the robust plan yielded a much lower dose) and about 25.4 (a voxel for which the nominal plan yielded a lower dose).

Figure 9(a) shows DVH plots for the dose distributions of figure 7, with robust solution shown as solid lines and nominal solution as dotted lines. DVH curves for normal region and PTV are not shown, since the definitions of these voxels differ between the robust and nominal cases. (The PTV voxels from the nominal formulation become normal voxels in the robust formulation, distorting both curves.) This figure shows clearly that, as in section 4.2,

Figure 10. Variances for CTV and critical voxels, ordered by total variance to each voxel. The three lines are the total variance (blue), organ movement (green) and dose calculation (red) components.

Figure 11. Spatial plot of variance contributions to CTV and critical voxels from (a) organ position uncertainty and (b) dose calculation uncertainty. (No variance calculations are performed for the normal tissue.)

the robust solution is better able to satisfy the lower bound on the CTV. We see also that the doses to parotids and spinal cord are generally lower for the robust case, as we hoped for in targeting the CTV more precisely in this case. Figure 9(b) shows DVH plots for 25 simulated treatment plans. This plot shows the same salient features as the expected-value plot, figure 9(a).

Figure 10 shows total variances and contributions from the organ position and dose calculation terms for voxels in the CTV and critical structures, ordered by total variance. We observe three distinct levels of variance. Approximately the first 900 voxels, which are those of the CTV, have a relatively high dose-calculation variance, corresponding to their generally high doses. The next approximately 700 voxels, corresponding to the parotids, have a somewhat lower dose-calculation variance, while the final approximately 400 voxels,

Figure 12. Probability of violating (a) the lower and (b) the upper bound constraints for the CTV voxels.

Table 4. Computational results for algorithm 1.

Iter(k)	Δ	$m(\Delta w^k)$	$\tau(w^k)$	s	$\ c_L(w^k)_+\ _\infty$	$\ c_U(w^k)_+\ _\infty$
0	30.0	33 939.10	56 936.75	0.403	0.16	0.12
1	45.0	33 866.78	40 573.42	0.165	0.08	0.04
2	11.5	33 999.47	40 573.42	0.162	0.08	0.04
3	17.2	34 102.10	39 563.47	0.138	0.07	0.03
4	7.4	34 293.54	39 563.47	0.133	0.07	0.03
5	11.1	34 334.85	36 131.17	0.050	0.10	0.02
6	5.5	34 432.19	36 131.17	0.047	0.10	0.02
7	8.3	34 319.21	35 732.93	0.040	0.06	0.01
8	4.2	34 476.41	35 732.93	0.035	0.06	0.01
9	6.2	34 518.82	35 227.47	0.020	0.08	0.01
10	3.1	34 583.72	35 227.47	0.018	0.08	0.01
11	4.7	34 536.09	34 977.89	0.013	0.08	0.01
12	2.3	34 612.13	34 977.89	0.010	0.08	0.01
13	3.5	34 615.03	34 848.02	0.007	0.07	0.01
14			34 848.02		0.07	0.01

corresponding to the spinal cord, have a negligible variance, corresponding to the very low dose to these voxels. Within the first two categories we see a tailing-off of the contribution from the organ position uncertainty, with the higher values corresponding to boundary voxels and the lower values corresponding to interior voxels of each region. These patterns can be clearly observed in the spatial plots of figure 11.

Figure 12 shows violation probabilities for lower and upper bounds on CTV voxels. For the lower bounds (figure 12(a)), we see higher probabilities clustered around the boundary of the CTV, with a few voxels in the interior at higher risk also. As we see in table 4, the constraints $c_L^i(x) \leq 0$ are not all satisfied at the solution; in fact the maximum lower bound violation probability over all CTV voxels is about 0.16. Figure 12(b) indicates that upper bound violations also are more likely at boundary voxels; the maximum probability of violation for one of these constraints is about 0.029.

Computational results for algorithm 1 applied to the data of this section are shown in table 4, which has the same format as table 3. The algorithm requires 14 iterations to meet the stopping criteria. The robust problem took nearly 2 min to set up and then solved in 3.8 min. The nominal problem took about 2.5 min to set up and solved in 44 seconds. This case solves much faster than the case studied in subsection 4.2, since there are only 3751 constraints in this problem compared to 16 489 in the previous data set.

5. Conclusions and future work

We have shown in this paper how two important sources of uncertainty in IMRT treatment planning problems—errors in dose matrix calculations and positional uncertainties for the regions—can be accounted for rigorously and tractably in a robust optimization formulation. Incorporation of dose calculation uncertainty into the model allows us to use dose matrices that are quite inexact (thus spending less time in the Monte Carlo dose calculation procedure), yet achieve a higher level of satisfaction of the important constraints in the model, though sometimes at the cost of higher doses to non-target tissues. We have also presented an algorithm for solving the robust formulation whose solve time is a modest multiple of that required by the corresponding nominal linear programming formulation, which ignores the uncertainty.

Future work will involve introducing robust formulation methodology to other, more elaborate formulations of the IMRT planning problem, and testing more extensively on a wider range of data sets. From an algorithmic point of view, we also wish to identify other robust linear programs from other applications for which algorithm 1 is a more efficient solution approach than general SOCP codes, and to develop this algorithm further.

Acknowledgments

We thank the anonymous referees for valuable comments which greatly improved the paper. We are grateful also to our colleague Robert Jeraj for helpful discussions. This research was supported in part by the US National Science Foundation under grants CCF-0113051, DMS-0427689, CCF-0430504, CTS-0456694, CNS-0540147 and SCI-0330538, and by the US Department of Energy under grant DE-FG02-04ER25627.

Appendix. Dose variance

We describe here how the variance of the total dose to a single voxel is calculated approximately under our assumptions. Recall that $A^{(q)}$ is the dose actually delivered on fraction q . This matrix takes one of the values A^k , $k = 1, 2, \dots, L$, where each A^k represents the dose matrix for scenario k and is random with elements distributed normally with expected value $\tilde{A}_{ij}^k$ and variance $(\sigma_{ij}^k)^2$. Moreover, the A_{ij}^k are independently distributed for $i = 1, 2, \dots, m$ and $j = 1, 2, \dots, n$. We use $\tilde{A}^{(q)}$ to denote $\tilde{A}^k$ when scenario k occurs on fraction q . In a complete treatment of N fractions, scenario k will occur approximately $p_k N$ times, a fact that we use at various places in the derivation. Recall that we have assumed also that the scenarios delivered at each fraction q are independent.

For the total dose delivered to voxel i , we have that

$$x_i = \sum_{q=1}^N \sum_{j=1}^n A_{ij}^{(q)} w_j$$

and $E(x_i)$ is given by (2.4). For the variance we have

$$\begin{aligned}
 \text{Var}(x_i) &= E\{(x_i - E(x_i))^2\} \\
 &= E\left\{\left(\sum_{q=1}^N A_{i\cdot}^{(q)} w - N \hat{A}_{i\cdot} w\right)^2\right\} \\
 &= E\left\{\left(\sum_{q=1}^N \sum_{j=1}^n (A_{ij}^{(q)} - \hat{A}_{ij}) w_j\right)^2\right\} \\
 &= E\left\{\sum_{q=1}^N \sum_{j=1}^n \sum_{\hat{q}=1}^N \sum_{\hat{j}=1}^n (A_{ij}^{(q)} - \hat{A}_{ij})(A_{i\hat{j}}^{(\hat{q})} - \hat{A}_{i\hat{j}}) w_j w_{\hat{j}}\right\} \\
 &= E\left\{\sum_{q=1}^N \sum_{j=1}^n \sum_{\hat{q}=1}^N \sum_{\hat{j}=1}^n (A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)} + \tilde{A}_{ij}^{(q)} - \hat{A}_{ij})(A_{i\hat{j}}^{(\hat{q})} - \tilde{A}_{i\hat{j}}^{(\hat{q})} + \tilde{A}_{i\hat{j}}^{(\hat{q})} - \hat{A}_{i\hat{j}}) w_j w_{\hat{j}}\right\} \\
 &= E\left\{\sum_{q=1}^N \sum_{j=1}^n \sum_{\hat{q}=1}^N \sum_{\hat{j}=1}^n (A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)})(A_{i\hat{j}}^{(\hat{q})} - \tilde{A}_{i\hat{j}}^{(\hat{q})}) w_j w_{\hat{j}}\right\} \\
 &\quad + 2E\left\{\sum_{q=1}^N \sum_{j=1}^n \sum_{\hat{q}=1}^N \sum_{\hat{j}=1}^n (A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)})(\tilde{A}_{i\hat{j}}^{(\hat{q})} - \hat{A}_{i\hat{j}}) w_j w_{\hat{j}}\right\} \\
 &\quad + E\left\{\sum_{q=1}^N \sum_{j=1}^n \sum_{\hat{q}=1}^N \sum_{\hat{j}=1}^n (\tilde{A}_{ij}^{(q)} - \hat{A}_{ij})(\tilde{A}_{i\hat{j}}^{(\hat{q})} - \hat{A}_{i\hat{j}}) w_j w_{\hat{j}}\right\}.
 \end{aligned}$$

We consider the three terms in the latter expression separately.

For the first term, we have that when *different* scenarios occur on fractions q and $\hat{q}$ that the random variables $(A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)})$ and $(A_{i\hat{j}}^{(\hat{q})} - \tilde{A}_{i\hat{j}}^{(\hat{q})})$ are independent and both have zero expected value. Hence, we have $E[(A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)})(A_{i\hat{j}}^{(\hat{q})} - \tilde{A}_{i\hat{j}}^{(\hat{q})})] = 0$ in these circumstances. When $j \neq \hat{j}$, we have by independence across beamlets that $E[(A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)})(A_{i\hat{j}}^{(\hat{q})} - \tilde{A}_{i\hat{j}}^{(\hat{q})})] = 0$. Therefore the only nonzero terms in the first term are those for which $j = \hat{j}$ and the same scenario occurs on fractions q and $\hat{q}$. Of the N^2 possible fraction pairs $(q, \hat{q})$, scenario k occurs on both fractions in approximately $p_k^2 N^2$ cases. Hence, we have for the first term that

$$\begin{aligned}
 &E\left\{\sum_{q=1}^N \sum_{j=1}^n \sum_{\hat{q}=1}^N \sum_{\hat{j}=1}^n (A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)})(A_{i\hat{j}}^{(\hat{q})} - \tilde{A}_{i\hat{j}}^{(\hat{q})}) w_j w_{\hat{j}}\right\} \\
 &= \sum_{q=1}^N \sum_{\hat{q}=1}^N \sum_{j=1}^n E[(A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)})(A_{ij}^{(\hat{q})} - \tilde{A}_{ij}^{(\hat{q})})] w_j^2 \\
 &\approx \sum_{k=1}^L \sum_{j=1}^n p_k^2 N^2 E[(A_{ij}^k - \tilde{A}_{ij}^k)^2] w_j^2
 \end{aligned}$$

$$\begin{aligned}
&= \sum_{k=1}^L \sum_{j=1}^n p_k^2 N^2 (\sigma_{ij}^k)^2 w_j^2 \\
&= \sum_{j=1}^n N^2 \tilde{\sigma}_{ij}^2 w_j^2,
\end{aligned} \tag{A.1}$$

where $\tilde{\sigma}_{ij}$ is defined in (2.6b).

For the second term, we have that $E(A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)}) = 0$ for all i, j, q . Similarly to the first term, we can assert that the variables $(A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)})$ and $(\tilde{A}_{i\hat{j}}^{(\hat{q})} - \hat{A}_{i\hat{j}})$ are independent when the scenarios at fractions q and $\hat{q}$ are different. By similar reasoning to the above, we have

$$\begin{aligned}
&E \left\{ \sum_{q=1}^N \sum_{j=1}^n \sum_{\hat{q}=1}^N \sum_{\hat{j}=1}^n (A_{ij}^{(q)} - \tilde{A}_{ij}^{(q)}) (\tilde{A}_{i\hat{j}}^{(\hat{q})} - \hat{A}_{i\hat{j}}) w_j w_{\hat{j}} \right\} \\
&\approx \sum_{k=1}^L (N p_k)^2 E \left\{ \sum_{j=1}^n \sum_{\hat{j}=1}^n (A_{ij}^k - \tilde{A}_{ij}^k) (\tilde{A}_{i\hat{j}}^k - \hat{A}_{i\hat{j}}) w_j w_{\hat{j}} \right\} \\
&= \sum_{k=1}^L (N p_k)^2 E \left\{ \sum_{j=1}^n (A_{ij}^k - \tilde{A}_{ij}^k) w_j \left(\sum_{\hat{j}=1}^n (\tilde{A}_{i\hat{j}}^k - \hat{A}_{i\hat{j}}) w_{\hat{j}} \right) \right\}.
\end{aligned}$$

In the last expression, the expression in parentheses is constant for each scenario k , so from $E(A_{ij}^k - \tilde{A}_{ij}^k) = 0$, we deduce that the entire expression is approximately zero.

We now examine the third term. This term is quite similar to the variance term considered in Chu *et al* (2005), since it is based on the same position uncertainty model and contains no dose calculation errors. Therefore, we simply adopt the rationale of Chu *et al* (2005), which is based on the central limit theorem, in computing this contribution to the variance. We conclude that

$$\begin{aligned}
&E \left\{ \sum_{q=1}^N \sum_{j=1}^n \sum_{\hat{q}=1}^N \sum_{\hat{j}=1}^n (\tilde{A}_{ij}^{(q)} - \hat{A}_{ij}) (\tilde{A}_{i\hat{j}}^{(\hat{q})} - \hat{A}_{i\hat{j}}) w_j w_{\hat{j}} \right\} \\
&\approx N [\tilde{A}_i w - e p^T \tilde{A}_i w]^T P [\tilde{A}_i w - e p^T \tilde{A}_i w] \\
&= \|R_i \tilde{A}_i w\|_2^2
\end{aligned} \tag{A.2}$$

where we refer to (2.6) and the comments following these formulae for the definitions of $\tilde{A}_i$, R_i , P and p .

By conjoining (A.1) with (A.2), we can express the variance of the total dose x_i as

$$\text{Var}(x_i) \approx \sum_{j=1}^n N^2 \tilde{\sigma}_{ij}^2 w_j^2 + \|R_i \tilde{A}_i w\|_2^2 = \left\| \begin{bmatrix} \Sigma_i \\ R_i \tilde{A}_i \end{bmatrix} w \right\|_2^2, \tag{A.3}$$

where Σ_i is defined in (2.6b).

References

- Alber M and Reemtsen R 2004 Intensity modulated radiation therapy planning by use of a Lagrangian barrier-penalty algorithm *Optim. Methods Softw.* at press
- Bahr G, Kereiakes J, Horwitz H, Finney R, Galvin J and Goode K 1968 The method of linear programming applied to radiation treatment planning *Radiology* **91** 686–93
- Baum C, Alber M, Birkner M and Nüsslin F 2006 Robust treatment planning for intensity modulated radiotherapy of prostate cancer based on coverage probabilities *Radiother. Oncol.* **78** 27–35

- Beckham W A, Keall P J and Siebers J V 2002 A fluence-convolution method to calculate radiation therapy dose distributions that incorporate random set-up error *Phys. Med. Biol.* **47** 3465–73
- Ben-Tal A and Nemirovski A 2000 Robust solutions of Linear Programming problems contaminated with uncertain data *Math. Prog. A* **88** 411–24
- Brooke A, Kendrick D, Meeraus A and Raman R 1998 *GAMS: A User's Guide* (Washington, DC: GAMS Development Corporation) <http://www.gams.com>
- Chan T C Y, Bortfeld T and Tsitsiklis J N 2006 A Robust approach to IMRT Optimization *Phys. Med. Biol.* **51** 2567–83
- Chu M, Zinchenko Y, Henderson S G and Sharpe M B 2005 Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty *Phys. Med. Biol.* **50** 5463–77
- Hodes L 1974 Semiautomatic optimization of external beam radiation treatment planning *Radiology* **110** 191–6
- Hou Q, Wang J, Chen Y and Galvin J M 2003 An optimization algorithm for intensity modulated radiotherapy - The simulated dynamics with dose-volume constraints *Med. Phys.* **30** 61–8
- Intensity Modulated Radiation Therapy Collaborative Working Group 2001 Intensity-modulated radiotherapy: current status and issues of interest *Int. J. Radiat. Oncol. Biol. Phys.* **51** 880–914
- Jeraj R and Keall P 1999 Monte Carlo-based inverse treatment planning *Phys. Med. Biol.* **44** 1885–96
- Klepper L Y 1966 Electronic computers and methods of linear programming in the choice of optimal conditions for radiation teletherapy *Medicinskaya Radiol.* **11** 8–15
- Kowalok M E 2004 Adjoint methods for external beam inverse treatment planning *PhD Thesis* Medical Physics Department, University of Wisconsin, Madison, WI
- Langer M, Brown R, Urie M, Leong J, Stracher M and Shapiro J 1990 Large scale optimization of beam weights under dose-volume restrictions *Int. J. Radiat. Oncol. Biol. Phys.* **18** 887–93
- Leong J 1987 Implementation of random positioning error in computerised radiation treatment planning systems as a result of fractionation *Phys. Med. Biol.* **32** 327–34
- Llacer J, Deasy J O, Bortfeld T R, Solberg T D and Promberger C 2003 Absence of multiple local minima effects in intensity modulated optimization with dose-volume constraints *Phys. Med. Biol.* **48** 183–210
- Lu W, Olivera G H, Chen M-L, Reckwerdt P J and Mackie T R 2005 Accurate convolution/superposition for multi-resolution dose calculation using cumulative tabulated kernels *Phys. Med. Biol.* **50** 655–80
- Ma C-M, Mok E, Kapur A, Pawlicki T, Findley D, Brain S, Forster K and Boyer A 1999 Clinical implementation of a Monte Carlo treatment planning system *Med. Phys.* **26** 2133–43
- Morrill S, Rosen I, Lane R and Belli J 1990 The influence of dose constraint point placement on optimized radiation therapy treatment planning *Int. J. Radiat. Oncol. Biol. Phys.* **19** 129–41
- Ólafsson A and Wright S J 2006 Linear programming formulations and algorithms for radiotherapy treatment planning *Optim. Methods Softw.* **21** 201–31
- Reemtsen R and Alber M 2004 Continuous optimization of beamlet intensities for photon and proton radiotherapy Tech. Rep., Brandenburgische Technische Universität Cottbus
- Rosen I I, Lane R G, Morrill S M and Belli J A 1991 Treatment plan optimization using linear programming *Med. Phys.* **18** 141–52
- Sempau J, Wilderman S J and Bielajew A F 2000 DPM, a fast, accurate Monte Carlo code optimized for photon and electron radiotherapy treatment planning dose calculations *Phys. Med. Biol.* **45** 2263–91
- Shepard D M, Ferris M C, Olivera G H and Mackie T R 1999 Optimizing the delivery of radiation therapy to cancer patients *SIAM Rev.* **41** 721–44
- Stapleton S, Zavgorodni S, Popescu I A and Beckham W A 2005 Implementation of random set-up errors in Monte Carlo calculated dynamic IMRT treatment plans *Phys. Med. Biol.* **50** 429–39
- 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
- van Herk M 2004 Errors and margins in radiotherapy *Semin. Radiat. Oncol.* **14** 52–64
- Wu C 2002 Treatment planning in adaptive radiotherapy *PhD Thesis* Medical Physics Department, University of Wisconsin, Madison, WI
- Zhang X, Liu H, Wang X, Dong L, Wu Q and Mohan R 2004 Speed and convergence properties of gradient algorithms for optimization of IMRT *Med. Phys.* **31** 1141–52