

Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty

Millie Chu†, Yuriy Zinchenko†, Shane G Henderson† and
Michael B Sharpe‡

† School of Operations Research & Industrial Engineering, Cornell University, Ithaca NY 14853, USA

‡ Princess Margaret Hospital, Toronto ON M5G 2M9, Canada

E-mail: mchu@orie.cornell.edu, yzinchen@orie.cornell.edu,
sgh9@cornell.edu, michael.sharpe@rmp.uhn.on.ca

Abstract. The recent development of Intensity Modulated Radiation Therapy (IMRT) allows the dose distribution to be tailored to match the tumour's shape and position, avoiding damage to healthy tissue to a greater extent than previously possible. Traditional treatment plans assume the target structure remains in a fixed location throughout treatment. However, many studies have shown that because of organ motion, inconsistencies in patient positioning over the weeks of treatment, etc., the tumour location is not stationary. We present a probabilistic model for the IMRT inverse problem and show that it is identical to using robust optimization techniques, under certain assumptions. For a sample prostate case, our computational results show that this method is computationally feasible and promising - compared to traditional methods, our model has the potential to find treatment plans that are more adept at sparing healthy tissue while maintaining the prescribed dose to the target.

Submitted to: *Phys. Med. Biol.*

1. Introduction

Over the past ten years, continuing technological developments have raised the accuracy and precision of radiation therapy delivery to new levels of sophistication. This progress is well-demonstrated in rapid adoption of technologies that support Intensity Modulated Radiation Therapy (IMRT). The development of IMRT has combined treatment optimization and computer-controlled treatment delivery to shape radiation dose distributions with high precision (IMRT Collaborative Working Group 2001). It is now possible to treat irregularly shaped target volumes aggressively with therapeutic doses, while sparing the surrounding healthy tissues as far as physically possible with photon radiotherapy.

However, high precision treatment delivery requires a high precision of targeting. Treatment planning is a strategy that uses Computed Tomography (CT) images to visualize the location and extent of the tumour and surrounding anatomic structures. Fractionation, another strategy that limits normal tissue damage, extends the treatment time over several weeks, dividing up the total dose into several smaller doses that are administered daily. Motion and deformation of the patient or of the inner organs during, or between, the treatment fractions, as well as any positioning uncertainties, could have a detrimental effect on local control of the disease, or on treatment toxicity.

Our contributions in this paper are:

- to demonstrate that current optimization methods for finding treatment plans are suitable for a fixed target but can perform poorly when motion occurs,
- to present a robust optimization model for the IMRT inverse problem,
- to show that this robust formulation is computationally tractable and produces plans that perform better than current methods when motion occurs, and
- to emphasize that performance measures used to evaluate the quality of a treatment plan need to address uncertainty and show potential outcomes.

We discuss the uncertainty inherent in this problem in section 1.1, and show an example in section 1.2 of how current methods can produce plans that are not robust under uncertainty. In section 2 we present our model formulation, derived from both a probabilistic and robust approach. In section 3 we address dose-volume constraints, and in section 4 we present computational results for a prostate case.

1.1. Sources of Uncertainty and Current Techniques

Immobilization devices are commonly used to help patients maintain consistent positioning, both during treatment and from fraction to fraction (Bentel 1999). Even with these aids, small discrepancies are inevitable because of, for example, periodic breathing and cardiac motion within a treatment; as well as changes in intra-abdominal pressure and weight changes over the protracted course of treatment.

Until recently, treatment planning has been performed on CT images that represent a static patient model, and uncertainties caused by movement and deformation have been addressed by adding safety margins to targets in order to assure that they receive the full prescribed dose. The safety margins are designed to accommodate systematic and random uncertainties that might be introduced following the acquisition of the initial planning data. Following the nomenclature and guidelines of the International Commission on Radiation Units and Measurements (ICRU 1993, ICRU 1999), a clinical target volume (CTV) is defined as the region to be irradiated, including the gross tumour volume (GTV) identified by diagnostic procedures, and regions suspected of harbouring undetectable disease. In order to address the uncertainty described above, the CTV is encompassed by a margin to form the planning target volume (PTV), which is to be treated with a high, uniform dose. Safety margins are sometimes also added to critical structures adjacent to the target to provide a margin of safety to assure the structures remain protected.

Alternative approaches model and consider the effect of movement on the dose distribution during the treatment planning process, with the aim of giving a more complete and accurate indication of the potential consequences of patient motion and setup uncertainty (Bortfeld *et al* 2002). Studies have been done to examine a convolution/smearing approximation of the dose distribution (Craig *et al* 2003a, Craig *et al* 2003b) and the effects of a blurred dose matrix on tumour control probability (van Herk *et al* 2002). In optimization models, approaches include incorporating the expectation (Li and Xing 2000) and variance (Unkelbach and Oelfke 2004, Unkelbach and Oelfke 2005) of dose delivered into the objective function, using an estimated probability distribution for structure locations.

With the advent of online imaging technologies, it is now possible to routinely acquire CT scans in the treatment setting. Using these images to locate the target can reduce patient setup errors, but patient and organ motion during treatment is inevitable.

1.2. A Small Example

Instead of representing the target region as a PTV, our model discards the safety margin and instead uses the CTV but allows for many possible shifts. We also remove any safety margins on healthy structures.

We use the conventional modeling approach of using voxels to measure the amounts of dose absorbed, measured in Gray (Gy). For simplicity, we use a deterministic dose calculation from a matrix of doses delivered to a water phantom, although in practice we would instead use dose calculations specific to the case. We assume that the cumulative dose delivered to a voxel over all fractions is linearly additive. Further, we assume that the beam angles have been pre-selected by an experienced planner, so our solutions are an assignment of intensities for individual beamlets.

To illustrate the potential downfall of using a safety margin to deal with uncertainty, consider the stylized example shown in figure 1.

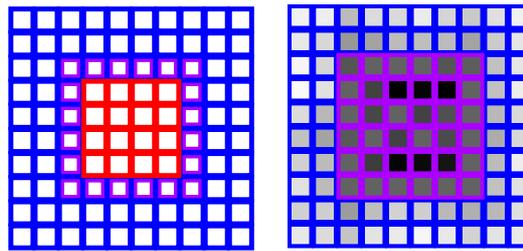


Figure 1. (Left) For a small example, the CTV is the center four-by-four square, surrounded by the PTV margin, one unit wide, and by the healthy structure, three units wide. Note the overlap between the PTV margin and the healthy structure. (Right) The solution of deterministic model is shown, where darker gray indicates higher dose delivered, lighter indicates lower.

Figure 1 shows the optimal solution (the “deterministic solution”) obtained by using a simple linear program that penalizes overdosing the healthy structure and penalizes underdosing the PTV.

Now suppose that a voxel’s location is not stationary throughout treatment, and our best estimate of its actual position during a given fraction is represented by five scenarios - staying in its original position or shifting 1 unit in each direction (right, left, up, down), with each scenario equally likely. Figure 2 shows how the deterministic plan performs when all voxels shift together. The CTV voxels remain within the PTV boundary for all scenarios, so adequate dose is delivered to the CTV. But because the deterministic plan delivers a high, uniform dose to the entire PTV, the healthy structure receives much of this high dose in most scenarios.

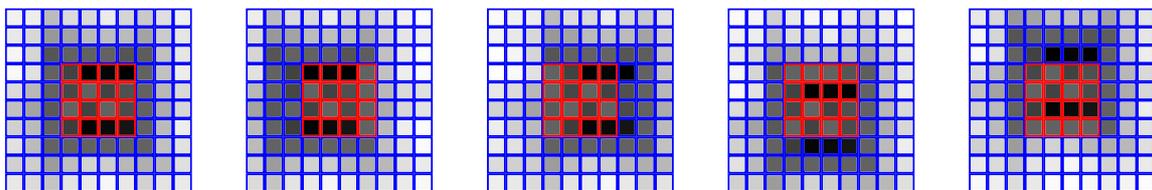


Figure 2. The optimal solution of deterministic model shown in five potential scenarios (no shift, right, left, up, down, respectively). The CTV voxels are the center four-by-four square, surrounded by the healthy structure. Darker gray indicates higher dose delivered, lighter indicates lower.

As an alternative to the deterministic model, we incorporate the scenario information into a robust model (see section 2 for details), so that the constraints now reflect voxel location uncertainty and therefore uncertainty in dose delivered. The optimal solution to this robust model is shown in figure 3. A sufficiently high dose is still delivered to the CTV in all scenarios, and the smoother dose distribution reduces the high dose to the healthy structure.

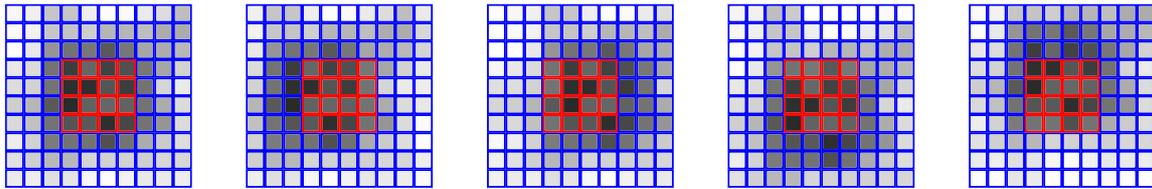


Figure 3. The optimal solution of robust model shown in five potential scenarios (no shift, right, left, up, down, respectively). The CTV voxels are the center four-by-four square, surrounded by the healthy structure. Darker gray indicates higher dose delivered, lighter indicates lower, on the same scale as figure 2.

2. The Formulation

The example in the previous section motivates us to look for a plan that is more robust to the various uncertainties in the problem. There are at least two ways to view this uncertainty. One interpretation is a probabilistic one, where we require that the plan be of high quality with high probability for any particular patient. This interpretation is intuitive, but relies on a normal approximation for the distribution of the cumulative dose delivered to a voxel. We present a second interpretation that does not rely on this assumption. Using a robust optimization approach, we require that the patient’s plan be of high quality so long as the problem parameters fall within some “reasonable” set. In section 2.1 and section 2.2 we describe these interpretations in turn, and under certain assumptions show that they lead to the same form of constraint. The formulation of the problem appears in section 2.3, minus dose-volume constraints which are discussed in section 3.

2.1. A Probabilistic Interpretation

Let $D_i(\mathbf{x})$ denote the total dose delivered to voxel i over the course of all N fractions, as a function of the beamlet intensity vector \mathbf{x} . We view $D_i(\mathbf{x})$ as a random variable because it depends on the doses delivered on each of the N fractions, and the dose per fraction depends on the location of the voxel during the fraction, which is random. More specifically, $D_i(\mathbf{x})$ is the sum of N random variables $D_{i\ell}(\mathbf{x})$, $\ell = 1, \dots, N$, where $D_{i\ell}(\mathbf{x})$ is the dose delivered to voxel i on the ℓ th fraction with beamlet intensity vector \mathbf{x} . In great generality, a sum of random variables is approximately normally distributed, as follows from a central limit theorem. We therefore assume that $D_i(\mathbf{x})$ is normally distributed. Denote the mean and variance of the total dose $D_i(\mathbf{x})$ by μ and σ^2 respectively. These parameters depend on i and \mathbf{x} , but for now we suppress that dependence in our notation.

Suppose that voxel i comes from some healthy structure H_k , so that we do not want the dose to voxel i to exceed some level m_k that we take to be structure-specific. (The level m_k could be taken to be voxel-specific, but we do not do so.) This desire can be naturally expressed by requiring that the dose exceeds m_k with probability at most δ , where δ is some small constant, e.g., 0.05. So we require that $P(D_i(\mathbf{x}) > m_k) \leq \delta$,

which is equivalent to saying that

$$\mathbb{P}\left(\frac{D_i(\mathbf{x}) - \mu}{\sigma} > \frac{m_k - \mu}{\sigma}\right) \leq \delta. \quad (1)$$

But $(D_i(\mathbf{x}) - \mu)/\sigma$ is normally distributed with mean 0 and variance 1, so (1) is equivalent to requiring that

$$\frac{m_k - \mu}{\sigma} \geq z_{1-\delta},$$

where $z_{1-\delta}$ is chosen so that $\mathbb{P}(Z > z_{1-\delta}) = \delta$, and Z is a normal random variable with mean 0 and variance 1. We can rearrange this equation to give the constraint

$$\sigma \leq \frac{m_k - \mu}{z_{1-\delta}}.$$

If voxel i comes from the CTV, then we do not want the dose to fall below a certain threshold, which we denote by m_{Tmin} . A similar argument using the fact that $z_\delta = -z_{1-\delta}$ shows that the analogous constraint is of the form

$$\sigma \leq \frac{\mu - m_{\text{Tmin}}}{z_{1-\delta}}. \quad (2)$$

Notice that $z_{1-\delta} > 0$ and we would expect that $m_{\text{Tmin}} < \mu$ in any reasonable solution, so that the right-hand side of (2) is positive, as it should be.

But how do we compute μ and σ^2 ?

Recall that μ and σ^2 are the mean and variance of $D_i(\mathbf{x})$, and are therefore functions of the beamlet intensity vector \mathbf{x} . We assume that the N random variables giving the dose per fraction are independent and identically distributed (i.i.d.). This assumption is a natural first step as it acknowledges the randomness in the problem, and leads to a computationally-feasible optimization problem. It is important to recognize that it is almost certainly violated to some degree. For example, recall that the treatment is delivered over several weeks, during which time the patient is likely to be losing weight as a side effect of the radiation. But if the patient's body geometry changes as time passes, then so will the doses received by various locations in the body, so that the assumption that the fraction doses are identically distributed is best viewed as an approximation.

The i.i.d. assumption ensures that $\mu = \text{NED}_{i1}(\mathbf{x})$ and $\sigma^2 = N\text{Var}D_{i1}(\mathbf{x})$, so the problem reduces to computing the mean and variance of the dose from a single fraction. Let \mathbf{Y}_i denote a random column vector, indexed by beamlets, representing the dose delivered to voxel i from each beamlet, when the beamlets have unit intensity. The dose to voxel i is the sum of the doses delivered from each beamlet, and therefore $D_{i1}(\mathbf{x}) = \mathbf{Y}_i^\top \mathbf{x}$, where \mathbf{v}^\top denotes the transpose of the vector \mathbf{v} . The variance can be rewritten as

$$\sigma^2 = N\text{Var}(\mathbf{Y}_i^\top \mathbf{x}) = N\mathbf{x}^\top \text{Cov}(\mathbf{Y}_i)\mathbf{x},$$

where $\text{Cov}(\mathbf{Y}_i)$ is a $B \times B$ matrix, and B is the number of beamlets. In practice, B is commonly on the order of one thousand, and there are tens of thousands of voxels, so storing a covariance matrix for each voxel requires too much memory for this formulation to be computationally tractable. To avoid the need for this storage and therefore make

the problem tractable, we instead adopt the following model of the random dose on a single fraction.

Suppose that on any single fraction, one of n possible scenarios s_1, \dots, s_n can occur with probabilities p_1, \dots, p_n respectively. Let \mathbf{a}_{ij} denote a column vector, indexed by beamlets, giving the deterministic dose delivered to voxel i in scenario j from each beamlet, when the beamlets have unit intensity. We can then write

$$D_{i\ell}(\mathbf{x}) = \mathbf{a}_{i,S(\ell)}^\top \mathbf{x},$$

where $S(\ell)$ is the index of the (random) scenario that occurs in fraction ℓ . Hence

$$\mathbb{E}D_{i1}(\mathbf{x}) = \mathbb{E}[\mathbf{a}_{i,S(1)}]^\top \mathbf{x} = \sum_{j=1}^n p_j \mathbf{a}_{ij}^\top \mathbf{x},$$

and

$$\text{Var}D_{i1}(\mathbf{x}) = \sum_{j=1}^n p_j [\mathbf{a}_{ij}^\top \mathbf{x} - \mathbb{E}D_{i1}(\mathbf{x})]^2.$$

These expressions can be written in a more useful form as follows. Let $\mathbf{p} = (p_1, \dots, p_n)$ be the column vector of scenario probabilities, and

$$A_i = \begin{pmatrix} \mathbf{a}_{i,1}^\top \\ \mathbf{a}_{i,2}^\top \\ \vdots \\ \mathbf{a}_{i,n}^\top \end{pmatrix}$$

be a matrix where the j th row contains the vector giving the dose to voxel i in scenario j . Then

$$\mathbb{E}D_{i1}(\mathbf{x}) = \mathbf{p}^\top A_i \mathbf{x}.$$

Let \mathbf{e} denote an $n \times 1$ (column) vector where each element equals 1, I denote the $n \times n$ identity matrix, and P denote the diagonal matrix where $P_{jj} = p_j$. Then

$$\begin{aligned} \text{Var}D_{i1}(\mathbf{x}) &= [A_i \mathbf{x} - \mathbf{e}(\mathbf{p}^\top A_i \mathbf{x})]^\top P [A_i \mathbf{x} - \mathbf{e}(\mathbf{p}^\top A_i \mathbf{x})] \\ &= [(I - \mathbf{e}\mathbf{p}^\top) A_i \mathbf{x}]^\top P [(I - \mathbf{e}\mathbf{p}^\top) A_i \mathbf{x}] \\ &= \|\mathbf{R}A_i \mathbf{x}\|^2, \end{aligned}$$

where $\|v\|^2 = v_1^2 + \dots + v_n^2$ denotes the square of the Euclidean norm, $\mathbf{R} = P^{1/2}(I - \mathbf{e}\mathbf{p}^\top)$, and $P^{1/2}$ is the diagonal matrix with $P_{jj}^{1/2} = \sqrt{p_j}$.

Putting it all together we see that the constraint (1) can be written as

$$\|\mathbf{R}A_i \mathbf{x}\| \leq \frac{m_k - N\mathbf{p}^\top A_i \mathbf{x}}{z_{1-\delta}\sqrt{N}}. \quad (3)$$

The constraint (2) can be expressed in almost exactly the same way.

There is one constraint of the form (3) for each voxel. We require any treatment plan to simultaneously satisfy all of these constraints. That does *not* mean that all voxels *simultaneously* receive an appropriate dose with high probability. Rather it means that each voxel, considered *one at a time*, has a high probability of receiving an appropriate dose. Put another way, we constrain the *marginal* distributions of the doses to the voxels, rather than the joint distribution of the (vector) dose to all voxels.

2.2. A Robust Interpretation

Recall that we want our approach not only to incorporate the uncertainty but also to be computationally tractable; that is, possible to implement for realistically-sized problems. In this section we

- introduce the Second Order Conic Programming Problem (SOCP),
- introduce the Robust Linear Programming Problem (Robust LP),
- demonstrate the connection between these two problems and highlight how they arise in our setting,
- give a formulation using Robust LP and show that under certain probabilistic assumptions the two formulations are indeed the same, and
- comment on the overall flexibility of the proposed approach.

For an excellent reference on robust LP, SOCP, and their interconnections, see Ben-Tal and Nemirovski (2001).

SOCP. Consider the conic programming problem

$$\begin{aligned} \min \mathbf{c}^\top \mathbf{x} \\ A\mathbf{x} = \mathbf{b} \\ \mathbf{x} \in K_2^\Pi \end{aligned}$$

with $\mathbf{c}, \mathbf{x} \in \mathfrak{R}^n$, $A \in \mathfrak{R}^{m \times n}$, $\mathbf{b} \in \mathfrak{R}^m$ and K_2^Π being a direct product of second order cones. Notice that at this point SOCP looks almost like LP, except for the nonnegativity constraints $\mathbf{x} \geq 0$, which have been replaced by $\mathbf{x} \in K_2^\Pi$. More precisely,

$$K_2^\Pi = \{\mathbf{x} \in \mathfrak{R}^n : (x_1, x_2, \dots, x_{k_1}) \in K_1, \\ (x_{k_1+1}, \dots, x_{k_2}) \in K_2, \dots, (x_{k_l+1}, \dots, x_n) \in K_{l+1}\}$$

where $1 \leq k_1 < k_2 < \dots < k_l < n$ and each $K_i, i = 1 \dots l + 1$ is a second order cone (SOC). A cone $K \subseteq \mathfrak{R}^{k+1}$ is a SOC if $K = \{(\mathbf{x}, t) \in \mathfrak{R}^k \times \mathfrak{R} : \|\mathbf{x}\| \leq t\}$, and if $k = 0$ this is a nonnegativity constraint, $t \geq 0$. Therefore, SOCP is a generalization of LP.

SOCP is a powerful modeling framework that is computationally tractable, both theoretically and practically. In fact, many existing solvers that are based on interior-point methods require a similar number of iterations to solve LP instances of a similar size. It should be noted, however, that the amount of work associated with each iteration is somewhat more demanding for SOCP.

Robust LP. Consider the LP

$$\begin{aligned} \min \mathbf{c}^\top \mathbf{x} \\ A\mathbf{x} \geq \mathbf{b} \end{aligned}$$

with $\mathbf{c}, \mathbf{x} \in \mathfrak{R}^n$, $A \in \mathfrak{R}^{m \times n}$, $\mathbf{b} \in \mathfrak{R}^m$. (Here, for convenience, we adopt a slightly different formulation that uses inequality constraints rather than equality constraints. Transformations are possible between the different forms so there is no loss of generality

in doing so.) In Robust LP, the underlying idea is that we do not know the matrix A with certainty, but have a good idea of its whereabouts. This knowledge is represented by requiring that $A \in \mathcal{U}$, where \mathcal{U} is usually termed the uncertainty set. Robust LP requires a solution \mathbf{x} to be feasible for *any* choice of $A \in \mathcal{U}$, i.e.,

$$\begin{aligned} \min \mathbf{c}^\top \mathbf{x} \\ A\mathbf{x} \geq \mathbf{b} \quad \forall A \in \mathcal{U}. \end{aligned}$$

We will show that this problem can be cast as a SOCP under some assumptions on the shape of the uncertainty set \mathcal{U} .

Robust LP and SOCP. The so-called ellipsoidal uncertainty set \mathcal{U} corresponds to a direct product of ellipsoids E_i for each constraint $\mathbf{a}_i^\top \mathbf{x} \geq b_i$ of the set of constraints $A\mathbf{x} \geq \mathbf{b}$, where \mathbf{a}_i^\top is the i th row of the matrix A , $i = 1, \dots, m$. We want to solve

$$\begin{aligned} \min \mathbf{c}^\top \mathbf{x} \\ \mathbf{a}_i^\top \mathbf{x} \geq b_i \quad \forall \mathbf{a}_i \in E_i, \text{ for } i = 1, \dots, m. \end{aligned}$$

The ellipsoid E_i can be represented as $\{\mathbf{a}_i(\mathbf{u}_i) = \mathbf{a}_i^* + W_i \mathbf{u}_i : \mathbf{u}_i^\top \mathbf{u}_i \leq 1\}$, where \mathbf{a}_i^* , \mathbf{u}_i and W_i are vectors and a matrix, respectively, of corresponding dimensions. We want the constraint $\mathbf{a}_i^\top \mathbf{x} \geq b_i$ to hold for any choice of $\mathbf{a}_i \in E_i$, that is

$$\begin{aligned} 0 &\leq \min\{\mathbf{a}_i(\mathbf{u}_i)^\top \mathbf{x} - b_i : \mathbf{u}_i^\top \mathbf{u}_i \leq 1\} \\ &= \mathbf{a}_i^{*\top} \mathbf{x} - b_i + \min\{\mathbf{u}_i^\top W_i^\top \mathbf{x} : \mathbf{u}_i^\top \mathbf{u}_i \leq 1\} \\ &= \mathbf{a}_i^{*\top} \mathbf{x} - b_i - \|W_i^\top \mathbf{x}\|, \end{aligned}$$

for $i = 1 \dots m$. Hence, the Robust LP can be written as

$$\begin{aligned} \min \mathbf{c}^\top \mathbf{x} \\ \|W_i^\top \mathbf{x}\| \leq \mathbf{a}_i^{*\top} \mathbf{x} - b_i, \quad \text{for } i = 1, \dots, m, \end{aligned}$$

which is an instance of SOCP.

Now that we have a sense of how Robust LP and SOCP are related, let us turn to a robust formulation of maximum/minimum dose constraints.

We still compute the dose to voxel i on fraction ℓ as a linear combination of the beamlet intensities \mathbf{x} , but now $D_{i\ell}(\mathbf{x}) = \mathbf{a}_i^\top \mathbf{x}$, where $\mathbf{a}_i \in E_i$, E_i is the ellipse $\{\mathbf{a}_i^* + W_i \mathbf{u}_i : \mathbf{u}_i^\top \mathbf{u}_i \leq 1\}$, $\mathbf{a}_i^{*\top} = \mathbf{p}^\top A_i$, and W_i is some appropriately-chosen matrix. The interpretation here is that the coefficients that determine the dose to voxel i are not known with certainty, but are presumed to lie in E_i . The ellipse E_i is centered on the expected dose coefficients \mathbf{a}_i^* . The question remains how to choose W_i .

A natural approach for selecting W_i is to appeal to a probabilistic model like the one in section 2.1. In that case, the appropriate choice turns out to be $W_i = z_{1-\delta} A_i^\top R^\top / \sqrt{N}$, using previous notation. For voxel i in healthy structure H_k , if we require the dose delivered in a single fraction $D_{i\ell}(\mathbf{x}) = \mathbf{a}_i^\top \mathbf{x}$ be at most m_k/N , for all possible dose coefficient vectors \mathbf{a}_i coming from E_i , then the cumulative dose over N fractions will be at most m_k , as desired. The resulting constraint is:

$$\frac{m_k}{N} \geq (\mathbf{a}_i^* + W_i \mathbf{u}_i)^\top \mathbf{x} = \mathbf{p}^\top A_i \mathbf{x} + \frac{z_{1-\delta}}{\sqrt{N}} \mathbf{u}_i^\top R A_i \mathbf{x},$$

for all \mathbf{u}_i satisfying $\mathbf{u}_i^\top \mathbf{u}_i \leq 1$, which is the same as

$$\begin{aligned} 0 &\leq m_k - N\mathbf{p}^\top A_i \mathbf{x} - z_{1-\delta} \sqrt{N} \mathbf{u}_i^\top R A_i \mathbf{x}, \quad \forall \mathbf{u}_i \text{ satisfying } \mathbf{u}_i^\top \mathbf{u}_i \leq 1 \\ &= m_k - N\mathbf{p}^\top A_i \mathbf{x} - z_{1-\delta} \sqrt{N} \max \{ \mathbf{u}_i^\top R A_i \mathbf{x} : \mathbf{u}_i^\top \mathbf{u}_i \leq 1 \} \\ &= m_k - N\mathbf{p}^\top A_i \mathbf{x} - z_{1-\delta} \sqrt{N} \|R A_i \mathbf{x}\|. \end{aligned}$$

Rearranging this inequality gives

$$\|R A_i \mathbf{x}\| \leq \frac{m_k - N\mathbf{p}^\top A_i \mathbf{x}}{z_{1-\delta} \sqrt{N}},$$

which is the same as (3).

Robust LP flexibility. We have demonstrated that imposing an ellipsoidal uncertainty set on the input data (the dose coefficients) gives rise to an ellipsoidal uncertainty set on the cumulative dose $\sum_{\ell=1}^N D_{iS(\ell)}(\mathbf{x})$ to a voxel, which then gives constraints that form part of a SOCP. It is possible to prescribe more sophisticated shapes for the uncertainty set. We do not do so here because the ellipsoidal form seems sufficient for our purposes.

2.3. An Initial Formulation

Many models have been proposed and discussed in the literature, and many others are used by commercial treatment planning systems. For an overview of formulations, see Langer *et al* (2003) and Shepard *et al* (1999). We ignore dose-volume constraints for now, and address them in section 3.

With no definitive model on hand, we will use the following formulation to demonstrate our approach. This formulation has a linear penalty objective function which is similar, in principle, to the model used in the ADAC Pinnacle³ treatment planning system (ADAC 2002). Let h be the number of healthy structures, indexed by k . The formulation is

Minimize

$$w_{T\min} r_{T\min} + w_{T\max} r_{T\max} + \sum_{j=1}^n w_{Tj} r_{Tj} + \sum_{k=1}^h w_k r_k$$

subject to:

$$\begin{aligned} \|R A_i \mathbf{x}\| &\leq \frac{N\mathbf{p}^\top A_i \mathbf{x} - u_{T\min}}{z_{1-\delta} \sqrt{N}} && \forall i \in \text{CTV} \\ m_{T\min} - u_{T\min} &\leq r_{T\min} \\ \|R A_i \mathbf{x}\| &\leq \frac{u_{T\max} - N\mathbf{p}^\top A_i \mathbf{x}}{z_{1-\delta} \sqrt{N}} && \forall i \in \text{CTV} \\ u_{T\max} - m_{T\max} &\leq r_{T\max} \\ \mathbf{a}_{i,j}^\top \mathbf{x} &\geq u_{Tj} && \forall i \in \text{CTV}, \text{ for } j = 1, \dots, n \\ m_T - u_{Tj} &\leq r_{Tj} && \text{for } j = 1, \dots, n \end{aligned}$$

$$\begin{aligned}
\|RA_i\mathbf{x}\| &\leq \frac{u_k - N\mathbf{p}^\top A_i\mathbf{x}}{z_{1-\delta}\sqrt{N}} && \forall i \in H_k, \text{ for } k = 1, \dots, h \\
u_k - m_k &\leq r_k && \text{for } k = 1, \dots, h \\
r_{\text{Tmin}}, r_{\text{Tmax}} &\geq 0 \\
r_{\text{Tj}} &\geq 0 && \text{for } j = 1, \dots, n \\
r_k &\geq 0 && \text{for } k = 1, \dots, h \\
\mathbf{x} &\geq \mathbf{0},
\end{aligned}$$

where w_{Tmin} , w_{Tmax} , w_{Tj} , and w_k are the penalty weights for failing to reach the minimum total dose m_{Tmin} for the CTV, exceeding the maximum total dose m_{Tmax} for the CTV, failing to reach the minimum dose m_{T} in scenario j for the CTV, and exceeding the maximum dose m_k for structure H_k , respectively.

In this penalty formulation, u_k is an intermediate variable replacing parameter m_k in (3), so u_k can be interpreted as the maximum dose to any voxel in H_k with probability $1 - \delta$ when \mathbf{x} is delivered. If $u_k > m_k$, then we penalize the excess dose, $u_k - m_k$, in the objective function, by some weight parameter $w_k > 0$ chosen by the planner. Otherwise, if $u_k \leq m_k$, there is no penalty. The excess dose is measured by imposing the linear constraints $u_k - m_k \leq r_k$ and $r_k \geq 0$. Minimizing the objective will cause $r_k = u_k - m_k > 0$ if $u_k > m_k$, and $r_k = 0$ when $u_k \leq m_k$. We have chosen to make the value u_k structure-specific rather than voxel-specific, since it seems sufficient for our purposes.

The maximum and minimum dose constraints for voxels in the CTV are implemented similarly and allow for some control over dose homogeneity within the target region. The minimum dose requirement for a CTV voxel per scenario allows some control over the minimum dose to the target per fraction.

In our implementation, we include additional small penalties to improve conditioning. For example, when $u_k > m_k$, a penalty with slope w_k is incurred in the objective function, and we further impose a penalty with slope ϵ for when $u_k < m_k$, where ϵ is small.

3. Dose-Volume Constraints

Solving the above formulation produces optimal solutions where most voxels in a healthy structure receive their maximum prescribed dose m_k , or very close to it. These treatment plans are not clinically acceptable, and physicians are accustomed to imposing dose-volume (DV) constraints, which are of the form “no more than $100v_k\%$ of healthy structure H_k may receive more than d_k Gy.” DV constraints are especially meaningful in parallel structures (organs in which the same function is performed throughout) since the organ’s functionality may be preserved in the portion that does not receive a high dose.

Accurately modeling a DV constraint would require a binary variable for each voxel to indicate whether or not its dose is above the threshold d_k . Doing so would add

thousands of binary variables to our model, making it computationally intractable. The alternative method we use here (which is similar to a method described in Ferris *et al* (2003) and Ferris *et al* (2004)) imposes the following constraint on the sum of the “excess” doses (expected doses above the threshold d_k) for structure H_k :

$$\sum_{i \in H_k} (N\mathbf{p}^\top A_i \mathbf{x} - d_k)_+ \leq g_k, \quad (4)$$

where $(\cdot)_+$ means positive part, and g_k is a parameter chosen by the planner. A standard linear programming trick ensures that this constraint can be introduced using linear equalities so that the overall optimization problem remains convex. By restricting the sum of the excess doses we achieve an effect similar to that of a DV constraint, while retaining computational tractability.

One possible value for g_k is $v_k |H_k| (m_k - d_k)$, which is the total amount of excess dose there would be if the fraction v_k of voxels exceeded the threshold d_k by receiving the maximum dose m_k . Values for g_k may also be chosen iteratively, by tightening (or loosening) this bound as solutions dictate.

Alternative methods for imposing dose-volume constraints have been proposed (Cotrutz and Xing 2002, Preciado-Walters 2003, Romeijn *et al* 2003), and further research may include evaluating such alternative approaches.

Building on our initial formulation in section 2.3, we now add our approximate dose-volume constraints. Let $\bar{h} \subset \{1, \dots, h\}$ be the indices of healthy structures that have a prescribed dose-volume restriction. For simplicity, we show the formulation where each structure in \bar{h} has just one dose-volume restriction, but it is straightforward to extend this formulation for multiple dose-volume restrictions on a structure. Continuing with the penalty model, let \bar{w}_k be the penalty weight for exceeding the upper bound g_k on excess doses. The formulation is

Minimize

$$w_{T_{\min}} r_{T_{\min}} + w_{T_{\max}} r_{T_{\max}} + \sum_{j=1}^n w_{T_j} r_{T_j} + \sum_{k=1}^h w_k r_k + \sum_{k \in \bar{h}} \bar{w}_k q_k$$

subject to:

$$\begin{aligned} \sum_{i \in H_k} (N\mathbf{p}^\top A_i \mathbf{x} - d_k)_+ &\leq f_k && \forall k \in \bar{h} \\ f_k - g_k &\leq q_k && \forall k \in \bar{h} \\ q_k &\geq 0 && \forall k \in \bar{h}, \end{aligned}$$

in addition to the constraints in section 2.3.

4. Computational Results

We tested this formulation on a sample prostate case, where the CTV is adjacent to two critical structures, and entirely surrounded by a healthy region (the “Unspecified”

Table 1. Distribution for voxel locations.

Scenario	Direction	Shift (mm)	Probability
1	None	0	0.250
2	Anterior	5	0.125
3	Posterior	3	0.125
4	Left	2	0.125
5	Right	2	0.125
6	Inferior	3	0.125
7	Superior	4	0.125

Table 2. Treatment parameters.

Structure	Prescribed doses
PTV/CTV	Uniform dose of 82.8 Gy
Bladder	Maximum dose of 81.0 Gy
Rectal solid	Maximum dose of 79.2 Gy
Unspecified	Maximum dose of 72.0 Gy
Left femur	Maximum dose of 50.0 Gy
Right femur	Maximum dose of 50.0 Gy
Bladder	No more than 50% may receive more than 60.0 Gy
Rectal solid	No more than 50% may receive more than 25.0 Gy
Rectal solid	No more than 30% may receive more than 50.0 Gy
Rectal solid	No more than 25% may receive more than 60.0 Gy
Rectal solid	No more than 15% may receive more than 73.8 Gy

structure), with $N = 45$ treatment fractions. Our selected distribution for voxel locations is given in table 1. The magnitude of the shifts was approximated from literature on observed prostate motion during the course of treatment (Langen and Jones 2001, Wong *et al* 2005), and we chose to use the same distribution for all voxels. Though setup uncertainties can be systematic in nature, the probabilistic model helps guard against bad outcomes in any one scenario. The distribution can be modified per patient, and we view this distribution as a reasonable first step and plan to study scenario and probability selection in more detail in future work. Values for all treatment parameters were taken from clinical protocols, shown in table 2, which are inherently conflicting due to overlap between the PTV and rectal solid, and between the PTV/CTV and bladder.

We used voxels of size 0.8 cm^3 , yielding about 6500 total voxels. In practice, 0.4 cm^3 sized voxels are used for this type of case, but such a large SOCP required too much memory. For simplicity, our solutions shown reflect the doses delivered to the 0.8 cm^3 sized voxels.

We chose $\delta = 0.05$ for the SOC constraints, with the interpretation being the maximum (or minimum) total dose constraints must hold with probability 0.95. The pre-selected beam angles are five co-planar beams, at gantry angles $0^\circ, 72^\circ, 144^\circ, 216^\circ, 288^\circ$,

with about 200 beamlets at each angle and 940 beamlets total. We used SeDuMi v.1.05 (Sturm 1999) to solve our SOCP to global optimality in approximately 3.3 hours in 97 iterations on a Dell PowerEdge 1750 with a single 3 GHz Xeon CPU and 4 GB RAM. For comparison, a problem on 1.0 cm³ sized voxels (about 3300 total voxels) solved in about one hour in 65 iterations. The optimal robust solution mostly penalizes the maximum dose to the bladder and approximate dose-volume constraints to the rectal solid and bladder.

For comparison, we solved the deterministic LP version of our formulation, where the PTV is the target region instead of the CTV, and all voxels are assumed to stay in their original positions for all N fractions. The approximate dose-volume constraints are similar to (4), and the maximum dose constraints are of the form

$$N\mathbf{a}_{i,1}^\top \mathbf{x} \leq u_k, \quad \forall i \in H_k$$

where $\mathbf{a}_{i,1}$ is the dose coefficient vector for voxel i in scenario s_1 (no shift), and u_k is penalized in the objective function for exceeding m_k . The minimum dose constraints for voxels in the PTV are similar. This LP was solved in about 17 minutes (67 iterations) also using SeDuMi on 0.8 cm³ sized voxels, and the resulting treatment plan was evaluated based on doses delivered to the CTV. The optimal deterministic solution mostly penalizes the maximum dose constraints to the rectal solid and unspecified region, and the approximate dose-volume constraints to the rectal solid.

A Dose-Volume Histogram (DVH) plots the fraction of each structure that receives at least some total dose. Physicians are accustomed to evaluating the quality of a treatment plan based on these traditional plots. Since these plots are based on the dose delivered to static voxels, they do not account for uncertainties and could be misleading. As a natural extension, in figure 4, we plot a ‘‘Dose - Expected Volume Histogram’’ (DEVH) which is like a DVH but, for each dose d , plots the *expected* percentage of a structure that receives at least d :

$$\text{DEVH}_k(d, \mathbf{x}) = \mathbb{E} \left(\frac{1}{|H_k|} \sum_{i \in H_k} \mathbf{1}_{\{D_i(\mathbf{x}) \geq d\}} \right) = \frac{1}{|H_k|} \sum_{i \in H_k} \mathbb{P}(D_i(\mathbf{x}) \geq d),$$

where voxel i belongs to structure H_k and $\mathbf{1}_{\{\cdot\}}$ is the indicator function. Using our assumptions and results from section 2.1, we can express this as:

$$\text{DEVH}_k(d, \mathbf{x}) = \frac{1}{|H_k|} \sum_{i \in H_k} \mathbb{P} \left(Z \geq \frac{d - N\mathbf{p}^\top A_i \mathbf{x}}{\sqrt{N} \|R A_i \mathbf{x}\|} \right),$$

where Z is a normal random variable with mean 0 and variance 1, and \mathbf{x} is the optimal set of beamlet intensities.

The plot in figure 4 shows the expected outcome, whereas the actual outcome may be significantly different. Evaluating the quality of a treatment plan should include consideration of the potential outcomes due to patient motion and setup uncertainties.

For the same prostate case, we simulated a treatment by randomly selecting 45 shifts representing possible positions of the patient at each fraction. We then used each model’s (robust and deterministic) optimal solution of beamlet intensities to calculate

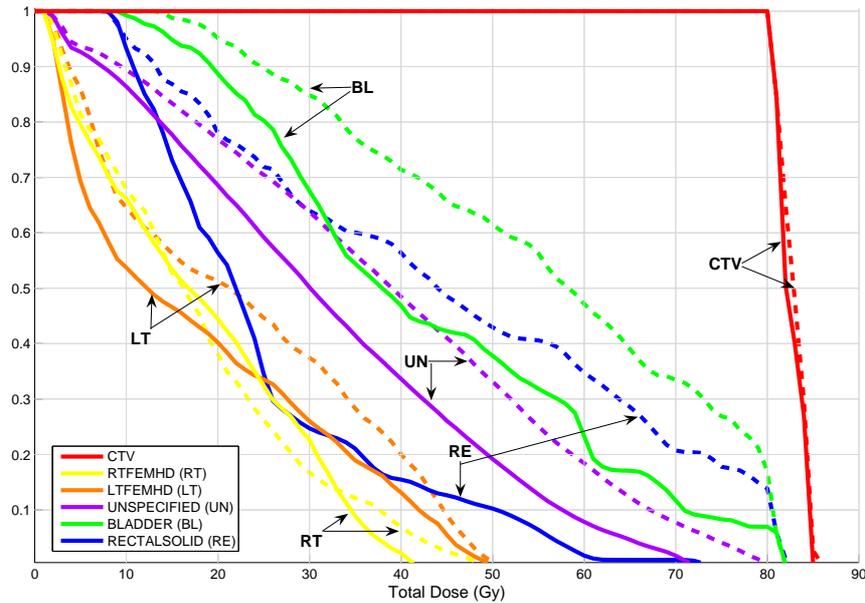


Figure 4. This Dose - Expected Volume Histogram (DEVH) plots the expected fraction of volume of each structure that will receive at least some dose. The solid (dashed) curves represent the robust (deterministic) solution.

cumulative doses delivered. We plotted the resulting DVH, and then simulated more treatments, overlaying all the curves on the same axes. The final plot, shown in figure 5, should convey a sense of the range of possible outcomes for each solution.

More specifically, figure 5 plots the DVH cloud for each of the 100 simulated treatments of $N = 45$ fractions each, sampling from the distribution given by s_1, \dots, s_7 and p_1, \dots, p_7 described above. For each fraction, we assume that all voxels shift together.

The robust and deterministic solutions achieve comparable dose homogeneity within the CTV, but the robust plan delivers less dose to every healthy structure, and especially, significantly less dose to the rectal solid. In all simulations, the deterministic solution violates all four of the prescribed dose-volume constraints for the rectal solid, and the robust solution satisfies all four.

5. Conclusions and Future Research

A robust optimization approach for finding IMRT treatment plans provides a more accurate representation of patient motion and setup uncertainties than traditional methods. For our test patient case, our model produces plans that are more adept at sparing healthy tissue while maintaining the prescribed dose to the target when shifts occur.

Despite the fact that we use larger voxels than the current state of the art in

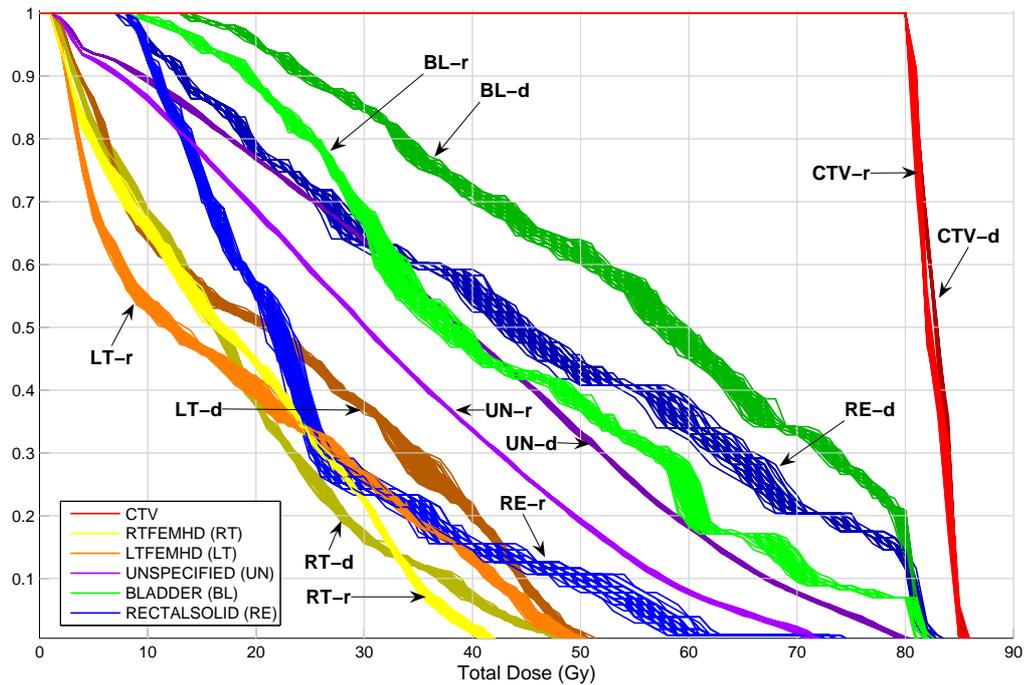


Figure 5. Dose-Volume Histogram cloud for solution to deterministic model ('-d') and robust model ('-r'), for 100 simulated treatments.

planning systems, we believe that the plans obtained using a robust formulation should be preferred to those obtained using smaller voxels but ignoring uncertainty, at least when the uncertainty is moderate to significant. We are currently investigating methods to reduce computational requirements for our formulation.

Other areas of future research include investigating appropriate probability distributions for voxel locations and alternative performance measures that convey the underlying uncertainty.

Acknowledgments

This work was partially supported by National Science Foundation Grant #DMI 0400287, and the Cornell Computational Optimization Project (CCOP). We thank the referees for helpful suggestions that improved the presentation.

References

- ADAC 2002 *P³IMRT User Guide. Inverse Planning and IMRT for Pinnacle³* (Milpitas: ADAC Laboratories)
- Ben-Tal A and Nemirovski A 2001 *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications* (Philadelphia: SIAM-MPS Series in Optimization)

- Bentel G C 1999 *Patient Positioning and Immobilization in Radiation Oncology* (New York: McGraw-Hill)
- Bortfeld T, Jokivarsi K, Goitein M, Kung J and Jiang S B 2002 Effects of intra-fraction motion on IMRT dose delivery: Statistical analysis and simulation *Phys. Med. Biol.* **47** 2203–20
- Cotrutz C and Xing L 2002 Using voxel-dependent importance factors for interactive DVH-based dose optimization *Phys. Med. Biol.* **47** 1659–69
- Craig T, Battista J and Van Dyk J 2003a Limitations of a convolution method for modeling geometric uncertainties in radiation therapy. I. The effect of shift invariance *Med. Phys.* **30** 2001–11
- 2003b Limitations of a convolution method for modeling geometric uncertainties in radiation therapy. II. The effect of a finite number of fractions *Med. Phys.* **30** 2012–20
- Ferris M C, Lim J and Shepard D M 2003 Radiosurgery treatment planning via nonlinear programming *Ann. Oper. Res.* **119** 247–60
- Ferris M C, Einarsson R, Jiang Z and Shepard D 2004 Sampling issues for optimization in radiotherapy Intensity Modulated Radiation Therapy Collaborative Working Group 2001 Intensity-modulated radiotherapy: current status and issues of interest *Int. J. Radiat. Oncol.* **51** 880–914
- International Commission on Radiation Units and Measurements 1993 *ICRU Report 50: Prescribing, Recording and Reporting Photon Beam Therapy* (Bethesda: International Commission on Radiation Units and Measurements)
- 1999 *ICRU Report 62: Prescribing, Recording and Reporting Photon Beam Therapy (Supplement to ICRU Report 50)* (Bethesda: International Commission on Radiation Units and Measurements)
- Langen K M and Jones D T L 2001 Organ motion and its management *Int. J. Radiat. Oncol.* **50** 265–78
- Langer M, Lee E K, Deasy J O, Rardin R L and Deye J A 2003 Operations research applied to radiotherapy, an NCI-NSF-sponsored workshop - February 7-9, 2002 *Int. J. Radiat. Oncol.* **57** 762-8
- Li J G and Xing L 2000 Inverse planning incorporating organ motion *Med. Phys.* **27** 1573–8
- Preciado-Walters F 2003 *Optimal External Radiation Therapy Planning for Cancer: A Mixed Integer Approach* (Ph.D. Thesis, Purdue University)
- Romeijn H E, Ahuja R K, Dempsey J F, Kumar A and Li J G 2003 A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning *Phys. Med. Biol.* **48** 3521–42
- 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
- Sturm J F 1999 Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones *Optim. Method Softw.* **11-12** 625–53
- 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
- 2005 Incorporating organ movements in inverse planning: assessing dose uncertainties by Bayesian inference *Phys. Med. Biol.* **50** 121–39
- van Herk M, Remeijer P and Lebesque J V 2002 Inclusion of geometric uncertainties in treatment plan evaluation *Int. J. Radiat. Oncol.* **52** 1407–22
- Wong J R, Grimm L, Uematsu M, Oren R, Cheng C W, Merrick W and Schiff P 2005 Image-guided radiotherapy for prostate cancer by CT-linear accelerator combination: prostate movements and dosimetric considerations *Int. J. Radiat. Oncol.* **61** 561–9