ROBUST INTENSITY MODULATED RADIATION THERAPY TREATMENT PLANNING

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Millie Chu

January 2006

ROBUST INTENSITY MODULATED RADIATION THERAPY TREATMENT PLANNING

Millie Chu, Ph.D.

Cornell University 2006

The recent development of Intensity Modulated Radiation Therapy (IMRT) provides a more sophisticated method for delivering radiation therapy for cancer treatment. Using IMRT, the dose distribution can be tailored to match the tumour's shape and position, thereby avoiding damage to healthy tissue to a greater extent than previously possible. Traditional planning for IMRT treatment assumes the target area remains in a fixed location throughout the weeks of treatment. However, many studies have shown that because of organ motion, inconsistencies in patient positioning, etc., the tumour location does not remain stationary. We present a probabilistic model for IMRT treatment planning and show that it is identical to using robust optimization techniques, under certain assumptions. Since this formulation is computationally more intensive, we propose two methods to reduce computational requirements. For a sample prostate case, our model is tractable and, compared to traditional methods, has the potential to find treatment plans that are more adept at sparing healthy tissue while maintaining the prescribed dose of radiation to the target.

BIOGRAPHICAL SKETCH

Millie Chu has been alive since August 29, 1975. Raised in the sunny suburbs of L.A., she enjoys cooking, musicals, and ballet. In 1997, she received her B.A. in Math from Wellesley College, where she was head of the yearbook staff. She worked as an actuarial analyst for three years before she decided to switch careers. In Fall 2000, she entered Cornell University to get her Ph.D in Operations Research, the "science of better". Upon completion of her degree, she will catch up on T.V.

ACKNOWLEDGEMENTS

Many people have helped me over the course of my research. I thank everyone!

I especially thank my advisor, Shane Henderson. His enthusiasm and support have made this work possible and my life as a graduate student enjoyable.

I wish to thank my other committee members, Stephen Vavasis and David Shmoys for their interest and helpful comments. I thank David for his encouragement during my early years as a graduate student.

This work could not have been done without the help of Michael B. Sharpe, Senior Physicist and Assistant Professor in the Department of Radiation Therapy Physics at Princess Margaret Hospital in Toronto. His guidance and mini-lessons in medical physics have been invaluable.

I thank Yuriy Zinchenko for all his helpful contributions and technical support. I wish him and Deniz a happy and healthy future together.

Also, this work was partially supported by National Science Foundation Grant #DMI 0400287, and the Cornell Computational Optimization Project (CCOP).

Thanks to all my friends, in particular: Anke, Davina, Frans, Ghosh, Iori, Mohan, Ranjith, Rif, Sam, Tim, Tuncay, and Vardges. Thanks to everyone at Ithaca Ballet for keeping me in shape. Thanks to Rachael Ray and Alton Brown for teaching me how to cook. And thanks to Stella's for the great coffee.

Finally, I thank my family - yes, I am finally done!

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

Chapter 1

Introduction

The American Cancer Society expects that nearly 1.4 million Americans will be diagnosed with some form of cancer in 2005 [2]. About half of these patients will undergo radiation therapy, some in conjunction with chemotherapy or surgery.

Radiation therapy uses ionizing radiation to kill cancerous cells and to shrink tumours. More than one hundred years of experience and technological innovation have firmly established radiation therapy as a proven means of treating many cancerous and non-cancerous diseases [7]. Medical linear accelerators are widely used to generate ionizing radiation, mainly in the form of high-energy x-rays that permit deep tissue penetration with very little energy deposited in the skin. Because radiation does not discriminate between normal and healthy tissue, strategic delivery is required to limit the adverse effects of radiation on normal tissues.

1.1 Treatment Planning Problem

Treatment planning is a process that uses Computed Tomography (CT) images to visualize the location and extent of the tumour and the surrounding anatomic structures. This information is used to design shielding that confines high-dose radiation, and to design plans that uniformly irradiate targeted tissues, while sparing normal tissues to reduce the toxic effects of treatment. Computer-aided visualization of anatomy and simulations of the effects of treatment are essential for maximizing the dose to targets without exposing normal tissues to potential injury.

Fractionation limits normal tissue damage by extending the treatment time over several weeks, and dividing up the total dose into several smaller doses, which

are administered daily. Fractionation gives normal tissues time to recover, and allows radioresistant tumour cells to progress to a more radiosensitive phase of their growth [25].

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). In addition to the benefit of sparing more normal tissue, IMRT provides the capacity to respond therapeutically to new information emerging from molecular biology and medical imaging.

IMRT combines treatment planning and computer-controlled treatment delivery to precisely shape the distribution of dose to encompass targeted tissues, while avoiding adjacent normal tissues [10, 24, 38]. Detailed reviews of IMRT principles and clinical guidelines are described elsewhere [19, 23, 64]. Briefly, dose distributions are formed by the strategic superposition of a series of beams intersecting the target from many directions. Several beam apertures are formed for each angle of incidence, using a multi-leaf collimator. During the optimization process, these apertures are idealized as narrow individual *beamlets*, or pencil beams. The radiation intensity within each beamlet is altered, or modulated, to achieve the prescribed doses for tissue elements it traverses. In the leaf-sequencing process, beamlets with common intensities are then combined to generate efficient apertures for treatment delivery.

Due to cost and time restrictions, the same treatment (beam angles and beamlet intensities) is administered at each treatment fraction, and the number of beam angles must be small, usually no more than 12. Experienced treatment planners and physicians are often able to pre-select effective beam angles, and in this work

we assume such a selection has been made. Studies have shown that beam angle optimization only leads to marginal improvement for the dose distribution for prostate cases, but can have a more significant effect on other cases [58, 51]. The pre-selection of beam angles reduces the treatment planning problem to that of finding beamlet intensities which will most likely ensure local tumour control (killing all viable tumour cells), yet refrain from adverse effects on normal tissues.

1.2 Uncertainties

The development of IMRT has provided tools to shape radiation dose distributions with high precision. It is now possible to treat irregularly shaped target volumes aggressively with therapeutic doses, while sparing the surrounding healthy tissues. However, high precision treatment delivery requires a high precision of targeting. Motion and deformation of the patient or of the inner organs during, or between, treatment fractions, as well as any positioning uncertainties, could have a detrimental effect on local control of the disease, or on treatment toxicity.

Immobilization devices are commonly used to help patients maintain consistent positioning, both during treatment and from fraction to fraction [6]. But even with the aid of immobilization devices, small discrepancies are inevitable because of, for example, periodic breathing and cardiac motion within a treatment, 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 regions as a way to ensure that they receive doses within the prescribed bounds. The safety margins are

designed to accommodate both systematic and random uncertainties that might be introduced following the acquisition of the initial planning data. A number of excellent review articles summarize the current state of motion management in radiotherapy [30, 31]. Following the nomenclature and guidelines of the International Commission on Radiation Units and Measurements (ICRU) [46, 47], 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 (see Figure 1.1). 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. Such regions are called planning risk volumes (PRVs).

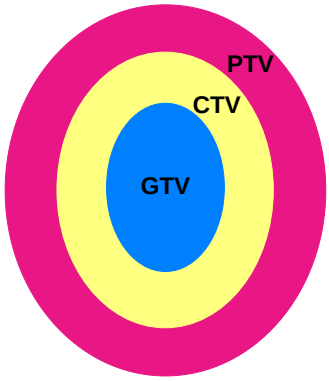

Figure 1.1: Safety margin recommended by ICRU.

However, using this approach, "good" solutions tend to be treatment plans where there is a sharp change in the dose distribution at the PTV boundary, which could be problematic for several reasons. First, if the safety margin encom-

passes any healthy tissue, it will be aggressively irradiated because it is part of the PTV. Second, during treatment, if some healthy tissue moves into the high dose area, severe side effects may occur. Moreover, it is possible to generate regions where PRVs and PTVs overlap, creating conflicting clinical objectives that must be resolved. And finally, if some malignant cells move into the surrounding low dose area, then the tumour will be underdosed and could persist. So a plan, that at first sight looks acceptable, may in fact be undesirable.

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. The treatment parameters can then be adjusted manually, or with the assistance of an algorithm to compensate for the changes in dose distribution caused by motion. The understanding and modeling of motion effects on dose distributions is an active area of inquiry. Studies have been done to determine these effects for a variety of treatment methods [8, 12, 29, 68], while others have examined a convolution/smearing approximation of the dose distribution [4, 13, 14, 41]. Biological studies have investigated these effects for structures with time-dependent geometry [48, 49], the implications of organ deformation [9], the effects of a blurred dose matrix on tumour control probability [63], and a model for tumour growth during treatment [65].

Optimization approaches to address uncertainties include incorporating the expectation of dose delivered into the objective function, using an estimated probability distribution for structure locations [37]. Physics-based models aim to optimize fluence profiles and minimize the variance of dose delivered [60, 61].

With the advent of online imaging technologies, it is now possible to routinely

acquire CT scans in the treatment setting [28, 42, 57, 66]. Using these images to locate the target can reduce patient setup errors, but patient motion during treatment is inevitable. Currently, because of time and personnel restrictions, these images are only used for correcting patient positioning, and not for online re-optimization of treatment plans.

Adaptive Radiation Therapy [40, 43, 52, 67] incorporates multiple patient images during treatment (e.g. daily CT scans) into a feedback algorithm to tune the sequence of treatment plans. Similar adaptive techniques rely on knowledge of actual amounts of dose absorbed during each treatment fraction to adjust subsequent fractions [22].

Our robust model, presented in [11], aims to create plans which remain acceptable over a broad range of likely circumstances. Instead of representing the target region as a PTV, our model discards the safety margin and 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* as sample points on a grid where we measure amounts of dose absorbed. Our voxels are selected from a uniform 3-dimensional grid, and for simplicity in notation, we refer to a voxel from a 0.8 cm $\times$ 0.8 cm $\times$ 0.8 cm grid as a "0.8 cm voxel," for example.

Dose is energy per unit mass, measured in Gray (abbreviated Gy, where 1 Gy = 1 joule per kilogram). The actual amounts of dose absorbed are unknown, as they depend on the random scattering of electrons and the resulting cellular damage caused by these displaced electrons. Commercial treatment planning systems employ Monte Carlo techniques to estimate absorbed doses specific to each patient, but in this work we use an approximate (deterministic) dose calculation from a

matrix of doses delivered to a water phantom and only address uncertainty in voxel location. We suspect that using a more accurate dose calculation would not greatly affect the relative performance of our robust model compared to a deterministic model. Further, we make the conventional assumption that the cumulative dose delivered to a voxel over all fractions is linearly additive.

The structure contours from the patient's initial planning CT scan determine the set of voxels for each structure and the "original" location of each voxel, which may or may not be its average location. The planning CT scan gives only one snapshot of structure contours, and therefore may be quite biased. However, additional scans are typically not available, and without further information, it seems reasonable to use original voxel locations as a starting point for potential movement (see Section 3.1 for details on scenario selection).

Since we are assuming beam angles have been pre-selected by an experienced planner, our solutions will be an assignment of intensities (or weights) for each beamlet, which would then be delivered to the patient at each treatment fraction.

1.3 Small Example

To illustrate the potential downfall of using a safety margin to deal with uncertainty, consider the stylized example shown in Figure 1.2. Voxels in the healthy structure are on the perimeter, three squares wide. The CTV voxels are the center 4×4 squares. The PTV margin is a one square border around the CTV, making the PTV a 6×6 square. Note the overlap between the PTV margin and the healthy structure.

Figure 1.2 shows the optimal plan (the "deterministic solution") obtained by using a simple linear program that penalizes overdosing the healthy structure and

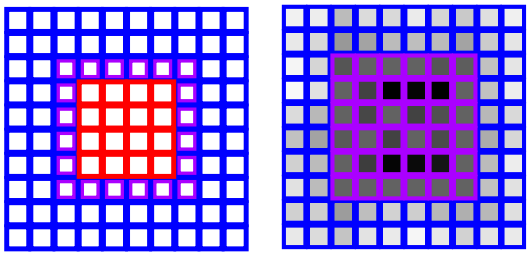

Figure 1.2: (Left) Voxel setup for small example. (Right) Treatment plan of deterministic model's optimal solution.

penalizes underdosing the PTV, where darker gray indicates higher dose delivered and lighter indicates lower.

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 one unit in each direction (right, left, up, down), with each scenario equally likely. Figure 1.3 shows how the deterministic plan performs when all voxels shift together for the five scenarios. 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.

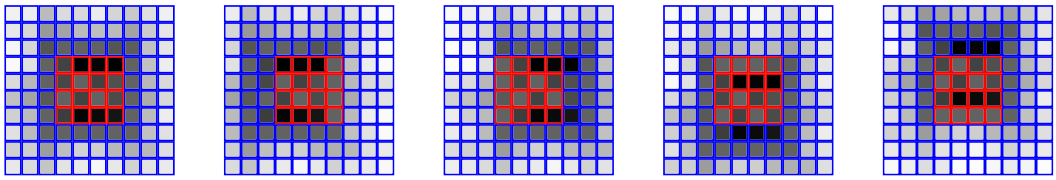

Figure 1.3: Small example: solution of deterministic model, in five potential scenarios.

As an alternative to the deterministic model, we incorporate the scenario information into a robust model (see Chapter 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 1.4, with doses on the same grayscale as in Figure 1.3. 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.

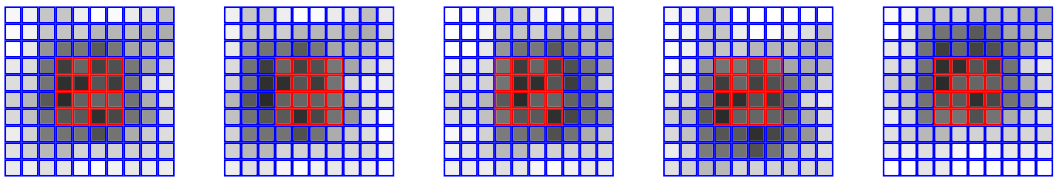

Figure 1.4: Small example: solution of robust model, in five potential scenarios.

1.4 Contributions

The contributions of this dissertation are:

- to demonstrate that current methods find treatment plans that are optimal for static cases but can perform poorly when motion occurs,
- to present a robust optimization model for IMRT treatment planning,
- to show that this robust formulation is computationally tractable and can find treatment plans that perform better than deterministic methods when motion occurs, and
- to emphasize that performance measures used by oncologists to evaluate the quality of a treatment plan need to address uncertainty and convey potential outcomes.

1.5 Outline of Thesis

The remainder of this thesis is organized as follows. In Chapter 2, we develop the formulation for robust IMRT treatment planning, from both a probabilistic approach (Section 2.1) and a robust optimization approach (Section 2.2). Computational results for a prostate case are presented in Chapter 3, with details of scenario selection in Section 3.1. Efforts to reduce computational requirements are described in Chapter 4, and concluding remarks are in Chapter 5. A summary of the computational implementation appears in the Appendix, and attached is a CD containing our MATLAB functions.

Chapter 2

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. 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 plan be of high quality so long as the problem parameters fall within some "reasonable" set. In Sections 2.1 and 2.2 we describe these interpretations in turn, and under certain assumptions show that they lead to the same form of constraint. A penalty formulation is presented in Section 2.3, and dose-volume constraints are discussed in Section 2.4. The full formulation of the problem appears in Section 2.5.

2.1 Probabilistic Approach

Let $D_i(\boldsymbol{x})$ denote the total dose delivered to voxel i over the course of all N fractions, as a function of the beamlet intensity vector $\boldsymbol{x}$. We view $D_i(\boldsymbol{x})$ 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, $D_i(\boldsymbol{x})$ is the sum of N random variables $D_{i\ell}(\boldsymbol{x})$, $\ell = 1, \ldots, N$, where $D_{i\ell}(\boldsymbol{x})$ is the dose delivered to voxel i on the ℓth fraction with beamlet intensity vector $\boldsymbol{x}$. In great generality, a sum of random variables is approximately normally distributed, as follows from a central limit theorem. We

therefore assume that $D_i(\boldsymbol{x})$ is normally distributed. Denote the mean and variance of the total dose $D_i(\boldsymbol{x})$ by μ and σ^2 respectively. These parameters depend on i and $\boldsymbol{x}$, but for now we suppress that dependence in our notation.

Suppose that voxel i comes from some healthy structure $\mathcal{H}_k$, 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 $\mathrm{P}(D_i(\boldsymbol{x}) > m_k) \leq \delta$, which is equivalent to saying that

$$\mathrm{P}\left(\frac{D_i(\boldsymbol{x}) - \mu}{\sigma} > \frac{m_k - \mu}{\sigma}\right) \leq \delta. \tag{2.1}$$

But $(D_i(\boldsymbol{x}) - \mu)/\sigma$ is normally distributed with mean 0 and variance 1, so (2.1) is equivalent to requiring that

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

where $z_{1-\delta}$ is chosen so that $\mathrm{P}(Z > z_{1-\delta}) = \delta$, 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

$$\sigma \leq \frac{\mu - m_{\mathrm{Tmin}}}{z_{1-\delta}}. \tag{2.2}$$

Notice that $z_{1-\delta} > 0$ and we would expect that $m_{\mathrm{Tmin}} < \mu$ in any reasonable solution, so that the right-hand side of (2.2) 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 $D_i(\boldsymbol{x})$, and are therefore functions of the beamlet intensity vector $\boldsymbol{x}$. 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 $\mu = N\mathrm{E}D_{i1}(\boldsymbol{x})$ and $\sigma^2 = N\mathrm{Var}D_{i1}(\boldsymbol{x})$, so the problem reduces to computing the mean and variance of the dose from a single fraction. Let $\boldsymbol{Y}_i$ 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 $D_{i1}(\boldsymbol{x}) = \boldsymbol{Y}_i^\top \boldsymbol{x}$, where $\boldsymbol{v}^\top$ denotes the transpose of the vector $\boldsymbol{v}$. The variance can be rewritten as

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

where $\mathrm{Cov}(\boldsymbol{Y}_i)$ 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 $s_1, \ldots, s_n$ can occur with probabilities $p_1, \ldots, p_n$ respectively. Let $\boldsymbol{a}_{ij}$ 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

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

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

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

and

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

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

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

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

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

Let $\boldsymbol{e}$ 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

$$\begin{aligned} \mathrm{Var}D_{i1}(\boldsymbol{x}) &= [A_i\boldsymbol{x} - \boldsymbol{e}\left(\boldsymbol{p}^{\top}A_i\boldsymbol{x}\right)]^{\top}P[A_i\boldsymbol{x} - \boldsymbol{e}\left(\boldsymbol{p}^{\top}A_i\boldsymbol{x}\right)] \\ &= [\left(I - \boldsymbol{e}\boldsymbol{p}^{\top}\right)A_i\boldsymbol{x}]^{\top}P[\left(I - \boldsymbol{e}\boldsymbol{p}^{\top}\right)A_i\boldsymbol{x}] \\ &= \|RA_i\boldsymbol{x}\|^2, \end{aligned}$$

where $\|v\|^2 = v_1^2 + \cdots + v_n^2$ denotes the square of the Euclidean norm, $R = P^{1/2}(I - \boldsymbol{e}\boldsymbol{p}^\top)$, 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 (2.1) can be written as

$$\|RA_i\boldsymbol{x}\| \leq \frac{m_k - N\boldsymbol{p}^\top A_i\boldsymbol{x}}{z_{1-\delta}\sqrt{N}}. \tag{2.3}$$

The constraint (2.2) can be expressed in almost exactly the same way.

There is one constraint of the form (2.3) 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 Robust Optimization Approach

Consider the SOCP problem [5]:

$$\begin{aligned} &\min \boldsymbol{c}^\top \boldsymbol{x} \\ &A\boldsymbol{x} = \boldsymbol{b} \\ &\boldsymbol{x} \in K_2^{\prod} \end{aligned}$$

where $\boldsymbol{c}, \boldsymbol{x} \in \Re^n$, $A \in \Re^{m\times n}$, $\boldsymbol{b} \in \Re^m$ and $K_2^{\prod}$ is a direct product of second order cones. Notice that SOCP almost looks like LP, except the nonnegativity constraints $\boldsymbol{x} \geq 0$ have been replaced by $\boldsymbol{x} \in K_2^{\prod}$. More precisely,

$$\begin{aligned} K_2^{\prod} \;=\; &\{\boldsymbol{x} \in \Re^n : (x_1, x_2, \ldots, x_{k_1}) \in K_1, (x_{k_1+1}, \ldots, x_{k_2}) \in K_2, \\ &\ldots, (x_{k_l+1}, \ldots, x_n) \in K_{l+1}\} \end{aligned}$$

where $1 \leq k_1 < k_2 < \cdots < k_l < n$ and for $i = 1, \ldots, l+1$, K_i is a second order cone (SOC). A cone $K \subseteq \Re^{k+1}$ is a SOC if $K = \{(\boldsymbol{x}, t) \in \Re^k \times \Re : \|\boldsymbol{x}\| \leq t\}$. 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. 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.

Consider the LP:

$$\min \boldsymbol{c}^\top \boldsymbol{x}$$

$$A\boldsymbol{x} \geq \boldsymbol{b}$$

where $\boldsymbol{c}, \boldsymbol{x} \in \Re^n, A \in \Re^{m \times n}$, and $\boldsymbol{b} \in \Re^m$. (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, developed by Ben-Tal and Nemirovski [5], 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 $A \in \mathcal{U}$, where $\mathcal{U}$ is usually termed the uncertainty set. Robust LP requires a solution $\boldsymbol{x}$ to be feasible for *any* choice of $A \in \mathcal{U}$, i.e.,

$$\min \boldsymbol{c}^\top \boldsymbol{x}$$

$$A\boldsymbol{x} \geq \boldsymbol{b} \qquad \forall A \in \mathcal{U}.$$

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

An ellipsoidal uncertainty set $\mathcal{U}$ corresponds to a direct product of ellipsoids E_i for each constraint $\boldsymbol{a}_i^\top \boldsymbol{x} \geq b_i$ in the set of constraints $A\boldsymbol{x} \geq \boldsymbol{b}$, where $\boldsymbol{a}_i^\top$ is the

ith row of the matrix A, for $i = 1, \ldots, m$. We want to solve:

$$\begin{aligned} &\min \boldsymbol{c}^\top \boldsymbol{x} \\ &\boldsymbol{a}_i^\top \boldsymbol{x} \geq b_i \qquad \forall \boldsymbol{a}_i \in E_i, \text{ for } i = 1, \ldots, m. \end{aligned}$$

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

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

for $i = 1, \ldots, m$. Hence, a Robust LP can be written as:

$$\begin{aligned} &\min \boldsymbol{c}^\top \boldsymbol{x} \\ &\left\|W_i^\top \boldsymbol{x}\right\| \leq {\boldsymbol{a}_i^*}^\top \boldsymbol{x} - b_i, \qquad \text{for } i = 1, \ldots, m, \end{aligned}$$

which is an instance of a 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 in IMRT treatment planning.

We still compute the dose to voxel i on fraction ℓ as a linear combination of the beamlet intensities $\boldsymbol{x}$, $D_{i\ell}(\boldsymbol{x}) = \boldsymbol{a}_i^\top \boldsymbol{x}$, but now we require $\boldsymbol{a}_i \in E_i$, where E_i is the ellipse $\{\boldsymbol{a}_i^* + W_i \boldsymbol{u}_i : \boldsymbol{u}_i^\top \boldsymbol{u}_i \leq 1\}$. 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. Naturally, we select ${\boldsymbol{a}_i^*}^\top = \boldsymbol{p}^\top A_i$, so that the ellipse E_i is centered at the expected dose coefficients for voxel i. The question remains on how to choose W_i.

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

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

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

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

Rearranging this inequality gives:

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

which is the same as (2.3).

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 $\sum_{\ell=1}^{N} D_{iS(\ell)}(\boldsymbol{x})$ 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, but we do not do so here because the ellipsoidal form seems sufficient for our purposes.

2.3 Penalty 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 [32, 33, 56]. Some proposed models have been formulated as linear programs [27], mixed-integer programs [35, 36, 50], nonlinear programs [21], and multicriteria models [18, 17, 26]. With no definitive model on hand, we will minimize a linear penalty objective function which is similar, in principle, to the model used in the ADAC Pinnacle^3 treatment planning system [1].

Our decision variables are the nonnegative beamlet intensities, $\boldsymbol{x}$. Let B be the number of beamlets available (from the pre-selected beam angles), so that $\boldsymbol{x}$ is a $B \times 1$ vector.

Let $\mathcal{H}_k$, for $k = 1, \ldots, h$, be the set of voxels for healthy structure k. Traditionally, treatment planners impose a maximum cumulative dose of m_k to every voxel in structure $\mathcal{H}_k$. In our formulation, we require the cumulative dose to be at most m_k with high probability. Let u_k be an intermediate variable replacing parameter m_k in (2.3), so for any voxel $i \in \mathcal{H}_k$ we have

$$\|RA_i\boldsymbol{x}\| \leq \frac{u_k - N\boldsymbol{p}^\top A_i\boldsymbol{x}}{z_{1-\delta}\sqrt{N}}, \tag{2.4}$$

where u_k can be interpreted as the maximum dose to any voxel $i \in \mathcal{H}_k$ with probability $1 - \delta$ when $\boldsymbol{x}$ is delivered. If $u_k > m_k$, then we penalize the excess dose, $u_k - m_k$, in the objective function, by some weight $w_k > 0$. Otherwise, if $u_k \leq m_k$, there is no penalty. The excess is measured by imposing the following linear constraints, for each healthy structure $k = 1, \ldots, h$:

$$u_k - m_k = r_k^+ - r_k^-, \tag{2.5}$$

and

$$r_k^+, r_k^- \geq 0. \tag{2.6}$$

The term $w_k r_k^+$ appears in the objective function, where parameter $w_k > 0$ is a penalty weight chosen by the planner. Minimizing the objective will cause $r_k^+ = u_k - m_k > 0$ if $u_k > m_k$, and $r_k^+ = 0$ when $u_k \leq m_k$. Note that we have chosen to make the value u_k structure-specific rather than voxel-specific, since it seems sufficient for our purposes.

For the target region, the (practically unobtainable) goal is for every CTV voxel to receive the exact same dose t. A homogeneous high dose to the entire CTV is essential for treatment since cold spots may allow some tumour cells to survive, while hot spots may cause unnecessary damage. Therefore, we impose both an upper bound m_{Tmax} and a lower bound m_{Tmin} to the cumulative dose delivered to any CTV voxel, where these bounds are close to t (typically within 5% to 10%).

The maximum CTV dose constraints have the same form as for the healthy structures. Let parameter m_{Tmax} be the desired maximum dose to the CTV, chosen by the planner, let $u_{\text{Tmax}}, r_{\text{Tmax}}^+, r_{\text{Tmax}}^-$ be three intermediate variables, and for any voxel $i \in \text{CTV}$, we restrict

$$\|RA_i\boldsymbol{x}\| \leq \frac{u_{\text{Tmax}} - N\boldsymbol{p}^\top A_i\boldsymbol{x}}{z_{1-\delta}\sqrt{N}}, \tag{2.7}$$

$$u_{\text{Tmax}} - m_{\text{Tmax}} = r_{\text{Tmax}}^+ - r_{\text{Tmax}}^-, \tag{2.8}$$

and

$$r_{\text{Tmax}}^+, r_{\text{Tmax}}^- \geq 0. \tag{2.9}$$

where $w_{\text{Tmax}} r_{\text{Tmax}}^+$ appears in the objective function, with parameter $w_{\text{Tmax}} > 0$.

The minimum CTV dose constraints are similar. Let $m_{\text{Tmin}}, u_{\text{Tmin}}, r_{\text{Tmin}}^+, r_{\text{Tmin}}^-$ be the minimum dose parameter and intermediate variables, respectively. For any

voxel $i \in$ CTV we also have:

$$\|RA_i\boldsymbol{x}\| \leq \frac{N\boldsymbol{p}^\top A_i\boldsymbol{x} - u_{\mathrm{Tmin}}}{z_{1-\delta}\sqrt{N}}, \tag{2.10}$$

$$m_{\mathrm{Tmin}} - u_{\mathrm{Tmin}} = r^+_{\mathrm{Tmin}} - r^-_{\mathrm{Tmin}}, \tag{2.11}$$

and

$$r^+_{\mathrm{Tmin}}, r^-_{\mathrm{Tmin}} \geq 0, \tag{2.12}$$

where $w_{\mathrm{Tmin}} r^+_{\mathrm{Tmin}}$ is in the objective function, with parameter $w_{\mathrm{Tmin}} > 0$.

Further, a minimum dose per fraction m_T is sometimes required for every voxel in the target to ensure that some lethal damage is achieved at every treatment fraction. Since the dose per fraction is represented by any one of our n scenarios, we set a lower bound on the dose per scenario. Let $u_{\mathrm{T}j}, r^+_{\mathrm{T}j}, r^-_{\mathrm{T}j}$ for scenario $j = 1, \ldots, n$ be intermediate variables, and for any voxel $i \in$ CTV, $j = 1, \ldots, n$, we impose the linear constraints

$$\boldsymbol{a}^\top_{i,j}\boldsymbol{x} \geq u_{\mathrm{T}j}, \tag{2.13}$$

$$m_{\mathrm{T}} - u_{\mathrm{T}j} = r^+_{\mathrm{T}j} - r^-_{\mathrm{T}j}, \tag{2.14}$$

and

$$r^+_{\mathrm{T}j}, r^-_{\mathrm{T}j} \geq 0, \tag{2.15}$$

where $w_{\mathrm{T}j} r^+_{\mathrm{T}j}$ appears in the objective function, with parameter $w_{\mathrm{T}j} > 0$. Note $u_{\mathrm{T}j}$ is the minimum (deterministic) dose delivered by $\boldsymbol{x}$ to the CTV in scenario j, and a penalty is incurred if this minimum dose is below m_{T}. Like the maximum and minimum dose constraints, this penalty is not incurred per voxel but rather for the entire CTV, for each scenario.

2.4 Dose-Volume Constraints

Solving the formulation thus far 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 $\mathcal{H}_k$ 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 [20, 21]) imposes the following constraint on the sum of the "excess" doses (expected doses above the threshold d_k) for structure $\mathcal{H}_k$:

$$\sum_{i \in \mathcal{H}_k} \left(N \boldsymbol{p}^\top A_i \boldsymbol{x} - d_k \right)_+ \leq g_k, \tag{2.16}$$

where $(\cdot)_+$ means positive part, and g_k is a parameter chosen by the planner. A standard linear programming trick (see below) 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.

Let $\overline{h} \subset \{1, \ldots, h\}$ be the indices of healthy structures that have a prescribed dose-volume restriction. For simplicity, we show the formulation where each structure in $\overline{h}$ has just one dose-volume restriction, but it is straightforward to extend

this formulation for multiple dose-volume restrictions on a structure. For each $k \in \overline{h}$, our DV constraints are as follows:

$$N\boldsymbol{p}^\top A_i \boldsymbol{x} - d_k = v_i^+ - v_i^-, \qquad \forall i \in \mathcal{H}_k, \tag{2.17}$$

$$v_i^+, v_i^- \geq 0, \qquad \forall i \in \mathcal{H}_k, \tag{2.18}$$

$$\sum_{i \in \mathcal{H}_k} v_i^+ - g_k = q_k^+ - q_k^-, \tag{2.19}$$

and

$$q_k^+, q_k^- \geq 0. \tag{2.20}$$

Continuing with the penalty objective function, let $\overline{w}_k > 0$ be the penalty weight for exceeding the upper bound g_k on excess doses in structure $\mathcal{H}_k$, and so the term $\overline{w}_k q_k^+$ appears in the objective.

One possible value for the upper bound g_k is $v_k|\mathcal{H}_k|(m_k - d_k)$, which is the total amount of "excess" 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 (see [50, 54]), and further research may include evaluating such alternative approaches.

2.5 Full Formulation

Minimize the sum of weighted penalties:

$$w_{\mathrm{Tmin}} r_{\mathrm{Tmin}}^+ + w_{\mathrm{Tmax}} r_{\mathrm{Tmax}}^+ + \sum_{j=1}^{n} w_{\mathrm{T}j} r_{\mathrm{T}j}^+ + \sum_{k=1}^{h} w_k r_k^+ + \sum_{k \in \overline{h}} \overline{w}_k q_k^+,$$

subject to:

maximum dose to healthy voxels, with high probability, (2.4) to (2.6):

$$\begin{aligned} \|RA_i\boldsymbol{x}\| &\leq \frac{u_k - N\boldsymbol{p}^\top A_i\boldsymbol{x}}{z_{1-\delta}\sqrt{N}}, && \forall i \in \mathcal{H}_k, \text{ for } k = 1, \ldots, h, \\ u_k - m_k &= r_k^+ - r_k^-, && \text{for } k = 1, \ldots, h, \\ r_k^+, r_k^- &\geq 0, && \text{for } k = 1, \ldots, h, \end{aligned}$$

maximum total dose to target voxels, with high probability, (2.7) to (2.9):

$$\begin{aligned} \|RA_i\boldsymbol{x}\| &\leq \frac{u_{\mathrm{Tmax}} - N\boldsymbol{p}^\top A_i\boldsymbol{x}}{z_{1-\delta}\sqrt{N}}, && \forall i \in \mathrm{CTV}, \\ u_{\mathrm{Tmax}} - m_{\mathrm{Tmax}} &= r_{\mathrm{Tmax}}^+ - r_{\mathrm{Tmax}}^-, \\ r_{\mathrm{Tmax}}^+, r_{\mathrm{Tmax}}^- &\geq 0, \end{aligned}$$

minimum total dose to target voxels, with high probability, (2.10) to (2.12):

$$\begin{aligned} \|RA_i\boldsymbol{x}\| &\leq \frac{N\boldsymbol{p}^\top A_i\boldsymbol{x} - u_{\mathrm{Tmin}}}{z_{1-\delta}\sqrt{N}}, && \forall i \in \mathrm{CTV}, \\ m_{\mathrm{Tmin}} - u_{\mathrm{Tmin}} &= r_{\mathrm{Tmin}}^+ - r_{\mathrm{Tmin}}^-, \\ r_{\mathrm{Tmin}}^+, r_{\mathrm{Tmin}}^- &\geq 0, \end{aligned}$$

minimum dose per scenario to target voxels, (2.13) to (2.15):

$$\begin{aligned} \boldsymbol{a}_{i,j}^\top\boldsymbol{x} &\geq u_{\mathrm{T}j}, && \forall i \in \mathrm{CTV}, \text{ for } j = 1, \ldots, n, \\ m_{\mathrm{T}} - u_{\mathrm{T}j} &= r_{\mathrm{T}j}^+ - r_{\mathrm{T}j}^-, && \text{for } j = 1, \ldots, n, \\ r_{\mathrm{T}j}^+, r_{\mathrm{T}j}^- &\geq 0, && \text{for } j = 1, \ldots, n, \end{aligned}$$

approximate dose-volume constraints, (2.17) to (2.20):

$$\begin{aligned} N\boldsymbol{p}^\top A_i\boldsymbol{x} - d_k &= v_i^+ - v_i^-, && \forall i \in \mathcal{H}_k, \forall k \in \overline{h}, \\ v_i^+, v_i^- &\geq 0, && \forall i \in \mathcal{H}_k, \forall k \in \overline{h}, \\ \sum_{i \in \mathcal{H}_k} v_i^+ - g_k &= q_k^+ - q_k^-, && \forall k \in \overline{h}, \\ q_k^+, q_k^- &\geq 0, && \forall k \in \overline{h}, \end{aligned}$$

and nonnegative beamlet intensities:

$$\boldsymbol{x} \geq \mathbf{0}.$$

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.

Chapter 3

Patient Case

In this chapter we apply our robust formulation to a test case, a patient with prostate cancer, and compare the resulting plan to one obtained from a deterministic formulation. For prostate cancer, the target region is the prostate gland, entirely surrounded by healthy tissue and adjacent to two critical structures, the bladder and rectal solid. Figure 3.1 shows a sample CT scan, with structures delineated, from a series of 81 CT scans taken at 3 mm intervals along the patient's superior-inferior (head-toe) axis. Note the overlap between the PTV and rectal solid, and between the PTV/CTV and bladder. Figure 3.2 shows the 3-d structure contours, excluding the Unspecified region, constructed from the series of CT scans. The Unspecified region does not overlap with any other structures. The PTV margin is 1 cm around the CTV, and for the deterministic setup, we choose not to use any safety margins around healthy structures. Table 3.1 shows the clinical protocols set by experienced treatment planners for this prostate case.

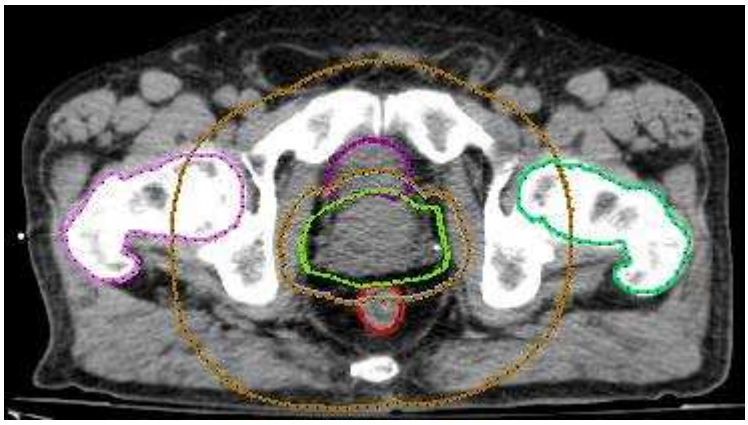

Figure 3.1: CT scan of patient displayed by CERR [15].

In Section 3.1, a distribution is selected for scenarios, and computational re-

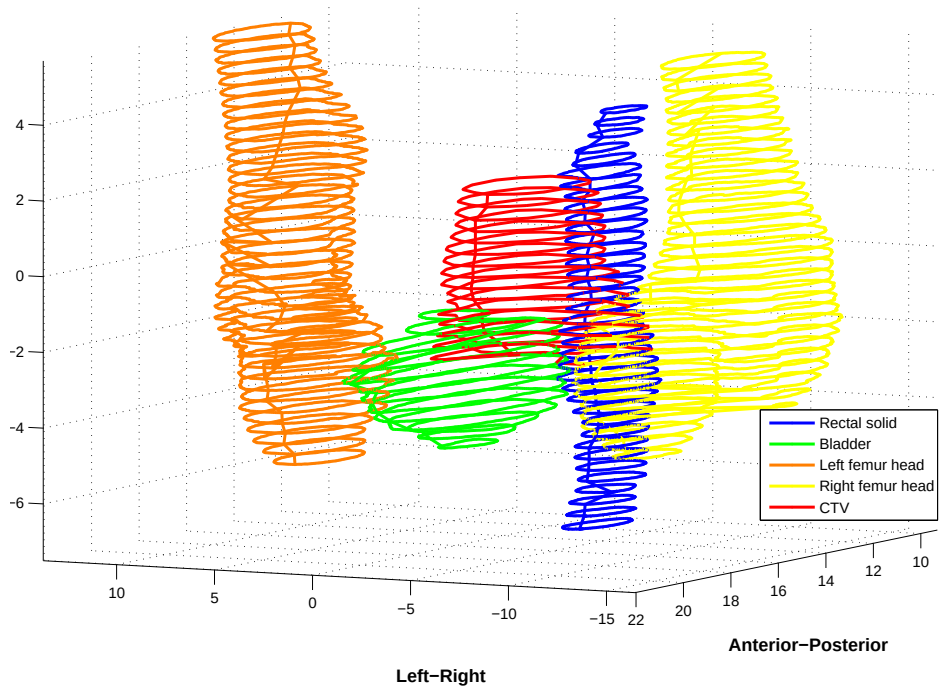

Figure 3.2: Structure contours for patient case. Unspecified region not shown.

sults are presented in Section 3.2. Methods for evaluating treatment plans are described in Sections 3.3 and 3.4, and computational results are compared for various parameter values in Section 3.5.

3.1 Distribution Selection for Scenarios

We construct a distribution to be used in our formulation from data published in the literature. Several studies have been done on prostate motion over the course of radiation therapy. The common approach to measuring motion is to acquire CT scans just prior to treatment (either at every treatment, or weekly), delineate structures, and compare their volume, shape, and position relative to

Table 3.1: Clinical protocols for prostate case.

Structure	**Dose Prescription**
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 head	Maximum dose of 50.0 Gy
Right femur head	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

bony anatomy. A reference table of prostate motion studies can be found in the review by Langen and Jones [31].

For studying changes in prostate volume, Mechalakos *et al* [45] took four CT scans for each of 50 prostate cancer patients undergoing radiation therapy. They observed significant trends in increasing bladder volume, but no trends in the prostate volume. Roeske *et al* [53] studied 10 patients by taking weekly CT scans and found that bladder volumes generally increased in half the patients, and generally decreased in the other half. But like Mechalakos *et al*, they found no significant trends in prostate volume over the course of treatment.

Shape variation was studied by Deurloo *et al* [16] in a study of 19 patients. They

concluded that the extent of prostate deformation during the course of radiotherapy is small, relative to organ motion.

For investigating rotations, van Herk *et al* [62] took repeat CT scans of 11 patients and measured the rotation of the prostate relative to the pelvic bone. They found that the prostate is fairly stable, except for rotations around the patient's left-right axis. The magnitude of this rotation was 4.0°, measured as one standard deviation.

A number of studies measure shifts of the prostate's center of mass during radiation therapy, usually measured along the patient's three anatomical axes (superior-inferior, anterior-posterior, and left-right). Wong *et al* [66] treated 108 prostate cancer patients, taking CT scans during the latter part of treatment. Along each axis, most shifts were less than 10 mm, and of these, a majority were less than 3 mm. Balter *et al* [3] studied 10 patients and concluded that displacements of the prostate were in a similar range, 1.7 mm to 4.5 mm.

Liu *et al* [39] studied 24 patients and found that shifts were greater in the anterior rather than posterior direction, and slightly greater in the superior rather than inferior direction. Lattanzi *et al* [34] performed a similar study with 6 patients and also found that shifts were greater in the anterior rather than posterior direction. Lattanzi *et al* found that shifts were slightly greater to the right rather than left, yet Roeske *et al* concluded that movement of the prostate in the left-right direction was insignificant. In their study of 10 patients, Roeske *et al* found that shifts ranged from 5.0 mm in the anterior direction to 5.3 mm posterior, and 6.3 mm inferior to 4.2 mm superior.

Based on these data, we choose to ignore changes in volume, shape, and rotations of the prostate. Our selected distribution has scenarios that correspond to

voxel location shifts as shown in Table 3.2. Although this distribution is based on motion of the prostate, we use this distribution for voxels in all structures, for simplicity. Future studies may include investigating structure-specific distributions.

Table 3.2: The scenario distribution.

Scenario	**Direction**	**Shift**	**Probability**
1	None	0 mm	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

In general, the distribution can be patient-specific. Results based on alternate distributions for all voxels appear in Section 3.5.

Though setup uncertainties can be systematic in nature, the probabilistic model can help guard against unfavorable outcomes in any particular scenario. If the target region is consistently located in one position, a deterministic solution may consistently overdose the healthy portion of the PTV, but a robust solution may recognize this as a possible scenario and therefore irradiate more sparingly.

3.2 Computational Results

For this prostate case, treatment planners pre-selected five coplanar beam angles on the plane perpendicular to the patient's superior-inferior axis. The beam angles

are equi-spaced at 0°, 72°, 144°, 216°, and 288°, where 0° corresponds to the beam pointing directly down towards the supine patient. From each beam angle, we selected beamlets which would deliver at least some significant amount of dose to the CTV, creating a set of 940 beamlets (approximately 200 from each angle).

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. Further study on selecting an appropriate value for δ appears in Section 3.5.

We used SeDuMi version 1.05 [55, 59] to solve our SOCP to global optimality. SeDuMi is a MATLAB [44] add-on that implements the self-dual embedding technique for solving optimization problems over symmetric cones, which includes SOCs. All problems were solved on a Dell PowerEdge 1750 with a single 3 GHz Xeon CPU and 4 GB RAM.

The A matrix for the SeDuMi input (for 1.0 cm voxels) contains 62.5% zeros with the largest entries being on the order of 100. After setting entries in A whose absolute values were less than 10^{-5} to zero, the resulting matrix had 70.6% zeros. The new solve time was 20% faster, with very minimal differences in the solution, so we adopted this convention of setting low entries in A to zero, as done in [20].

In our formulation, we have a maximum dose constraint for all voxels and a minimum dose constraint for CTV voxels, so the number of SOCs in our formulation is $V + T$, where V is the total number of voxels and T is the number of voxels in the CTV. With 7 scenarios and 940 beamlets, we were able to solve problems on 0.8 cm voxels in about 3.3 hours. For a fixed set of parameters, beamlets, and the scenarios in Table 3.2, the CPU solve times and number of iterations by voxel size are shown in Table 3.3 and plotted in Figure 3.3. We observe a slight increase in solve time versus number of SOC constraints as $V + T$ increases. In practice,

Table 3.3: Solve times for robust formulation.

Voxel Size	V	T	**Iterations**	**Solve Time**
2.0 cm	411	10	31	3 min
1.6	826	25	39	8
1.2	1934	56	61	30
1.0	3307	88	65	62
0.8	6569	191	97	199
0.6	15,417	449	(Out of memory)	

since treatment plans must be found within a few hours or overnight, and often require adjustments requested by the physician, the solve times for the 1.0 cm or 0.8 cm sized voxel instances are not unreasonable.

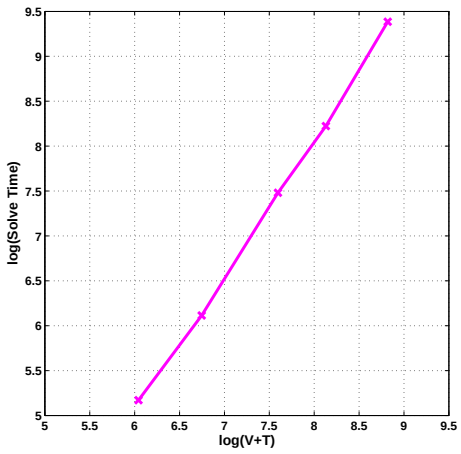

Figure 3.3: Log-log plot of solve time versus number of SOC constraints.

For comparison, we solved the deterministic LP formulation on SeDuMi. Like

the robust formulation, the objective is to minimize the sum of penalties,

$$w_{\text{Tmin}} r^+_{\text{Tmin}} + w_{\text{Tmax}} r^+_{\text{Tmax}} + \sum_{k=1}^{h} w_k r^+_k + \sum_{k \in \overline{h}} \overline{w}_k q^+_k,$$

subject to:

maximum total dose to healthy voxels:

$$\begin{aligned} N \boldsymbol{a}_{i,1}^\top \boldsymbol{x} &\leq u_k, && \forall i \in \mathcal{H}_k, \text{ for } k = 1, \ldots, h, \\ u_k - m_k &= r^+_k - r^-_k, && \text{for } k = 1, \ldots, h, \\ r^+_k, r^-_k &\geq 0, && \text{for } k = 1, \ldots, h, \end{aligned}$$

maximum total dose to target voxels:

$$\begin{aligned} N \boldsymbol{a}_{i,1}^\top \boldsymbol{x} &\leq u_{\text{Tmax}}, && \forall i \in \text{CTV}, \\ u_{\text{Tmax}} - m_{\text{Tmax}} &= r^+_{\text{Tmax}} - r^-_{\text{Tmax}}, \\ r^+_{\text{Tmax}}, r^-_{\text{Tmax}} &\geq 0, \end{aligned}$$

minimum total dose to target voxels:

$$\begin{aligned} N \boldsymbol{a}_{i,1}^\top \boldsymbol{x} &\geq u_{\text{Tmin}}, && \forall i \in \text{CTV}, \\ m_{\text{Tmin}} - u_{\text{Tmin}} &= r^+_{\text{Tmin}} - r^-_{\text{Tmin}}, \\ r^+_{\text{Tmin}}, r^-_{\text{Tmin}} &\geq 0, \end{aligned}$$

approximate dose-volume constraints:

$$\begin{aligned} N \boldsymbol{a}_{i,1}^\top \boldsymbol{x} - d_k &= v^+_i - v^-_i, && \forall i \in \mathcal{H}_k, \forall k \in \overline{h}, \\ v^+_i, v^-_i &\geq 0, && \forall i \in \mathcal{H}_k, \forall k \in \overline{h}, \\ \sum_{i \in \mathcal{H}_k} v^+_i - g_k &= q^+_k - q^-_k, && \forall k \in \overline{h}, \\ q^+_k, q^-_k &\geq 0, && \forall k \in \overline{h}, \end{aligned}$$

and nonnegative beamlet intensities:

$$\boldsymbol{x} \geq \boldsymbol{0},$$

where $\boldsymbol{a}_{i,1}$ is the dose-coefficient vector for voxel i in its original CT scan location. Table 3.4 shows the solve times for various voxel sizes.

Table 3.4: Solve times for deterministic formulation.

Voxel Size	V	T	**Iterations**	**Solve Time**
2.0 cm	430	29	32	1 min
1.6	855	54	33	1
1.2	2010	132	60	4
1.0	3429	210	58	8
0.8	6796	418	67	17
0.6	15,989	1021	129	100
0.4	54,156	3423	150*	430*

* Maximum number of iterations reached

3.3 Dose-Volume Histograms

In order to evaluate the therapeutic quality of a solution of beamlet intensities, physicians rely on clinical guidelines, expertise, and intuition. Physicians are accustomed to examining Dose-Volume Histograms (DVHs), a cumulative histogram of the dose delivered to each structure, similar in nature to dose-volume constraints (see Section 2.4). For a set of beamlet intensities $\boldsymbol{x}$, a DVH plots the fraction of each structure $\mathcal{H}_k$ that receives at least some total dose d:

$$\mathrm{DVH}_k(d, \boldsymbol{x}) = \frac{1}{|\mathcal{H}_k|} \sum_{i \in \mathcal{H}_k} \mathbf{1}\{D_i(\boldsymbol{x}) \geq d\}, \tag{3.1}$$

where $\mathbf{1}_{\{\cdot\}}$ is the indicator function. Traditional DVHs are based on the dose delivered to static voxels, located as in the original planning CT scan, but in our setting $D_i(\boldsymbol{x})$ is a random variable.

In Figure 3.4, we present a Dose-Expected Volume Histogram (DEVH) which

is a DVH, except it plots the expected fraction of a structure that receives at least d. Taking the expectation of (3.1), we have

$$\mathrm{DEVH}_k(d, \boldsymbol{x}) = \mathrm{E}\left(\frac{1}{|\mathcal{H}_k|}\sum_{i\in\mathcal{H}_k}\mathbf{1}\{D_i(\boldsymbol{x}) \geq d\}\right) = \frac{1}{|\mathcal{H}_k|}\sum_{i\in\mathcal{H}_k}\mathrm{P}\left(D_i(\boldsymbol{x}) \geq d\right).$$

Using our results from Section 2.1, under the assumption that $D_i(\boldsymbol{x})$ has a normal distribution with mean $N\boldsymbol{p}^\top A_i\boldsymbol{x}$ and variance $N\|RA_i\boldsymbol{x}\|^2$, we can write

$$\mathrm{DEVH}_k(d, \boldsymbol{x}) = \frac{1}{|\mathcal{H}_k|}\sum_{i\in\mathcal{H}_k}\mathrm{P}\left(Z \geq \frac{d - N\boldsymbol{p}^\top A_i\boldsymbol{x}}{\sqrt{N}\|RA_i\boldsymbol{x}\|}\right), \tag{3.2}$$

where Z is a normal random variable with mean 0 and variance 1.

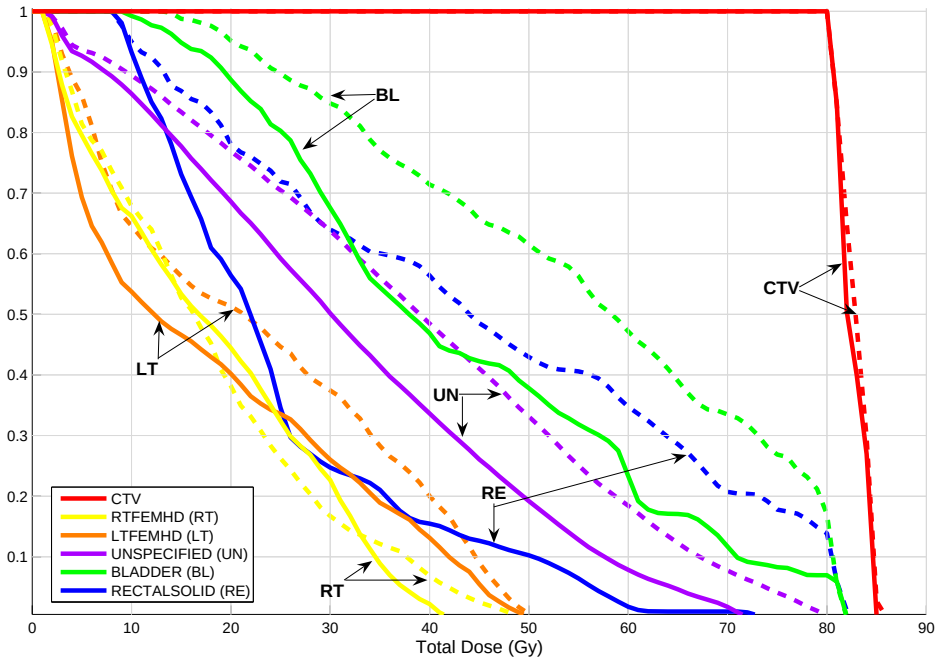

Figure 3.4: Dose-Expected Volume Histogram (DEVH) for robust solution (solid curves) and deterministic solution (dashed curves) for 0.8 cm voxels.

Figure 3.4 shows that both solutions achieve a high, homogeneous dose to the CTV. However, the solid curves are mostly below the dashed curves, and hence

the robust solution delivers considerably less dose to the healthy structures. In the penalty objective function, the optimal robust solution mostly penalizes the minimum dose to the CTV and the approximate dose-volume constraints. 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.

Figure 3.4 is a straightforward extension of a DVH, showing the expected outcome, but the actual outcome may be significantly different. Also, under the normal distribution assumption, a DEVH depends on the selected distribution via $\boldsymbol{p}, R$, and A_i, as in (3.2).

For the prostate case, we simulated a treatment by randomly selecting N shifts from the n scenarios representing possible positions of the patient at each treatment fraction. We then used each model's (robust and deterministic) optimal solution of beamlet intensities to calculate 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 3.5, conveys a sense of the range of possible outcomes associated with each solution.

More specifically, Figure 3.5 plots the DVH cloud based on 100 simulated treatments of $N = 45$ fractions each, sampling from the distribution given in Table 3.2. For each fraction, we assume that all voxels shift together.

The plots in Figures 3.4 and 3.5 show that both solutions deliver comparable dose to the CTV, but the robust solution delivers less dose to every healthy structure. Most significantly, the robust solution satisfies all four dose-volume constraints to the rectal solid, while the deterministic solution violates all four. Since the deterministic solution must deliver a high, uniform dose to the entire PTV, it

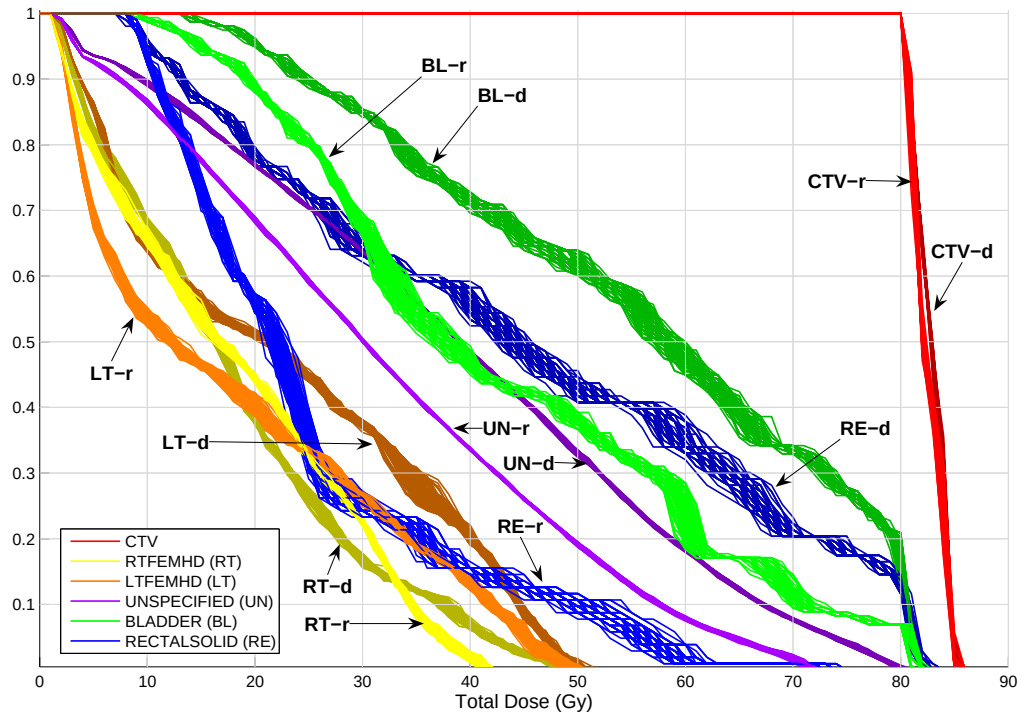

Figure 3.5: Dose-Volume Histogram cloud for robust solution ('-r') and deterministic solution ('-d'), for 100 simulated treatments.

delivers more dose to healthy regions than the robust solution.

3.4 Other Performance Measures

In addition to Dose-Volume Histograms, we can also plot doses delivered to patient slices for analyzing solution quality. Though these figures do not explicitly convey the implications of voxel location uncertainty, they may provide insight on the robust formulation's performance. Figure 3.6 displays the dose delivered to one patient plane, the same slice as shown in Figure 3.1, and Figure 3.7 shows the surface plots.

These plots clearly show that the deterministic solution delivers a high dose to a much larger region covering the entire PTV. Since the PTV overlaps the rectal solid and bladder, both of these healthy regions receive high dose. The robust solution more carefully avoids these regions and delivers high dose to other parts immediately surrounding the CTV that don't overlap with the rectal solid or bladder. Therefore, it cleverly falls short of delivering high dose to parts of the target if an anterior or posterior shift occurs, and yet succeeds in the case of a left or right shift, so that the overall dose to the CTV is sufficiently high, with high probability.

3.5 Varying δ and n

Our maximum and minimum dose constraints must be satisfied with probability $1 - \delta$, and for our results thus far, we have set $\delta = 0.05$. If we vary δ and keep all other parameters the same, with 7 scenarios as in Table 3.2, and 1.0 cm voxels, the solution times are shown in Table 3.5, and the DEVHs are shown in Figure 3.8.

Table 3.5: Solve times for various δ.

δ	**Iterations**	**Solve Time**
0.010	60	50 min
0.025	62	52
0.050	65	62
0.100	51	44
0.250	53	48

There is no discernible pattern for the solve times. The DEVH curves are

nearly identical for the CTV, and they are similar in the tail end for healthy structures. A possible explanation for this is that our penalty model only penalizes the maximum dose to healthy structures, which is comparable for all solutions, and all dose-volume constraints are met by all solutions.

Since the effect of varying δ is minimal and unclear, our choice of $\delta = 0.05$ seems reasonable in terms of performance, though perhaps we could achieve faster solve times and similar solutions for other values of δ.

Next, we investigate the effect of varying the number of scenarios n, and the scenario distribution. In our SOC constraints, the vector $RA_i\boldsymbol{x}$ inside the norm has dimension n, so increasing n increases the dimension of each second order cone.

For a distribution with 27 scenarios, we selected a shift to each corner, face, and edge of a 1 cm^3 cube centered at a voxel's original CT scan location. Then for each n, we selected n symmetric shifts from these 27 shifts, with each of the n shifts being equally likely. The mean shifts are all $(0, 0, 0)$, but the variance varies. We continue to use 1.0 cm voxels and $\delta = 0.05$.

Table 3.6: Scenarios and solve times for various n.

n	**Shifts, on 1 cm^3 cube**	**Iterations**	**Solve Time**
5	Center, 4 symmetric shifts	55	49 min
7	Center, 6 faces	54	51
9	Center, 8 corners	63	77
13	Center, 12 edges	65	114
15	Center, 6 faces, 8 corners	60	123
19	Center, 6 faces, 12 edges	(Out of memory)	

Table 3.6 summarizes the solve times for various n, and Figure 3.9 shows the DEVH for each solution. These problems were solved with n scenarios, and evaluated with 27 equally likely shifts. We observe that all solutions have similar quality, so our choice of $n = 7$ scenarios seems reasonable.

Further, we compare the performance of the robust and deterministic solutions on a different scenario distribution than the one used in the robust formulation. Figure 3.10 is DVH cloud of the same setup as in Figure 3.5 except that the scenarios have been simulated from the 27 equally likely shifts described above, instead of the distribution in Table 3.2. The robust solution delivers much less dose to the bladder and rectal solid, similar to its performance in Figure 3.5. However, its dose to the CTV is not as homogeneous as before. Most likely this is caused by the fact that the magnitudes of the 7 shifts from Table 3.2 are all less than or equal to those of the 27 shifts. This suggests that the scenarios in the robust formulation must carefully be chosen to include the range of possible shifts by CTV voxels in order to ensure that every CTV voxel receives high dose.

For comparison, instead of using the distribution in Table 3.2, we solve a robust formulation on 0.8 cm voxels with 7 equally likely scenarios (center and 6 shifts to each face of a 1 cm^3 cube), and evaluate it on 27 scenarios. The dose homogeneity in the CTV is closer to that of the deterministic solution on 0.8 cm voxels, as shown in Figure 3.11. We suspect that the robust solution on 0.8 cm voxels solved on 9 shifts (center and 8 shifts to each corner of a 1 cm^3 cube) would perform even better when evaluated on 27 shifts, but this problem requires too much memory.

Instead, we can solve these problems on 1.0 cm voxels and compare their performance. Figure 3.12 plots the DEVH of four solutions solved on 1.0 cm voxels: a deterministic formulation, a robust formulation solved with the 7 scenarios in

Table 3.2, with 7 equally likely symmetric scenarios (center and 6 faces of a 1 cm^3 cube), and with 9 equally likely symmetric scenarios (center and 8 corners of a 1 cm^3 cube). These solutions were all evaluated on the 27 scenarios. We observe that of the three robust solutions, the 9 scenario solution achieves the best CTV dose homogeneity, followed by the solution with 7 symmetric scenarios, and then by the other 7 scenario solution. Hence, if the 27 scenarios are reasonable realizations of the patient's position, then the 9 scenario solution is preferable, but has a longer solve time and more computational requirements.

In summary, for this prostate case, the robust formulation is computationally tractable. The robust solution is favorable to the deterministic one since it achieves superior sparing of the healthy tissue while maintaining a high dose to the CTV, as the DEVH and DVH clouds show. The robust solution tailors the dose distribution to match the CTV's potential motion, while the deterministic solution conforms to the PTV's fixed position. The robust solution is sensitive to the scenario distribution, so it should encompass all possible shifts. However, selecting too many shifts will cause computational problems. Overall, compared to deterministic solutions, we believe the extra computational requirements are worthwhile for a robust solution's quality, and further, we investigate methods to reduce these requirements in the next chapter.

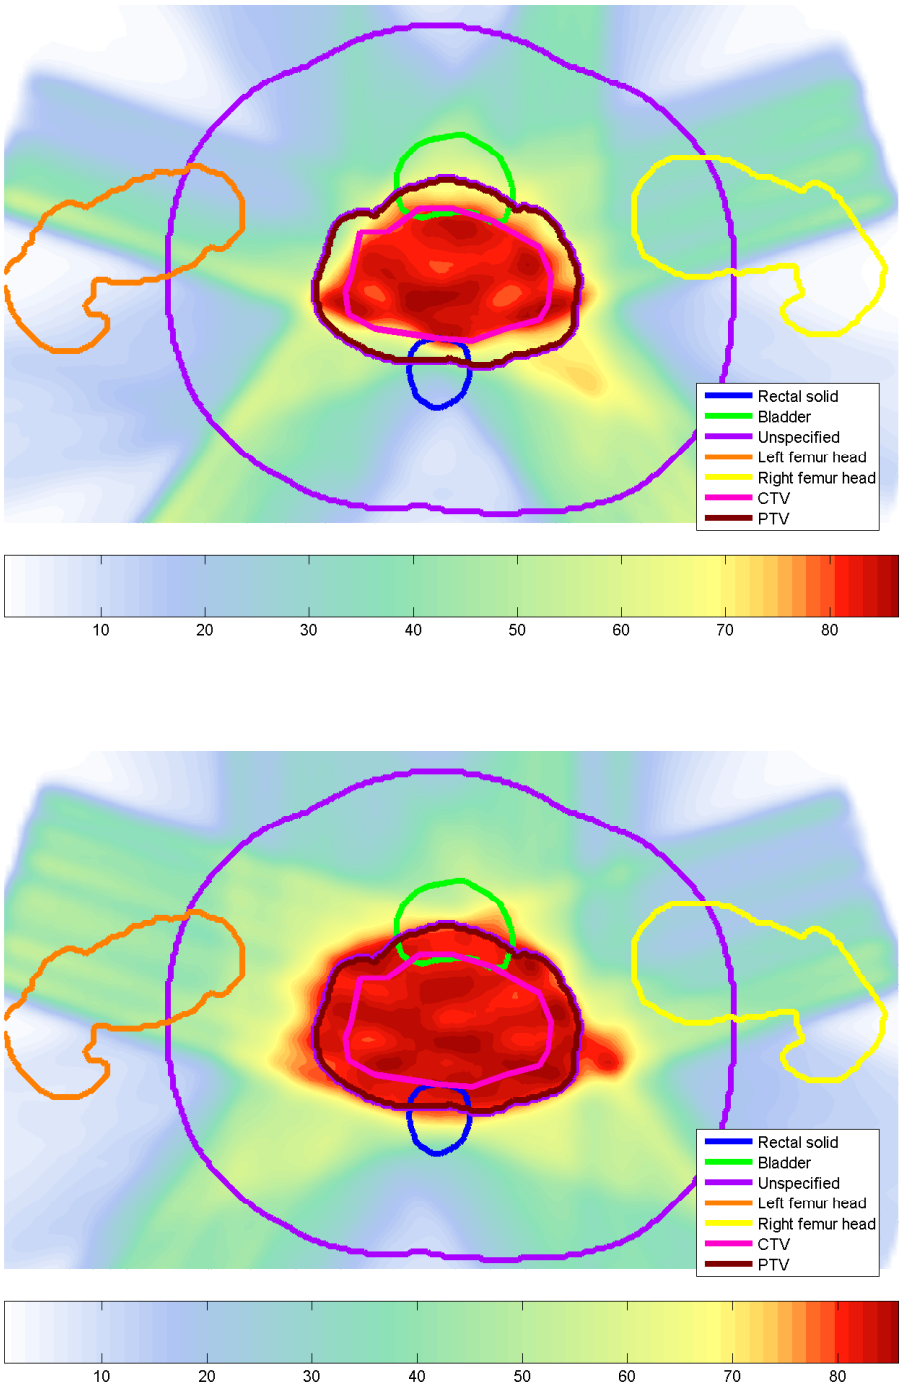

Figure 3.6: 2-D plot of dose delivered to one patient plane by the robust solution (top) and deterministic solution (bottom), with CT scan structure contours overlaid. Color scale indicates dose delivered (Gy).

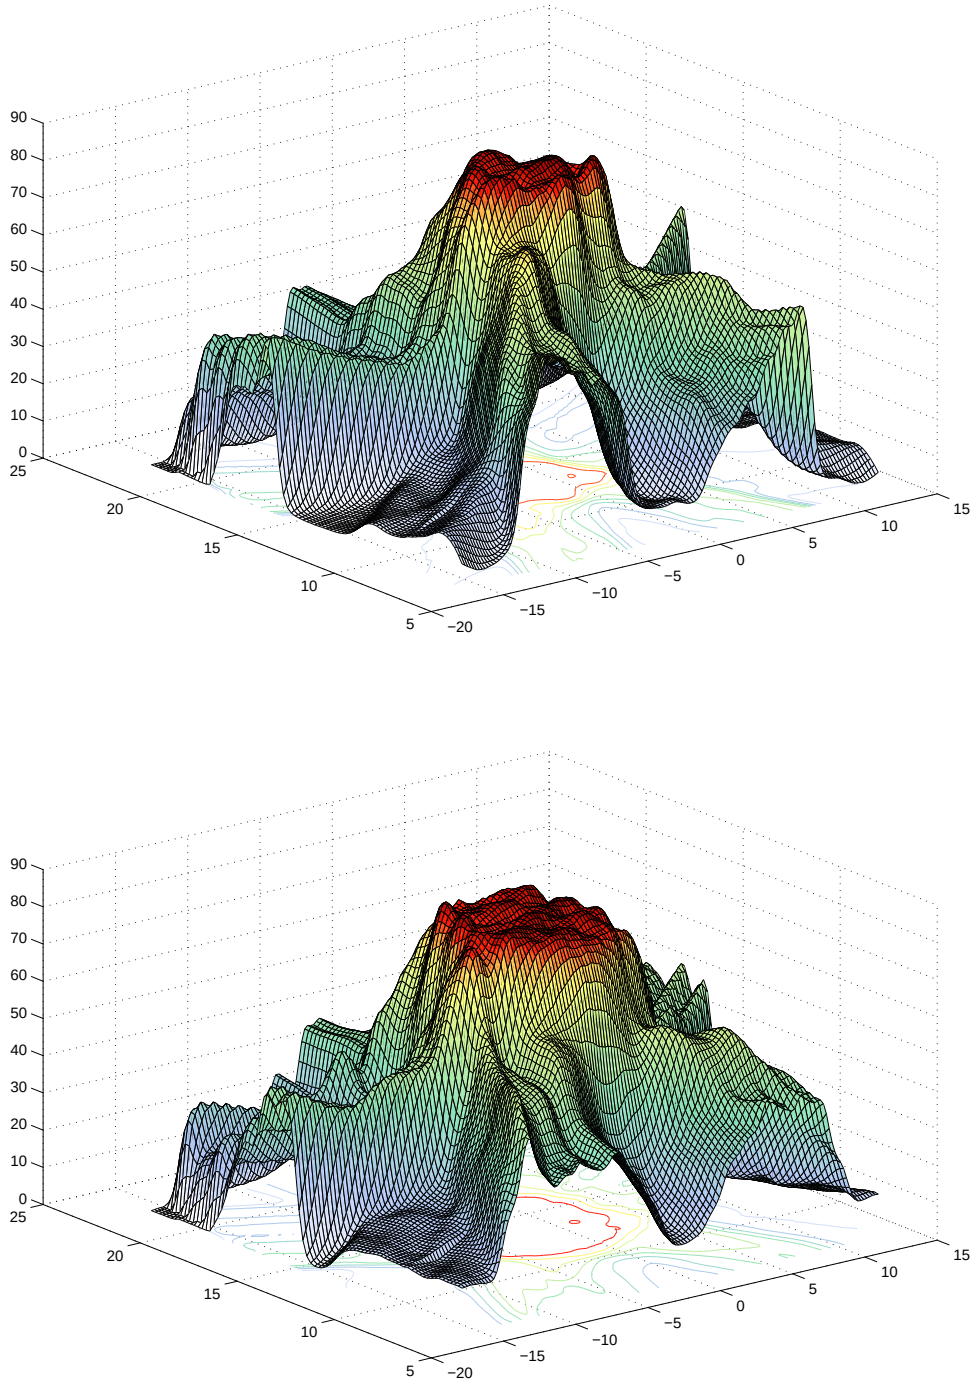

Figure 3.7: 3-D plot of dose distributions to one patient plane by the robust solution (top) and deterministic solution (bottom). Dose delivered (Gy) is on the vertical axis, same color scale as in Figure 3.6.

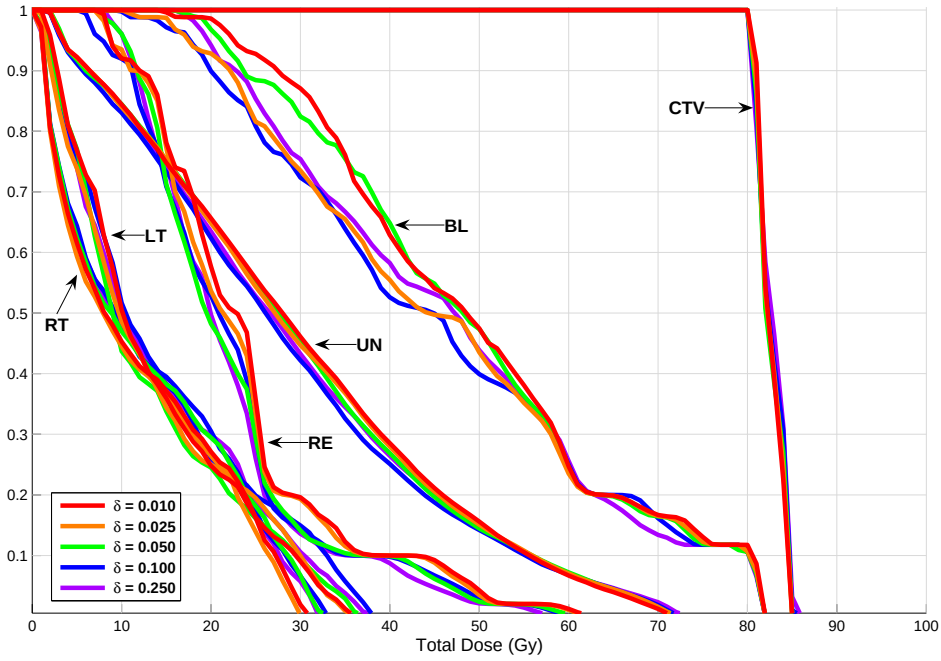

Figure 3.8: DEVH for robust solutions with varying δ, solved and evaluated on 1.0 cm voxels.

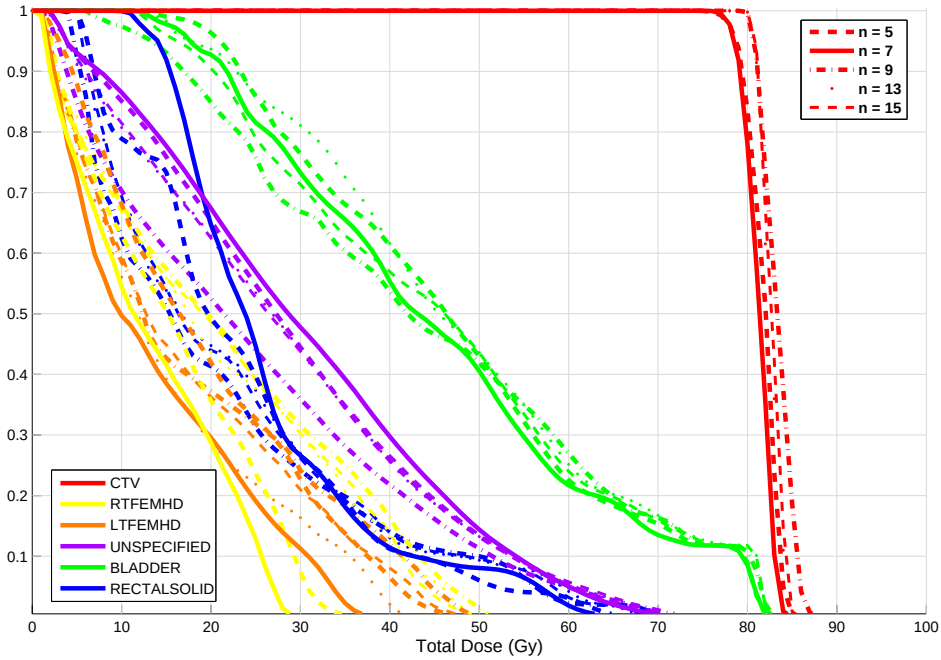

Figure 3.9: DEVH for robust solutions with various n as in Table 3.6, based on 27 equally likely shifts, solved and evaluated on 1.0 cm voxels.

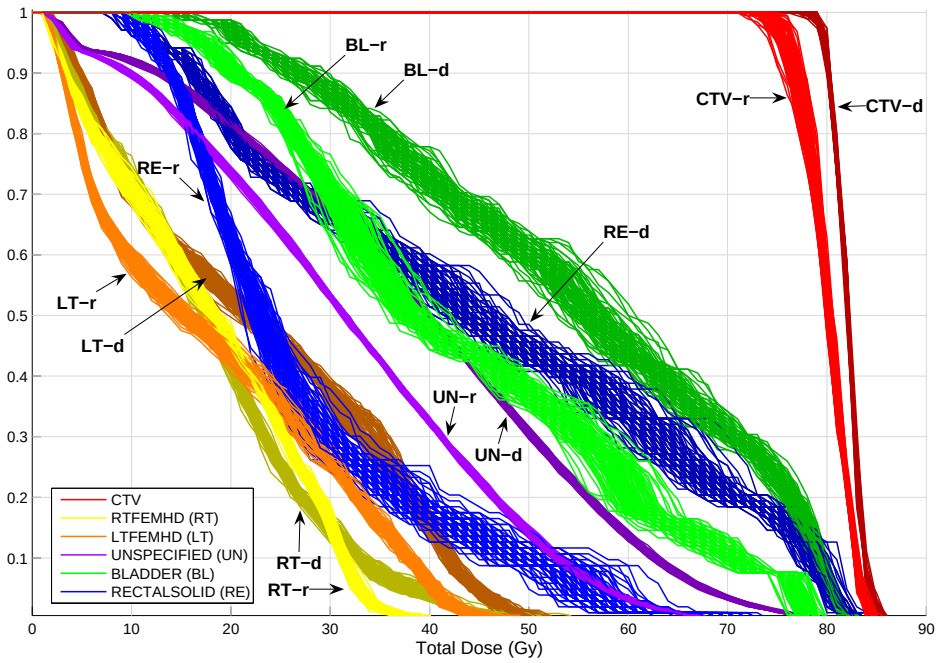

Figure 3.10: DVH cloud for 0.8 cm robust ('-r') and deterministic ('-d') solutions, based on 100 simulated treatments sampling from 27 equally likely shifts. Robust solutions solved on the 7 scenario distribution in Table 3.2.

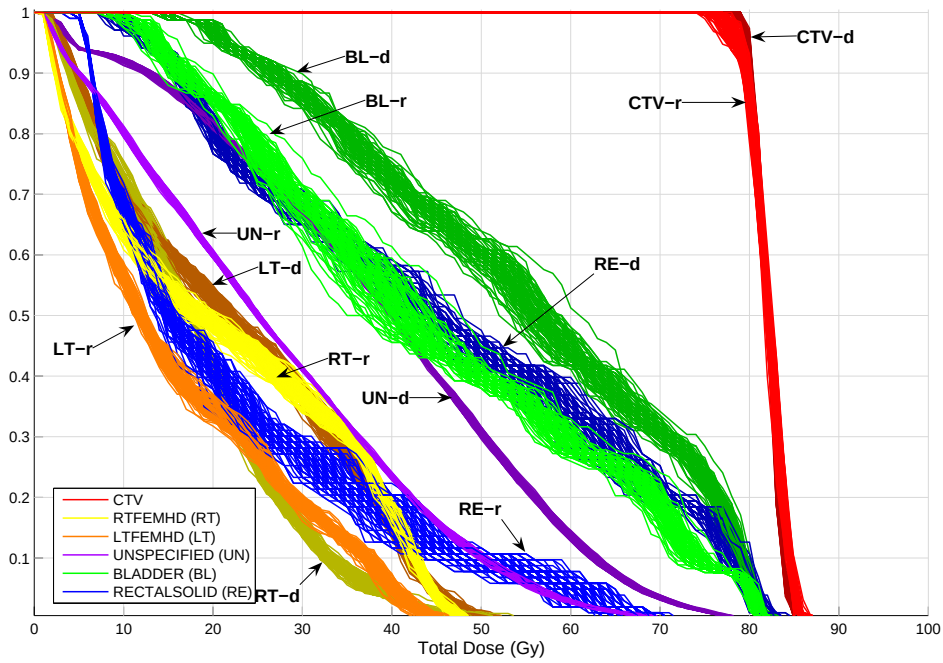

Figure 3.11: DVH cloud for 0.8 cm robust ('-r') and deterministic ('-d') solutions, based on 100 simulated treatments sampling from 27 equally likely shifts. Robust solution solved on 7 equally likely symmetric scenarios (center and 6 faces of a cube).

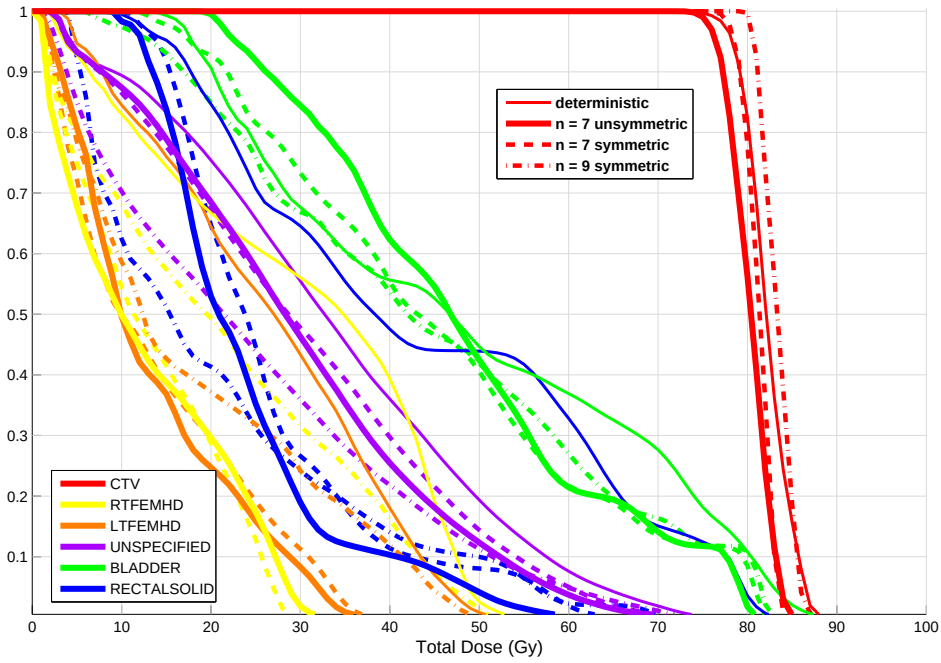

Figure 3.12: DEVH for 1.0 cm robust and deterministic solutions, evaluated on 27 equally likely shifts. Robust solutions solved on 7 equally likely symmetric scenarios (center and 6 faces of a cube), 7 scenarios from Table 3.2, and 9 equally likely symmetric scenarios (center and 8 corners of a cube).

Chapter 4

Reducing Computational Requirements

In the previous chapter, treatment plans were evaluated based on doses delivered to the same sized voxels (0.8 cm or 1.0 cm) as used for the problem input. However, current clinical treatment planning systems use 0.4 cm voxels for optimizing and evaluating plans for prostate cases, and 0.4 cm problem instances are too large for us to solve (see Tables 3.3 and 3.4). But since we calculate dose delivered by interpolating from a matrix associated with a 0.2 cm grid, we can calculate the doses delivered to 0.4 cm voxels to evaluate any solution of beamlet intensities.

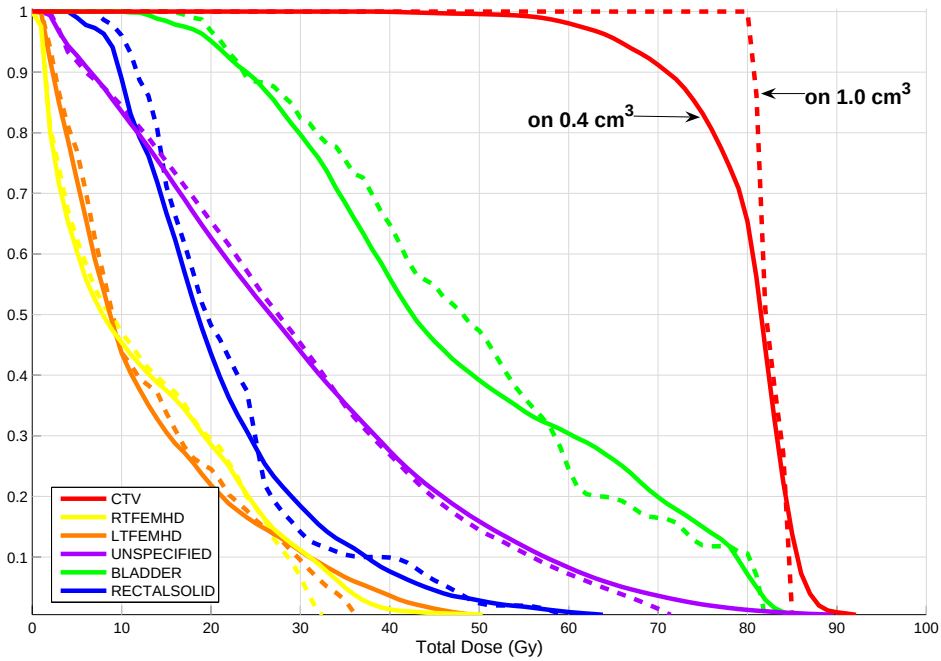

Figure 4.1: DEVH of robust solutions solved on 1.0 cm voxels, evaluated on 0.4 cm voxels (solid curves) and on 1.0 cm (dashed curves).

Figure 4.1 shows the DEVH of the optimal robust solution for the 1.0 cm voxel

input, when evaluated on 0.4 cm versus 1.0 cm voxels. The dose to the CTV on 0.4 cm voxels is clearly not as homogeneous as on 1.0 cm voxels, and the doses to healthy regions seems comparable, though the unspecified region and femur heads receive slightly more dose when viewed on 0.4 cm voxels. We observe a similar picture for the 0.8 cm robust solution on 0.4 cm voxels.

Since CTV dose homogeneity has high priority, we explore two methods that maintain computational tractability for finding better robust solutions on 0.4 cm voxels. Recall that 0.8 cm voxels are points on a $0.8 \times 0.8 \times 0.8$ cm grid, where each voxel represents a cube of volume 0.8^3 cm^3. The dose delivered to that cube is measured at one point, the center. We investigate taking the mean of doses delivered to several points within each cube in Section 4.1 and sampling extra voxels near the boundary of a structure in Section 4.2.

Also, the entire grid of original potential voxel locations can be shifted prior to determining the set of voxels for each structure. Doing so creates slightly different problems of approximately the same size. When evaluted on fine voxels, the quality of the solutions may be different, and the methods proposed in this chapter may reduce those differences.

4.1 Dose Aggregation

The robust formulation for the 1.0 cm voxel input fails to capture doses delivered between voxels, as seen in Figure 4.1. This shortcoming motivates us to explore a method of dose aggregation in which we replace the dose coefficients for a particular voxel with the mean of the dose coefficients of nearby points. One could also use the maximum or minimum dose coefficients of nearby points within the volume represented by each voxel. For a conservative approach, one could use the

maximum for all maximum dose constraints, and the minimum for all minimum dose constraints. In this work we use the mean, and in our formulation, we replace $D_i(\boldsymbol{x})$ with an aggregate dose, $\frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} D_j(\boldsymbol{x})$, where $\mathcal{N}_i$ is some set of neighbors of voxel i. Using our scenario model from Section 2.1, the expected aggregate dose is

$$\mathrm{E}\left(\frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} D_j(\boldsymbol{x})\right) = \frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} N\boldsymbol{p}^\top A_j\boldsymbol{x},$$

and the variance is

$$\mathrm{Var}\left(\frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} D_j(\boldsymbol{x})\right) = \frac{1}{|\mathcal{N}_i|^2}\mathrm{Var}\left(\sum_{j\in\mathcal{N}_i}\sum_{\ell=1}^{N} D_{j\ell}(\boldsymbol{x})\right).$$

In Section 2.1, we assumed that the N random variables giving the dose per fraction to any particular voxel are i.i.d. Following the same reasoning, we assume that the sums of doses per fraction are also i.i.d. random variables, and we get

$$\begin{aligned}\mathrm{Var}\left(\frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} D_j(\boldsymbol{x})\right) &= \frac{1}{|\mathcal{N}_i|^2}N\mathrm{Var}\left(\sum_{j\in\mathcal{N}_i} D_{j1}(\boldsymbol{x})\right)\\ &= \frac{1}{|\mathcal{N}_i|^2}N\left\|R\left(\sum_{j\in\mathcal{N}_i} A_j\right)\boldsymbol{x}\right\|^2,\end{aligned}$$

following the derivation in Section 2.1.

A maximum dose constraint for voxel $i\in\mathcal{H}_k$ with dose aggregation becomes

$$\frac{1}{|\mathcal{N}_i|}\sqrt{N}\left\|R\left(\sum_{j\in\mathcal{N}_i} A_j\right)\boldsymbol{x}\right\| \le \frac{u_k - \frac{1}{|\mathcal{N}_i|}N\boldsymbol{p}^\top\left(\sum_{j\in\mathcal{N}_i} A_j\right)\boldsymbol{x}}{z_{1-\delta}}$$

or

$$\left\|R\left(\frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} A_j\right)\boldsymbol{x}\right\| \le \frac{u_k - N\boldsymbol{p}^\top\left(\frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} A_j\right)\boldsymbol{x}}{z_{1-\delta}\sqrt{N}},$$

which is the same as (2.4) except the dose coefficient matrix A_i has been replaced with $\frac{1}{|\mathcal{N}_i|}\sum_{j\in\mathcal{N}_i} A_j$.

We restrict voxel aggregation to voxels $i \in \mathcal{H}_k$ which belong to the *interior* of $\mathcal{H}_k$, defined as $\mathcal{I}_k \subset \mathcal{H}_k$, where $i \in \mathcal{I}_k$ if and only if every $j \in \mathcal{N}_i$ is within the original CT scan contours of structure $\mathcal{H}_k$. For each voxel i not in the interior, we continue to use dose coefficients A_i as in our original formulation. This prevents dose aggregation from aggregating doses from different structures.

We select $\mathcal{N}_i$ to be the set containing i and the 8 corners of a cube centered at voxel i, where the cube's edge length is half the grid size of the original voxels. By selecting such points, $\mathcal{N}_{i_1}$ and $\mathcal{N}_{i_2}$ are disjoint for any two distinct voxels i_1 and i_2.

Figure 4.2 shows the DEVHs for a 1.0 cm voxel input, with and without dose aggregation for all voxels, when evaluated on 0.4 cm voxels. The solution using dose aggregation delivers more dose to the CTV and less dose to the bladder, but also delivers slightly more dose to all other healthy structures.

Note that dose aggregation does not change the number of SOC constraints. The CPU solve time for the formulation with dose aggregation was slightly faster (2873 seconds, versus 3727) and took fewer iterations (53 iterations, versus 65). Out of the 3307 voxels, 2691 (about 81%) of them lie in the interior of their structure and had their doses aggregated.

Since dose aggregation captures some measure of the dose delivered between voxels, we explore the possibility of increasing the voxel grid size in the unspecified region. The majority of voxels (about 82%) are in the unspecified region, so reducing the number of these voxels reduces the number of SOC constraints and speeds up the solution time considerably. In Figure 4.3, the voxel size for the unspecified region is varied, while a 1.0 cm grid is used for all other structures. The computational details are summarized in Table 4.1. There is a tradeoff for faster solve time and less dose to the unspecified region, and we observe that the

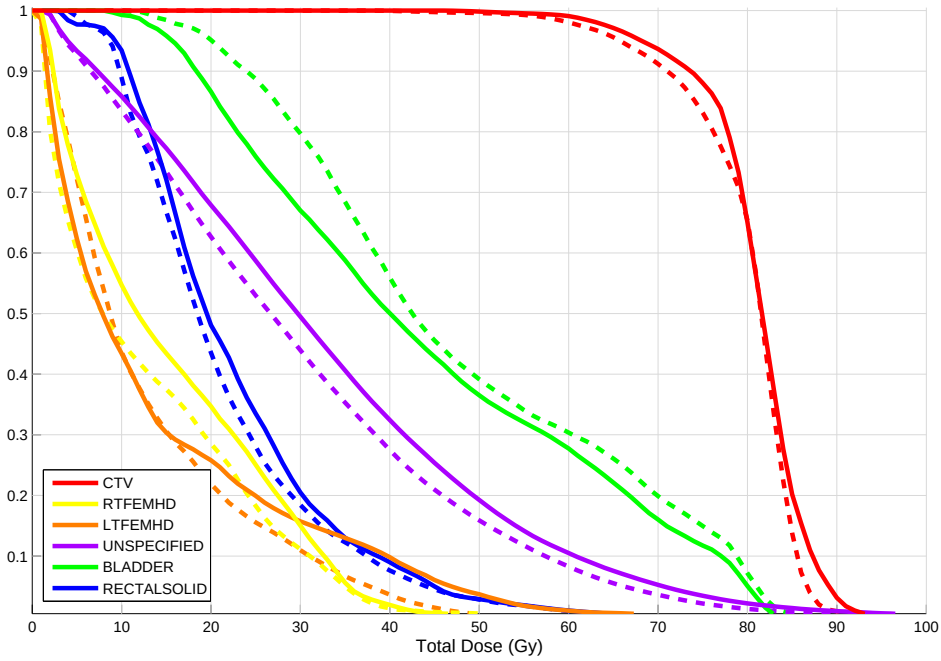

Figure 4.2: DEVH of robust solution solved on 1.0 cm voxels, original solution (dashed) and with dose aggregation (solid), all evaluated on 0.4 cm voxels.

solution with 1.4 cm voxels gives a comparable solution to 1.0 cm, but solves in half the time. Therefore, we proceed with 1.4 cm voxels in the unspecified region and 1.0 cm in all other regions, with dose aggregation for all voxels.

Moreover, on a 1.4 cm grid, we have 1018 voxels in the unspecified region, which is about 5.6 times as many than the next largest structure (the right femur head, with 181 voxels). Ferris *et al* [20] also studied the effects of coarser voxels in the unspecified region, and they concluded that a factor of five times as many voxels as the next largest structure is "more than adequate" for the unspecified region, and our results support this finding.

With the computational savings from having less voxels in the unspecified re-

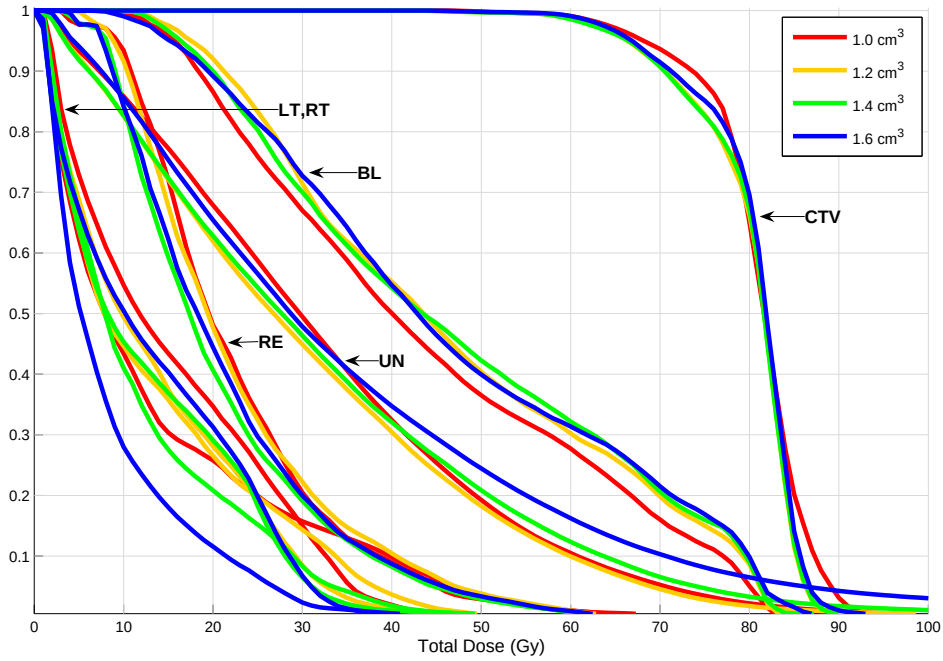

Figure 4.3: DEVH of robust solutions with dose aggregation, solved on 1.0 cm voxels, except 1.0 cm (red), 1.2 cm (yellow), 1.4 cm (green), and 1.6 cm (blue) voxels in the unspecified region, all evaluated on 0.4 cm voxels.

gion, we can now accommodate more voxels in other regions. Most importantly, to improve dose homogeneity, CTV voxels can be on a finer grid. Figure 4.4 shows the DEVH for solutions on 1.0 cm voxels, except for the unspecified region on 1.4 cm voxels, for various sized CTV voxels, with dose aggregation for all voxels. Computational results are summarized in Table 4.2. The solution with 0.8 cm voxels in the CTV achieves better dose homogeneity than the 1.0 cm voxels, and solves in about 35 minutes. This solution also delivers slightly more dose to the femur heads, rectal solid, and the unspecified region, but this compromise seems worthwhile. Therefore, we proceed with 0.8 cm voxels in the CTV.

Table 4.1: Solve times for various sized voxels in the unspecified region (UN), with all other voxels 1.0 cm, all voxels with dose aggregation.

UN Voxel Size	V	**UN Voxels**	**Iterations**	**Solve Time**
1.0 cm	3307	2725	53	48 min
1.2	2165	1583	55	33
1.4	1600	1018	52	23
1.6	1263	681	52	19

Table 4.2: Solve times for various sized voxels in the CTV, with all other voxels of size 1.0 cm except 1.4 cm in the unspecified region, and all voxels with dose aggregation.

CTV Voxel Size	V	T	**Iterations**	**Solve Time**
1.0 cm	1600	88	52	23 min
0.8	1703	191	63	35
0.6	1961	449	75	64
0.4	3057	1545	98	231

4.2 Boundary Sampling

We now explore the impact of voxels on a non-uniformly spaced grid. Our boundary sampling method selects a higher density of voxels near the border of a structure, with the hope that carefully constraining the doses delivered on the boundary will influence the interior doses. Also, boundary sampling in the CTV will give us extra voxels near the target boundary, a critical area where we expect the greatest change in the dose distribution.

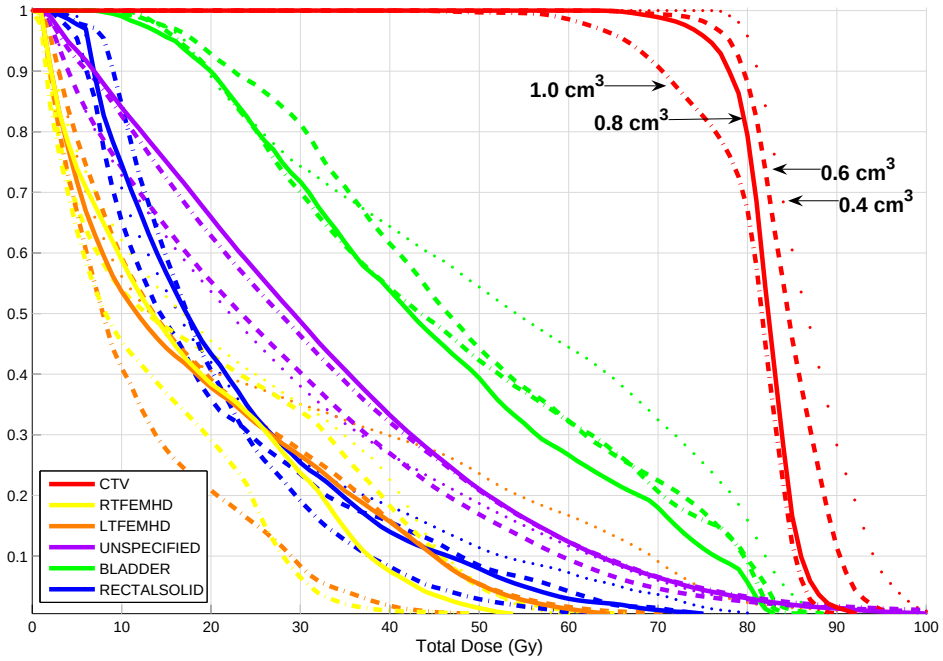

Figure 4.4: DEVH of robust solutions with dose aggregation, solved on 1.0 cm voxels except 1.4 cm in the unspecified region, and 1.0 cm (dash-dotted), 0.8 cm (solid), 0.6 cm (dashed), or 0.4 cm (dotted) CTV voxels, all evaluated on 0.4 cm voxels.

Using our definition of interior voxels from the previous section, let the *boundary* voxels of structure $\mathcal{H}_k$ be defined as $\mathcal{B}_k = \mathcal{H}_k \backslash \mathcal{I}_k$, i.e., all voxels which are originally in the structure but not in the interior. Then, for any boundary voxel $i \in \mathcal{B}_k$, only a subset of its neighbors $\mathcal{N}_i$ lie within the structure's CT scan contours. Let $\bar{\mathcal{N}}_i$ be the subset of $\mathcal{N}_i$ that lies within the structure's contours, and take the new set of voxels to be

$$\bar{\mathcal{H}}_k := \mathcal{I}_k \bigcup \left(\bigcup_{i \in \mathcal{B}_k} \bar{\mathcal{N}}_i \right).$$

Then for each voxel in $\bar{\mathcal{H}}_k$, the same maximum or minimum dose constraints apply, as in Section 2.5.

However, a weighting adjustment is required in our approximate dose-volume constraints to account for the non-uniformly spaced voxels. We have chosen $\mathcal{N}_i$ to consist of i and the 8 voxels at the corners of a cube centered at voxel i, so $|\mathcal{N}_i| = 9$, and there are now 9 voxels in the space previously represented by one voxel. For simplicity, we assume each of these "sub-voxels" represents an equal amount of volume, $s^3/9$ where s is the original voxel grid size. Consequently, the approximate total volume of structure $\mathcal{H}_k$ is

$$s^3|\mathcal{I}_k| + \frac{s^3}{9}\sum_{i\in\mathcal{B}_k}|\bar{\mathcal{N}}_i|. \tag{4.1}$$

Recall our approximate dose-volume constraint (2.16), where we set the upper bound g_k equal to $v_k|\mathcal{H}_k|(m_k - d_k)$, so that we have

$$\sum_{i\in\mathcal{H}_k}\left(N\boldsymbol{p}^\top A_i\boldsymbol{x} - d_k\right)_+ \leq g_k = v_k|\mathcal{H}_k|(m_k - d_k).$$

To measure the excess dose for each voxel $i \in \bar{\mathcal{H}}_k$, we now weight $\left(N\boldsymbol{p}^\top A_i\boldsymbol{x} - d_k\right)_+$ by the voxel's volume and replace $|\mathcal{H}_k|$ with (4.1), so that our approximate dose-volume constraint becomes

$$\begin{aligned} s^3\sum_{i\in\mathcal{I}_k}\left(N\boldsymbol{p}^\top A_i\boldsymbol{x} - d_k\right)_+ + \frac{s^3}{9}\sum_{i\in\mathcal{B}_k}\sum_{j\in\bar{\mathcal{N}}_i}\left(N\boldsymbol{p}^\top A_j\boldsymbol{x} - d_k\right)_+ \\ \leq v_k\left(s^3|\mathcal{I}_k| + \frac{s^3}{9}\sum_{i\in\mathcal{B}_k}|\bar{\mathcal{N}}_i|\right)(m_k - d_k)\,, \end{aligned}$$

or

$$\begin{aligned} \sum_{i\in\mathcal{I}_k}\left(N\boldsymbol{p}^\top A_i\boldsymbol{x} - d_k\right)_+ + \frac{1}{9}\sum_{i\in\mathcal{B}_k}\sum_{j\in\bar{\mathcal{N}}_i}\left(N\boldsymbol{p}^\top A_j\boldsymbol{x} - d_k\right)_+ \\ \leq v_k\left(|\mathcal{I}_k| + \frac{1}{9}\sum_{i\in\mathcal{B}_k}|\bar{\mathcal{N}}_i|\right)(m_k - d_k)\,. \end{aligned}$$

The voxels in $\bar{\mathcal{H}}_k$ are on a non-uniform grid, giving a biased view of the dose delivered to a structure. But any error from this bias, and from the approximate

nature of our dose-volume constraints, should be negligible if we evaluate plans on 0.4 cm uniform voxels. Actually, regardless of the set of voxels used for the maximum dose constraints, we could apply (2.16) to voxels on a 0.4 cm (or some fine, uniform) grid, which would add more linear constraints to our formulation. But in all of our experiments, all dose-volume requirements are met, so the approximate dose-volume constraints appear to be satisfactory.

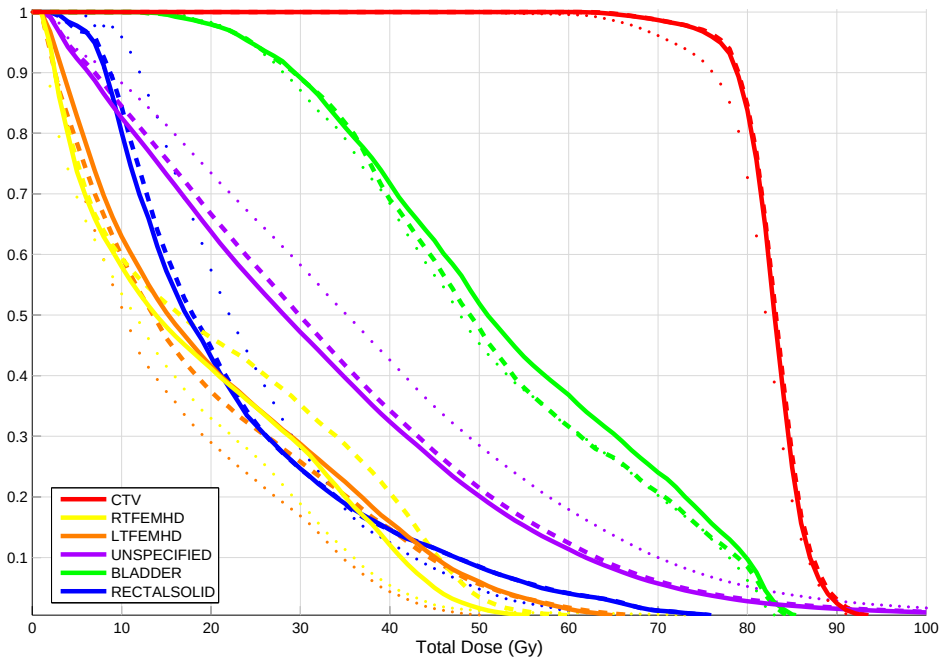

Figure 4.5: DEVH of robust solutions with dose aggregation, solved on voxels of size 1.0 cm except 1.4 cm unspecified region, and 0.8 cm (solid) and 1.0 cm (dotted) CTV voxels with boundary sampling, and 0.8 cm CTV, 1.0 cm bladder with boundary sampling (dashed), all evaluated on 0.4 cm voxels.

Building on our setup from the previous section, we start with 0.8 cm voxels in the CTV and add extra boundary voxels. From the original 191 voxels, 120 are in the interior, and we add 339 more voxels near the boundary, for a total of

530 voxels. The solve time nearly doubles (63 minutes versus 35), and we do not observe much improvement in CTV dose homogeneity compared to the previous solution without boundary sampling. The DEVH of this solution is shown in Figure 4.5 (solid curves), along with the solution from adding boundary voxels in both the CTV and bladder (dashed curves). Very little difference is observed between these two solutions. Also shown (dotted curves) is the solution for the same problem represented by the solid curves, except that we use 1.0 cm CTV voxels with boundary sampling. As expected, the dose homogeneity in the CTV is not as good. Not shown is the solution from adding boundary voxels in the CTV, bladder, and rectal solid, which delivers too much dose to the CTV, femur heads, and unspecified region to be clinically acceptable. This finding of extra boundary voxels not significantly improving the solution quality is consistent with the results in [20].

In conclusion, in Figure 4.6, the DEVH of the solution without boundary sampling is compared to the original robust solutions on 0.8 cm and 1.0 cm voxels from Section 3.2. Among these three solutions, the solution on 0.8 cm voxels (dashed curves) is of the highest quality, but requires the longest solve time, exceeding the others by a factor of three. For the other two solutions, the one using a mixed setup of voxels (solid curves) solves in half the time and delivers comparable dose to the CTV, rectal solid, and bladder (the three regions with highest priority), but delivers more dose to the unspecified region and femur heads.

We observe the typical tradeoff between solution quality and solve time. Overall, we recommend that if time allows, the original robust formulation should be solved on fine voxels. If time is critical, then the robust formulation should be solved on fine CTV voxels and coarser voxels in other structures, particularly the

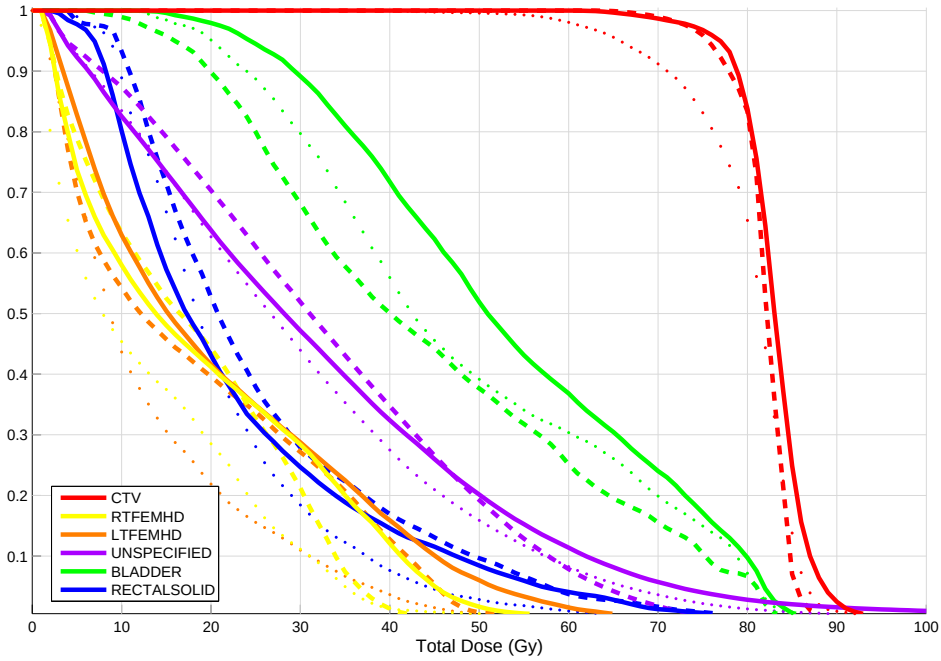

Figure 4.6: DEVH of robust solutions, all evaluated on 0.4 cm voxels. Solid curves: 1.0 cm voxels except 1.4 cm unspecified and 0.8 cm uniform CTV, all voxels with dose aggregation (solve time: 35 minutes). Dotted curves: original robust solution on 1.0 cm voxels (solve time: 62 minutes). Dashed curves: original robust solution on 0.8 cm voxels (solve time: 199 minutes).

unspecified region, with dose aggregation. This method is faster and gives better solutions than simply using coarser voxels in all structures.

Chapter 5

Conclusions

The current method of addressing uncertainty by adding a safety margin around the target region results in treatment plans that irradiate a larger region than necessary, and consequently overdose healthy tissue. This approach ensures that the tumour receives proper dosage, but at the high cost of causing potential damage to nearby organs.

Our computational results show that under uncertainty, IMRT treatment plans derived from a robust formulation should be clinically preferable to those from a deterministic formulation with the safety margin. The benefit of a robust plan is the superior organ sparing, while maintaining an adequately high dose to the tumour.

The drawback of using a robust formulation is the computational requirement. Very fine voxels make it too large, and we run out of memory. With larger voxels, we achieved reasonable solve times (a few hours). We observed that these treatment plans performed well when we looked at the doses delivered to the voxels used in the problem setup, but that some quality is lost when the plan is evaluated on fine voxels.

We proposed methods to reduce computational requirements, yielding faster solve times but with some tradeoff in solution quality. For a prostate case, we observe that dose aggregation helps to achieve better dose homogeneity in the CTV, but at the cost of more dose to healthy structures. Also, extra voxels near the CTV boundary do not improve solution quality significantly. Dose aggregation and varying inter-structure voxel sizes are only recommended when a fast solution

is required, and otherwise, we recommend solving the original robust formulation on the finest sized voxels possible.

Our results are based on experiments with one prostate case, so clearly, further experiments are needed to test our methods. It would be interesting to try other prostate cases, as well as other types of cancers. These cases may provide valuable insight into the effects and performance of the robust model.

Further, more clinical data on organ and tumour motion during IMRT treatment would help us to refine the scenario selection. The shifts are especially critical for ensuring no part of the CTV is under-dosed.

In conclusion, we hope that the promising results from this work will encourage further clinical study and that robust optimization will someday become incorporated into a commercial treatment planning system.

In order to achieve this goal, we believe graphical visualization of the solution, with the potential consequences of motion, is key. We have presented two types of plots, a Dose - Expected Volume Histogram and a Dose - Volume Histogram cloud, both of which acknowledge some uncertainty in voxel location. A sensible, and relatively easily accomplished, feature to add to current treatment planning systems would be an option to evaluate plans under simple shifts of the patient. A more ambitious feature would be an interactive movie, showing doses delivered under potential motion, with an interface that allows the user to select distributions of organ motion.

Appendix A

Computational Implementation

All computations were performed in the MATLAB environment [44]. In this appendix, we first summarize the steps for importing patient and dose data. Then, we describe the MATLAB functions written to create problem inputs, followed by functions for graphing data and solutions.

Our patient case and dose data were exported from the commercial treatment planning system ADAC Pinnacle^3 [1] into RTOG format, and then converted into a MATLAB data file using CERR (A Computational Environment for Radiotherapy Research). CERR was developed by the Advanced Radiotherapy Treatment Planning Group of Washington University and provides "a software platform for developing and sharing research results in radiation therapy treatment planning" [15]. CERR is a free download for academic use and runs in MATLAB. CT scans and dose distributions can also be displayed by CERR, as shown in Figure 3.1.

CERR saves the imported data as a cell array called `planC`. The size of `planC` is approximately 20 MB for the patient data, and 70 MB for the dose data. The MATLAB commands

```
>> load planC     % the patient data
>> structures = planC{planC{end}.structures};
>> clear planC; save structures
```

will save the relevant portions of the patient data in the 3.5 MB struct array `structures`. No structure contours were included for the CTV, so we created our own by subtracting a 1 cm margin from the PTV, and appended it to `structures`.

After using CERR to import the data, we programmed all of the following MATLAB functions, except for `sedumi`. All of our functions are on the attached CD.

To get voxels from `structures`, the function `getVoxelParameters` requires the indices `s` of the desired structures, the voxel size `voxelSize`, and the index `t` of the target structure. For example,

```
>> load structures
>> voxelParamS = getVoxelParameters(structures, s, voxelSize, t);
>> voxelParamS = removeSomeUN(voxelParamS);
>> clear structures; save voxelParamS
```

will create a struct array `voxelParamS` which contains all parameters relating to voxels and structures, e.g. voxel locations, scenario distribution, and dose prescriptions. The function `removeSomeUN` removes the voxels in the Unspecified region that overlap with other healthy structures.

The structures of interest (and their indices) for our prostate case are the rectal solid (3), bladder (5), unspecified region (10), left femur head (11), right femur head (12), CTV (13), and PTV (7).

For the dose data, first run the above commands for voxels with `s = [3 5 10 11 12 13]`, `voxelSize = 0.2`, and `t = 13`. This will make the dose matrix correspond to points on the finest grid possible, 0.2 cm, and cover the appropriate ranges. Then, the commands

```
>> load planC     % the dose data
>> dose = planC{planC{end}.dose};
>> clear planC; save dose
```

```
>> doseParamS = getDoseParameters(dose, voxelParamS);
>> doseParamS = setBeamlets(doseParamS);
>> save doseParamS
```

will create a 1 x 5 struct array `doseParamS`, containing a dose coefficient matrix for each specified beam angle. In the function `getDoseParameters`, the original `dose` data (8 MB) is used to calculate dose coefficients for the range and size of voxels in `voxelParamS`. The function `setBeamlets` turns a set of $B = 940$ beamlets "on" (by default, all are off), and this set can be modified. Only beamlets which are on will be used in the optimization.

The function `getRobust` creates the SeDuMi input variables from `voxelParamS` and `doseParamS`. Optimization parameters such as probability δ, penalty slope ϵ, approximate dose-volume constraint bounds g_k, and penalty weights in the objective function are set in this function.

As an example, the following commands will create a robust problem instance on 1.0 cm voxels, with the formulation as in Section 2.5:

```
>> voxelSize = 1;          % voxel size, in cm
>> t = 13;                 % index of CTV in structures
>> s = [3 5 10 11 12 13];  % indices of ROIs in structures
>> load structures
>> voxelParamS = getVoxelParameters(structures, s, voxelSize, t);
>> voxelParamS = removeSomeUN(voxelParamS);
>> clear structures; load doseParamS
>> [At, b, c, K] = getRobust(voxelParamS, doseParamS);
```

The variables `At, b, c, K` are the input variables for the solver SeDuMi. Se-

DuMi is a free download [55] and more details can be found in [59]. Use the command

```
>> [x, y, info] = sedumi(At,b,c,K);
```

to solve the problem with SeDuMi. The optimal solution `y` contains the optimal beamlet intensities as its first B entries, where B is the number of beamlets used. Often, some entries are slightly negative, so we set these values equal to zero. To calculate the doses delivered by the optimal solution, the commands

```
>> int = y(1:doseParamS(1).totalBeamlets);
>> int = (int>0).*int;
>> bigA = getA(voxelParamS, doseParamS);
>> doseParamS = assignInt(doseParamS, int);
>> voxelParamS = getDoseDelivered(voxelParamS, doseParamS, bigA);
```

saves the doses delivered in `voxelParamS`. The function `getA` outputs a dose coefficient matrix `bigA` for each voxel in each scenario from each beamlet. The solution of beamlet intensities `int` is stored in `doseParamS` by the function `assignInt`. Then, the function `getDoseDelivered` calculates the doses delivered and stores them in `voxelParamS`.

Commonly, we wish to calculate doses delivered to 0.4 cm uniformly spaced voxels instead of the voxels used in the problem input. To save time, compute and save `voxelParamS` and `bigA` corresponding to 0.4 cm voxels and 7 scenarios as in Table 3.2. Then, for a solution of beamlet intensities `int`, call the function

```
>> voxelParamS = getDoseDeliveredOn4mm(int);
```

to get doses delivered to 0.4 cm voxels.

To perform dose aggregation in all structures, as described in Section 4.1, use the commands

```
>> voxelParamS = checkInterior(structures, voxelParamS);
>> S = length(voxelParamS.roiS);
>> for k = 1:S
>> voxelParamS.roiS(k).aggregateInterior = 1;
>> end
```

after `removeSomeUN` and before `getRobust`. The function `checkInterior` will find the subset of voxels that lie in the interior of each structure. To aggregate doses in some but not all structures, simply set `voxelParamS.roiS(k).aggregateInterior` to 1 for dose aggregation in structure `k` and to 0 (default) for no dose aggregation, after running `checkInterior`.

To get extra boundary voxels, as described in Section 4.2, use

```
>> voxelParamS = checkIntExt(structures, voxelParamS);
```

after `removeSomeUN` and before `getRobust`. The function `checkIntExt` will append boundary voxels to the voxels in `voxelParamS`, so it should only be run for the structures in which boundary voxels are desired. For example, to get boundary voxels in the CTV only (the CTV corresponds to `voxelParamS.roiS(6)`), the commands

```
>> CTV = voxelParamS;
>> CTV.roiS([1:5]) = [ ];
>> CTV = checkIntExt(structures, CTV);
>> voxelParamS.roiS(6) = CTV.roiS(1);
```

will append the boundary CTV voxels in `voxelParamS`.

To set up a problem instance in the deterministic formulation, use the commands

```
>> voxelSize = 1;          % voxel size, in cm
>> t = 7;                  % index of PTV in structures
>> s = [3 5 10 11 12 7];   % indices of ROIs in structures
>> load structures
>> voxelParamS = getVoxelParameters(structures, s, voxelSize, t);
>> voxelParamS = removeSomeUN(voxelParamS);
>> load doseParamS
>> [At, b, c, K] = getDet(voxelParamS, doseParamS);
```

and then proceed by calling SeDuMi as for the robust case.

After running `getDoseDelivered` to get doses delivered by a solution of beamlet intensities `int`, we have written several functions that plot figures for plan evaluation.

The function `DEVH(voxelParamS)` plots a DEVH according to the doses and scenarios in `voxelParamS`. The DEVH of more than one solution can be plotted in the same figure, and the line style or colors should be changed in `DEVH` between each call.

To plot DVH clouds, e.g. Figure 3.5, the commands

```
>> N = 45;             % number of treatment fractions
>> bigNumber = 100;    % number of simulated treatments
>> R = rand(N, bigNumber);
```

```
>> plotDVHcloud(voxelParamS, R)
```

will overlay a DVH for `bigNumber` simulated treatments of `N` fractions each. For consistency, the same `R` should be used for the robust and deterministic DVH clouds.

For visualization of the structure contours,

- `plot2dContour` plots 2-dimensional structure contours for a patient slice,
- `plot2dContourByScenario` applies a shift to the contours, and
- `plotStructure` plots 3-dimensional structure contours, as in Figure 3.2.

For pictures of the dose data,

- `plotBeamlet` plots the dose delivered along a beamlet's path for one patient slice,
- `plot3Beamlet` plots the 3-dimensional picture, and
- `plot3ContoursAndBeamlets` adds structure contours.

For analyzing solutions, the function `plotDoseSurface` plots the 3-d surface plot of dose delivered to one patient slice, as in Figure 3.7. To get the 2-d picture as in Figure 3.6, use `plotDoseSurface` but set the options `view([180 270])` and `shading interp`. To view a solution of beamlet intensities `int`, the function `plotIntensityMap` plots each beamlet's intensity, indicated by color.

BIBLIOGRAPHY

[1] ADAC. P^{3}IMRT User Guide. Inverse Planning and IMRT for Pinnacle^3, January 2002.

[2] American Cancer Society. Cancer facts and figures 2005, http://www.cancer.org.

[3] James M. Balter, Howard M. Sandler, Kwok Lam, Robert L. Bree, Allen S. Lichter, and Randall K. Ten Haken. Measurement of prostate movement over the course of routine radiotherapy using implanted markers. *Int. J. Radiation Oncology Biol. Phys.*, 31(1):113–118, 1995.

[4] W. A. Beckham, P. J. Keall, and J. V. Siebers. A fluence-convolution method to calculate radiation therapy dose distributions that incorporate random set-up error. *Phys. Med. Biol.*, 47(19):3465–3473, 2002.

[5] A. Ben-Tal and A. Nemirovski. *Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications*. SIAM-MPS Series in Optimization, Philadelphia, PA, 2001.

[6] Gunilla Carleson Bentel. *Patient Positioning and Immobilization in Radiation Oncology*. McGraw-Hill, New York, NY, 1999.

[7] Jacques Bernier, Eric J. Hall, and Amato Giaccia. Radiation oncology: a century of achievements. *Nature Reviews Cancer*, 4(9):737–747, 2004.

[8] Thomas Bortfeld, Kimmo Jokivarsi, Michael Goitein, Jong Kung, and Steve B. Jiang. Effects of intra-fraction motion on IMRT dose delivery: statistical analysis and simulation. *Phys. Med. Biol.*, 47(13):2203–2220, 2002.

[9] K. K. Brock, D. L. McShan, R. K. Ten Haken, S. J. Hollister, L. A. Dawson, and J. M. Balter. Inclusion of organ deformation in dose calculations. *Med. Phys.*, 30(3):290–295, 2003.

[10] M. Carol, W. H. Grant, D. Pavord, P. Eddy, H. S. Targovnik, B. Butler, S. Woo, J. Figura, V. Onufrey, R. Grossman, and R. Selkar. Initial clinical experience with the Peacock intensity modulation of a 3-D conformal radiation therapy system. *Sterotactic and Functional Neurosurgery*, 66(1-3):30–34, 1996.

[11] Millie Chu, Yuriy Zinchenko, Shane G. Henderson, and Michael B. Sharpe. Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty. *Phys. Med. Biol.*, 50(23):5463–5477, 2005.

[12] C. S. Chui, E. Yorke, and L. Hong. The effects of intra-fraction organ motion on the delivery of intensity-modulated field with a multileaf collimator. *Med. Phys.*, 30(7):1736–1746, 2003.

[13] T. Craig, J. Battista, and J. Van Dyk. Limitations of a convolution method for modeling geometric uncertainties in radiation therapy. I. The effect of shift invariance. *Med. Phys.*, 30(8):2001–2011, 2003.

[14] T. Craig, J. Battista, and J. Van Dyk. Limitations of a convolution method for modeling geometric uncertainties in radiation therapy. II. The effect of a finite number of fractions. *Med. Phys.*, 30(8):2012–2020, 2003.

[15] Joseph Deasy, Konstantine Zakarian, Issam El Naqa, James Alaly, and Patricia E. Lindsay. Cerr: A computational environment for radiotherapy research. http://radium.wustl.edu/CERR/about.php, January 2005.

[16] Kirsten E. I. Deurloo, Roel J. H. M. Steenbakkers, Lambert J. Zijp, Josien A. de Bois, Peter J. C. M. Nowak, Coen R. N. Rasch, and Marcel van Herk. Quantification of shape variation of prostate and seminal vesicles during external beam radiotherapy. *Int. J. Radiation Oncology Biol. Phys.*, 61(1):228–238, 2005.

[17] Matthias Ehrgott. An optimisation model for intensity modulated radiation therapy. *Proceedings of the 37th Annual Conference of the Operational Research Society of New Zealand*, pages 23–31, 2002.

[18] Matthias Ehrgott and Mena Burjony. Radiation therapy planning by multicriteria optimisation. *Proceedings of the 36th Annual Conference of the Operational Research Society of New Zealand*, pages 244–253, 2001.

[19] G. A. Ezzell, J. M. Galvin, D. Low, J. R. Palta, I. Rosen, M. B. Sharpe, P. Xia, Y. Xiao, L. Xing, and C. X. Yu. Guidance document on delivery, treatment planning, and clinical implementation of IMRT: Report of the IMRT subcommittee of the AAPM radiation therapy committee. *Med. Phys.*, 30(8):2089–2115, 2003.

[20] Michael C. Ferris, Rikhardur Einarsson, Ziping Jiang, and David Shepard. Sampling issues for optimization in radiotherapy. Technical Report, Computer Sciences Department, University of Wisconsin. 2004.

[21] Michael C. Ferris, Jinho Lim, and David M. Shepard. Radiosurgery treatment planning via nonlinear programming. *Annals of Operations Research*, 119(1–4):247–260, 2003.

[22] Michael C. Ferris and Meta M. Voelker. Neuro-dynamic programming for radiation treatment planning. Numerical Analysis Group Research Report NA-02/06, 2002.

[23] James M. Galvin, Gary Ezzell, Avraham Eisbrauch, Cedric Yu, Brian Butler, Ying Xiao, Isaac Rosen, Julian Rosenman, Michael Sharpe, Lei Xing, Ping Xia, Tony Lomax, Daniel A. Low, and Jatinder Palta. Implementing IMRT

in clinical practice: a joint document of the American Society for Therapeutic Radiology and Oncology and the American Association of Physicists in Medicine. *Int. J. Radiation Oncology Biol. Phys.*, 58(5):1616–1634, 2004.

[24] Intensity Modulated Radiation Therapy Collaborative Working Group. Intensity-modulated radiotherapy: current status and issues of interest. *Int. J. Radiation Oncology Biol. Phys.*, 51(4):880–914, 2001.

[25] Eric J. Hall. *Radiobiology for the Radiologist.* J. B. Lippincott Company, 1994.

[26] H. W. Hamacher and K. H. Küfer. Inverse radiation therapy planning a multiple objective optimization approach. *Discrete Applied Mathematics*, 118:145–161, 2002.

[27] Allen Holder. Radiotherapy treatment design and linear programming. *Operations Research and Health Care: A Handbook of Methods and Applications (International Series in Operations Research and Management Science)*, pages 741–774, 2004.

[28] David A. Jaffray, Jeffrey H. Siewerdsen, John W. Wong, and Alvaro A. Martinez. Flat-panel cone-beam computed tomography for image-guided radiation therapy. *Int. J. Radiation Oncology Biol. Phys.*, 53(5):1337–1349, 2002.

[29] Steve B. Jiang, Cynthia Pope, Khaled M. Al Jarrah, Jong H. Kung, Thomas Bortfeld, and George T. Y. Chen. An experimental investigation on intrafractional organ motion effects in lung IMRT treatments. *Phys. Med. Biol.*, 48(12):1773–1784, 2003.

[30] P. Keall. 4-Dimensional computed tomography imaging and treatment planning. *Seminars in Radiation Oncology*, 14(1):81–90, 2004.

[31] K. M. Langen and D. T. L. Jones. Organ motion and its management. *Int. J. Radiation Oncology Biol. Phys.*, 50(1):265–278, 2001.

[32] Mark Langer, Eva K. Lee, Joseph O. Deasy, Ronald L. Rardin, and James A. Deye. Operations research applied to radiotherapy, an NCI-NSF-sponsored workshop - February 7-9, 2002. *Int. J. Radiation Oncology Biol. Phys.*, 57(3):762–768, 2003.

[33] Mark Langer and Joseph Leong. Optimization of beam weights under dose-volume restrictions. *Int. J. Radiation Oncology Biol. Phys.*, 13:1255–1260, 1987.

[34] Joseph Lattanzi, Shawn McNeely, Alexandra Hanlon, Indra Das, Timothy E. Schultheiss, and Gerald E. Hanks. Daily CT localization for correcting portal errors in the treatment of prostate cancer. *Int. J. Radiation Oncology Biol. Phys.*, 41(5):1079–1086, 1998.

[35] Eva K. Lee, Tim Fox, and Ian Crocker. Optimization of radiosurgery treatment planning via mixed integer programming. *Med. Phys.*, 27(5):995–1004, 2000.

[36] Eva K. Lee, Tim Fox, and Ian Crocker. Integer programming applied to intensity-modulated radiation therapy treatment planning. *Annals of Operations Research*, 119(1–4):165–181, 2003.

[37] Jonathan G. Li and Lei Xing. Inverse planning incorporating organ motion. *Med. Phys.*, 27(7):1573–1578, 2000.

[38] C. Clifton Ling, Chandra Burman, Chen S. Chui, Gerald J. Kutcher, Stephen A. Leibel, Thomas LoSasso, Radhe Mohan, Thomas Bortfeld, Larry Reinstein, Spiridon Spirou, X. H. Wang, Quiwen Wu, Michael Zelefsky, and Zvi Fuks. Conformal radiation treatment of prostate cancer using inversely-planned intensity-modulated photon beams produced with dynamic multileaf collimation. *Int. J. Radiation Oncology Biol. Phys.*, 35(4):721–730, 1996.

[39] Yu-Ming Liu, Stella Ling, K. M. Langen, Katsuto Shinohara, Vivian Weinberg, Jean Pouliot, and Mack Roach III. Prostate movement during simulation resulting from retrograde urethrogram compared with "natural" prostate movement. *Int. J. Radiation Oncology Biol. Phys.*, 60(2):470–475, 2004.

[40] Johan Löf, Bengt K. Lind, and Anders Brahme. An adaptive control algorithm for optimization of intensity modulated radiotherapy considering uncertainties in beam profiles, patient set-up and internal organ motion. *Phys. Med. Biol.*, 43(6):1605–1628, 1998.

[41] Anthony E. Lujan, Edward W. Larsen, James M. Balter, and Randall K. Ten Haken. A method for incorporating organ motion due to breathing into 3D dose calculations. *Med. Phys.*, 26(5):715–720, 1999.

[42] Thomas Rockwell Mackie, Jeff Kapatoes, Ken Ruchala, Weiguo Lu, Chuan Wu, Gustavo Olivera, Lisa Forrest, Wolfgang Tome, Jim Welsh, Robert Jeraj, Paul Harari, Paul Reckwerdt, Bhudatt Paliwal, Mark Ritter, Harry Keller, Jack Fowler, and Minesh Mehta. Image guidance for precise conformal radiotherapy. *Int. J. Radiation Oncology Biol. Phys.*, 56(1):89–105, 2003.

[43] Alvaro A. Martinez, Di Yan, David Lockman, Donald Brabbins, Kamal Kota, Michael Sharpe, David A. Jaffray, Frank Vicini, and John Wong. Improvement in dose escalation using the process of adaptive radiotherapy combined with three-dimensional conformal or intensity-modulated beams for prostate cancer. *Int. J. Radiation Oncology Biol. Phys.*, 50(5):1226–1234, 2001.

[44] MATLAB Version 7.0.0.19901 (R14). The MathWorks, May 2004.

[45] James G. Mechalakos, Gikas S. Mageras, Michael J. Zelefsky, Olga Lyass, Marcel van Herk, Hanne M. Kooy, Steven A. Leibel, and C. Clifton Ling. Time trends in organ position and volume in patients receiving prostate three-dimensional conformal radiotherapy. *Radiotherapy and Oncology*, 62:261–265, 2002.

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

[47] International Commission on Radiation Units and Measurements. ICRU Report 62: Prescribing, recording and reporting photon beam therapy (supplement to ICRU Report 50), 1999.

[48] H. Paganetti, H. Jiang, and A. Trofimov. 4D Monte Carlo simulation of proton beam scanning: modelling of variations in time and space to study the interplay between scanning pattern and time-dependent patient geometry. *Phys. Med. Biol.*, 50(5):983–990, 2005.

[49] Harald Paganetti, Hongyu Jiang, Judith A. Adams, George T. Chen, and Eike Rietzel. Monte Carlo simulations with time-dependent geometries to investigate effects of organ motion with high temporal resolution. *Int. J. Radiation Oncology Biol. Phys.*, 60(3):942–950, 2004.

[50] Felisa Preciado-Walters, Ronald Rardin, Mark Langer, and Van Thai. A coupled column generation, mixed-integer approach to optimal planning of intensity modulated radiation therapy for cancer. *Mathematical Programming*, 101(2):319–338, 2004.

[51] Andrei Pugachev, Jonathan G. Li, Arthur L. Boyer, Steven L. Hancock, Quynh-Thu Le, Sarah S. Donaldson, and Lei Xing. Role of beam orientation optimization in intensity-modulated radiation therapy. *Int. J. Radiation Oncology Biol. Phys.*, 50(1):551–560, 2001.

[52] Henrik Rehbinder, Camilla Forsgren, and Johan Löf. Adaptive radiation therapy for compensation of errors in patient setup and treatment delivery. *Med. Phys.*, 31(12):3363–3371, 2004.

[53] John C. Roeske, Jeffrey D. Forman, Carmen F. Mesina, Tony He, Charles A. Pelizzari, Ernesto Fontenla, S. Vijayakumar, and George T. Y. Chen. Evaluation of changes in the size and location of the prostate, seminal vesicles, bladder, and rectum during a course of external beam radiation therapy. *Int. J. Radiation Oncology Biol. Phys.*, 33:1321–1329, 1995.

[54] H. Edwin Romeijn, Ravindra K. Ahuja, James F. Dempsey, Arvind Kumar, and Jonathan G. Li. A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning. *Phys. Med. Biol.*, 48:3521–3542, 2003.

[55] SeDuMi. McMaster University, Advanced Optimization Laboratory. http://sedumi.mcmaster.ca/, 2005.

[56] David M. Shepard, Michael C. Ferris, Gustavo H. Olivera, and T. Rockwell Mackie. Optimizing the delivery of radiation therapy to cancer patients. *SIAM Review*, 41(4):721–744, 1999.

[57] Almon S. Shiu, Eric L. Chang, Jin-Song Ye, MingFwu Lii, Laurence D. Rhines, Ehud Mendel, Jeffrey Weinberg, Sanjay Singh, Moshe H. Maor, Radhe Mohan, and James D. Cox. Near simultaneous computed tomography image-guided stereotactic spinal radiotherapy: An emerging paradigm for achieving true stereotaxy. *Int. J. Radiation Oncology Biol. Phys.*, 57(3):605–613, 2003.

[58] J. Stein, R. Mohan, X. H. Wang, T. Bortfeld, Q. Wu, K. Preiser, C. C. Ling, and W. Schlegel. Number and orientations of beams in intensity-modulated radiation treatments. *Med. Phys.*, 24(2):149–160, 1997.

[59] Jos F. Sturm. Using SeDuMi 1.02, a MATLAB toolbox for optimization over symmetric cones. *Optimization Methods and Software. Special issue on Interior Point Methods*, 11-12:625–653, 1999.

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

[61] J. Unkelbach and U. Oelfke. Incorporating organ movements in inverse planning: assessing dose uncertainties by Bayesian inference. *Phys. Med. Biol.*, 50(1):121–139, 2005.

[62] Marcel van Herk, Allison Bruce, Guus Kroes, Tarek Shouman, Adriaan Touw, and Joos V. Lebesque. Quantification of organ motion during conformal radiotherapy of the prostate by three dimensional image registration. *Int. J. Radiation Oncology Biol. Phys.*, 33(5):1311–1320, 1995.

[63] Marcel van Herk, Peter Remeijer, and Joos V. Lebesque. Inclusion of geometric uncertainties in treatment plan evaluation. *Int. J. Radiation Oncology Biol. Phys.*, 52(5):1407–1422, 2002.

[64] S. Webb. *Intensity Modulated Radiation Therapy (Series in Medical Physics)*. The Institute of Physics, Bristol, UK, 2001.

[65] Lawrence M. Wein, Jonathan E. Cohen, and Joseph T. Wu. Dynamic optimization of a linear-quadratic model with incomplete repair and volume-dependent sensitivity and repopulation. *Int. J. Radiation Oncology Biol. Phys.*, 47(4):1073–1083, 2000.

[66] James R. Wong, Lisa Grimm, Minoru Uematsu, Reva Oren, Chee Wai Cheng, Scott Merrick, and Peter Schiff. Image-guided radiotherapy for prostate cancer by CT-linear accelerator combination: prostate movements and dosimetric considerations. *Int. J. Radiation Oncology Biol. Phys.*, 61(2):561–569, 2005.

[67] Di Yan, Frank Vicini, John Wong, and Alvaro Martinez. Adaptive radiation therapy. *Phys. Med. Biol.*, 42(1):123–132, 1997.

[68] Cedric X. Yu, David A. Jaffray, and John W. Wong. The effects of intra-fraction organ motion on the delivery of dynamic intensity modulation. *Phys. Med. Biol.*, 43(1):91–104, 1998.