

Stochastic Programming for Off-line Adaptive Radiotherapy

MUSTAFA Y. SIR^{a,*}, MARINA A. EPELMAN^b, and
STEPHEN M. POLLOCK^b

^a *Department of Industrial and Manufacturing Systems Engineering, University of Missouri, Columbia, MO 65211-2200, USA*

^b *Department of Industrial and Operations Engineering, The University of Michigan, Ann Arbor, MI 48109-2117, USA*

In intensity-modulated radiotherapy (IMRT), a treatment is designed to deliver high radiation doses to tumors, while avoiding the healthy tissue. Optimization-based treatment planning often produces sharp dose gradients between tumors and healthy tissue. Random shifts during treatment can cause significant differences between the dose in the “optimized” plan and the actual dose delivered to a patient. An IMRT treatment plan is delivered as a series of small daily dosages, or fractions, over a period of time (typically 35 days). It has recently become technically possible to measure variations in patient setup and the delivered doses after each fraction. We develop an optimization framework, which exploits the dynamic nature of radiotherapy and information gathering by adapting the treatment plan in response to temporal variations measured during the treatment course of an individual patient. The resulting (suboptimal) control policies, which re-optimize before each fraction, include two approximate dynamic programming schemes: certainty equivalent control (CEC) and open-loop feedback control (OLFC). Computational experiments show that resulting individualized adaptive radiotherapy plans promise to provide a considerable improvement compared to non-adaptive treatment plans, while remaining computationally feasible to implement.

Keywords: adaptive radiotherapy, patient setup error, stochastic programming

* Corresponding author.

1. Introduction

1.1. Background

In 2007, an estimated 1.5 million people were diagnosed with cancer in the United States (Pickle et al., 2007). More than half of cancer patients in the US undergo radiotherapy (the use of radiation as a means for treating cancer) at some point during the course of their disease (Perez and Brady, 1998).

In radiotherapy, radiation is delivered to cancerous regions in order to damage the DNA of the cells in the area being treated, interfering with their ability to divide and grow. As a result, radiation can be used to kill cancerous tumor cells and stop them from regenerating. Since both cancer and healthy cells are affected by radiation (see Stone et al., 2003, for effects of radiotherapy on normal tissue), any treatment plan should be designed in such a way that the radiation dose delivered to the tumor(s) is high enough to stop the cancer cells from regenerating, while simultaneously avoiding the delivery of excessive doses of radiation to surrounding healthy tissue. A “treatment protocol,” prescribed by a radiation oncologist, is a description of the desired dose to be delivered to the various regions of interest.

Intensity-Modulated Radiation Therapy (IMRT), one of the most advanced and common forms of (external beam) radiotherapy, allows a significant amount of control over the characteristics of the radiation delivered by providing flexibility in the shapes of treatment beams, time of exposures, etc. (Webb, 2005). One beam can be represented by a decomposition into a grid of many “beamlets,” each of which can have a different intensity. As a result, radiation treatment plans can potentially conform much more precisely to the treatment protocol compared to traditional conformal radiotherapy techniques. Treatment planning, then, refers to activities involved in finding beamlet intensities which deliver doses of radiation that are as close to those specified by the treatment protocol as possible. These activities usually include repeated solutions of an optimization model embedded in specialized software, which, given a representation of patient anatomy via a collection of voxels (3-dimensional rectangular volume elements) and the treatment protocol, produces a solution in the form of beamlet intensities and resulting dose distribution, *i.e.*, dose that would be delivered to each voxel. In recent years, optimization models used in this step of treatment planning have been an active area of research, and a wide variety of sophisticated models have been proposed (for an overview see, for instance, Shepard et al., 1999; Romeijn

et al., 2003; Lee et al., 2003; Küfer et al., 2003; Kessler et al., 2005); proprietary models are embedded in a number of commercial treatment planning systems (see Petric et al., 2005).

In clinical practice, once the treatment plan is decided on, it is delivered over a period of time as a series of small dosages called “fractions.” The number of fractions in a treatment is determined by the radiation oncologist (typically 5 treatments per week for a period of 4-9 weeks), and is part of the treatment protocol. Fractionation is used to increase tumor control probability and decrease damage to the healthy tissue surrounding the tumor by giving it time to recover (Thames and Hendry, 1987).

1.2. Uncertainty

Many aspects of radiotherapy are subject to uncertainty. The most common uncertainty in radiotherapy is the position of the target(s) (*i.e.*, the tumor(s) and other regions, such as lymph nodes, to be treated), and surrounding patient anatomy, with respect to the treatment beams. This uncertainty leads to delivery errors, *i.e.*, differences between the dose distribution as calculated by a treatment plan and the actual dose distribution delivered to a patient. These errors can be classified as preparation (“systematic”) and execution (“random”) errors. Systematic errors include those occurring due to organ displacement on images and scans taken during the treatment planning session, tissue delineation (identification of voxels belonging to the targets and other organs and tissues), and equipment calibration. Random errors include those caused by movement during radiation delivery, day-to-day variation in the patient setup and equipment, and inherent inaccuracy of the delivery system (van Herk, 2004; Keller et al., 2003). Most existing optimization models used in treatment planning do not account for the uncertainties leading to these errors and typically produce plans containing sharp dose gradients between a tumor and its neighboring healthy tissue. As a result, it is possible, in some fractions, for the tumor to be in a low-dose radiation field while the healthy tissue regions may be exposed to a high-dose radiation field, resulting in complications and possibly failure of the treatment. Therefore, an “optimized” plan that does not take uncertainties into account may easily lose its advantage when it is subject to random shifts relative to a “nominal” position of the target and surrounding patient anatomy (*e.g.*, a static position on the CT scanner) assumed in the treatment planning (Bortfeld

et al., 2004; van Herk, 2004; Lujan et al., 1999).

Immobilization devices are commonly used to reduce uncertainty by helping patients maintain consistent positioning, both during treatment and from fraction to fraction (Bentel, 1999). Moreover, active breathing control devices are used to reduce the effects of respiratory movement (Wong et al., 1999; Koshani et al., 2006). However, even with these aids, discrepancies between the nominal and actual position of the target and surrounding patient anatomy are inevitable.

The traditional treatment planning approach to dealing with uncertainty in radiotherapy is to expand the clinical target volume (CTV) by some amount (called a “margin”) to form a planning target volume (PTV) (ICRU, 1993, 1999). The width of the margin is typically set to be some number k times the standard deviation of the patient setup error and/or the internal organ motion (van Herk, 2004; van Herk et al., 2000; Sir et al., 2006).

More recent approaches used in practice to account for geometrical uncertainties include convolution-based methods (Lujan et al., 1999; Chetty et al., 2003; Beckham et al., 2002) and Multiple Instance of Geometry Approximation (MIGA) (McShan et al., 2006). In convolution-based methods, uncertainties caused by daily setup procedures at the beginning of each treatment fraction, as well as the inter- and intra-fraction motion of the internal organs, are incorporated in dose calculations to calculate the average delivered dose values. MIGA, on the other hand, approximates the distribution of the random patient position by a distribution of a discrete random variable, and performs a dose deposition calculation for each of the possible discrete variable realizations, referred to as “setup instances” (see Section 2.2 for further details on MIGA).

1.3. Adaptive Treatment Planning

As mentioned above, uncertainties in IMRT can lead to significant differences between the dose distribution calculated by an “optimal” treatment plan and the actual dose distribution delivered to a patient. With advanced technology (Schlegel et al., 2006), however, the position of the target and surrounding patient anatomy with respect to treatment beams can be measured during the treatment process using imaging devices such as electronic portal imaging device (EPID) (Herman et al., 2001), ultrasound (Keller et al., 2004), and CT scanner (Wu et al., 2006). Thus, fraction-to-fraction setup variation, long-term organ deformation, organ filling, patient respiration, and/or dose response can be mon-

itored and the treatment plan can be adapted to these measurements during the treatment course of an individual patient.

Many techniques, referred to as Adaptive Radiation Therapy (ART), have been developed to exploit the additional information provided by the imaging devices (Yan et al., 1997a, 2000; Martinez et al., 2000; Löf et al., 1995; Keller et al., 2003; Ferris and Voelker, 2004). In general, Yan (2006) describes ART as a system “designed to systematically manage treatment feedback, planning, and adjustment in response to temporal variations occurring during the treatment course.”

In this paper, we concentrate on “off-line” ART, where measurements from imaging devices become available *after* the delivery of a fraction (Rehbinder et al., 2004; de la Zerda et al., 2007; Ferris and Voelker, 2004). Based on these measurements, treatment is then adjusted for future fractions. Online ART, on the other hand, attempts to adapt the treatment to variations during the delivery of a fraction (Mestrovic et al., 2007; Thongphiew et al., 2008). Unlike online ART, in off-line ART, a more thorough and careful analysis of measurement data can be performed after the delivery. Another advantage of off-line ART is that it does not affect the duration of a fraction, and therefore the patient throughput, since the treatment is not modified during the delivery. In the remainder of this paper, we will use the term ART to mean off-line ART.

There have been two main directions in ART research. Most of the attention has been concentrated on the sequential refinement of knowledge about an individual patient by taking measurements of the patient’s fraction-to-fraction setup variation during the course of the treatment. These methods use well-established techniques such as Kalman filtering (Keller et al., 2003; Yan et al., 1997b) and Bayesian statistical methods (Lam et al., 2005; Sir, 2007) to “update” knowledge for correcting for systematic setup error (Keller et al., 2003; Bortfeld et al., 2002) and create individualized treatment margins (Yan et al., 2000; Keller et al., 2004).

Imaging devices can also be used for reconstructing dose distributions actually *delivered* in a particular fraction (Partridge et al., 2002; McNutt et al., 1997), and therefore highlighting where the delivered treatment may not be as it was planned. This knowledge can then be exploited using a dynamic programming (DP), or stochastic control, framework. Ferris and Voelker (2004) formulate a DP model that attempts to compensate, over time, for movement of the patient and error in the delivery process. In this model, the target dose distribution for the remaining fractions of the treatment course is adjusted after each fraction,

taking into account the total dose that has already been delivered to each voxel (in effect, they modify the treatment protocol to be adhered to in planning for future fractions, but not the treatment plan, *i.e.*, the beamlet intensities).¹ Due to the size and complexity of the DP formulation, they resort to neuro-dynamic programming techniques to derive effective rule-of-thumb policies for modifying the target dose distribution. Wu et al. (2002) proposed a similar re-optimization framework for ART.

Rehbinder et al. (2004), on the other hand, formulate ART as a linear-quadratic (LQ) control problem which has a simple analytical recursive solution, in the form of the discrete-time Ricatti equation (Bertsekas, 2005). However, they do not constrain the beamlet intensities to be non-negative nor consider any additional constraints that may be required by the treatment protocol and/or the type of delivery machine. With these constraints, there is no analytical solution to the LQ control problem (Sir, 2007, Chapter 4).

In this paper, we develop a similar optimization framework that exploits the dynamic nature of radiotherapy and information gathering by adapting the treatment plan in response to variations measured during the treatment course of an individual patient, while incorporating, when appropriate, non-negativity and other constraints. We use and compare two approximate dynamic programming schemes, namely certainty equivalent control (CEC) and open-loop feedback control (OLFC), to find beamlet intensities for each fraction of the treatment. Depending on the scheme used, the optimization problem for finding the beamlet intensities can be posed as a linear programming or stochastic programming problem, the latter solved by a cutting plane algorithm. We analyze the structure of optimization problems arising in computing these policies when arbitrary convex cost functions and convex constraints are used. We then specialize our analysis to the case when the objective and constraint functions are based on equivalent uniform dose (EUD) functions commonly used in assessment of IMRT treatment plans.

The paper is organized as follows: In Section 2 we give the basics of modeling radiation delivery and uncertainty in patient positioning. In Section 3 we describe the objective and constraint functions we use in optimization problems

¹ “Target dose,” *i.e.*, the dose distribution treatment planners aim to deliver in the remaining fractions of the treatment, should not be confused with the target volume, *i.e.*, the tumor(s) and other tissues to be treated to a high dose of radiation.

for finding beamlet intensities. We then discuss how approximate dynamic programming schemes CEC and OLFC are utilized in adaptive radiation treatment planning. In Sections 4 and 5 we formulate optimization problems for CEC and OLFC, respectively, and discuss solution techniques for the resulting problems. Finally, Section 6 presents several numerical examples comparing the performance of different control policies via simulation.

2. Modeling radiation delivery and uncertainty in fractionated IMRT

2.1. Setup position in a fraction

In every fraction, the position of the patient (referred to as the “setup position”) is characterized by the three-dimensional random vector

$$\mathbf{X} \equiv \begin{bmatrix} X \\ Y \\ Z \end{bmatrix}$$

representing the position of an anatomical reference point with respect to an arbitrary fixed origin in the room coordinate system. A realization \mathbf{x} of the random vector corresponds to a shift of the whole set of voxels rigidly, without any rotation or organ deformation, so that the anatomical reference point is located at \mathbf{x} .

\mathbf{X} is typically assumed to have a three-dimensional Normal distribution with probability density function (pdf) given by

$$f_{\mathbf{X}}(\mathbf{x}; \mathbf{m}, \mathbf{v}) = (2\pi)^{-\frac{3}{2}} |\mathbf{v}|^{-\frac{1}{2}} e^{-\frac{1}{2}(\mathbf{x}-\mathbf{m})^T \mathbf{v}^{-1} (\mathbf{x}-\mathbf{m})}. \quad (1)$$

where $\mathbf{m} \in \mathbb{R}^3$ is the mean vector and \mathbf{v} is the 3×3 covariance matrix (Lam et al., 2005; Schewe et al., 1996).

For simplicity, in the remainder of this paper, we assume that there is no systematic error (*i.e.*, $\mathbf{m} = [0\ 0\ 0]^T$). This is not a very strong assumption as most clinics use daily localization and correction procedures primarily aimed at eliminating systematic errors (Lam et al., 2007). Although these procedures along with immobilization and active breathing control devices help reducing random errors, discrepancies between the nominal and actual position of the patient anatomy cannot be eliminated. Sir (2007) discusses applications of Bayesian statistical methods in the presence of both systematic and random errors.

When there are N remaining fractions in the treatment, the setup positions $\mathbf{X}_1, \mathbf{X}_2, \dots, \mathbf{X}_N$ are independent and identically distributed (iid), where \mathbf{X}_n is the setup position in the n^{th} upcoming fraction.

As implicitly assumed above, we consider only random setup errors with rigid body motion in this paper. In theory, one can also include rotations and organ deformations in the model of uncertainty. However, this might significantly complicate the problem because of

- (i) the complexity (or even unavailability) of appropriate stochastic models for changes in patient anatomy due to organ deformation and weight loss; and
- (ii) the fact that realizations of such changes are likely not iid over the course of the treatment.

Conceptually, the approximate dynamic programming framework described in the remainder of this paper can be extended to address these issues, possibly at the expense of substantially increased computation times.

2.2. Dose models, dose deposition matrix and MIGA

Recall that in IMRT optimization the tissue subjected to radiation is represented by a set of voxels, and radiation beams are represented by a set of beamlets. Using physics principles, a dose deposition matrix can be generated by computing how a unit-intensity beamlet will deposit energy to voxels as it travels through the body of the patient. In particular, with v the number of voxels and b the number of beamlets, $\mathbf{D} \in \mathbb{R}_+^{v \times b}$ is the dose deposition matrix with elements d_{ij} defined as the dose delivered to voxel i by beamlet j at unit intensity. Then if $\mathbf{w} \in \mathbb{R}_+^b$ is the vector of beamlet intensities used in a fraction, the vector of doses delivered to the patient is $\mathbf{D}\mathbf{w} \in \mathbb{R}_+^v$. We assume that dose is additive, *i.e.*, dose delivered to a voxel in the course of the entire treatment is computed as the sum of doses delivered to this voxel in each fraction of the treatment. Although dose additivity is a common assumption (Chu et al., 2005; Ferris and Voelker, 2004), other dose models have been proposed to reflect the biological effects of fluctuations in dose delivery over fractions (Bortfeld and Paganetti, 2006). Many aspects of the analysis we present can be extended to these dose models. These extensions, however, fall outside the scope of this paper.

The dose deposition matrix depends on the patient's position with respect to the beams. We therefore define the mapping $\mathbf{D}(\mathbf{x}) : \mathbb{R}^3 \rightarrow \mathbb{R}_+^{v \times b}$ as the dose deposition matrix when the anatomical reference point on the patient is at room

coordinate \mathbf{x} . When the position of the anatomical reference point is the random variable \mathbf{X} , the dose deposition matrix is the random matrix $\mathbf{D}(\mathbf{X})$. Since the calculation of $\mathbf{D}(\cdot)$ is usually a subroutine in the treatment planning system and \mathbf{X} is typically assumed to have a continuous distribution, there is no convenient way to use $\mathbf{D}(\mathbf{X})$ directly in subsequent analysis. Therefore, an approximation for the distribution of $\mathbf{D}(\mathbf{X})$ is in order.

One method, known as the “multiple instance of geometry approximation” (MIGA) (McShan et al., 2006), is to approximate the distribution of setup position in each fraction with a discrete distribution over “setup instances.” Suppose that there are K setup instances used in the approximation, each of which is assigned probability p_k , $k = 1, \dots, K$, of occurring. The k^{th} setup instance is characterized by $\bar{\mathbf{x}}_k \in \mathbb{R}^3$ and corresponds to a shift of the voxels, so that the anatomical reference point is located at $\bar{\mathbf{x}}_k$. Figure 1 shows an example of a MIGA distribution for $K = 7$. The parameters of MIGA are selected to approximate the true distribution of \mathbf{X} (1).

A dose deposition matrix calculation can be performed for each setup instance of MIGA, storing the resulting matrices $\mathbf{D}(\bar{\mathbf{x}}_1), \mathbf{D}(\bar{\mathbf{x}}_2), \dots, \mathbf{D}(\bar{\mathbf{x}}_K)$. Thus the random dose deposition matrix $\mathbf{D}(\mathbf{X})$ is approximated by another random matrix $\hat{\mathbf{D}}$, having the probability mass function (pmf)

$$\hat{\mathbf{D}} = \begin{cases} \mathbf{D}(\bar{\mathbf{x}}_1), & \text{with probability } p_1, \\ \mathbf{D}(\bar{\mathbf{x}}_2), & \text{with probability } p_2, \\ \vdots & \vdots \\ \mathbf{D}(\bar{\mathbf{x}}_K), & \text{with probability } p_K. \end{cases} \quad (2)$$

3. Optimization models in adaptive treatment planning

As mentioned in the introduction, in the process of IMRT treatment planning, an optimization problem is often solved to obtain a treatment plan that optimizes a cost function that serves as a performance measure for comparing feasible plans that meet the constraints of the treatment protocol. A wide variety of optimization models have been proposed in medical and operations research literature. In this section, we will outline the structure of the objective function and constraints used in our models, and conclude the section with an outline of two approximate dynamic programming schemes for ART.




Figure 1. Multiple Instance of Geometry Approximation. The centers of the circles represent the position of each setup instance; the size of the circles and the numbers inside are the probabilities for each setup instance.

3.1. Objective function

Let $\mathbf{d} \in \mathbb{R}_+^v$ be the total dose delivered at the end of the treatment course. The cost function we will use to evaluate treatment plans is of the form

$$f(\mathbf{d}) \equiv \sum_{l=1}^L r_l f_l(\mathbf{d}),$$

where L is the number of treatment criteria considered in the treatment plan, $f_l(\mathbf{d})$, $l = 1, 2, \dots, L$, are convex functions of \mathbf{d} , and $r_l \geq 0$, $l = 1, 2, \dots, L$. Each function $f_l(\mathbf{d})$ represents a criterion used as a measure of performance for a single treatment region. Note that there might be several criteria for each treatment region. Therefore, L might be greater than the number of treatment regions in a treatment plan.

A possible functional form for $f_l(\mathbf{d})$ is defined through an Equivalent Uniform Dose (EUD) function which summarizes a heterogenous dose distribution over a specific treatment region (a target or a healthy organ or tissue) by a single value: the uniform dose that would have the same biological effect on the region. The generalized EUD formula (Niemierko, 1999) is

$$\text{EUD}_a(\mathbf{d}) = \left(\frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} (d_i)^a \right)^{1/a},$$

where \mathcal{S} is the set of voxels of the treatment region, \mathbf{d} is the delivered dose vector,

and a is a region-specific parameter. We make the following observations (Choi and Deasy, 2002):

- As $a \rightarrow \infty$, $\text{EUD}_a(\mathbf{d})$ approaches $\max_{i \in \mathcal{S}} d_i$. This function is suitable for a region of healthy organ tissue with a *serial* structure, in which damage to any part of the organ results in the failure of the entire organ.
- When $a = 1$, $\text{EUD}_a(\mathbf{d})$ is equal to the average dose delivered to all voxels in the treatment region, *i.e.*, $\frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} d_i$. This function is suitable for a region of healthy organ tissue with a *parallel* structure, in which even if some part of the organ is damaged, the undamaged portion still performs the functions of the organ.
- As $a \rightarrow -\infty$, $\text{EUD}_a(\mathbf{d})$ approaches $\min_{i \in \mathcal{S}} d_i$. This function can be utilized in cost functions for tumors.

Notice that $\text{EUD}_a(\cdot)$ defined over $\mathbb{R}_+^{|\mathcal{S}|}$ is convex when $a > 1$, concave when $a < 1$, and linear when $a = 1$; see Choi and Deasy (2002).

An alternative expression for EUD, referred to as “linear EUD” and denoted in the literature by αEUD , was presented in Thieke et al. (2002). αEUD is a convex combination of appropriate generalized EUD functions of Niemierko (1999), namely:

- For tumors,

$$\alpha\text{EUD}(\mathbf{d}) = \alpha \min_{i \in \mathcal{S}} d_i + (1 - \alpha) \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} d_i. \quad (3)$$

- For organs and healthy tissue,

$$\alpha\text{EUD}(\mathbf{d}) = \alpha \max_{i \in \mathcal{S}} d_i + (1 - \alpha) \frac{1}{|\mathcal{S}|} \sum_{i \in \mathcal{S}} d_i. \quad (4)$$

In the above, $\alpha \in [0, 1]$ is a region-specific parameter. Note that the function in (3) is concave in \mathbf{d} , while the function in (4) is convex. These formulas are attractive because they give rise to linear re-optimization problems as demonstrated in the rest of this paper (see also (Craft et al., 2005)).

3.2. Constraints

Constraints typically imposed in treatment planning include:

- Nonnegativity constraints for beamlet intensities:

$$\mathbf{w} \geq 0. \quad (5)$$

- General constraints of the form

$$g_m(\mathbf{d}) \leq 0, \quad m = 1, 2, \dots, M, \quad (6)$$

where the convex functions $g_m(\mathbf{d})$ represent some measures of performance relevant to the outcome of the treatment but not incorporated into the cost function.

In general, we will consider three types of treatment regions: the tumor(s), referred to as the clinical target volume (CTV), organ(s)-at-risk (OAR), and healthy (also called normal) tissue. We denote the set of voxels in the CTV by \mathcal{S}_t , the set of voxels in the OAR by \mathcal{S}_o , and the set of healthy tissue voxels by \mathcal{S}_h . $\mathcal{S} \equiv \mathcal{S}_t \cup \mathcal{S}_o \cup \mathcal{S}_h$ denotes the set of all voxels.

The constraints of (6) include lower and upper bound constraints:

$$l_t \leq d_i \leq u_t, \quad i \in \mathcal{S}_t, \quad (7)$$

$$d_i \leq u_o, \quad i \in \mathcal{S}_o, \quad (8)$$

$$d_i \leq u_h, \quad i \in \mathcal{S}_h, \quad (9)$$

where u_t , u_o and u_h are the upper bounds on the total dose to be delivered to voxels in the CTV, OAR and healthy tissue, respectively, and l_t is the lower bound on the dose to be delivered to voxels in the CTV. In addition, constraints on linear EUDs in different treatment regions are included:

$$\alpha\text{EUD}_t(\mathbf{d}) \equiv \alpha_t \min_{i \in \mathcal{S}_t} d_i + (1 - \alpha_t) \frac{1}{|\mathcal{S}_t|} \sum_{i \in \mathcal{S}_t} d_i \geq L_t, \quad (10)$$

$$\alpha\text{EUD}_o(\mathbf{d}) \equiv \alpha_o \max_{i \in \mathcal{S}_o} d_i + (1 - \alpha_o) \frac{1}{|\mathcal{S}_o|} \sum_{i \in \mathcal{S}_o} d_i \leq U_o, \quad (11)$$

$$\alpha\text{EUD}_h(\mathbf{d}) \equiv \alpha_h \max_{i \in \mathcal{S}_h} d_i + (1 - \alpha_h) \frac{1}{|\mathcal{S}_h|} \sum_{i \in \mathcal{S}_h} d_i \leq U_h, \quad (12)$$

where L_t is the lower bound on αEUD in the CTV, and U_o and U_h are upper bounds on αEUD in the OAR and healthy tissue, respectively.

3.3. Approximate dynamic programming for adaptive radiation treatment

Suppose there are N fractions remaining in the treatment, and $\widehat{\mathbf{d}}$ is the total dose delivered in the prior fractions (we set $\widehat{\mathbf{d}} = 0$ at the beginning of the treatment). Let $\mathbf{w}_1, \dots, \mathbf{w}_N$ be the beamlet weight vectors used in the remaining

fractions. Then the (random) dose delivered at the end of the treatment is given by

$$\mathbf{d} = \hat{\mathbf{d}} + \sum_{n=1}^N \mathbf{D}(\mathbf{X}_n) \mathbf{w}_n.$$

The Adaptive Radiation Therapy problem (with N remaining fractions) can then be stated as the problem of minimizing the expectation

$$E_{\mathbf{x}_1, \dots, \mathbf{x}_N} \left[f \left(\hat{\mathbf{d}} + \sum_{n=1}^N \mathbf{D}(\mathbf{X}_n) \mathbf{w}_n \right) \right]$$

over $\mathbf{w}_1, \dots, \mathbf{w}_N$, subject to the constraints (5) and (6) with the qualification that \mathbf{w}_n is selected with the knowledge of realizations $\mathbf{x}_1, \dots, \mathbf{x}_{n-1}$, and hence the dose $\hat{\mathbf{d}} + \mathbf{D}(\mathbf{x}_1) \mathbf{w}_1 + \dots + \mathbf{D}(\mathbf{x}_{n-1}) \mathbf{w}_{n-1}$.

Two challenges are presented by the above optimization problem. Firstly, the presence of constraints (6) makes the standard dynamic programming algorithm inapplicable. Without these constraints, the dynamic programming algorithm is theoretically applicable, but becomes computationally intractable. To address these challenges we apply and compare the performance of two techniques of approximate dynamic programming: certainty equivalent control and open-loop feedback control (see, for example, Bertsekas, 2005, for details).

- *Certainty equivalent control (CEC):*

For the purposes of optimization, future stochastic fraction-to-fraction setup positions are replaced by a nominal deterministic one (*e.g.*, the MIGA setup instance corresponding to the most likely position of the patient, or to the expected value of the random variable \mathbf{X}). The expression for the total dose delivered is simplified to

$$\mathbf{d} = \hat{\mathbf{d}} + \sum_{n=1}^N \mathbf{D}_{\text{nom}} \mathbf{w} = \hat{\mathbf{d}} + N \mathbf{D}_{\text{nom}} \mathbf{w}.$$

where \mathbf{D}_{nom} is the dose deposition matrix at the nominal setup position.

The optimization problem for determining the beamlet intensities in this scheme becomes a deterministic optimization problem; the details of its formulation are described in Section 4.

- *Open-loop feedback control (OLFC):*

The stochastic nature of the doses to be delivered in future fractions is incorporated into the optimization, but it is assumed that no further measurements

will be taken, and therefore no further adjustment will be made to the plan. The resulting treatment plan design problem becomes that of minimizing the *expected* cost for the remaining fractions of the treatment course, and can be solved using stochastic programming techniques. The expression for the total dose delivered becomes

$$\mathbf{d} = \hat{\mathbf{d}} + \sum_{n=1}^N \mathbf{D}(\mathbf{X}_n) \mathbf{w},$$

or, using MIGA,

$$\mathbf{d} = \hat{\mathbf{d}} + \sum_{n=1}^N \hat{\mathbf{D}}_n \mathbf{w},$$

where $\hat{\mathbf{D}}_n$ is the dose deposition matrix in the n^{th} fraction with pmf given by (2). The optimization problem for determining the beamlet intensities in this scheme becomes a stochastic programming problem; the details of its formulation are described in Section 5.

The adaptive radiation treatment planning proceeds as follows: with N fractions remaining, the dose-to-date $\hat{\mathbf{d}}$ is observed/calculated, and the optimization problem for determining beamlet intensities is formulated in accordance to the approximate dynamic programming technique employed. The beamlet intensities obtained in the optimization are used for treatment in the next fraction, during which the setup position of the patient is observed and the dose-to-date vector is updated. The process is then repeated until the end of the treatment is reached. We will compare the performance of the CEC and OLFC schemes used within ART framework in Section 6.

4. The re-optimization problem for CEC

In CEC, at every future fraction the dose deposition matrix is assumed to be \mathbf{D}_{nom} . The resulting re-optimization problem becomes a deterministic open-loop control problem from the present fraction to the end of the treatment course:

$$\min_{\mathbf{w}, \mathbf{d}} f(\mathbf{d}) \equiv \sum_{l=1}^L r_l f_l(\mathbf{d}) \tag{13}$$

$$\text{subject to } \mathbf{d} = \hat{\mathbf{d}} + N \mathbf{D}_{\text{nom}} \mathbf{w} \tag{14}$$

$$g_m(\mathbf{d}) \leq 0, \quad m = 1, 2, \dots, M \tag{15}$$

$$\mathbf{w} \geq \mathbf{0} \quad (16)$$

Since the functions $f_l(\mathbf{d})$ and $g_m(\mathbf{d})$ are convex in \mathbf{d} and \mathbf{d} is an affine function of \mathbf{w} , the problem (13)-(16) is a convex optimization problem.

In the numerical examples that follow, we use EUD-based cost function

$$f(\mathbf{d}) = -r_t \alpha \text{EUD}_t(\mathbf{d}) + r_o \alpha \text{EUD}_o(\mathbf{d}) + r_h \alpha \text{EUD}_h(\mathbf{d}) \quad (17)$$

and constraints given in equations (7)-(12). Using standard transformation techniques, the resulting optimization problem can be rewritten as

$$\min_{\mathbf{d}, \mathbf{w}, y_t, y_o, y_h, \mu_t, \mu_o, \mu_h} \left[-r_t (\alpha_t y_t + (1 - \alpha_t) \mu_t) + r_o (\alpha_o y_o + (1 - \alpha_o) \mu_o) + r_h (\alpha_h y_h + (1 - \alpha_h) \mu_h) \right] \quad (18)$$

$$\text{subject to } \mathbf{d} = \widehat{\mathbf{d}} + N\mathbf{D}_{\text{nom}} \mathbf{w} \quad (19)$$

$$\mathbf{w} \geq \mathbf{0} \quad (20)$$

CTV constraints:

$$y_t \leq d_i, \quad i \in \mathcal{S}_t \quad (21)$$

$$\mu_t = \frac{1}{|\mathcal{S}_t|} \sum_{i \in \mathcal{S}_t} d_i \quad (22)$$

$$y_t \geq l_t \quad (23)$$

$$d_i \leq u_t, \quad i \in \mathcal{S}_t \quad (24)$$

$$\alpha_t y_t + (1 - \alpha_t) \mu_t \geq L_t \quad (25)$$

OAR constraints:

$$y_o \geq d_i, \quad i \in \mathcal{S}_o \quad (26)$$

$$\mu_o = \frac{1}{|\mathcal{S}_o|} \sum_{i \in \mathcal{S}_o} d_i \quad (27)$$

$$y_o \leq u_o \quad (28)$$

$$\alpha_o y_o + (1 - \alpha_o) \mu_o \leq U_o \quad (29)$$

Healthy tissue constraints:

$$y_h \geq d_i, \quad i \in \mathcal{S}_h \quad (30)$$

$$\mu_h = \frac{1}{|\mathcal{S}_h|} \sum_{i \in \mathcal{S}_h} d_i \quad (31)$$

$$y_h \leq u_h \quad (32)$$

$$\alpha_h y_h + (1 - \alpha_h) \mu_h \leq U_h \quad (33)$$

Note that the optimization problem (18)-(33) is a linear program (LP), which can be solved efficiently.

5. The re-optimization problem for OLFC

In OLFC, the stochastic nature of the doses delivered in future fractions is incorporated into the re-optimization. However, it is assumed that no further measurements of the delivered dose will be taken, and therefore no further adjustment will be made to the beamlet intensities (Sir, 2007, Chapter 4). The goal is to find the beamlet intensity vector which minimizes expected cost from the present fraction to the end of the treatment course and does not violate the constraints under any realization of the dose deposition matrices. This section discusses the resulting re-optimization problem when MIGA with K setup instances is used as the approximate model of uncertainty and how it can be solved using stochastic programming techniques (Birge and Louveaux, 1997; Ruszczyński and Shapiro, 2003; Diwekar et al., 2002).

5.1. Scenarios and the re-optimization problem formulation

To derive the stochastic programming formulation for OLFC, we introduce the concept of a scenario: a scenario represents a sequence of setup positions of the patient in all (remaining) fractions. Using a MIGA model with K setup instances, with N remaining fractions, the number of scenarios is K^N (see Figure 2).

In view of the dose additivity assumption, to calculate the dose delivered to the patient in a particular scenario, it is sufficient to specify the number of times each MIGA setup instance occurs in this scenario. We thus refine the concept of a scenario: a scenario ξ is characterized by K nonnegative integers $N_1(\xi), N_2(\xi), \dots, N_K(\xi)$ such that $\sum_{k=1}^K N_k(\xi) = N$, where $N_k(\xi)$ is the number of times the setup instance k occurs in the remaining N fractions. The probability of this scenario occurring is given by

$$q(\xi) = N! \prod_{k=1}^K \frac{p_k^{N_k(\xi)}}{N_k(\xi)!}.$$

Let Ω be the set of all scenarios defined this way. Table 1 shows elements of Ω for $K = 3$ and $N = 5$. It can be shown that $|\Omega| = \frac{(N+K-1)!}{(K-1)!N!}$, growing significantly slower than K^N as a function of either K or N .



 Figure 2. Illustration of the scenario tree with three setup instances ($K = 3$).

 Table 1
 List of scenarios Ω for $K = 3$ and $N = 5$.

	Scenario index, ξ																				
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20	21
N_1	0	0	0	0	0	0	1	1	1	1	1	2	2	2	2	3	3	3	4	4	5
N_2	0	1	2	3	4	5	0	1	2	3	4	0	1	2	3	0	1	2	0	1	0
N_3	5	4	3	2	1	0	4	3	2	1	0	3	2	1	0	2	1	0	1	0	0

We denote by $\overline{\mathbf{D}}(\xi)$ the *total* dose deposition matrix under scenario ξ :

$$\overline{\mathbf{D}}(\xi) \equiv \sum_{k=1}^K N_k(\xi) \mathbf{D}(\overline{\mathbf{x}}_k),$$

and for every $\xi \in \Omega$, we define

- $\tilde{f}_l(\mathbf{w}; \xi) \equiv f_l(\hat{\mathbf{d}} + \overline{\mathbf{D}}(\xi)\mathbf{w})$, $l = 1, 2, \dots, L$,
- $\tilde{f}(\mathbf{w}; \xi) = \sum_{l=1}^L r_l \tilde{f}_l(\mathbf{w}; \xi)$, and
- $\tilde{g}_m(\mathbf{w}; \xi) \equiv g_m(\hat{\mathbf{d}} + \overline{\mathbf{D}}(\xi)\mathbf{w})$, $m = 1, 2, \dots, M$.

Using the above notation, the re-optimization model for OLFC is a stochastic programming problem given by

$$\min_{\mathbf{w}} \sum_{\xi \in \Omega} q(\xi) \tilde{f}(\mathbf{w}; \xi) \quad (34)$$

$$\text{subject to } \tilde{g}_m(\mathbf{w}; \xi) \leq 0, \quad \forall m, \quad \forall \xi \quad (35)$$

$$\mathbf{w} \geq \mathbf{0}. \quad (36)$$

Note that the functions $\tilde{f}(\mathbf{w}; \xi)$ and $\tilde{g}_m(\mathbf{w}; \xi)$ are convex functions of \mathbf{w} , and thus (34)-(36) is a convex optimization problem.

The number of constraints in (35) can be significantly reduced by considering only *basic* scenarios. We define a *basic* scenario ξ_k^* , $k = 1, 2, \dots, K$, to be one in which the patient will be in setup instance k in all of the remaining fractions of the treatment course (*i.e.*, $N_k(\xi_k^*) = N$, $N_j(\xi_k^*) = 0$ for $j \neq k$). Under this scenario the total dose deposition matrix is given by

$$\bar{\mathbf{D}}(\xi_k^*) \equiv N\mathbf{D}(\bar{\mathbf{x}}_k).$$

For example, in Table 1, the basic scenarios are 1, 6 and 21. Proposition 5.1 shows that problem (34)-(36) is equivalent to

$$\min_{\mathbf{w}} \sum_{\xi \in \Omega} q(\xi) \tilde{f}(\mathbf{w}; \xi) \quad (37)$$

$$\text{subject to } \tilde{g}_m(\mathbf{w}; \xi_k^*) \leq 0, \quad \forall m, \quad \forall k \quad (38)$$

$$\mathbf{w} \geq \mathbf{0}. \quad (39)$$

Proposition 5.1. Optimization problems (34)-(36) and (37)-(39) are equivalent.

Proof of Proposition 5.1. The optimization problems have identical objective functions, and the feasible region of the former is easily seen to be contained in the feasible region of the latter.

Suppose that $\mathbf{w} \geq \mathbf{0}$ satisfies (38). Consider an arbitrary scenario $\xi \in \Omega$ which is characterized by K nonnegative integers N_1, N_2, \dots, N_K such that $\sum_{k=1}^K N_k = N$. Under this scenario, the total dose delivered at the end of the treatment course is equal to

$$\hat{\mathbf{d}} + \sum_{k=1}^K N_k \mathbf{D}(\bar{\mathbf{x}}_k) \mathbf{w}.$$

Then, for any m ,

$$\begin{aligned}
\tilde{g}_m(\mathbf{w}; \xi) &= g_m \left(\hat{\mathbf{d}} + \sum_{k=1}^K N_k \mathbf{D}(\bar{\mathbf{x}}_k) \mathbf{w} \right) \\
&= g_m \left(\sum_{k=1}^K \frac{N_k}{N} \left(\hat{\mathbf{d}} + N \mathbf{D}(\bar{\mathbf{x}}_k) \mathbf{w} \right) \right) \\
&\leq \sum_{k=1}^K \frac{N_k}{N} g_m \left(\hat{\mathbf{d}} + N \mathbf{D}(\bar{\mathbf{x}}_k) \mathbf{w} \right) \tag{40} \\
&= \sum_{k=1}^K \frac{N_k}{N} \tilde{g}_m(\mathbf{w}; \xi_k^*) \leq 0,
\end{aligned}$$

where (40) follows by convexity of g_m . Thus \mathbf{w} satisfies (35). \square

As mentioned above, the framework presented in this paper can be easily extended to more realistic dose models, which take into account the biological effects of fluctuations in dose delivery over fractions. For example, it can be readily shown that the above proposition is still valid if the additive dose model is replaced with a linear-quadratic dose model (Bortfeld and Paganetti, 2006; Thames and Hendry, 1987) under the mild assumptions that

- the fractions are temporally well separated;
- the constraint functions g_m are convex functions of the biologic effect vector and monotonic in each argument (*e.g.*, penalizing the deviations beyond a certain threshold).

5.2. Cutting plane algorithm

The problem in (37)-(39) is a convex optimization problem. The difficulty in solving this problem lies in the functional form of the objective function, which can become intractable when the number of scenarios is large. Therefore, we will solve this problem by employing a cutting plane approach. In the remainder of this section, we will closely follow the terminology and notation of Ruszczyński and Shapiro (2003). They present a general cutting plane algorithm similar to Kelley's cutting plane method (Kelley, 1960) in which a convex optimization problem is approximated by a *master program* where

- the objective function is a convex piece-wise linear function which is a lower bound on the objective function of the original problem, and

- the feasible region is a polyhedron which contains the feasible region of the original problem.

The master problem is solved using linear programming techniques, and its solution is used to generate *cuts*, *i.e.*, linear inequality constraints that improve the approximation by either tightening the lower bound on the objective function, or reducing the size of the polyhedron containing the feasible region. The cuts are added to the master problem, and the process is repeated. In solving the problem (37)-(39), only objective function cuts will be generated, since the feasible region will be represented exactly in the master problem.

We denote the set of all subgradients of a convex function $h : \mathbb{R}^n \rightarrow \mathbb{R}$ at a point $\mathbf{x}_0 \in \mathbb{R}^n$ (also called the subdifferential) by $\partial h(\mathbf{x}_0)$, that is,

$$\mathbf{z} \in \partial h(\mathbf{x}_0) \Leftrightarrow h(\mathbf{x}) - h(\mathbf{x}_0) \geq \mathbf{z}^\top(\mathbf{x} - \mathbf{x}_0), \forall \mathbf{x} \in \mathbb{R}^n.$$

For example, if $h(\cdot)$ is differentiable, its gradient at \mathbf{x}_0 is a unique element of the subdifferential, *i.e.*, $\partial h(\mathbf{x}_0) = \{\nabla h(\mathbf{x}_0)\}$.

Suppose that \mathbf{w}' satisfies (38)-(39) and let $\mathbf{z}_l(\mathbf{w}', \xi) \in \partial \tilde{f}_l(\mathbf{w}'; \xi)$. Then, from the definition of subgradient, for any \mathbf{w} ,

$$\sum_{l=1}^L r_l \tilde{f}_l(\mathbf{w}; \xi) \geq \sum_{l=1}^L r_l \tilde{f}_l(\mathbf{w}'; \xi) + \sum_{l=1}^L [r_l \mathbf{z}_l(\mathbf{w}'; \xi)]^\top (\mathbf{w} - \mathbf{w}'),$$

i.e., $\mathbf{z}(\mathbf{w}'; \xi) \equiv \sum_{l=1}^L r_l \mathbf{z}_l(\mathbf{w}'; \xi) \in \partial \tilde{f}(\mathbf{w}; \xi)$. Therefore, for any \mathbf{w} ,

$$\sum_{\xi \in \Omega} q(\xi) \tilde{f}(\mathbf{w}; \xi) \geq \gamma(\mathbf{w}') + \mathbf{z}(\mathbf{w}')^\top \mathbf{w}, \quad (41)$$

where

$$\mathbf{z}(\mathbf{w}') \equiv \sum_{\xi \in \Omega} q(\xi) \mathbf{z}(\mathbf{w}'; \xi), \quad (42)$$

$$\gamma(\mathbf{w}') \equiv \sum_{\xi \in \Omega} q(\xi) \tilde{f}(\mathbf{w}'; \xi) - \mathbf{z}(\mathbf{w}')^\top \mathbf{w}'. \quad (43)$$

Suppose \mathbf{w}^j , $j = 1, \dots, \bar{j}$ are iterates of the algorithm. Inequality (41) allows us to construct the following linear programming formulation of a master program, in which we approximate the objective function with a convex piecewise linear function:

$$\min_{\mathbf{w}, \theta} \theta \quad (44)$$

$$\text{subject to } \mathbf{z}(\mathbf{w}^j)^\top \mathbf{w} + \gamma(\mathbf{w}^j) \leq \theta, \quad j = 1, \dots, \bar{j}, \quad (45)$$

$$\tilde{g}_m(\mathbf{w}; \xi_k^*) \leq 0, \quad \forall m, \quad \forall k, \quad (46)$$

$$\mathbf{w} \geq \mathbf{0}. \quad (47)$$

The cutting plane algorithm for problem (37)-(39) can then be summarized as follows:

Cutting Plane Algorithm

Initialization: Set $j = 1$, $\theta^1 = -\infty$, and let \mathbf{w}^1 be an arbitrary feasible solution of (37)-(39) (e.g., $\mathbf{w}^1 = \mathbf{0}$).

Iteration $j \geq 1$:

Step 1 If $\sum_{\xi \in \Omega} q(\xi) \tilde{f}(\mathbf{w}^j; \xi) = \theta^j$, stop; \mathbf{w}^j is an optimal solution of (37)-(39).

Step 2 Compute $\mathbf{z}^j \equiv \mathbf{z}(\mathbf{w}^j)$ and $\gamma^j \equiv \gamma(\mathbf{w}^j)$ via (42) and (43) and add the cut

$$(\mathbf{z}^j)^\top \mathbf{w} + \gamma^j \leq \theta$$

to the master program (44)-(47).

Step 3 Find a solution $(\mathbf{w}^{j+1}, \theta^{j+1})$ of the master problem. Set $j \leftarrow j + 1$, and go to **Step 1**.

5.3. Special case: stochastic re-optimization model based on linear EUD functions

We now restrict the OLFC re-optimization problem and the cutting plane algorithm of Subsections 5.1 and 5.2 to the case when the cost function and the constraints are as described in Subsections 3.1 and 3.2.

Let us define

$$\begin{aligned} \tilde{f}_t(\mathbf{w}; \xi) &\equiv -\alpha \text{EUD}_t(\hat{\mathbf{d}} + \overline{\mathbf{D}}(\xi) \mathbf{w}), \\ \tilde{f}_o(\mathbf{w}; \xi) &\equiv \alpha \text{EUD}_o(\hat{\mathbf{d}} + \overline{\mathbf{D}}(\xi) \mathbf{w}), \\ \tilde{f}_h(\mathbf{w}; \xi) &\equiv \alpha \text{EUD}_h(\hat{\mathbf{d}} + \overline{\mathbf{D}}(\xi) \mathbf{w}), \end{aligned}$$

where the linear EUD functions are as defined in (10)-(12). The stochastic re-optimization model is then as follows:

$$\min_{\mathbf{w}, \mathbf{d}(\xi_1^*), \dots, \mathbf{d}(\xi_K^*)} \sum_{\xi \in \Omega} q(\xi) \left[r_t \tilde{f}_t(\mathbf{w}; \xi) + r_o \tilde{f}_o(\mathbf{w}; \xi) + r_h \tilde{f}_h(\mathbf{w}; \xi) \right]$$

$$\begin{aligned}
& \text{subject to } \mathbf{d}(\xi_k^*) = \widehat{\mathbf{d}} + \overline{\mathbf{D}}(\xi_k^*)\mathbf{w}, \quad \forall k \\
& l_t \leq d_i(\xi_k^*) \leq u_t, \quad \forall i \in \mathcal{S}_t, \forall k \\
& d_i(\xi_k^*) \leq u_o, \quad \forall i \in \mathcal{S}_o, \forall k \\
& d_i(\xi_k^*) \leq u_h, \quad \forall i \in \mathcal{S}_h, \forall k \\
& \alpha \text{EUD}_t(\mathbf{d}(\xi_k^*)) \geq L_t, \quad \forall k \\
& \alpha \text{EUD}_o(\mathbf{d}(\xi_k^*)) \leq U_o, \quad \forall k \\
& \alpha \text{EUD}_h(\mathbf{d}(\xi_k^*)) \leq U_h, \quad \forall k \\
& \mathbf{w} \geq \mathbf{0}.
\end{aligned}$$

To apply the subgradient algorithm to the above problem we need to derive the expression for the subgradient of the objective function at the point \mathbf{w} . For $\xi \in \Omega$, let

$$\begin{aligned}
i_t & \in \operatorname{argmin}_{i \in \mathcal{S}_t} \left\{ \widehat{d}_i + \sum_{k=1}^K N_k(\xi) (\mathbf{d}_{i.}^k)^\top \mathbf{w} \right\}, \\
i_o & \in \operatorname{argmax}_{i \in \mathcal{S}_o} \left\{ \widehat{d}_i + \sum_{k=1}^K N_k(\xi) (\mathbf{d}_{i.}^k)^\top \mathbf{w} \right\}, \\
i_h & \in \operatorname{argmax}_{i \in \mathcal{S}_h} \left\{ \widehat{d}_i + \sum_{k=1}^K N_k(\xi) (\mathbf{d}_{i.}^k)^\top \mathbf{w} \right\},
\end{aligned}$$

where $(\mathbf{d}_{i.}^k)^\top$ denotes the i^{th} row of the matrix $\mathbf{D}(\mathbf{x}_k)$.

Then, it can be readily shown that

$$\begin{aligned}
\mathbf{z}_t(\mathbf{w}; \xi) & \equiv -\alpha_t \left[\sum_{k=1}^K N_k(\xi) \mathbf{d}_{i_t.}^k \right] - (1 - \alpha_t) \left[\frac{1}{|\mathcal{S}_t|} \sum_{i \in \mathcal{S}_t} \sum_{k=1}^K N_k(\xi) \mathbf{d}_{i.}^k \right], \\
\mathbf{z}_o(\mathbf{w}; \xi) & \equiv \alpha_o \left[\sum_{k=1}^K N_k(\xi) \mathbf{d}_{i_o.}^k \right] + (1 - \alpha_o) \left[\frac{1}{|\mathcal{S}_o|} \sum_{i \in \mathcal{S}_o} \sum_{k=1}^K N_k(\xi) \mathbf{d}_{i.}^k \right], \\
\mathbf{z}_h(\mathbf{w}; \xi) & \equiv \alpha_h \left[\sum_{k=1}^K N_k(\xi) \mathbf{d}_{i_h.}^k \right] + (1 - \alpha_h) \left[\frac{1}{|\mathcal{S}_h|} \sum_{i \in \mathcal{S}_h} \sum_{k=1}^K N_k(\xi) \mathbf{d}_{i.}^k \right]
\end{aligned}$$

are subgradients of $f_t(\mathbf{w}; \xi)$, $f_o(\mathbf{w}; \xi)$ and $f_h(\mathbf{w}; \xi)$, respectively, and $\mathbf{z}(\mathbf{w}; \xi) \equiv r_t \mathbf{z}_t(\mathbf{w}; \xi) + r_o \mathbf{z}_o(\mathbf{w}; \xi) + r_h \mathbf{z}_h(\mathbf{w}; \xi)$ is a subgradient of the objective function.

Finally, using standard transformation techniques and introducing auxiliary variables $y_t(\xi_k^*)$, $y_o(\xi_k^*)$ and $y_h(\xi_k^*)$, the master problem of the cutting plane algorithm can be written as the following LP:

$$\min \theta \tag{48}$$

$$\text{subject to } (\mathbf{z}^j)^\top \mathbf{w} + \gamma^j \leq \theta, \quad \forall j \tag{49}$$

$$\mathbf{d}(\xi_k^*) = \hat{\mathbf{d}} + \overline{\mathbf{D}}(\xi_k^*) \mathbf{w}, \quad \forall k \tag{50}$$

$$y_t(\xi_k^*) \leq d_i(\xi_k^*), \quad \forall i \in \mathcal{S}_t, \forall k \tag{51}$$

$$l_t \leq y_t(\xi_k^*), \quad \forall k \tag{52}$$

$$d_i(\xi_k^*) \leq u_t, \quad \forall i \in \mathcal{S}_t, \forall k \tag{53}$$

$$\alpha_t y_t(\xi_k^*) + (1 - \alpha_t) \frac{1}{|\mathcal{S}_t|} \sum_{i \in \mathcal{S}_t} d_i(\xi_k^*) \geq L_t, \quad \forall k \tag{54}$$

$$y_o(\xi_k^*) \geq d_i(\xi_k^*), \quad \forall i \in \mathcal{S}_o, \forall k \tag{55}$$

$$y_o(\xi_k^*) \leq u_o, \quad \forall k \tag{56}$$

$$\alpha_o y_o(\xi_k^*) + (1 - \alpha_o) \frac{1}{|\mathcal{S}_o|} \sum_{i \in \mathcal{S}_o} d_i(\xi_k^*) \leq U_o, \quad \forall k \tag{57}$$

$$y_h(\xi_k^*) \geq d_i(\xi_k^*), \quad \forall i \in \mathcal{S}_h, \forall k \tag{58}$$

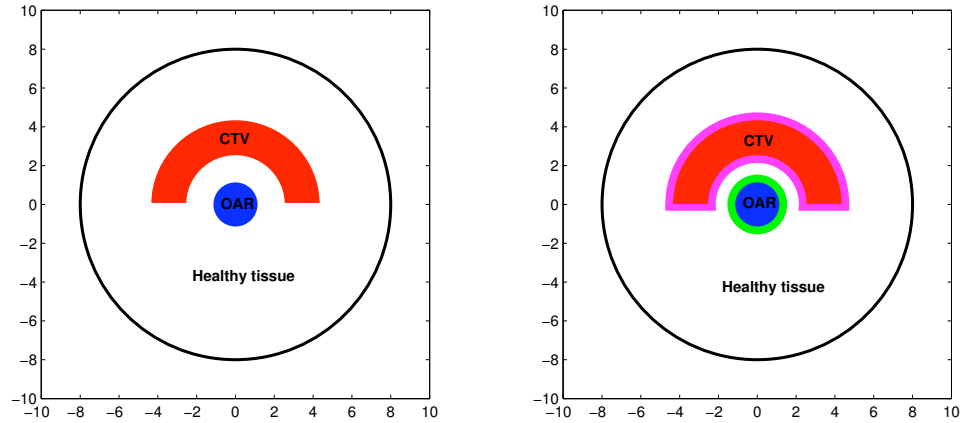
$$y_h(\xi_k^*) \leq u_h, \quad \forall k \tag{59}$$

$$\alpha_h y_h(\xi_k^*) + (1 - \alpha_h) \frac{1}{|\mathcal{S}_h|} \sum_{i \in \mathcal{S}_h} d_i(\xi_k^*) \leq U_h, \quad \forall k \tag{60}$$

$$\mathbf{w} \geq \mathbf{0}. \tag{61}$$

6. Numerical examples

It is of obvious interest to compare the performance of CEC and OLFC, the most practical way to do so being via simulation. It would have been clearly preferable to test the performance of these policies on three-dimensional real patient cases to validate the applicability of our approach to actual clinical use. However, since our main goal in this paper is to clearly show the benefits of adaptive radiation therapy without cluttering the results with complex geometries and trade-offs that generally exist in three-dimensional real patient cases, we use as an example a simple two-dimensional geometry. It consists of a horseshoe-shaped CTV and circular OAR surrounded by circular healthy tissue (see Figure 3(a)). The radius of the region under consideration, R_p , is set to 8 cm. The



(a) Geometry without a margin, used for OLFC. (b) Geometry with a margin, used for CEC.

Figure 3. Simple geometries with and without a margin. The radius of the area being treated (R_p) is 8 cm. The radius of the circular OAR, R_1 , is 1.1 cm. The inner and outer radii of the horseshoe-shaped CTV are $R_2 = 2.5$ cm and $R_3 = 4.3$ cm, respectively.

radius of the circular OAR, R_1 , is 1.1 cm. The inner and outer radii of the horseshoe-shaped CTV are $R_2 = 2.5$ cm and $R_3 = 4.3$ cm, respectively.

There are 5025 voxels, each of which is a $0.2 \text{ cm} \times 0.2 \text{ cm}$ square. 5 treatment beams are used at angles of 15° , 90° , 165° , 225° , and 315° . There are 100 beamlets (20 per beam), each having a width of 0.5 cm. Dose matrix calculation is based on a triple Gaussian pencil beam model for photon beams developed by Ulmer and Harder (1995).

The distribution of random setup position in each fraction is assumed to be a two-dimensional Normal distribution with a zero mean vector (*i.e.*, there is no systematic error) and covariance matrix

$$\mathbf{v} = \begin{bmatrix} 0.4 & 0 \\ 0 & 0.4 \end{bmatrix},$$

where the units of setup position are centimeters. In other words, the standard deviations along x- and y-axes are 0.63. Using a popular margin recipe of $2.5\Sigma + 0.7\sigma$ (van Herk et al., 2000), where Σ and σ represent the standard deviations of systematic and random errors, respectively, this kind of random setup error would be accounted for by expanding the CTV by a margin of $2.5 \times 0 + 0.7 \times 0.63 \approx 0.4$ cm (recall that, in this paper, we assume that there are no systematic errors) in clinical practice.




Figure 4. MIGA with 5 setup instances. Each point represent the position of an instance. The center instance is the nominal position of the patient. The instances around the center represent extreme positions (0.4 cm off of the nominal position along each axis direction) that the patient might take in different fractions. The probabilities of instances are given along side the points.

A MIGA model with 5 setup instances, shown in Figure 4, is used as the *approximate* model of uncertainty for the OLFC re-optimization problem. The setup instances include the nominal position of the patient (the point in the center) and four others generated by translating the whole patient by the size of the margin found above (*i.e.*, 0.4 cm) along each axis (McShan et al., 2006). The probabilities of instances are given along side the points.

The constraints in the OLFC re-optimization problem of Section 5 are enforced for every setup instance in the MIGA model (see Proposition 5.1). Since the MIGA setup instances in Figure 4 translate the patient 0.4 cm in each axis direction, in order to make a comparison between CEC and OLFC more meaningful, we will include such a margin to expand both the CTV and the OAR for the CEC re-optimization (see Figure 3(b)).

Table 2 shows the bounds used in the constraints of the treatment plan. It should be noted that the OLFC (48)-(61) only enforces total dose and EUD constraints for all MIGA scenarios. Under the *true* (Normal) distribution of the random setup position, there is still a non-zero probability of violating some of these constraints. For this reason, the positions and probabilities of the MIGA setup instances should be selected to capture enough of the “spread” of the underlying distribution to enforce the constraints for the most probable error realizations. This is also true for the CEC (18)-(33), where margins replace MIGA to account for uncertainty. Our choice of MIGA parameters and margin size used

Table 2

The bounds used in the constraints of the treatment plan.

Quantity	Lower bound	Upper bound
CTV dose	$l_t = 95$	$u_t = 120$
CTV EUD	$L_t = 95$	-
OAR dose	-	$u_o = 120$
OAR EUD	-	$U_o = 120$
Healthy tissue dose	-	$u_h = 110$
Healthy tissue EUD	-	$U_h = 105$

in the CEC re-optimization are based on values used in the literature (McShan et al., 2006; van Herk et al., 2000). Optimal selection of these parameters is an ongoing research project and beyond the scope of this paper.

Table 3 shows the parameters of the linear EUD and relative weight coefficients.

Table 3

Linear EUD parameters and relative weights for different treatment regions.

Treatment region	α	r
CTV	$\alpha_t = 0.8$	$r_t = 0$
OAR	$\alpha_o = 0.8$	$r_o = 10$
Healthy tissue	$\alpha_h = 0.5$	$r_h = 1$

Each treatment plan was evaluated against a common set of simulated errors for 30 replications of treatment courses, each consisting of 10 fractions. The setup errors in these simulations were generated from the “true” two-dimensional normal distribution given above.

The CEC re-optimization problem and the subproblems of the cutting plane algorithm for the OLFC re-optimization were solved using CPLEX 11 on a 2.6 GHz Sun workstation with 2 GB of memory. While the CEC re-optimization problem, and each of the subproblems of the cutting plane algorithm, took less than a second of CPU time, the entire cutting plane algorithm to solve an OLFC re-optimization problem with 10 fractions and 5 MIGA setup instances (*i.e.*, a stochastic programming problem with 1001 scenarios) took 820 seconds of CPU time. It should be noted that the running time of the cutting plane algorithm can be improved substantially by using more efficient data structures and by

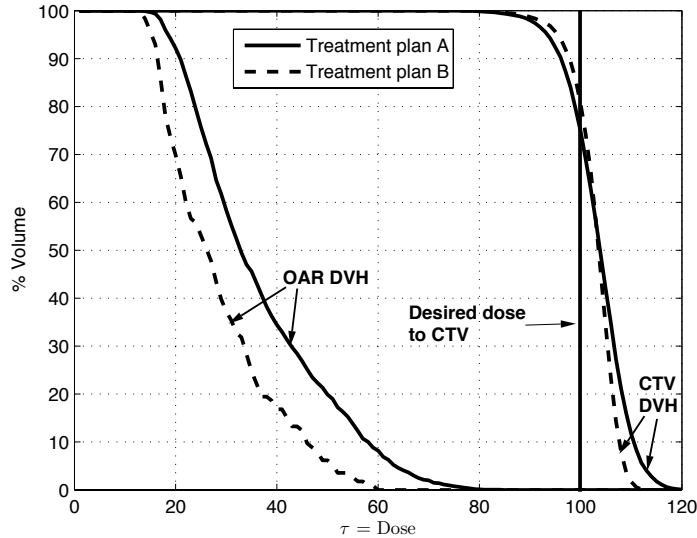


Figure 5. Illustration of how dose-volume histograms are used in comparing and evaluating treatment plans: two treatment plans, A and B, are compared using DVHs.

parallelizing many aspects of the subgradient generation. Implementing a multi-cut version (Birge and Louveaux, 1997) of the cutting plane algorithm could also improve the running times. Finally, sampling techniques such as sample average approximation (Ruszczynski and Shapiro, 2003, Chapter 6) might be used to handle large number of scenarios for real three-dimensional patient cases.

To compare the quality of treatment plans various treatment metrics can be used, such as the value of the resulting EUD as well as the minimum, mean, and maximum dose to different treatment regions. However, physicians usually rely on graphical tools to give a more holistic picture of the treatment plan. The Dose-Volume Histogram (DVH) is the most common of such tools.

A DVH plots the percentage of a treatment region \mathcal{S}_r that receives at least some dose τ . This curve is readily obtained, given delivered dose $\mathbf{d} \in \mathbb{R}_+^u$ to a patient, by using

$$\text{DVH}_r(\tau; \mathbf{d}) = \frac{1}{|\mathcal{S}_r|} \sum_{i \in \mathcal{S}_r} I_{[d_i \geq \tau]},$$

where $I_{[\cdot]}$ is the indicator function.

Figure 5 compares the DVHs for two treatment plans, A and B. As can be seen, the CTV DVH for treatment plan B is steeper than that for treatment plan

A; the dose delivered to CTV by treatment B is more homogenous and produces fewer cold spots. Moreover, the OAR DVH for treatment plan B is further to the left compared to that for treatment plan A, which means that treatment plan B delivers less radiation to OAR than treatment plan A.

Based on the above observations, most physicians would consider treatment plan B to be superior to treatment plan A. However, it should be noted that comparisons of realistic treatment plans are never as clear-cut, since usually each plan has both advantageous and disadvantageous aspects. To make comparisons simple in the numerical experiments below, more importance is given to CTV coverage and OAR-sparing. The sparing of healthy tissue is of secondary importance. Note that this is consistent with our choice of constraints and objective function, parameters of which are specified in Table 2 and 3, respectively.

In the following sections, the performance of several adaptive radiation therapy policies is compared using various treatment metrics and DVHs. All metrics and DVHs reported in the following sections are calculated based on the original structures, CTV, OAR, and healthy tissue, without margins.

6.1. Adaptive vs. Non-adaptive

As discussed in Section 1.3, in both CEC and OLFC, a re-optimization is performed before each fraction, *adapting* the treatment plan to the most current information (*i.e.*, the total dose delivered in previous fractions). In practice, however, only one optimization is performed before the treatment starts and the resulting beamlet intensity vector is used throughout the treatment course (we refer to such treatment schemes as *non-adaptive*). In this section, we compare the *adaptive* schemes of CEC and OLFC to their corresponding *non-adaptive* schemes.

Figure 6 compares the performance of adaptive and non-adaptive scheme of CEC, both of which use a margin size of 0.4 cm; Table 4 lists summary statistics of several important metrics along with the result of significance tests on the difference in the mean values of these metrics at significance level 0.05. As can be seen from Table 4, the adaptive scheme of CEC covers the CTV better than the non-adaptive scheme. However, the adaptive scheme also delivers significantly higher amounts of radiation dose to the healthy tissue and the CTV than the non-adaptive scheme. As can be seen from Figure 6, in certain simulated treatment courses, some CTV voxels receive more than 130 gray of radiation dose (recall

Table 4

Comparison of basic metrics for adaptive and non-adaptive schemes of CEC. Both policies use a margin size of 0.4 cm. Comparisons are based on 30 replications of 10-fraction treatment. The difference in the mean values of the metrics are evaluated using significance tests at level 0.05.

	Statistic	Adaptive	Non-adaptive	Significance Test
Max total OAR dose	Min	40.04	41.92	NSE
	Avg	49.82	50.61	
	Max	58.91	62.95	
	Std	4.73	5.31	
Avg total OAR dose	Min	25.14	29.50	NSE
	Mean	34.08	34.40	
	Max	41.81	39.06*	
	Std	4.36	2.47	
OAR total EUD	Min	37.06*	39.80	NSE
	Mean	46.67	47.36	
	Max	54.50	56.94	
	Std	4.56	4.48	
Max total healthy tissue dose	Min	112.69	105.20	Non-adaptive
	Mean	116.96	109.34	
	Max	127.75	113.92	
	Std	2.83	2.33	
Avg total healthy tissue dose	Min	58.12	56.61*	Non-adaptive
	Mean	60.57	57.07	
	Max	62.62	57.65	
	Std	1.06	0.22	
Healthy tissue total EUD	Min	86.53	81.16	Non-adaptive
	Avg	88.76	83.21	
	Max	94.69	85.78	
	Std	1.61	1.17	
Min total CTV dose	Min	87.75	85.12	Adaptive
	Avg	94.94	91.82	
	Max	99.35	98.27	
	Std	2.24	3.65	
CTV total EUD	Min	92.22	88.94	Adaptive
	Avg	98.05	94.48	
	Max	101.86	99.98	
	Std	1.86	3.05	

NSE = No Strong Evidence that the mean values of the metrics are different.

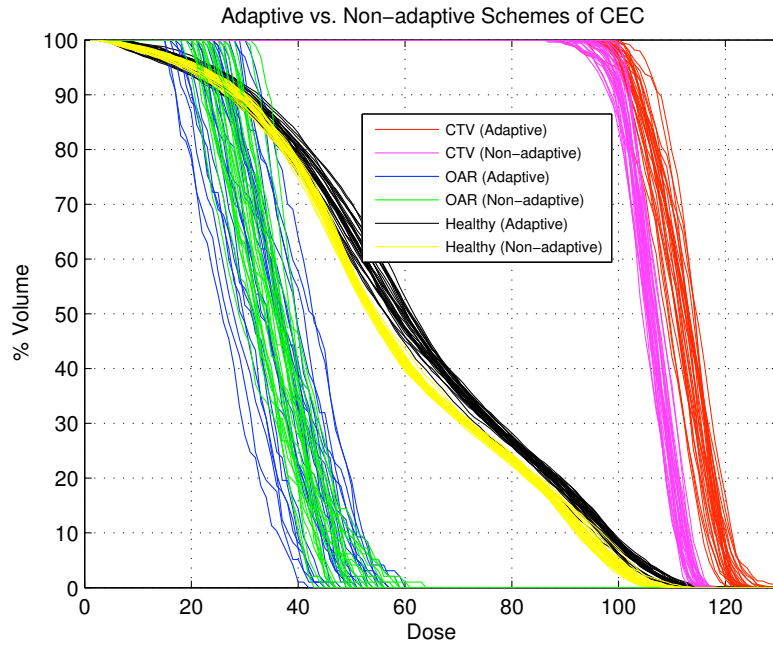


Figure 6. Comparison of DVHs, resulting from 30 replications of 10-fraction treatment, for adaptive and non-adaptive schemes of CEC. Both policies use a margin size of 0.4 cm.

that the upper bound on CTV dose was set to 120 gray). Table 4 shows that the adaptive scheme of CEC delivers as much as 127.75 gray to the healthy tissue, which is only 113.92 grays for the non-adaptive scheme (recall that the upper bound on healthy tissue dose was set to 110 gray).

Table 4 also shows that both the adaptive and non-adaptive schemes achieve similar levels of OAR-sparing.

In conclusion, the adaptive scheme of CEC does not provide a significant improvement over the non-adaptive scheme. Even though CTV coverage is seemingly better for the adaptive scheme, it delivers significantly higher amount of radiation to the patient than the non-adaptive scheme.

One possible explanation of this is that the adaptive scheme of CEC is a reactive policy, which tries to compensate dose delivery errors in previous fractions at every re-optimization. However, since its re-optimization model is based on the assumption that the patient will be at the nominal position in future fractions, its reaction to delivery errors in previous fractions can be excessive. The non-adaptive scheme of CEC, on the other hand, does not suffer from this and,

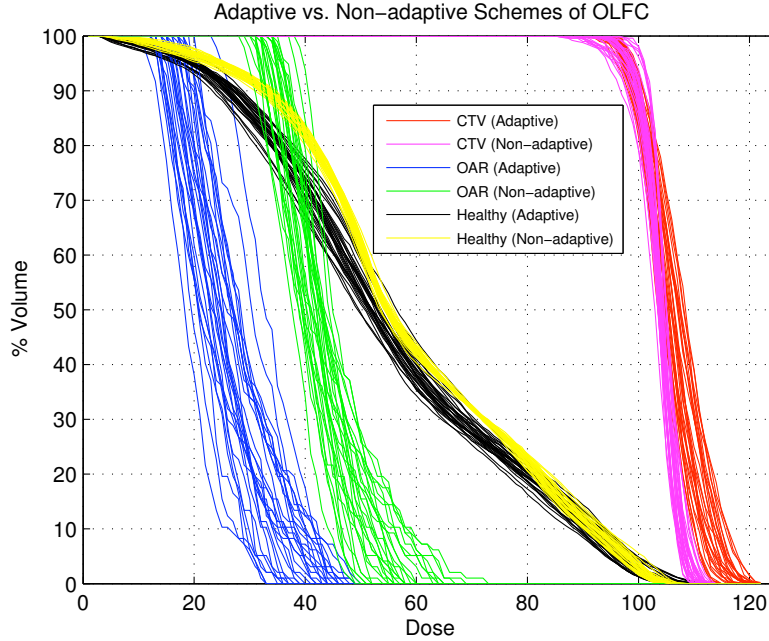


Figure 7. Comparison of DVHs, resulting from 30 replications of 10-fraction treatment, for adaptive and non-adaptive schemes of OLFC. Both policies use the MIGA model shown in Figure 4.

the use of margins allows it to handle uncertainty in patient setup fairly well.

Table 5 compares the performance of the adaptive and non-adaptive schemes of OLFC. As can be seen, the adaptive scheme of OLFC delivers significantly less radiation dose to the OAR than the non-adaptive scheme. Figure 7 shows this even more clearly.

Table 5 also shows that, although the adaptive scheme of OLFC results in a lower average healthy tissue dose, its max healthy tissue dose are higher than the non-adaptive scheme. The resulting healthy tissue EUD values are similar for both policies.

Figure 7 show that the adaptive scheme of OLFC does a better job in covering the CTV than the non-adaptive OLFC, without creating hot spots (in almost every simulated treatment courses, the adaptive scheme delivers less than the prescribed upper bound of 120 gray to the CTV voxels).

In conclusion, the adaptive scheme of OLFC results in better CTV coverage and OAR-sparing than the non-adaptive scheme. This result is not surprising because, unlike the adaptive scheme of CEC, the adaptive scheme of OLFC not

Table 5

Comparison of basic metrics for adaptive and non-adaptive schemes of OLFC. Both policies use the MIGA model shown in Figure 4. Comparisons are based on 30 replications of 10-fraction treatment. The difference in the mean values of the metrics are evaluated using significance tests at level 0.05.

	Statistic	Adaptive	Non-adaptive	Significance Test
Max total OAR dose	Min	31.06	47.40	Adaptive
	Avg	40.30	56.55	
	Max	48.01	71.22	
	Std	5.20	6.15	
Avg total OAR dose	Min	18.10	36.65	Adaptive
	Avg	24.60	41.53	
	Max	32.41	45.27	
	Std	3.27	2.20	
OAR total EUD	Min	28.92	45.25	Adaptive
	Avg	37.16	53.55	
	Max	43.92	65.47	
	Std	4.62	5.21	
Max total healthy tissue dose	Min	104.19	102.91	Non-adaptive
	Avg	108.01	105.44	
	Max	114.40	107.97	
	Std	2.62	1.43	
Avg total healthy tissue dose	Min	52.63	57.92	Adaptive
	Avg	54.68	58.32	
	Max	58.04	58.75	
	Std	1.26	0.17	
Healthy tissue total EUD	Min	78.64	80.63	NSE
	Avg	81.34	81.88	
	Max	85.64	83.03	
	Std	1.60	0.70	
Min total CTV dose	Min	88.33	84.76	Adaptive
	Avg	93.20	90.69	
	Max	95.82	95.70	
	Std	1.84	3.05	
CTV total EUD	Min	91.90	88.24	Adaptive
	Avg	95.58	93.14	
	Max	97.48	97.29	
	Std	1.52	2.53	

NSE = No Strong Evidence that the mean values of the metrics are different.

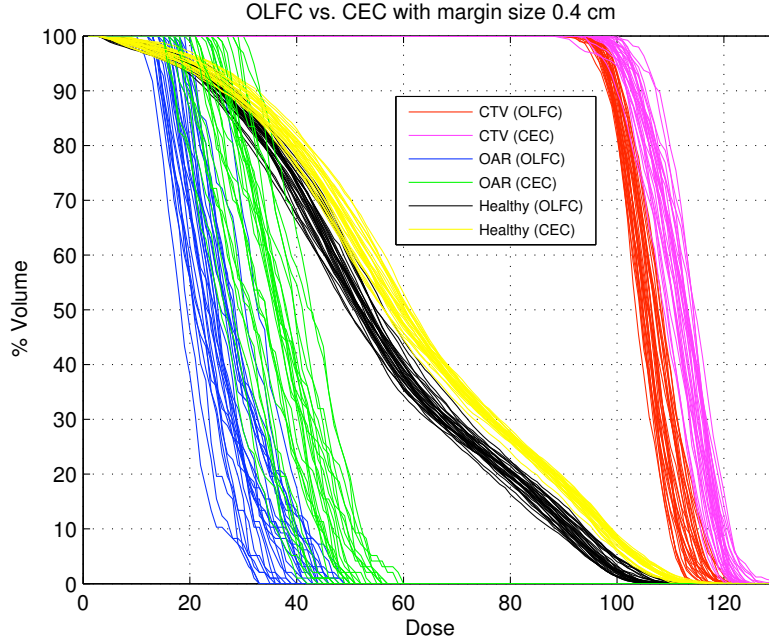


Figure 8. Comparison of DVHs, resulting from 30 replications of 10-fraction treatment, for OLFC and CEC. The OLFC re-optimization uses the MIGA model shown in Figure 4. The CEC re-optimization uses a margin of 0.4 cm.

only exploits the delivery error information from previous fractions, but also incorporates into its re-optimization model the more sophisticated MIGA model of uncertainty, making it able to account for random setup variations in future fractions.

6.2. OLFC vs. CEC with margin size of 0.4 cm

In this section, the performance of adaptive schemes of CEC and OLFC are compared. For the CEC re-optimization, a margin of 0.4 cm is used. For the OLFC re-optimization, the MIGA model with 5 setup instances, shown in Figure 4, is used as the model of uncertainty.

Figure 8 compares DVHs for CEC and OLFC policies; Table 6 lists summary statistics of several important metrics along with the result of significance tests on the difference in the mean values of these metrics at significance level 0.05.

As can be seen from Table 6 and Figure 8, OLFC spares the OAR and healthy tissue better than CEC.

Table 6

Comparison of basic metrics for OLFC and CEC. The OLFC re-optimization uses the MIGA model shown in Figure 4. The CEC re-optimization uses a margin of 0.4 cm. Comparisons are based on 30 replications of 10-fraction treatment. The difference in the mean values of the metrics are evaluated using significance tests at level 0.05.

	Statistic	OLFC	CEC	Significance Test
Max total OAR dose	Min	31.06	40.04	OLFC
	Avg	39.96	49.82	
	Max	48.01	58.91	
	Std	5.17	4.73	
Avg total OAR dose	Min	18.10	25.14	OLFC
	Avg	24.43	34.08	
	Max	32.41	41.81	
	Std	3.32	4.36	
OAR total EUD	Min	28.92	37.06	OLFC
	Avg	36.85	46.67	
	Max	43.92	54.50	
	Std	4.60	4.56	
Max total healthy tissue dose	Min	104.19	112.69	OLFC
	Avg	107.79	116.96	
	Max	114.40	127.75	
	Std	2.50	2.83	
Avg total healthy tissue dose	Min	52.63	58.12	OLFC
	Avg	54.51	60.57	
	Max	56.61	62.62	
	Std	1.10	1.06	
Healthy tissue total EUD	Min	78.64	86.53	OLFC
	Avg	81.15	88.76	
	Max	84.50	94.69	
	Std	1.41	1.61	
Min total CTV dose	Min	88.33	87.75	CEC
	Avg	93.18	94.94	
	Max	95.82	99.35	
	Std	1.88	2.24	
CTV total EUD	Min	91.90	92.22	CEC
	Avg	95.56	98.05	
	Max	97.48	101.86	
	Std	1.56	1.86	

In most of the simulated treatment courses, OLFC delivers a slightly lower dose to some of the CTV voxels than the prescribed minimum dose of 95 gray, as the CTV DVHs in Figure 8 shows (also see the summary statistics for min total CTV dose in Table 6). CEC, on the other hand, delivers a dose lower than the prescribed minimum dose to some of the CTV voxels in only half of the simulated treatment courses. As can be seen from Table 6, in terms of EUD (recall that lower bound on EUD in the CTV was set to 95 gray), CEC also seems to cover the CTV better than OLFC. This is possibly a direct consequence of significantly higher amount of radiation delivered to the CTV by CEC compared to OLFC as Figure 8 show. These DVHs show that CEC violates the upper bound of 120 gray on CTV dose in every simulated treatment course (some voxels in certain simulated treatment courses even receive more than 130 grays), whereas OLFC does not violate this constraint for most simulated treatment courses. As a consequence, OLFC delivers a more uniform dose distribution to the CTV (*i.e.*, steeper DVHs), which is often a desired property in a treatment plan, compared to CEC.

The summary statistics for max total healthy tissue dose in Table 6 shows that CEC creates hot spots in healthy tissue region (it delivers as much as 127.75 gray to the healthy tissue in some simulated treatment courses).

As discussed in previous section, CEC delivers excessive amount of radiation dose both to the CTV and other treatment regions since it tries to compensate dose delivery errors in previous fractions at every re-optimization based on the assumption that the patient will be at the nominal position in future fractions. The resulting dose distributions are therefore not able to account for random setup variations in future fractions. OLFC, on the other hand, not only exploits the dose delivery information from previous fractions, but also uses MIGA as an approximate model of uncertainty. Therefore, OLFC is able to spare OAR and healthy tissue while maintaining a sufficient dose level in CTV.

As mentioned above, in most of the simulated treatment courses, OLFC delivers a slightly lower dose to some of the CTV voxels than the prescribed minimum dose of 95 grays. Since the dose and EUD constraints are enforced for every setup instance, this shows that the setup instances in the MIGA model shown in Figure 4 do not capture enough of the “spread” of the underlying distribution the random setup position. We believe that improving the MIGA approximation by increasing the number of setup instances would improve the CTV coverage significantly without sacrificing OAR and healthy tissue sparing.

7. Conclusions and future work

As we have illustrated, an optimization framework, exploiting the dynamic nature of radiotherapy to adapt a treatment plan, can be readily developed. Computationally feasible control policies can be obtained using techniques of approximate dynamic programming such as CEC and OLFC. The resulting control policies can be formulated for arbitrary convex cost functions and convex constraints, and solution methods can be developed for specific forms of cost functions and constraints.

Numerical experiments show that this approach results in better CTV coverage as well as OAR and healthy tissue sparing as long as the the model of uncertainty used by the control policies is a good approximation of the *true* distribution of random setup position. At the beginning of a treatment course, however, no information is available about a particular patient except for the clinical data collected from the patient population. Sir (2007) discusses how Bayesian statistical methods can be used to refine population-based statistics and estimate the parameters of the true distribution of random setup position as a sequence of measurements of patient setup is taken during the course of the treatment. Sir (2007) also discusses applications of Bayesian statistical methods to deal with systematic errors within the framework described in this paper.

Numerical experiments also show that adapting the treatment plans to the most current information (i.e., the total dose delivered in previous fractions) does not provide any significant improvement over the non-adaptive treatment schemes used in practice if the uncertainty in future fractions is not adequately accounted for.

There are several important directions for future research, the most critical being a demonstration of the applicability of our methods to clinical data. However, we strongly believe that, once the technical details of data collection and computer implementation for three-dimensional real patient cases are resolved, the associated improvements in the resulting treatment plans would still reflect the results obtained for the two-dimensional simple patient representation. The pursuit of such validation analysis is in progress, and future papers will report on this.

Extension of our results to incorporate uncertainties in biological response as well as predictive models of tumor growth and shrinkage is also of great interest. Finally, since the performance of OLFC clearly depends on the quality

of the MIGA used, further study is needed on how to select the positions and probabilities for the setup instances, especially in the presence of heterogeneous tissue densities.

Acknowledgements

The authors are grateful to Drs Kwok Lam, Randall Ten Haken, Benedick Fraass, Daniel McShan, Ken Jee, and Marc Kessler from the Department of Radiation Oncology at the University of Michigan for insightful discussions of this work.

This work was supported in part by National Institute of Health (NIH) Grant P01-CA59827 and by National Science Foundation (NFS) Grant CCF-0306240.

References

- W Beckham, P Keall, and J Siebers. A fluence-convolution method to calculate radiation therapy dose distributions that incorporate random set-up error. *Phys. Med. Biol.*, 47:3465–3473, 2002.
- G C Bentel. *Patient Positioning and Immobilization in Radiation Oncology*. McGraw-Hill, 1999.
- D P Bertsekas. *Dynamic Programming and Optimal Control*, volume 1. Athena Scientific, Belmont, Massachusetts, 3 edition, 2005.
- J R Birge and F Louveaux. *Introduction to stochastic programming*. Springer, New York, 1997.
- T Bortfeld and H Paganetti. The biologic relevance of daily dose variations in adaptive treatment planning. *Int. J. Radiation Oncology Biol. Phys.*, 65:899–906, 2006.
- T Bortfeld, M van Herk, and S Jiang. When should systematic patient positioning errors in radiotherapy be corrected? *Phys. Med. Biol.*, 47:N297–N302, 2002.
- T Bortfeld, S Jiang, and E Rietzel. Effects of motion on the total dose distribution. *Semin. Radiat. Oncol.*, 14:41–51, 2004.
- I Chetty, M Rosu, N Tyagi, L H Marsh, D L McShan, J M Balter, B A Fraass, and R K Ten Haken. A fluence convolution method to account for respiratory motion in three-dimensional dose calculations for liver: A monte carlo study. *Med. Phys.*, 30:1776–17, 2003.

- B Choi and J O Deasy. The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning. *Phys. Med. Biol.*, 47:3579–3589, 2002.
- M Chu, Y Zinchenko, S G Henderson, and M B Sharpe. Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty. *Phys. Med. Biol.*, 50:5463–5477, 2005.
- D Craft, T Halabi, and T Bortfeld. Exploration of tradeoffs in intensity-modulated radiotherapy. *Phys. Med. Biol.*, 50:5857–5868, 2005.
- A de la Zerda, B Armbruster, and L Xing. Formulating adaptive radiation therapy (art) treatment planning into a closed-loop control framework. *Phys. Med. Biol.*, 52:4137–4153, 2007.
- U Diwekar, V Rico-Ramirez, K J Kim, and K Sahin. Optimization under uncertainty. *SIAG/OPT Views-and-News*, 13(1), 2002.
- M Ferris and M Voelker. Fractionation in radiation treatment planning. *Math. Program.*, Ser. B 101:387–413, 2004.
- M G Herman, J M Balter, D A Jaffray, K P McGee, P Munro, S Shalev, M van Herk, and J W Wong. Clinical use of electronic portal imaging: Report of AAPM radiation therapy committee task group 58. *Med. Phys.*, 28:712–737, 2001.
- ICRU. Prescribing, recording and reporting photon beam therapy ICRU Report 50. Technical Report 50, International Commission on Radiation Units and Measurements (ICRU), Bethesda, MD, 1993.
- ICRU. Prescribing, recording and reporting photon beam therapy (Supplement to ICRU Report 50),. Technical Report 62, International Commission on Radiation Units and Measurements (ICRU), Bethesda, MD, 1999.
- H Keller, M Ritter, and T Mackie. Optimal stochastic correction strategies for rigid-body target motion. *Int. J. Radiation Oncology Biol. Phys.*, 55:261–270, 2003.
- H Keller, W Tome, M A Ritter, and T R Mackie. Design of adaptive treatment margins for non-negligible measurement uncertainty: application to ultrasound-guided prostate radiation therapy. *Phys. Med. Biol.*, 49:69–86, 2004.
- J E Kelley. The cutting-plane method for solving convex programs. *J. Soc. for Indust. Appl. Math.*, 8:703–712, 1960.
- M L Kessler, D L Mcshan, M A Epelman, K A Vineberg, A Eisbruch, T S Lawrence, and B A Fraass. Costlets: a generalized approach to cost functions

- for automated optimization of IMRT treatment plans. *Optimization Eng.*, 6: 421–448, 2005.
- R Koshani, J M Balter, J A Hayman, G T Henning, and M van Herk. Short-term and long-term reproducibility of lung tumor position using active breathing control (ABC). *Int. J. Radiation Oncology Biol. Phys.*, 65:1553–1559, 2006.
- K A Küfer, A Scherrer, and M Monz. Intensity-modulated radiotherapy - a large scale multi-criteria programming problem. *OR Spectrum*, 25:223–249, 2003.
- K L Lam, R K Ten Haken, D Litzenberg, J M Balter, and S M Pollock. An application of bayesian statistical methods to adaptive radiotherapy. *Phys. Med. Biol.*, 50:3849–3858, 2005.
- K L Lam, J M Balter, and R K Ten Haken. Effect of daily localization and correction on the setup uncertainty: dependences on the measurement uncertainty, re-positioning uncertainty and action level. *Phys. Med. Biol.*, 52:6575–6587, 2007.
- E Lee, T Fox, and I Crocker. Integer programming applied to intensity-modulated radiation therapy treatment planning. *Annals of Operations Research*, pages 165–181, 2003.
- J Löf, B K Lind, and A Brahme. Optimal radiation beam profiles considering the stochastic process of patient positioning in fractionated radiation therapy. *Inverse Problems*, 11:1189–1209, 1995.
- A E Lujan, R K Ten Haken, E W Larsen, and J M Balter. Quantization of setup uncertainties in 3-D dose calculations. *Med. Phys.*, 26:2397–2402, 1999.
- A Martinez, D Yan, D Lockman, D Brabbins, K Kota, M Sharpe, D A Jaffray, F Vicini, and J Wong. Improvement in dose escalation using the process of adaptive radiotherapy combined with three-dimensional conformal or intensity-modulated beams for prostate cancer. *IntRad*, 48:289–302, 2000.
- T R McNutt, T R Mackie, and B R Paliwal. Analysis and convergence of the iterative convolution/superposition dose reconstruction technique for multiple treatment beams and tomotherapy. *Med. Phys.*, 24:1465–1476, 1997.
- D L McShan, M L Kessler, K Vineberg, and B A Fraass. Inverse plan optimization accounting for random geometric uncertainties with a multiple instance geometry approximation (MIGA). *Med. Phys.*, 33:1510–1521, 2006.
- A Mestrovic, M Milette, A Nichol, B G Clark, and K Otto. Direct aperture optimization for online adaptive radiation therapy. *Med. Phys.*, 34:1631–1646, 2007.
- A Niemierko. A generalized concept of equivalent uniform dose (EUD) (abstract).

- Med. Phys.*, 26:1100, 1999.
- M Partridge, M Ebert, and B M Hesse. IMRT verification by three-dimensional dose reconstruction from portal beam measurements. *Med. Phys.*, 29:1847–1858, 2002.
- C A Perez and L W Brady. *Principles and Practice of Radiotherapy*. Lippincott-Raven, Philadelphia, PA, 1998.
- M P Petric, B G Clark, and J L Robar. A comparison of two commercial treatment-planning systems to IMRT. *J. Appl. Clin. Med. Phys.*, 6:63–80, 2005.
- L W Pickle, Y Hao, A Jemal, Z Zou, R C Tiwari, E Ward, M Hachey, H L Howe, and E J Feuer. A new method of estimating united states and state-level cancer incidence counts for the current calendar year. *CA Cancer J Clin*, 57:30–42, 2007.
- H Rehbinder, C Forsgren, and J Löf. Adaptive radiation therapy for compensation of errors in patient setup and treatment delivery. *Med. Phys.*, 31:3363–3371, 2004.
- H E Romeijn, R K Ahuja, J F Dempsey, A Kumar, and J 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.
- A Ruszczyński and A Shapiro. *Handbooks in operations research and management science: Stochastic Programming*, volume 10. Elsevier, 2003.
- J E Schewe, J M Balter, K L Lam, and R K Ten Haken. Measurement of patient setup errors using port films and a computer-aided graphical alignment tool. *Med. Dosim.*, 21:97–104, 1996.
- W Schlegel, T Bortfeld, and A L Grosu. *New Technologies in Radiation Oncology*. Medical Radiology / Radiation Oncology. Springer, 2006.
- D M Shepard, M C Ferris, G H Olivera, and T R Mackie. Optimizing the delivery of radiation therapy to cancer patients. *SIAM Review*, 41:721–744, 1999.
- M Y Sir. *Optimization of radiotherapy considering uncertainties caused by daily setup procedures and organ motion*. PhD thesis, The University of Michigan, Ann Arbor, MI, 2007.
- M Y Sir, S M Pollock, M A Epelman, K L Lam, and R K Ten Haken. Ideal spatial radiotherapy dose distributions subject to positional uncertainties. *Phys. Med. Biol.*, 51:6329–6347, 2006.
- H B Stone, C N Coleman, M S Anscher, and W H McBride. Effects of radiation on normal tissue: consequences and mechanisms. *Lancet Oncol*, 4:529–536,

2003.

- H D Thames and J H Hendry. *Fractionation in radiotherapy*. Taylor and Francis, London, New York, 1987.
- C Thieke, T Bortfeld, and K Küfer. Characterization of dose distributions through the max and mean dose concept. *Acta Oncologica*, 41:158 – 161, 2002.
- D Thongphiew, V Chankong, F Yin, and Q J Wu. An online adaptive radiation therapy system for intensity modulated radiation therapy: An application of multi-objective optimization. *JIMO*, 4:453 – 475, 2008.
- W Ulmer and D Harder. A triple gaussian pencil beam model for photon beam treatment planning. *Z. Med. Phys.*, 5:25–30, 1995.
- M van Herk. Errors and margins in radiotherapy. *Semin. Radiat.Oncol.*, 14: 52–64, 2004.
- M van Herk, P Remeijer, C Rasch, and J V Lebesque. The probability of correct target dosage: dose-population histograms for deriving treatment margins in radiotherapy. *Int. J. Radiation Oncology Biol. Phys.*, 47:1121–1135, 2000.
- S Webb. *Contemporary IMRT: Developing Physics and Clinical Implementation*. Institute of Physics Publishing, Bristol and Philadelphia, 2005.
- J W Wong, M B Sharpe, D A Jaffray, V R Kini, J M Roberson, J S Stromberg, and A Martinez. The use of active breathing control (ABC) to reduce margin for breathing motion. *Int. J. Radiation Oncology Biol. Phys.*, 44:911–999, 1999.
- C Wu, R Jeraj, G Olivera, and T Mackie. Re-optimization in adaptive radiotherapy. *Phys. Med. Biol.*, 47:3181–3195, 2002.
- Q Wu, J Liang, and D Yan. Application of dose compensation in image-guided radiotherapy of prostate cancer. *Phys. Med. Biol.*, 51:1405–1419, 2006.
- D Yan. “Image-Guided/Adaptive Radiotherapy” in *New Technologies in Radiation Oncology*, chapter 4, pages 321–336. Medical Radiology. Springer Berlin Heidelberg, 2006.
- D Yan, F Vicini, J Wong, and A Martinez. Adaptive radiation therapy. *Phys. Med. Biol.*, 42:123–132, 1997a.
- D Yan, J Wong, F Vicini F, J Michalski, C Pan, A Frazier, E Horwitz, and A Martinez. Adaptive modification of treatment planning to minimize the deleterious effects of treatment setup errors. *Int. J. Radiation Oncology Biol. Phys.*, pages 197–206, 1997b.
- D Yan, D Lockman, D Brabbins, L Tyburski, and A Martinez. An off-line strategy for constructing a patient-specific planning target volume in adaptive treatment process for prostate cancer. *Int. J. Radiation Oncology Biol. Phys.*, 48:289–302,

2000.