



Innovative Applications of O.R.

## Adaptive and robust radiation therapy optimization for lung cancer

Timothy C.Y. Chan<sup>a</sup>, Velibor V. Mišić<sup>b,\*</sup><sup>a</sup> Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, ON M5S 3G8, Canada<sup>b</sup> Operations Research Center, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Building E40, Cambridge, MA 02139, United States

## ARTICLE INFO

## Article history:

Received 23 May 2012

Accepted 3 June 2013

Available online 11 June 2013

## Keywords:

OR in health services

Intensity-modulated radiation therapy

Adaptive radiation therapy

Robust optimization

Linear programming

## ABSTRACT

A previous approach to robust intensity-modulated radiation therapy (IMRT) treatment planning for moving tumors in the lung involves solving a single planning problem before the start of treatment and using the resulting solution in all of the subsequent treatment sessions. In this paper, we develop an adaptive robust optimization approach to IMRT treatment planning for lung cancer, where information gathered in prior treatment sessions is used to update the uncertainty set and guide the reoptimization of the treatment for the next session. Such an approach allows for the estimate of the uncertain effect to improve as the treatment goes on and represents a generalization of existing robust optimization and adaptive radiation therapy methodologies. Our method is computationally tractable, as it involves solving a sequence of linear optimization problems. We present computational results for a lung cancer patient case and show that using our adaptive robust method, it is possible to attain an improvement over the traditional robust approach in both tumor coverage and organ sparing simultaneously. We also prove that under certain conditions our adaptive robust method is asymptotically optimal, which provides insight into the performance observed in our computational study. The essence of our method – solving a sequence of single-stage robust optimization problems, with the uncertainty set updated each time – can potentially be applied to other problems that involve multi-stage decisions to be made under uncertainty.

© 2013 Elsevier B.V. All rights reserved.

## 1. Introduction

Lung cancer is the leading cause of death due to cancer in North America, killing an estimated 180,000 people in 2010 (American Cancer Society, 2010; Canadian Cancer Society's Steering Committee, 2010) and accounting for over 25% of all cancer deaths. Lung cancer is often treated using radiation therapy (Toschi, Cappuzzo, & Janne, 2007). Of the different types of radiation therapy, one of the most commonly used in practice for treating cancer in general is intensity-modulated radiation therapy (IMRT) (Mell, Mehrotra, & Mundt, 2005). In an IMRT treatment, the patient is irradiated from multiple beams, each of which is decomposed into a large number of small beamlets. The beamlet intensities can be controlled through the use of a multileaf collimator (MLC) that moves metal leaves in and out of the beam field in order to block certain parts of the beam. By appropriately setting the beamlet intensities, the volume that is irradiated can be made to closely conform to the shape of the target. The basic problem in planning an IMRT treatment is to determine how the beamlet intensities or weights should be set so that the target receives an adequate dose while

the healthy tissue receives a minimal dose. This is known as the beamlet weight optimization problem or the fluence map optimization problem. Since the inception of IMRT, much research has focused on modeling and solving this problem as a mathematical program (see Romeijn & Dempsey, 2008 for a comprehensive overview).

In practice, the beamlet weight optimization problem is complicated by the presence of uncertainties, such as those arising from errors in beam positioning and patient placement, internal organ motion during treatment, and changes in organ position between treatment sessions. All of these factors affect the relative position of the tumor with respect to the beams, which in turn affects how much dose is deposited in the tumor and the healthy tissue. For tumors in the lung, the most significant uncertainty arises from breathing motion. During treatment, the patient is constantly breathing, and the tumor moves with the expansion and contraction of the patient's lungs. Furthermore, the patient's breathing pattern during treatment is not known exactly beforehand and can vary from day to day. If a treatment is planned with a specific breathing pattern in mind but a different pattern is realized during treatment, the tumor may end up being underdosed and the quality of the treatment may thus be greatly compromised (Lujan, Baltzer, & Ten Haken, 2003; Sheng et al., 2006). At the same time, if the treatment is designed to deliver the prescription dose to the tumor

\* Corresponding author. Tel.: +1 781 492 4276.

E-mail addresses: [tcychan@mie.utoronto.ca](mailto:tcychan@mie.utoronto.ca) (T.C.Y. Chan), [vmisic@mit.edu](mailto:vmisic@mit.edu) (V.V. Mišić).

under any breathing pattern, regardless of how unlikely some of those patterns may be, then an unnecessarily high amount of damage will be done to the healthy tissue. There exist approaches that aim to mitigate the effect of uncertainty by only activating the beam when the patient is in a certain phase of the breathing cycle (Ohara et al., 1989), but these approaches have the disadvantage of prolonging the treatment time, leading to reduced throughput in the treatment center.

One methodology that allows for tumor coverage under uncertainty to be balanced with healthy tissue sparing, without extending the treatment time, is robust optimization. There are many robust optimization approaches to IMRT treatment planning. In this paper, we build on the method developed in Bortfeld, Chan, Trofimov, and Tsitsiklis (2008) and Chan, Bortfeld, and Tsitsiklis (2006), which is specifically designed to manage the instantaneous breathing motion uncertainty that may be realized during treatment sessions. Although robust optimization has been shown to be a valuable methodology for managing uncertainty due to breathing motion, it can be further improved by leveraging the fact that treatments are typically fractionated. Fractionation refers to the practice of dividing up the prescription dose and delivering small amounts in daily treatment sessions over multiple weeks, in order to exploit the more capable repair mechanism that healthy tissue has over cancerous tissue (Thames & Hendry, 1987).

To date, the typical approach to fractionated treatment planning and delivery has been to determine what the beamlet intensities should be to deliver a certain prescription dose, scale them by  $1/n$  where  $n$  is the number of fractions, and use the resulting intensities in every treatment fraction. This, however, does not take advantage of the dynamics that fractionation injects into IMRT treatment planning and the underlying uncertainty. For example, it may be the case that the pre-treatment estimates of the patient's breathing pattern are erroneous, and the breathing patterns that are planned for differ from those realized over the treatment. In this case, the wrong beamlet intensities would be used for every fraction, without the possibility of correcting them. On the other hand, the beamlet intensities may be optimized in such a way that the prescription dose is delivered under a very large range of breathing patterns, when in actuality the patient's breathing pattern does not vary significantly from fraction to fraction. In this case, there would be no way to correct the conservatism of the treatment, and the patient's healthy tissue will incur more damage than necessary. Another example comes from the possibility that the patient's breathing pattern may change over the course of treatment as the patient becomes more comfortable in the treatment room or as the evolution of the underlying disease affects the patient's respiratory ability. As a result, latter treatment sessions may deliver insufficient dose to the tumor and more dose than anticipated to healthy tissue. These issues motivate the incorporation of adaptive re-optimization into robust optimization models.

To deal with fractionation in more general IMRT treatment planning contexts, much research has been conducted in the medical physics community in what is known as off-line adaptive radiation therapy (ART). In off-line ART, after the current fraction is delivered, new information, typically obtained from imaging devices, is fed back into the planning process to design the beamlet intensities of the next fraction (Yan, Vicini, Wong, & Martinez, 1997a). To date, off-line ART and robust optimization have largely been considered independently of each other.

In this paper, we propose a novel method that generalizes the framework of off-line adaptive radiation therapy and the robust optimization methodology of Bortfeld et al. (2008) and Chan et al. (2006) to treat lung cancer in the presence of breathing motion uncertainty. Our method is based on adaptively updating an uncertainty set that models the breathing motion uncertainty

and then solving a robust optimization problem with the new uncertainty set prior to each fraction. As a result, our method is able to combat both the instantaneous uncertainty realized in each fraction as well as changes in this uncertainty that occur from day to day. The clinical value of this method is that it has the potential to improve upon the non-adaptive robust optimization method in both tumor coverage and healthy tissue sparing. This is significant because patient survival rates have been shown to improve when the tumor dose is increased (see Perez, Bauer, Edelstein, Gillespie, & Birch, 1986). Unfortunately, dose escalation to the tumor is limited by the radiation tolerance of healthy tissue, particularly the lungs (Cox et al., 1990). If it is possible to increase tumor dose and reduce healthy tissue dose from the levels achieved by non-adaptive robust optimization, further dose escalation will become viable, and improvements in patient survival rate may be realized.

Our specific contributions in this paper are as follows:

1. We develop a method that combines robust optimization with adaptive radiation therapy and suitably generalizes both frameworks. This method has modest computational requirements and can be easily generalized to other types of cancer and other types of uncertainty.
2. Using real patient data, we demonstrate computationally that this method generates treatments that improve upon treatments obtained by non-adaptive robust methods in both tumor coverage and lung dose. We show that the quality of the final treatment does not vary significantly with the choice of initial uncertainty set. We also show that the achieved performance is comparable to that of two "prescient" solutions, which correctly anticipate the patient's breathing over the entire treatment.
3. We prove an asymptotic optimality result: if the patient breathing pattern converges over the treatment course, then the dose distribution produced by our adaptive robust method converges to an "ideal" set of dose distributions that exhibit no tumor underdose or overdose and low healthy tissue dose. We use this result to explain multiple phenomena observed in our computational study.

The rest of the paper is organized as follows. In Section 2, we briefly review some of the relevant literature in robust optimization and off-line ART. In Section 3, we provide an overview of the non-adaptive robust optimization method on which this paper builds. In Section 4, we provide a detailed description of the proposed adaptive robust method. In Section 5, we provide results from our computational study with clinical lung cancer patient data. In Section 6, we conduct a theoretical analysis of the adaptive robust method. Finally, in Section 7, we provide some conclusions and directions for future work.

## 2. Literature review

Robust optimization has been extensively developed in the OR community (see Ben-Tal, El Ghaoui, & Nemirovski, 2009 and Bertsimas, Brown, & Caramanis, 2011), and has received increasing focus in the last decade as a methodology for IMRT treatment planning under uncertainty. In Chu, Zinchenko, Henderson, and Sharpe (2005), tumor position uncertainty due to interfraction motion (organ position changes that occur between treatment sessions) and setup error was modeled using an ellipsoidal uncertainty set and the authors showed that the resulting robust solution achieves the same level of tumor coverage as a clinical margin solution but with less healthy tissue dose. In Ólafsson and Wright (2006), the uncertainty due to dose calculation error and interfraction organ motion was also modeled using an ellipsoidal uncertainty set,

and the authors showed for a nasopharyngeal case that a robust solution achieves better tumor coverage than the nominal solution (one which assumes a dose matrix known with certainty) and leads to better organ sparing than the clinically prescribed margin solution. In Chan et al. (2006) and Bortfeld et al. (2008), uncertainty due to intrafraction breathing motion was modeled using a polyhedral uncertainty set, and it was shown that the robust solution provides improved tumor coverage over the nominal solution, while providing tumor coverage comparable to the margin solution at reduced healthy tissue dose. All of the robust IMRT optimization methods described above are non-adaptive.

With regard to off-line ART, a variety of procedures have been proposed for how to treat a patient based on information obtained in the most recent fraction. Many studies have proposed using daily computed tomography (CT) images as a form of feedback. For example, Mohan et al. (2005), Lu et al. (2006), Wu, Liang, and Yan (2006) and Wu et al. (2008) all studied schemes where daily CT images are used to periodically assess the treatment plan for errors in dose delivery due to organ deformation and intrafraction motion (organ motion that occurs during a treatment session), and to aid either the re-adjustment or complete re-optimization of the treatment for subsequent fractions. Other imaging methods have also been proposed for off-line ART. For instance, de la Zerda, Armbruster, and Xing (2007) proposed algorithms for generating new treatment plans from fraction to fraction in response to changes in the patient geometry detected from cone beam CT (CBCT) images, as well as to the cumulative delivered dose. Another example is Yan et al. (1998), who tested an ART method where the daily setup error is measured using portal imaging and is used to correct the treatment if the error becomes sufficiently large.

Some ART studies have also incorporated probabilistic models of uncertainty into their proposed planning methods. Yan et al. (1997a) studied an ART method where the daily treatment target and beam placement variation are measured and used to modify the treatment dose and field margin in each fraction under the assumption that the error is normally distributed. Löf, Lind, and Brahme (1998) studied an ART scheme where the daily setup position is measured in every fraction, but assumed that the probabilistic dynamics of the internal organ motion are known beforehand. Reh binder, Forsgren, and Löf (2004) applied linear quadratic control to ART planning and showed that using this form of adaptation removes the need for a margin, allowing for safer dose escalation. The problem of ART can also be seen as a problem of sequential decision making under uncertainty and as such, several studies have considered dynamic programming techniques. For example, Sir, Epelman, and Pollock (2012) showed that an adaptive open-loop feedback control (OLFC) policy achieves better performance than a non-adaptive OLFC policy. Ferris and Voelker (2004) and Deng and Ferris (2008) applied neuro-dynamic programming (NDP) to ART and similarly showed that their NDP policy offers an improvement over a constant, non-adaptive policy.

The limitation of the majority of algorithms for off-line ART is in how they have approached uncertainty. Many studies have not explicitly included uncertain effects into their treatment planning process. Of the studies that have, most have assumed that realizations of the uncertain effect from fraction to fraction are independent and identically distributed – although this may be appropriate for setup error and patient positioning, it may not be suitable for other types of uncertainty. The most significant assumption present in studies that have incorporated probabilistic models of uncertainty is that the underlying probability distribution of the uncertain effect is known precisely before the start of the treatment. In actuality, the distribution which underlies an uncertain effect for a particular patient in a particular treatment is never known precisely, and as already discussed, the quality of

a treatment designed with a specific probability distribution in mind can deteriorate significantly if a different distribution is realized.

The fundamental difference between the approach proposed in this paper and the majority of prior ART work lies in the assumptions that are made about the uncertainty. In our approach, we do not assume that the nature of the uncertainty is known precisely *a priori* nor that it needs to stay constant from fraction to fraction. There has been some prior research into using previous measurements of uncertainty to update an estimate of the uncertain effect, using such methods as Kalman filtering (Yan et al., 1997b; Keller, Ritter, & Mackie, 2003) and Bayesian updating (Lam, Haken, Litzenberg, Balter, & Pollock, 2005; Sir, 2007). Very few studies have considered these methods with treatment re-planning (Yan et al., 1997b & Sir, 2007 are exceptions) and these studies assume that the uncertain effect follows a particular distribution. No studies have combined such estimation methods with the kind of robust optimization approach we consider.

### 3. Static robust IMRT optimization

In this section, we review the robust optimization approach of Chan et al. (2006) and Bortfeld et al. (2008). In their approach, the planner decides on an uncertainty set before treatment and solves the robust optimization problem corresponding to that uncertainty set. The optimal solution of the robust problem is a vector that specifies the intensity of each beamlet in the ensemble of beams to be used for treatment. These optimized beamlet intensities are robust to the uncertain effect whenever it takes on values from the uncertainty set: whenever the uncertain effect realizes a value inside the uncertainty set while the patient is being irradiated according to the optimized intensities, the dose to every part of the tumor will be within the prescribed bounds. We will refer to this approach as the *static* robust optimization approach, as the beamlet intensity vector is fixed and does not change over the course of the treatment (in each fraction,  $1/n$ th of the intensity vector is delivered, where  $n$  is the number of fractions). The patient's breathing motion is modeled using a probability mass function (PMF). The breathing motion PMF for a patient in a particular fraction specifies the proportion of time the patient spends in each of a finite number of breathing motion states during that fraction. Each breathing motion state corresponds to a particular displacement or snapshot of the patient geometry during the patient's respiratory cycle. The uncertainty set is then a set of breathing motion PMFs that could be realized during treatment. We note here that in the method of Bortfeld et al. (2008), the set of breathing motion states is determined from pre-treatment breathing-correlated 4D computed tomography (4D-CT) images. This assumes that the amplitude of the tumor motion observed in these pre-treatment images represents the full extent of the tumor motion, and that the amplitude of the motion that will be observed during the ensuing treatment sessions will not exceed these limits. This potential issue can be circumvented by having the patient take deep breaths during the pre-treatment 4D-CT imaging session, allowing the treatment planner to more accurately construct the set of breathing motion states.

We now define the static robust IMRT optimization problem. Let  $\mathcal{B}$  be the set of beamlets and let  $w_b$  denote the intensity of beamlet  $b \in \mathcal{B}$ ; the  $w_b$  variables are the decision variables of the robust optimization problem. The patient geometry (the healthy organs and tumor tissue) is divided into three dimensional voxels, or volume elements. Let  $\mathcal{V}$  denote the set of all voxels and  $\mathcal{T}$  denote the set of tumor voxels. Let  $X$  denote the finite set of breathing motion states. Each breathing motion state is associated with a particular displacement of the patient geometry (the voxels), so to each

beamlet  $b \in \mathcal{B}$ , breathing motion state  $x \in X$  and voxel  $v \in \mathcal{V}$ , we associate a dose deposition coefficient  $\Delta_{v,x,b}$ . This coefficient specifies the amount of dose deposited in voxel  $v$  when the patient is in breathing motion state  $x$  and the beamlet  $b$  is at unit intensity. Let  $p(x)$  be the probability of the patient being in motion state  $x$  at any given time during a treatment session. The dose that is delivered to voxel  $v$  is then given by

$$\sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b,$$

which is the sum of the doses to voxel  $v$  that would be delivered under each breathing motion state weighted by the corresponding proportions of time spent in those states. Due to the use of beamlet intensities, this model does not account for the interplay effect between the motion of the MLC leaves and the tumor motion (see, for example, Yu, Jaffray, & Wong (1998) and Bortfeld, Jokivarsi, Goitein, Kung, & Jiang (2002)); the impact of such an effect, however, is small in practice (Bortfeld, Jiang, & Rietzel, 2004). For each tumor voxel  $v \in \mathcal{T}$ , let  $\theta_v$  be the prescribed minimum dose and let  $\gamma\theta_v$  be the prescribed maximum dose, where  $\gamma \geq 1$ .

Let  $\mathcal{P}$  be the set of all PMFs on the finite set  $X$ , defined as

$$\mathcal{P} = \left\{ \mathbf{p} \in \mathbb{R}^{|X|} \mid \text{For each } x \in X, p(x) \geq 0; \sum_{x \in X} p(x) = 1 \right\}.$$

The set  $\mathcal{P}$  defines the  $(|X| - 1)$ -dimensional unit simplex. Let  $P \subseteq \mathcal{P}$  be the uncertainty set, which is a bounded polyhedron defined by a lower bound vector  $\ell$  and an upper bound vector  $\mathbf{u}$ :

$$P = \{ \mathbf{p} \in \mathcal{P} \mid \text{For each } x \in X, \ell(x) \leq p(x) \leq u(x) \}.$$

With these definitions, the static robust problem from Bortfeld et al. (2008) is

$$\begin{aligned} & \text{minimize} && \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{p}(x) w_b \\ & \text{subject to} && \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b \geq \theta_v, \quad \forall v \in \mathcal{T}, \quad \mathbf{p} \in P, \\ & && \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b \leq \gamma\theta_v, \quad \forall v \in \mathcal{T}, \quad \mathbf{p} \in P, \\ & && w_b \geq 0, \quad \forall b \in \mathcal{B}, \end{aligned} \quad (1)$$

where  $\bar{\mathbf{p}}$  is some nominal PMF chosen by the treatment planner before the start of treatment as approximately representative of the patient's overall breathing pattern. The optimal solution of (1) is a beamlet intensity vector  $\mathbf{w}^*$  that meets minimum and maximum tumor dose requirements no matter which breathing motion PMF  $\mathbf{p} \in P$  is realized, while ensuring the lowest possible total dose to the patient under the nominal PMF. In the case that  $\mathbf{p} \notin P$ , it was shown empirically in Bortfeld et al. (2008) that  $\mathbf{w}^*$  remains feasible with high probability. Due to the linearity of this model and the use of polyhedral uncertainty sets, this model is capable of accommodating robust versions of many other types of clinically relevant constraints, such as partial volume constraints based on conditional-value-at-risk (see Romeijn, Ahuja, Dempsey, & Kumar, 2006) or constraints involving  $\alpha$ EUD (Thieke, Bortfeld, & Küfer, 2002), which is a linear approximation to the popular equivalent uniform dose (EUD) metric. Our goal in this paper is to elucidate the potential benefits of a framework that combines adaptation and robustness, not to design clinic-ready solutions. We therefore consider the stylized model above, which only enforces lower and upper dose bounds on the tumor.

There are two special cases of (1) that will be referred to later. The first is the case when  $P$  consists of a single PMF. The resulting problem is called the *nominal problem*; when  $P = \{\mathbf{p}\}$ , we refer to problem (1) as the *nominal problem with respect to  $\mathbf{p}$* . The solution of the nominal problem is the least conservative treatment we can deliver, as we are assuming that the PMF that will be realized dur-

ing treatment will be the exact PMF that we are guarding against. The second is the case when  $P$  consists of all possible PMFs on  $X$ , that is,  $P = \mathcal{P}$ ; this problem is called the *margin problem*. The solution of the margin problem is the most conservative treatment that we can deliver, because we assume that any PMF is possible. These two types of formulations represent two extremes of a continuum of robustness (Chan et al., 2006), with the nominal problem being the least robust and the margin problem being the most robust. The robust problem with a general uncertainty set (which is neither a singleton nor  $\mathcal{P}$ ) is somewhere in between. With increasing levels of robustness, characterized by larger uncertainty sets, the level of tumor coverage increases, but at the cost of increased dose to healthy tissue.

#### 4. Adaptive robust IMRT optimization

In the adaptive robust approach that we propose, we allow the beamlet intensities to be re-optimized from fraction to fraction in the following way. On a given day, the patient is irradiated with a beamlet intensity vector obtained by solving (1) with the current uncertainty set as input. During or immediately after the fraction is delivered, the patient's breathing motion PMF is measured. This measurement, together with the current uncertainty set, is used to generate a new uncertainty set. The robust problem is then solved with this new uncertainty set, leading to a new beamlet intensity vector to be used on the next day. This process is repeated for each fraction until the end of the treatment. This procedure is presented as Algorithm 1.

Although in theory one could measure the patient's PMF before the delivery of a fraction and plan the fraction using the resulting PMF, this approach would constrain the amount of time that would be available for treatment planning and quality assurance, as the patient would have to wait between measurement and fraction delivery. By measuring the patient's PMF during or after the delivery of the fraction, this temporal constraint can be avoided by having treatment planning and quality assurance done offline between treatment fractions, which would be more desirable from an operational point of view. Furthermore, we will show in Section 6 that these two methods have the same asymptotic performance.

#### Algorithm 1. Adaptive robust optimization method

---

**Require:** Total number of fractions  $n$ , initial uncertainty set  $P^1$

- 1: Initialize  $i = 1$
- 2: Solve problem (1) with  $P = P^1$  to obtain a beamlet intensity vector  $\mathbf{w}^1$
- 3: Deliver  $\mathbf{w}^1/n$  to the patient
- 4: **while**  $i \leq n - 1$  **do**
- 5: Measure the patient's breathing motion and construct the PMF  $\mathbf{p}^i$
- 6: Generate the new uncertainty set  $P^{i+1}$  from  $P^i$ ,  $\mathbf{p}^i$
- 7: Solve problem (1) with  $P = P^{i+1}$  to obtain a new beamlet intensity vector  $\mathbf{w}^{i+1}$
- 8: Deliver  $\mathbf{w}^{i+1}/n$  to the patient
- 9: Set  $i = i + 1$
- 10: **end while**

---

##### 4.1. Update algorithms

One major step in the adaptive robust approach specified as Algorithm 1 is the generation of the new uncertainty set from the old uncertainty set and the most recent measurement of the patient's breathing motion PMF. We refer to any procedure that

can be used to perform this step as an *uncertainty set update algorithm*. In our computational study we consider two such update algorithms, exponential smoothing and running average, though our theoretical analysis generalizes to a larger class of update algorithms.

Recall that we define an uncertainty set  $P$  in terms of a lower bound vector  $\ell$  and an upper bound vector  $\mathbf{u}$ . Therefore, the uncertainty set for fraction  $i$ , denoted  $P^i$ , is defined by lower and upper bound vectors  $\ell^i$  and  $\mathbf{u}^i$ , respectively. The update algorithms we present will determine lower and upper bound vectors  $\ell^{i+1}$ ,  $\mathbf{u}^{i+1}$  from the most recent vectors  $\ell^i$ ,  $\mathbf{u}^i$  and the most recent PMF  $\mathbf{p}^i$ .

The *exponential smoothing* update algorithm generates a new uncertainty set by taking a convex combination of the most recent lower and upper bound vectors with the most recent PMF. Specifically, given  $\ell^i$ ,  $\mathbf{u}^i$  and  $\mathbf{p}^i$  for day  $i$ ,  $\ell^{i+1}$  and  $\mathbf{u}^{i+1}$  are defined as

$$\ell^{i+1}(x) = (1 - \alpha)\ell^i(x) + \alpha p^i(x), \quad (2)$$

$$\mathbf{u}^{i+1}(x) = (1 - \alpha)\mathbf{u}^i(x) + \alpha \mathbf{p}^i(x), \quad (3)$$

for each  $x \in X$ , where  $\alpha \in [0, 1]$ .

The *running average* update algorithm generates a new uncertainty set by averaging the PMFs that have been observed so far, together with the very first lower and upper bound vectors. Specifically, given  $\ell^1$ ,  $\mathbf{u}^1$  and  $\mathbf{p}^1, \dots, \mathbf{p}^i$  on day  $i$ ,  $\ell^{i+1}$  and  $\mathbf{u}^{i+1}$  are defined as

$$\ell^{i+1}(x) = \frac{1}{i+1} \left( \ell^1(x) + \sum_{j=1}^i p^j(x) \right),$$

$$\mathbf{u}^{i+1}(x) = \frac{1}{i+1} \left( \mathbf{u}^1(x) + \sum_{j=1}^i \mathbf{p}^j(x) \right).$$

for each  $x \in X$ . These expressions can be rewritten as functions of only  $\ell^i(x)$ ,  $\mathbf{u}^i(x)$  and  $\mathbf{p}^i(x)$ :

$$\ell^{i+1}(x) = \frac{i}{i+1} \ell^i(x) + \frac{1}{i+1} p^i(x), \quad (4)$$

$$\mathbf{u}^{i+1}(x) = \frac{i}{i+1} \mathbf{u}^i(x) + \frac{1}{i+1} \mathbf{p}^i(x). \quad (5)$$

#### 4.2. Prescient solutions

Because the update algorithms are myopic and are not guaranteed to correctly anticipate the patient's breathing motion PMF in every fraction, it is important to be able to measure the loss of optimality with respect to a "prescient" solution – one that would be generated if we knew  $\mathbf{p}^1, \dots, \mathbf{p}^n$  ahead of time.

We consider two kinds of prescient solutions. The first type of prescient solution that we define is the *daily prescient* solution. This solution is obtained by setting each  $\mathbf{w}^i$  to the solution of problem (1) with  $P = P^i = \{\mathbf{p}^i\}$ . It is clear that using this solution, the dose to each tumor voxel  $v$  is between  $\theta_v/n$  and  $\gamma\theta_v/n$  in each fraction, which ensures that the cumulative dose to every tumor voxel  $v$  is within the prescribed bounds  $\theta_v$  and  $\gamma\theta_v$ . Furthermore, because each  $\mathbf{w}^i$  is obtained from a nominal problem, the overall dose delivered to the patient should be low compared to other treatments that also deliver the minimum required dose to every tumor voxel.

The second type of prescient solution that we define is the *average prescient* solution. This solution is obtained by solving the nominal problem with  $P = \{\mathbf{p}_{\text{avg}}\}$ , where  $\mathbf{p}_{\text{avg}} = 1/n \cdot \sum_{i=1}^n \mathbf{p}^i$  and setting each  $\mathbf{w}^i$  to the corresponding optimal solution  $\mathbf{w}$  of this problem. It can be easily verified that under this solution, the final dose to every tumor voxel is within the prescribed bounds for that voxel. Furthermore, because the intensity vector of every fraction is obtained from a nominal problem, the overall dose delivered to the patient should be low compared to other treatments that also deliver the minimum required dose to every tumor voxel. While each

tumor voxel  $v \in \mathcal{T}$  is not guaranteed to receive  $\theta_v/n$  in each fraction, the fact that  $\mathbf{w}^i = \mathbf{w}^j$  for any two fractions  $i$  and  $j$  ensures that the dose delivered by this solution in any fraction cannot be too low.

## 5. Computational study

### 5.1. Background

In order to compare our results with previous static robust optimization results, the patient geometry and beam geometry used in this paper were exactly the same as in Bortfeld et al. (2008). The patient geometry consisted of 110,275 voxels, of which 5495 were tumor voxels, and each voxel was of size  $2.93 \times 2.5 \times 2.93$  mm. The set of tumor voxels for this patient case corresponded to the clinical target volume (CTV). The robust problem solved in each fraction of the adaptive method had 122,515 variables and 65,940 constraints. The parameter  $\gamma$  was set to 1.1, while the minimum dose  $\theta_v$  was set to 72 gray for every tumor voxel  $v$ . The set  $X$  consisted of five motion states. The nominal PMF used in the objective function was the same nominal PMF used previously. The adaptive robust and the static robust methods were tested using two different PMF sequences derived from clinical lung cancer patient data obtained at Massachusetts General Hospital; details on these PMF sequences can be found in Section A of the Online Supplement.

The two PMF sequences consisted of the first 30 PMFs from the two experiments, respectively, used in Bortfeld et al. (2008). This number of PMFs was chosen to simulate a realistic 6-week, Monday-to-Friday treatment course. In the discussion that follows, we will focus on the first PMF sequence, as the results obtained from the second PMF sequence were qualitatively very similar; for completeness, results for the second PMF sequence are provided in Section B of the Online Supplement. For the adaptive robust method, we tested the exponential smoothing update algorithm with  $\alpha$  equal to 0.1, 0.5, 0.9 and 1, and the running average update algorithm. The specific values of  $\alpha$  for the exponential smoothing method were chosen to demonstrate the performance of the approach at varying speeds of adaptation (as  $\alpha$  increases, the uncertainty set is adjusted more rapidly to follow the realized PMFs). Three different initial uncertainty sets were used for the adaptive and static robust methods. The first type of uncertainty set was a nominal uncertainty set, with  $\ell$  and  $\mathbf{u}$  both equal to the nominal PMF. The second uncertainty set was the margin uncertainty set ( $P = \mathcal{P}$ , the entire unit simplex) – for this patient case, the margin uncertainty set treatment corresponds to treating the internal target volume (ITV), which is the volume obtained by taking the union of the tumor volume at all of the different phases of the breathing cycle. The third uncertainty set was a robust (intermediately-sized) uncertainty set, which was neither a singleton nor the whole space of PMFs  $\mathcal{P}$ ; for each PMF sequence we used the corresponding robust uncertainty set from Bortfeld et al. (2008). In addition to the adaptive and static robust methods, both prescient solutions were calculated for each of the two PMF sequences.

The experiments were performed on a Dell workstation with a quad-core Intel Xeon 2.67 gigahertz processor and 6 gigabytes of memory. All of the steps of the adaptive robust method – with the exception of the step where the robust problem is solved to obtain  $\mathbf{w}^{i+1}$  – were implemented in MATLAB (The MathWorks, Inc., Natick, Massachusetts). The robust problem in each fraction was solved using CPLEX 12.1 (IBM Corp., Armonk, NY) via AMPL (AMPL Optimization LLC, Albuquerque, NM) and took no more than 20 minutes to solve in the majority of cases using CPLEX 12.1's barrier method.

After the adaptive and static robust methods were executed with each of the PMF sequences described above, the resulting treatments were evaluated by calculating the final dose distribution. The final dose distribution is a vector  $\mathbf{d} \in \mathbb{R}^{|\mathcal{V}|}$  of doses delivered to the voxels in the patient geometry, defined component-wise as

$$d_v = \sum_{i=1}^n \sum_{x \in \mathcal{X}} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n}.$$

From the final dose distribution, we calculated three key statistics: the minimum tumor dose,  $d_T^{\min} = \min_{v \in \mathcal{T}} d_v$ ; the mean left lung dose,  $d_L^{\text{mean}} = \sum_{v \in \mathcal{L}} d_v / |\mathcal{L}|$ , where  $\mathcal{L}$  is the set of voxels belonging to the left lung; and the mean normal tissue dose,  $d_N^{\text{mean}} = \sum_{v \in \mathcal{N}} d_v / |\mathcal{N}|$ , where  $\mathcal{N}$  is the set of normal tissue (non-tumor) voxels. In addition to these statistics, we also compute two other metrics. The mean left lung dose-scaled (MLLD-scaled) tumor dose of a treatment is the minimum dose to the tumor when the dose to the patient is scaled so that the mean left lung dose of the treatment matches the mean left lung dose of the static margin treatment. Similarly, the left lung V20-scaled (LLV20-scaled) tumor dose of a treatment is the minimum dose to the tumor when the dose to the patient is scaled so that the lung V20 value (the volume of the lung receiving 20 gray or more) of the treatment matches the lung V20 value of the static margin treatment.

## 5.2. Results

Table 1 displays dose statistics for the first PMF sequence. In Figs. 1 and 2, we plot the minimum tumor dose as a percentage of 72 gray versus the mean left lung dose as a percentage of the static margin solution for each method applied to the first PMF sequence. Fig. 1 shows all of the treatments, while Fig. 2 zooms in on the adaptive and prescient treatments. On these plots, points obtained by the same method (static, exponential smoothing with parameter  $\alpha$  or running average) from different initial uncertainty sets are joined together by straight lines to illustrate the approximate Pareto frontiers generated by the corresponding method.

In Fig. 3, we also provide dose volume histograms (DVHs) for two treatments obtained from the first patient sequence—the exponential smoothing algorithm with smoothing factor 0.5 and robust uncertainty set (implementation (ES (0.5),R)) and the static method with the same uncertainty set (implementation (S,R)). A DVH is a plot for a particular structure which indicates, for a given level of dose, what volume of the structure receives at least that dose. Ideally, the healthy organ DVHs would be step functions from 100% to 0% at 0 gray (i.e., no dose is delivered to the healthy organs), and the tumor DVH would be a step function from 100% to 0% at 72 gray (i.e., the whole tumor volume receives exactly 72 gray).

## 5.3. Discussion

The computational results generate three important insights. First, the adaptive robust method generally dominates the static robust method, generating solutions that can simultaneously improve on tumor coverage and healthy tissue dose. As an example of this, compare the dose statistics of the static and the exponential smoothing with  $\alpha = 0.5$  methods for the first PMF sequence in Table 1. For the margin (M) uncertainty set, moving from the static (implementation (S,M)) to the exponential smoothing with  $\alpha = 0.5$  (implementation (ES (0.5),M)) treatment results in a reduction in the mean left lung dose of 12.73% of the static margin dose, with virtually no change in tumor coverage. For the robust (R) uncertainty set, moving from the static (implementation (S,R)) to the exponential smoothing with  $\alpha = 0.5$  (implementation (ES (0.5),R)) treatment results in an increase in minimum tumor dose

of 0.85% of the prescribed dose, and a decrease in the mean left lung dose of 2.65% of the static margin dose – both objectives are improved. For the nominal (N) uncertainty set, moving from the static (implementation (S,N)) to the exponential smoothing with  $\alpha = 0.5$  (implementation (ES (0.5),N)) treatment results in a significant increase in the minimum tumor dose of 6.23% of the prescription dose with only a marginal increase in mean left lung dose of 0.91% of the static margin dose. Clinically, the marginal increase in lung dose (0.18 gray) would be acceptable given the increase in minimum tumor dose (4.49 gray). The dominance of the adaptive robust method over the static robust method is also evident in Figs. 1 and 2. On these plots, we can see that the frontiers of all of the adaptive methods are generally below and to the left of the static frontier. The DVHs in Fig. 3 further support this insight. From Fig. 3, we can see that the left lung, esophagus and heart curves for the adaptive treatment are mostly to the left of the corresponding curves for the static treatment, indicating an overall reduction of dose in these structures. The maximum spinal cord dose is slightly higher in the adaptive treatment, but is still clinically acceptable. The homogeneity of the tumor dose is slightly better in the static treatment, as can be seen in the sharper decline of the static treatment tumor dose-volume curve. However, the adaptive treatment has less tumor underdose and still satisfies the maximum tumor dose constraint. Again, Fig. 3 supports the observation that the adaptive robust method generally leads to simultaneous improvement in tumor coverage and healthy tissue sparing.

Alternatively, we can analyze these simultaneous improvements from an isodose perspective, where the dose delivered to the patient under each treatment is normalized to match the static margin treatment in certain lung dose metrics and the resulting (scaled) tumor doses are then compared. From this perspective, our results show that adaptation allows for significant dose escalation beyond the existing static robust approach. For example, if the prescription dose for the static robust treatment (implementation (S,R)) for the first PMF sequence is scaled up so that the mean left lung dose matches that of the static margin treatment (implementation (S,M)), the resulting minimum tumor dose (the MLLD-scaled tumor dose in Table 1) is 80.01 gray. Doing the same for the exponential smoothing treatment with  $\alpha = 0.5$  and the robust uncertainty set (implementation (ES (0.5),R)) for the first PMF sequence, the MLLD-scaled tumor dose is 83.18 gray – an increase of 3.17 gray ( $\approx 4\%$  of the original prescription dose of 72 gray) from the static robust value. Using the estimated relationship between the 5-year local tumor control rate and tumor dose from Kong et al. (2005) for non-small cell lung cancer (NSCLC), this difference of 3.17 gray translates to a 4% increase in local tumor control. In addition to the mean lung dose, dose can also be escalated according to the commonly used lung V20 metric, which is defined as the volume of the lung tissue that receives a dose of 20 gray or higher; the minimum tumor dose that results from this type of escalation is given under the column labeled “LLV20-scaled tumor dose” in Table 1. If the prescription dose is escalated so that the (ES (0.5),R) and (S,R) treatments for the first PMF sequence have the same left lung V20 as the (S,M) treatment, then the adaptive robust solution leads to an additional minimum tumor dose escalation of 6.40 gray ( $\approx 9\%$  of the original prescription dose of 72 gray) beyond that of the static robust approach. Using the relationship from Kong et al. (2005), this difference of 6.40 gray translates approximately to an 8% increase in 5-year local tumor control. Given that the 5-year control rates for tumor doses in the 63–69 gray and 74–84 gray ranges are on the order of 12% and 35% respectively (Kong et al., 2005), the tumor dose escalation and the associated gains in local control rate using both the mean left lung dose and left lung V20 methods that we have described have the potential to be significant.

**Table 1**  
Dose statistics for the first PMF sequence.

Implementation <sup>a</sup>	Min. tumor dose		Mean lung dose		Mean n. tissue dose		MLLD-scaled tumor dose		LLV20-scaled tumor dose	
	Gy	% <sup>b</sup>	Gy	% <sup>c</sup>	Gy	% <sup>d</sup>	Gy	% <sup>e</sup>	Gy	% <sup>f</sup>
(S,N)	67.25	93.40	17.39	85.26	9.04	88.99	78.87	109.48	82.14	114.02
(ES (0.1),N)	70.79	98.32	17.50	85.85	9.05	89.11	82.46	114.46	88.42	122.73
(ES (0.5),N)	71.74	99.63	17.57	86.16	9.06	89.13	83.26	115.57	90.45	125.55
(ES (0.9),N)	71.83	99.77	17.57	86.17	9.06	89.13	83.36	115.71	90.71	125.92
(ES (1),N)	71.84	99.77	17.57	86.16	9.06	89.13	83.38	115.74	90.78	126.01
(RA,N)	71.47	99.26	17.52	85.91	9.06	89.12	83.19	115.47	90.10	125.06
(S,R)	71.37	99.13	18.19	89.20	9.44	92.93	80.01	111.06	83.88	116.43
(ES (0.1),R)	71.80	99.72	17.83	87.43	9.19	90.46	82.12	113.98	87.55	121.53
(ES (0.5),R)	71.99	99.98	17.65	86.55	9.09	89.50	83.18	115.46	90.28	125.32
(ES (0.9),R)	71.98	99.97	17.60	86.33	9.07	89.28	83.38	115.73	90.73	125.95
(ES (1),R)	71.97	99.95	17.59	86.29	9.07	89.26	83.40	115.77	90.80	126.03
(RA,R)	71.88	99.84	17.66	86.61	9.12	89.71	82.99	115.20	89.63	124.41
(S,M)	72.04	100.06	20.39	100.00	10.16	100.00	72.04	100.00	72.04	100.00
(ES (0.1),M)	72.05	100.07	18.31	89.78	9.44	92.91	80.25	111.40	82.45	114.45
(ES (0.5),M)	72.03	100.04	17.80	87.27	9.16	90.10	82.53	114.56	88.86	123.35
(ES (0.9),M)	72.02	100.02	17.71	86.87	9.11	89.67	82.90	115.07	89.87	124.75
(ES (1),M)	71.99	99.98	17.67	86.65	9.09	89.49	83.08	115.32	90.32	125.37
(RA,M)	72.03	100.04	17.94	87.97	9.23	90.86	81.88	113.66	86.36	119.88
(DLYP)	72.00	100.00	17.57	86.18	9.06	89.13	83.55	115.97	91.16	126.54
(AVGP)	72.00	100.00	17.54	86.02	9.06	89.13	83.70	116.19	91.59	127.13

<sup>a</sup> Under "Implementation", the first term indicates the type of method: static (S), exponential smoothing with parameter  $\alpha$  (ES ( $\alpha$ )), running average (RA), daily prescient (DLYP) and average prescient (AVGP). For the non-prescient methods, the second term indicates which of the uncertainty sets described in Section 5.1 was used as the initial uncertainty set: nominal (N), robust (R) and margin (M).

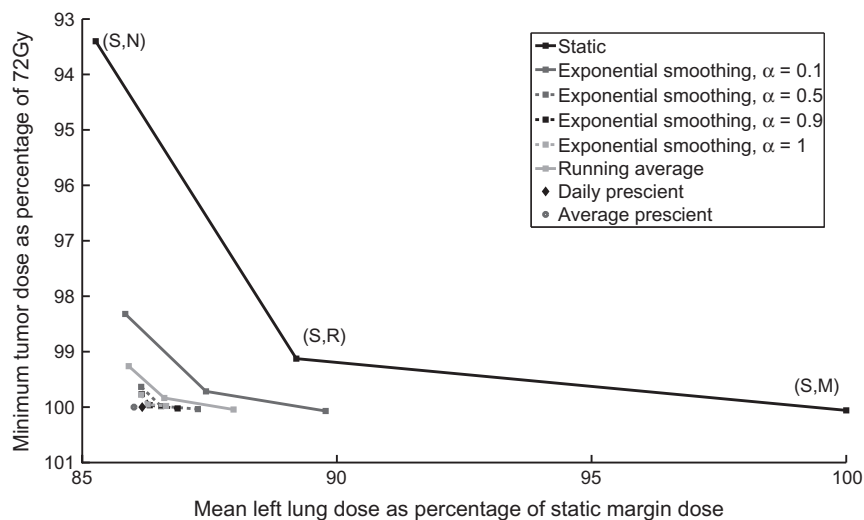
<sup>b</sup> The percentage under "Min. tumor dose" is the minimum tumor dose as a percentage of the prescription dose (72 gray).

<sup>c</sup> The percentage under "Mean lung dose" is the mean left lung dose as a percentage of the mean left lung dose delivered in the static margin treatment, indicated as implementation (S,M).

<sup>d</sup> The percentage under "Mean n. tissue dose" is defined analogously to the percentage under "Mean lung dose".

<sup>e</sup> The percentage under "MLLD-scaled tumor dose" is the minimum tumor dose as a percentage of the static margin minimum tumor dose when the treatment in question is scaled to have the same mean left lung dose as the static margin treatment.

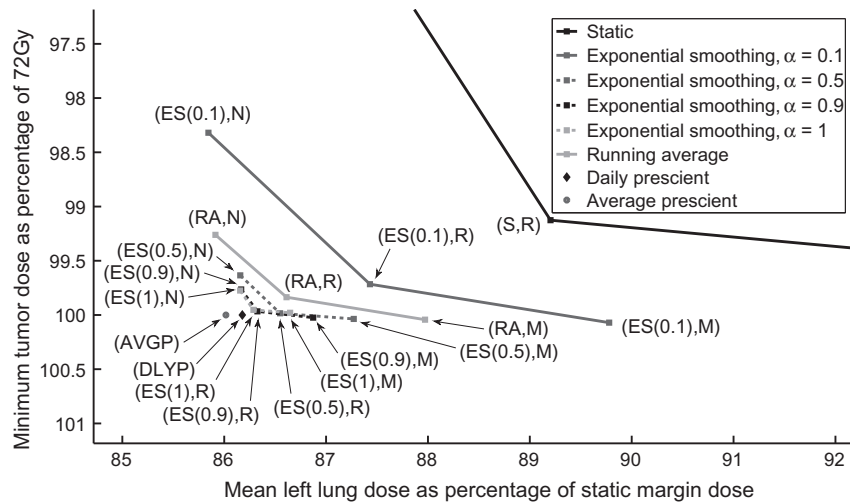
<sup>f</sup> The percentage under "LLV20-scaled tumor dose" is the minimum tumor dose as a percentage of the static margin minimum tumor dose when the treatment in question is scaled to have the same left lung V20 as the static margin treatment.



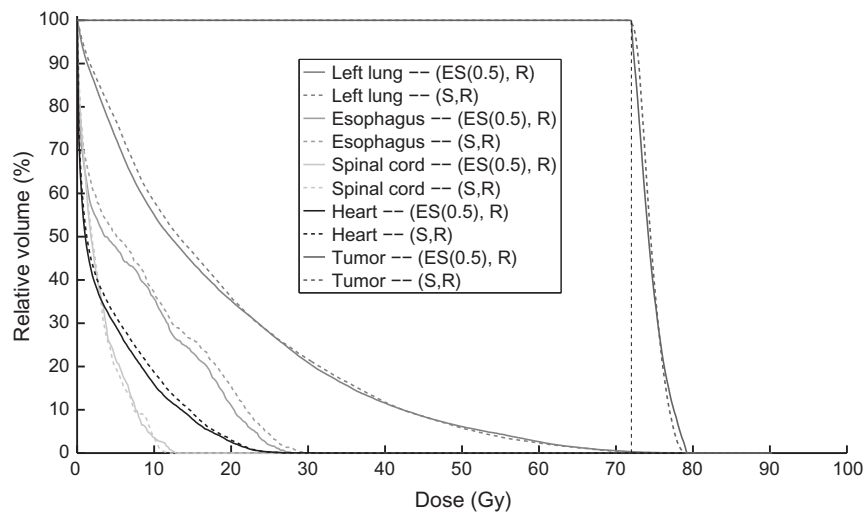
**Fig. 1.** Plot of minimum tumor dose versus mean left lung dose for the different implementations applied to the first PMF sequence (full). The points in the bottom left part of the plot are the adaptive robust and prescient treatments. Point labels follow the notation of Table 1.

The second insight that can be drawn is that even though the adaptive robust method is myopic and does not know what PMFs will be realized in future fractions, its performance for this patient case comes close to the performance of solutions that correctly anticipate future PMFs. This is evident in the fact that the treatments obtained from the adaptive robust method are very close in tumor coverage and healthy tissue dose to the two types of prescient solutions. For example, for the first PMF sequence, when the robust (R) uncertainty set is chosen and exponential smoothing

with  $\alpha = 0.9$  is applied (implementation (ES (0.9),R) in Table 1), the minimum tumor dose as a percentage is 99.97% and the mean left lung dose as a percentage is 86.33%. For both the daily and average prescient treatments, the minimum tumor dose as a percentage is 100%; for the daily prescient treatment, the mean left lung dose as a percentage is 86.18%, whereas for the average prescient treatment, the mean left lung dose as a percentage is 86.02%. Even at slower rates of adaptation, the adaptive robust treatments perform similarly to the prescient solutions. For in-



**Fig. 2.** Plot of minimum tumor dose versus mean left lung dose for the different implementations applied to the first PMF sequence (zoomed-in). Point labels follow the notation of Table 1.



**Fig. 3.** Dose volume histogram comparing the treatments obtained by exponential smoothing with  $\alpha = 0.5$  started from the robust uncertainty set (implementation (ES (0.5),R)) and the static method with the robust uncertainty set (implementation (S,R)) with the first PMF sequence.

stance, using exponential smoothing with  $\alpha = 0.1$  and starting with the robust (R) uncertainty set (implementation (ES (0.1),R)), the minimum tumor dose as a percentage of 72 gray is within 0.5% of the prescient solutions in minimum tumor dose and within 1.5% of the prescient solutions in mean left lung dose. Although our present computational results provide a specific example of this phenomenon, and our analysis in Section 6 provides theoretical justification for this phenomenon under quite general circumstances, further testing on other patient cases is needed to verify the degree to which this insight holds more broadly in practice.

The third insight is that the final dose distribution from any of the adaptive robust methods is far less sensitive to the choice of initial uncertainty set compared to the static method. For instance, for the first PMF sequence, consider the treatments that use exponential smoothing with  $\alpha = 0.9$  (implementations (ES (0.9),N), (ES (0.9),R) and (ES (0.9),M)). For the nominal (N), robust (R) and margin (M) uncertainty sets, the minimum tumor doses achieved using exponential smoothing with  $\alpha = 0.9$  are 71.83, 71.98 and 72.02 gray respectively, while the mean left lung doses are 17.57, 17.60 and 17.71 gray respectively. Visually, this insight is evident in the fact that the adaptive method frontiers are much smaller

than the static frontier on both Figs. 1. This insight is significant because it indicates that the success of this method is not contingent on having accurate knowledge of what the PMFs will be before the treatment begins. From a clinical standpoint, this eliminates an important barrier to clinical implementation, as the treatment planner may not have enough detailed knowledge of a particular patient's breathing patterns before the start of treatment to create a single uncertainty set suitable for the entire treatment course. Moreover, the method proposed in Bortfeld et al. (2008) for the design of uncertainty sets requires historical breathing data from prior patients that is similar to the breathing data of the current patient, which may not always be available.

Given the third insight, it may seem that adding adaptation diminishes the value of robust optimization, as the adaptive treatments that begin with a nominal uncertainty set have similar performance to the adaptive treatments that begin with the other two larger uncertainty sets. Furthermore, the treatments from exponential smoothing with  $\alpha = 1$  ((ES (1),N), (ES (1),R) and (ES (1),M)), for which the uncertainty set from fraction 2 on is always a singleton, perform almost as well as the prescient solutions with respect to the final dose distribution. Although this may be true for



the PMF sequences that we study here, it may not be true in general, as the PMF sequences that we have used here are well-behaved and relatively stable. For instance, sequences that exhibit more variability may cause poor performance if there is too much adaptation and too little robustness (i.e.,  $\alpha$  is set close to one and the initial uncertainty set is small); we explore such a sequence in Section C of the Online Supplement. Similarly, sequences that exhibit transient behavior and drift outside of the initial uncertainty set towards a new value could also necessitate the use of a large uncertainty set early on while the lower and upper bound vectors are still “homing in” on the PMFs. The value of the methodology that we have described in this paper is that it allows the treatment planner to balance robustness, which provides protection during the early fractions of the treatment, and adaptation, which provides protection in the later fractions of the treatment. Thus, this methodology is flexible enough to be applied to settings where the patient’s breathing will be very stable over the course of the treatment (such as the PMF sequences that we study here) and settings where the patient’s breathing could be highly variable early in the treatment or even throughout the entire treatment (such as the PMF sequence studied in Section C of the Online Supplement). Lastly, our computational results correspond to a standard treatment schedule of 30 fractions. When the number of fractions is reduced, as in hypofractionated treatments, a larger proportion of the total prescription dose needs to be delivered in each fraction. In such treatments, the consequences of selecting an inappropriate uncertainty set in each fraction therefore become more severe. Furthermore, the errors in delivered dose made in the early fractions can no longer be “averaged out” due to a large number of fractions. We expect that in such hypofractionated treatments, the quality of the treatments resulting from our method will be more strongly tied to the choice of initial uncertainty set. Thus, the performance of the adaptive robust method (using a non-singleton uncertainty set) will be further differentiated in a positive way from the performance of a nominal (no robustness) adaptive method.

### 6. Asymptotic optimality of the adaptive robust method

In Section 5.3, we observed that the adaptive robust treatments are (a) generally better than the static robust treatments in both lung dose and tumor coverage, (b) very close in quality to both prescient solutions, and (c) relatively insensitive to the choice of initial uncertainty set. In this section, we describe our core theoretical result, which we use to explain these observations. Our result is as follows: if the sequence of breathing motion PMFs  $(\mathbf{p}^i)_{i=1}^\infty$  converges to  $\mathbf{p}^*$  and the uncertainty set update algorithm belongs to a special class of update algorithms, then as the number of fractions  $n$  tends to infinity, the dose distribution eventually enters the epsilon neighborhood of a set  $\mathbf{D}$  of “ideal” dose distributions. The set  $\mathbf{D}$  is the set of dose distributions obtained when any optimal solution to the nominal problem with respect to  $\mathbf{p}^*$  is delivered while the patient is actually breathing according to  $\mathbf{p}^*$ . The proof of this statement, along with the assumptions and the auxiliary results leading to it, is given in Section D of the Online Supplement.

We define the epsilon neighborhood  $U(V, \epsilon)$  of a subset  $V \subseteq \mathbb{R}^n$  as

$$U(V, \epsilon) = \bigcup_{x \in V} B(x, \epsilon),$$

where  $B(x, \epsilon)$  is the open ball of radius  $\epsilon$  about  $x$  ( $B(x, \epsilon) = \{x' \in \mathbb{R}^n \mid \|x - x'\| < \epsilon\}$ ). The norm  $\|\cdot\|$  is the 1-norm on the appropriate finite dimensional real vector space, although technically any  $p$ -norm can be used.

We say that an uncertainty set update algorithm is a *convex-convergent update algorithm* if it satisfies two conditions:

1. For every  $i \in \mathbb{Z}_+$  (the set of positive integers),  $\ell^{i+1}$  and  $\mathbf{u}^{i+1}$  can be written as convex combinations of  $\mathbf{p}^i$  with  $\ell^i$  and  $\mathbf{u}^i$  respectively; that is, there exists an  $\alpha_i \in [0, 1]$  such that

$$\begin{aligned} \ell^{i+1} &= (1 - \alpha_i)\ell^i + \alpha_i\mathbf{p}^i, \\ \mathbf{u}^{i+1} &= (1 - \alpha_i)\mathbf{u}^i + \alpha_i\mathbf{p}^i. \end{aligned}$$

2. If  $\mathbf{p}^i \rightarrow \mathbf{p}^*$  as  $i \rightarrow \infty$ , then  $\ell^i \rightarrow \mathbf{p}^*$  and  $\mathbf{u}^i \rightarrow \mathbf{p}^*$  as  $i \rightarrow \infty$ .

It can be easily verified that both the exponential smoothing and running average update algorithms are convex-convergent.

Let  $\mathbf{w}^*(\ell, \mathbf{u})$  denote the set of optimal solutions for the robust problem (1) with its uncertainty set  $P$  defined by lower and upper bound vectors  $\ell$  and  $\mathbf{u}$ . Let  $\mathbf{w}^*(\mathbf{p}^*)$  denote the set of optimal solutions for the nominal problem with respect to  $\mathbf{p}^*$ .

Define the set  $\mathbf{D}$ , the set of dose distributions when  $\mathbf{p}^*$  is realized while  $\mathbf{w} \in \mathbf{w}^*(\mathbf{p}^*)$  is being delivered, as

$$\mathbf{D} = \{\mathbf{d} \in \mathbb{R}^{|\mathcal{V}|} \mid \mathbf{d} = \Delta\mathbf{p}^*\mathbf{w} \text{ for some } \mathbf{w} \in \mathbf{w}^*(\mathbf{p}^*)\}$$

where the product  $\Delta\mathbf{p}\mathbf{w}$  is defined as

$$\Delta\mathbf{p}\mathbf{w} = \left[ \sum_{x \in X} \sum_{b \in B} \Delta_{v,x,b} p(x) w_b \right]_{v \in \mathcal{V}}.$$

For an  $n$  fraction treatment during which the intensity vectors were  $\mathbf{w}^1, \dots, \mathbf{w}^n$  and the realized PMFs were  $\mathbf{p}^1, \dots, \mathbf{p}^n$ , the final dose distribution is given by  $1/n \cdot \sum_{i=1}^n \Delta\mathbf{p}^i\mathbf{w}^i$ . Our core result can be stated as Theorem 1.

**Theorem 1.** (Adaptive robust dose convergence) *Let  $(\mathbf{p}^i)_{i=1}^\infty$  be a sequence of PMFs that converges to  $\mathbf{p}^*$ . Let  $(\ell^i)_{i=1}^\infty$  and  $(\mathbf{u}^i)_{i=1}^\infty$  be lower and upper bound sequences generated from  $(\mathbf{p}^i)_{i=1}^\infty$  by a convex-convergent update algorithm. For each  $i \in \mathbb{Z}_+$ , let  $\mathbf{w}^i \in \mathbf{w}^*(\ell^i, \mathbf{u}^i)$ . Then for every  $\epsilon > 0$ , there exists an  $N \in \mathbb{Z}_+$  such that for all  $n > N$ ,*

$$\frac{1}{n} \sum_{i=1}^n \Delta\mathbf{p}^i\mathbf{w}^i \in U(\mathbf{D}, \epsilon).$$

The theorem states that the dose distribution obtained via any convex-convergent update algorithm approaches a set of optimal dose distributions  $\mathbf{D}$ . The intuition behind this result is that as  $i \rightarrow \infty$ ,  $\mathbf{w}^i$  is increasingly similar to some solution in  $\mathbf{w}^*(\mathbf{p}^*)$ , and  $\mathbf{p}^i$  is increasingly similar to  $\mathbf{p}^*$ . Therefore, as  $n$  tends to infinity, each term in the tail of the sum  $\sum_{i=1}^n \Delta\mathbf{p}^i\mathbf{w}^i$  becomes increasingly similar to some  $\mathbf{d} \in \mathbf{D}$ .

This set  $\mathbf{D}$  of optimal dose distributions is the same set that is approached by both the daily prescient and average prescient dose distributions.

**Theorem 2.** *Let  $(\mathbf{p}^i)_{i=1}^\infty$  be a sequence of PMFs that converges to  $\mathbf{p}^*$ .*

- (a) (Daily prescient dose convergence.) *For each  $i \in \mathbb{Z}_+$ , let  $\mathbf{w}^i \in \mathbf{w}^*(\mathbf{p}^i, \mathbf{p}^i)$ . Then for every  $\epsilon > 0$ , there exists an  $N \in \mathbb{Z}_+$  such that for all  $n > N$ ,*

$$\frac{1}{n} \sum_{i=1}^n \Delta\mathbf{p}^i\mathbf{w}^i \in U(\mathbf{D}, \epsilon).$$

- (b) (Average prescient dose convergence.) *For each  $n \in \mathbb{Z}_+$ , let  $\mathbf{w}^n \in \mathbf{w}^*(1/n \cdot \sum_{i=1}^n \mathbf{p}^i, 1/n \cdot \sum_{i=1}^n \mathbf{p}^i)$ . Then for every  $\epsilon > 0$ , there exists an  $N \in \mathbb{Z}_+$  such that for all  $n > N$ ,*

$$\frac{1}{n} \sum_{i=1}^n \Delta\mathbf{p}^i\mathbf{w}^n \in U(\mathbf{D}, \epsilon).$$

When the nominal problem with respect to  $\mathbf{p}^*$  has a unique optimal solution, the difference between the adaptive robust and either prescient dose distribution tends to zero.

**Corollary 1.** Let  $(\mathbf{p}^i)_{i=1}^\infty$  be a sequence of PMFs that converges to  $\mathbf{p}^*$ . For each  $i \in \mathbb{Z}_+$ , let  $\mathbf{w}_{AR}^i \in \mathbf{w}^*(\ell^i, \mathbf{u}^i)$  (the adaptive robust intensity vector for fraction  $i$ ), where  $(\ell^i)_{i=1}^\infty$  and  $(\mathbf{u}^i)_{i=1}^\infty$  are obtained by a convex-convergent update algorithm, and let  $\mathbf{w}_{DP}^i \in \mathbf{w}^*(\mathbf{p}^i, \mathbf{p}^i)$  (the daily prescient intensity vector for fraction  $i$ ). For each  $n \in \mathbb{Z}_+$ , let  $\mathbf{w}_{AP}^n \in \mathbf{w}^*(1/n \cdot \sum_{i=1}^n \mathbf{p}^i, 1/n \cdot \sum_{i=1}^n \mathbf{p}^i)$  (the average prescient intensity vector for fraction  $n$ ). If the nominal problem with respect to  $\mathbf{p}^*$  has a unique optimal solution (i.e.,  $\mathbf{w}^*(\mathbf{p}^*)$  is a singleton), then

$$\lim_{n \rightarrow \infty} \left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}_{AR}^i - \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}_{DP}^i \right\| = \lim_{n \rightarrow \infty} \left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}_{AR}^i - \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}_{AP}^n \right\| = 0.$$

### 6.1. Three theoretical insights

These results directly relate to the three insights of Section 5.3. With regard to the first insight – that adaptive robust treatments generally dominate static robust treatments – suppose that  $\mathbf{w}$  is a solution of the robust problem with uncertainty set  $P$ . The static robust dose distribution after  $n$  fractions is given by

$$\frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w} = \Delta \sum_{i=1}^n \frac{\mathbf{p}^i}{n} \mathbf{w}.$$

If  $(\mathbf{p}^i)_{i=1}^\infty$  converges to  $\mathbf{p}^*$ , then this dose distribution converges to  $\Delta \mathbf{p}^* \mathbf{w}$ . Whether or not this dose distribution is satisfactory depends on whether or not  $\mathbf{p}^*$  is in  $P$ . If  $\mathbf{p}^* \in P$ , then every tumor voxel will receive sufficient dose in the limiting dose distribution. However, if  $\mathbf{p}^* \notin P$ , then there is no guarantee that every tumor voxel will sufficient dose. In contrast to the static method, the adaptive robust method always leads to sufficient tumor dose in the limit. The reason for this is that every dose distribution  $\mathbf{d}$  in  $\mathbf{D}$  has the property that all tumor voxels receive the minimum prescription dose. This property is a consequence of the fact that  $\mathbf{d} = \Delta \mathbf{p}^* \mathbf{w}$  for some intensity vector  $\mathbf{w}$  corresponding to the singleton uncertainty set  $\{\mathbf{p}^*\}$ ; by the definition of problem (1), it is straightforward to see that the dose to tumor voxel  $v$  satisfies

$$d_v = \sum_{x \in X} \sum_{b \in B} \Delta_{v,x,b} \mathbf{p}^*(x) \mathbf{w}_b \geq \theta_v.$$

Furthermore, observe that every beamlet intensity vector in  $\mathbf{w}^*(\mathbf{p}^*)$  is an optimal solution of the nominal problem with respect to  $\mathbf{p}^*$ . As we discussed earlier, solutions to nominal problems are ones that generally result in the lowest level of normal tissue damage. Therefore, not only do the dose distributions in  $\mathbf{D}$  meet the tumor dose requirements, but they are also likely to lead to low levels of healthy tissue dose. In contrast, a static robust treatment will generally not possess both of these qualities simultaneously. If the treatment planner selects  $P$  to be a singleton (i.e.  $\mathbf{w}$  is the solution of the nominal problem with a specific PMF), then the treatment is likely to result in low healthy tissue dose, but the treatment planner is taking a major risk by assuming what  $\mathbf{p}^*$  will be. If the treatment planner selects a larger  $P$ , then there is a better chance that  $\mathbf{p}^*$  will be contained in  $P$ , but at the cost of increased healthy tissue dose.

With regard to the second insight – that the adaptive robust and prescient treatments become very similar – we know that both the average and daily prescient dose distributions also approach  $\mathbf{D}$ . When the nominal problem with respect to  $\mathbf{p}^*$  has a unique optimal solution, the difference between the adaptive robust dose distribution and either prescient dose distribution must tend to zero

by Corollary 1. In the case that there are multiple optimal solutions to the nominal problem with respect to  $\mathbf{p}^*$ , it seems reasonable to expect that the dose distributions that are achieved when these solutions are delivered and the patient breathes according to  $\mathbf{p}^*$  should not differ significantly. Therefore, the difference between either prescient dose distribution and the adaptive robust dose distribution should also become small in the limit.

With regard to the third insight – that the adaptive robust method is relatively insensitive to the choice of initial uncertainty set – we can see that in the proof of Theorem 1, the choice of  $\ell^1$  and  $\mathbf{u}^1$  is not important: the dose distribution always approaches the set  $\mathbf{D}$ , no matter what  $\ell^1$  and  $\mathbf{u}^1$  are. Again, if  $\mathbf{D}$  is a singleton, then any two dose distributions that are obtained using the same adaptive method but different  $\ell^1$  and  $\mathbf{u}^1$  vectors will converge to the single distribution in  $\mathbf{D}$ . Since the two dose distributions converge to the same limiting dose distribution, the difference between those dose distributions must tend to zero. In practice, we will still see some dependence on the initial uncertainty set due to the finite number of fractions, but with a large number of fractions, the effect of the initial uncertainty set will become diminished.

We note that although the theoretical results we have developed are helpful in understanding the performance of our methods, it is unlikely that a patient PMF sequence in real life will converge to a single limiting PMF. A more plausible scenario may be that after a transient period, during which the patient is being acclimated to the treatment, the PMF sequence converges to a set of PMFs, and on each day the patient’s PMF varies within this set, without converging to a single limiting PMF. Our results cannot directly answer this question without modifications. However, our core theoretical result (Theorem 1) relies on the fact that for every  $\epsilon$ , which represents the error between the set of ideal dose distributions  $\mathbf{D}$  corresponding to a limiting  $\mathbf{p}^*$  and the realized dose distribution  $1/n \cdot \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i$ , we are able to get  $\mathbf{p}^i$  within a suitable distance of  $\mathbf{p}^*$ . In the case that we can only say that  $\mathbf{p}^i$  gets arbitrarily close to a set of limiting PMFs, we expect that Theorem 1 can be modified to hold for  $\epsilon \geq \epsilon_0$ , where  $\epsilon_0 > 0$  depends on the size of the set of limiting PMFs, as opposed to any arbitrary  $\epsilon > 0$ . We expect that as the size of the limiting set of PMFs decreases, the error between the realized dose distribution and an ideal set of distributions (obtained by applying the daily prescient method to the sequence) will also decrease; Theorem 1 would then present an extreme case where the limiting set of PMFs is a singleton. The real patient PMF sequences used in Section 5 are relatively stable, which is why our present theoretical results agree with our computational insights. We also expect that when the PMFs converge only to a set, the performance may be more sensitive to the choice of the update algorithm, and how the convergence to a set is preserved (or not preserved) for the sequence of lower and upper bound vectors generated by the update algorithm. In Section C of the Online Supplement, we provide an example of such a PMF sequence that does not stabilize. For this sequence, we indeed find that the performance gap between the exponential smoothing and prescient solutions is generally quite large; however, the running average algorithm is still able to achieve good performance.

## 7. Conclusions

In this study, we developed an optimization method for fractionated IMRT treatment planning that combats breathing motion uncertainty by generalizing the methodologies of adaptive radiation therapy and robust optimization. We studied the effectiveness of our method on two PMF sequences obtained from clinical lung cancer patients and showed that our adaptive robust optimization method improves upon the static method in both tumor coverage and healthy tissue sparing simultaneously. We also studied the

effectiveness of our method from a theoretical standpoint and proved that the adaptive robust dose distribution approaches a limiting set of optimal dose distributions. Even though this adaptive approach is myopic, it performs well when the PMF sequence converges, and for some uncertainty set update algorithms, performs well even when the PMF sequence does not converge. Also, as we have shown empirically and theoretically, our method does not require perfectly accurate information to be available before the start of treatment. Overall, the clinical value of this method is that it allows for the tumor dose to be safely escalated without leading to additional healthy tissue toxicity, which may ultimately improve the rate of patient survival.

Future work should consider what results can be obtained when the convergence assumption is relaxed and the PMFs are instead drawn from a distribution on the probability simplex. The adaptive robust method we have studied can also be modified by changing the right hand side of problem (1) after each fraction, in order for the treatment to account for the dose delivered in prior fractions and to compensate for daily underdosages and overdosages in prior fractions. With such a method, we hypothesize that the daily dose delivered to the tumor will become increasingly inhomogeneous as the treatment progresses, with some parts of the tumor receiving significantly less than the daily prescribed amount. This seems to be the case in some preliminary computational tests and is the subject of ongoing work.

Finally, we note that although we have specifically studied fractionated IMRT planning under breathing motion uncertainty, the methodology and analysis we have presented here is of value to other practical problems outside of the domain of healthcare. Specifically, this method can be applied to problems involving sequential decision making under uncertainty, when new measurements of the uncertain effect are obtained over the planning horizon. This approach may be particularly attractive when it is computationally inexpensive to solve the robust optimization problem with the updated uncertainty set at each stage. Overall, the empirical performance and theoretical properties we have shown in this paper suggest that this type of adaptive robust optimization method could be a viable approach for other multi-stage decision making problems under uncertainty.

## Acknowledgements

The authors thank the referees for their careful reading of the paper and for their thoughtful comments which have helped to improve the paper. The authors thank Thomas Bortfeld of Massachusetts General Hospital for providing the patient data used in Section 5. The work of both authors was partially supported by the Natural Sciences and Engineering Research Council of Canada (NSERC) and the Canadian Institutes of Health Research (CIHR). The work of the second author was partially supported by an NSERC CGS-M award.

## Appendix A. Supplementary material

Supplementary data associated with this article can be found, in the online version, at <http://dx.doi.org/10.1016/j.ejor.2013.06.003>.

## References

American Cancer Society, 2010. Cancer Facts and Figs. 2010. American Cancer Society, Atlanta.

Ben-Tal, A., El Ghaoui, L., Nemirovski, A., 2009. Robust Optimization. Princeton University Press.

Bertsimas, D., Brown, D., Caramanis, C., 2011. Theory and applications of robust optimization. *SIAM Review* 53, 464–501.

Bortfeld, T., Chan, T., Trofimov, A., Tsitsiklis, J., 2008. Robust management of motion uncertainty in intensity-modulated radiation therapy. *Operations Research* 56, 1461–1473.

Bortfeld, T., Jiang, S., Rietzel, E., 2004. Effects of motion on the total dose distribution. *Seminars in Radiation Oncology* 14, 41–51.

Bortfeld, T., Jokivarsi, K., Goitein, M., Kung, J., Jiang, S., 2002. Effects of intra-fraction motion on IMRT dose delivery: statistical analysis and simulation. *Physics in Medicine and Biology* 47, 2203–2220.

Canadian Cancer Society's Steering Committee, 2010. Canadian Cancer Statistics 2010. Canadian Cancer Society, Toronto.

Chan, T., Bortfeld, T., Tsitsiklis, J., 2006. A robust approach to IMRT optimization. *Physics in Medicine and Biology* 51, 2567–2583.

Chu, M., Zinchenko, Y., Henderson, S., Sharpe, M., 2005. Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty. *Physics in Medicine and Biology* 50, 5463–5477.

Cox, J., Azarnia, N., Byhardt, R., Shin, K., Emami, B., Pajak, T., 1990. A randomized phase I/II trial of hyperfractionated radiation therapy with total doses of 60.0 Gy–79.2 Gy: possible survival benefit with greater than or equal to 69.6 Gy in favorable patients with Radiation Therapy Oncology Group stage III non-small-cell lung carcinoma: report of Radiation Therapy Oncology Group 83-11.. *Journal of Clinical Oncology* 8, 1543–1555.

Deng, G., Ferris, M., 2008. Neuro-dynamic programming for fractionated radiotherapy planning. In: Alves, C., Pardalos, P., Vicente, L. (Eds.), *Optimization in Medicine*. Springer Verlag, pp. 47–70.

Ferris, M., Voelker, M., 2004. Fractionation in radiation treatment planning. *Mathematical Programming* 101, 387–413.

Keller, H., Ritter, M., Mackie, T., 2003. Optimal stochastic correction strategies for rigid-body target motion. *International Journal of Radiation Oncology, Biology, Physics* 55, 261–270.

Kong, F., Ten Haken, R., Schipper, M., Sullivan, M., Chen, M., Lopez, C., Kalemkerian, G., Hayman, J., 2005. High-dose radiation improved local tumor control and overall survival in patients with inoperable/unresectable non-small-cell lung cancer: long-term results of a radiation dose escalation study. *International Journal of Radiation Oncology, Biology, Physics* 63, 324–333.

Lam, K., Haken, R., Litzenberg, D., Balter, J., Pollock, S., 2005. An application of Bayesian statistical methods to adaptive radiotherapy. *Physics in Medicine and Biology* 50, 3849–3858.

Löf, J., Lind, B., Brahme, A., 1998. An adaptive control algorithm for optimization of intensity modulated radiotherapy considering uncertainties in beam profiles, patient set-up and internal organ motion. *Physics in Medicine and Biology* 43, 1605–1628.

Lu, W., Olivera, G., Chen, Q., Ruchala, K., Haimerl, J., Meeks, S., Langen, K., Kupelian, P., 2006. Deformable registration of the planning image (kVCT) and the daily images (MVCT) for adaptive radiation therapy. *Physics in Medicine and Biology* 51, 4357–4374.

Lujan, A., Balter, J., Ten Haken, R., 2003. A method for incorporating organ motion due to breathing into 3D dose calculations in the liver: sensitivity to variations in motion. *Medical Physics* 30, 2643–2649.

Mell, L., Mehrotra, A., Mundt, A., 2005. Intensity-modulated radiation therapy use in the US, 2004. *Cancer* 104, 1296–1303.

Mohan, R., Zhang, X., Wang, H., Kang, Y., Wang, X., Liu, H., Ang, K., Kuban, D., Dong, L., 2005. Use of deformed intensity distributions for on-line modification of image-guided IMRT to account for interfractional anatomic changes. *International Journal of Radiation Oncology, Biology, Physics* 61, 1258–1266.

Ohara, K., Okumura, T., Akisada, M., Inada, T., Mori, T., Yokota, H., Calaguas, M., 1989. Irradiation synchronized with respiration gate. *International Journal of Radiation Oncology, Biology, Physics* 17, 853–857.

Ólafsson, A., Wright, S., 2006. Efficient schemes for robust IMRT treatment planning. *Physics in Medicine and Biology* 51, 5621–5642.

Perez, C., Bauer, M., Edelstein, S., Gillespie, B., Birch, R., 1986. Impact of tumor control on survival in carcinoma of the lung treated with irradiation. *International Journal of Radiation Oncology, Biology, Physics* 12, 539–547.

Rehbinder, H., Forsgren, C., Löf, J., 2004. Adaptive radiation therapy for compensation of errors in patient setup and treatment delivery. *Medical Physics* 31, 3363–3371.

Romeijn, H., Ahuja, R., Dempsey, J., Kumar, A., 2006. A new linear programming approach to radiation therapy treatment planning problems. *Operations Research* 54, 201.

Romeijn, H., Dempsey, J., 2008. Intensity modulated radiation therapy treatment plan optimization. *Top* 16, 215–243.

Sheng, K., Cai, J., Brookeman, J., Molloy, J., Christopher, J., Read, P., 2006. A computer simulated phantom study of tomotherapy dose optimization based on probability density functions (PDF) and potential errors caused by low reproducibility of PDF. *Medical Physics* 33, 3321–3326.

Sir, M., 2007. Optimization of Radiotherapy Considering Uncertainties Caused by Daily Setup Procedures and Organ Motion. Ph.D. Thesis, University of Michigan.

Sir, M., Epelman, M., Pollock, S., 2012. Stochastic programming for off-line adaptive radiotherapy. *Annals of Operations Research* 196, 767–797.

Thames, H., Hendry, J., 1987. Fractionation in Radiotherapy. Taylor and Francis, London.

Thieke, C., Bortfeld, T., Küfer, K., 2002. Characterization of dose distributions through the max and mean dose concept. *Acta Oncologica* 41, 158–161.

Toschi, L., Cappuzzo, F., Janne, P., 2007. Evolution and future perspectives in the treatment of locally advanced non-small cell lung cancer. *Annals of Oncology* 18, ix150–ix155.

Wu, Q., Liang, J., Yan, D., 2006. Application of dose compensation in image-guided radiotherapy of prostate cancer. *Physics in Medicine and Biology* 51, 1405–1419.

- Wu, Q., Thongphiew, D., Wang, Z., Mathayomchan, B., Chankong, V., Yoo, S., Lee, W., Yin, F., 2008. On-line re-optimization of prostate IMRT plans for adaptive radiation therapy. *Physics in Medicine and Biology* 53, 673–691.
- Yan, D., Vicini, F., Wong, J., Martinez, A., 1997a. Adaptive radiation therapy. *Physics in Medicine and Biology* 42, 123–132.
- Yan, D., Wong, J., Vicini, F., Michalski, J., Pan, C., Frazier, A., Horwitz, E., Martinez, A., 1997b. Adaptive modification of treatment planning to minimize the deleterious effects of treatment setup errors. *International Journal of Radiation Oncology, Biology, Physics* 38, 197–206.
- Yan, D., Ziaja, E., Jaffray, D., Wong, J., Brabbins, D., Vicini, F., Martinez, A., 1998. The use of adaptive radiation therapy to reduce setup error: a prospective clinical study. *International Journal of Radiation Oncology, Biology, Physics* 41, 715–720.
- Yu, C., Jaffray, D., Wong, J., 1998. The effects of intra-fraction organ motion on the delivery of dynamic intensity modulation. *Physics in Medicine and Biology* 43, 91–104.
- de la Zerda, A., Armbruster, B., Xing, L., 2007. Formulating adaptive radiation therapy (ART) treatment planning into a closed-loop control framework. *Physics in Medicine and Biology* 52, 4137–4153.