

Formulating adaptive radiation therapy (ART) treatment planning into a closed-loop control framework*

Adam de la Zerda¹, Benjamin Armbruster² and Lei Xing³

¹ Department of Electrical Engineering, Stanford University, Stanford, CA 94305-9505, USA

² Department of Management Science and Engineering, Stanford University, Stanford, CA 94305-4026, USA

³ Department of Radiation Oncology, Stanford University, Stanford, CA 94305-5847, USA

E-mail: lei@reyes.stanford.edu

Received 17 January 2007, in final form 25 April 2007

Published 14 June 2007

Online at stacks.iop.org/PMB/52/4137

Abstract

While ART has been studied for years, the specific quantitative implementation details have not. In order for this new scheme of radiation therapy (RT) to reach its potential, an effective ART treatment planning strategy capable of taking into account the dose delivery history and the patient's on-treatment geometric model must be in place. This paper performs a theoretical study of dynamic closed-loop control algorithms for ART and compares their utility with data from phantom and clinical cases. We developed two classes of algorithms: those *Adapting to Changing Geometry* and those *Adapting to Geometry and Delivered Dose*. The former class takes into account organ deformations found just before treatment. The latter class optimizes the dose distribution accumulated over the entire course of treatment by adapting at each fraction, not only to the information just before treatment about organ deformations but also to the dose delivery history. We showcase two algorithms in the class of those *Adapting to Geometry and Delivered Dose*. A comparison of the approaches indicates that certain closed-loop ART algorithms may significantly improve the current practice. We anticipate that improvements in imaging, dose verification and reporting will further increase the importance of adaptive algorithms.

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

1. Introduction

Current IMRT treatment plan optimization and dose delivery are two decoupled steps (AAPM IMRT Sub-committee 2003, Webb 2001). In each fraction, the patient geometry is hardly the

* Presented at ASTRO Annual Meeting, 2006, Philadelphia, PA.

same as in the pre-treatment CT simulation. A commonly used method to take the uncertainty into account is to add a safety margin, whose size is based on population statistics, to the target and sensitive structures (Balter *et al* 1995, van Herk 2006). This significantly compromises the success of radiation therapy.

Recently, cone-beam computed tomography (CBCT) integrated with a medical linear accelerator has become available and promises to improve the situation. CBCT provides a valuable 3D (or even possibly 4D) geometric model of the patient in the treatment position and allows for verification of the delivered dose distribution (Xing *et al* 2006). This not only affords an opportunity for on-line correction of patient setup error and inter-fraction rigid motion (Oldham *et al* 2005, Pouliot *et al* 2005), but also allows dose recalculation and adaptive radiation therapy (ART) (Langen *et al* 2005, Yang *et al* 2007), which uses the volumetric information to adjust the treatment plan each fraction to the updated patient anatomy and positioning. A significant promise of ART is the optimal compensation of uncertainties including organ deformation and inter-fraction organ motion as well as dosimetric errors incurred in previous fractions (Trofimov *et al* 2005, Wu *et al* 2006, Mohan *et al* 2005, de la Zerda *et al* 2006, Olivera *et al* 2006). To realize ART clinically and maximally exploit the potential of this new form of image guided radiation therapy (IGRT), a robust inverse planning strategy for ART must be in place.

The purpose of this work is to develop dynamic closed-loop control strategies for ART inverse planning and to demonstrate their utility with data from phantom and clinical cases. Closed-loop control algorithms are a general tool for dealing with time-dependent systems (Widrow and Stearns 1985, Widrow and Walach 1995) and are used to solve a variety of problems including automobile cruise control, supply chain optimization, computer chip design and spaceship navigation. The algorithms in all these applications share the same basic closed-loop control framework of repeated re-evaluation and re-planning. ART is a natural application for closed-loop control because CBCT provides frequently updated system information. To meet different clinical requirements, we investigate a few variants of these algorithms.

The next section starts with a brief review of the foundations of closed-loop control. Section 2 continues by presenting a formalism for ART planning, the intensity-modulated radiation therapy (IMRT) approach, and our new ART algorithms. In section 3, we demonstrate the significance of our new approach with phantom and patient case studies. Section 4 takes a broader perspective to discuss this work and directions for future research. We conclude in section 5.

2. Methods and materials

This section starts with an overview of closed-loop control and then later describes various closed-loop control algorithms for ART before summarizing these algorithms in table 2.

2.1. Foundations of closed-loop control

The use of feedback makes closed-loop control unique (Widrow and Stearns 1985, Widrow and Walach 1995, Yan *et al* 1997). In open-loop control, figure 1(a), some *input* is fed to the *controller* which then decides on an *action*. The controller then performs that action on the *system* (the system being the physical or biological system we want to control), effectively putting the system into a new state (which hopefully is 'better' in some sense). We could simulate an open-loop control context by connecting the output of the algorithm implemented

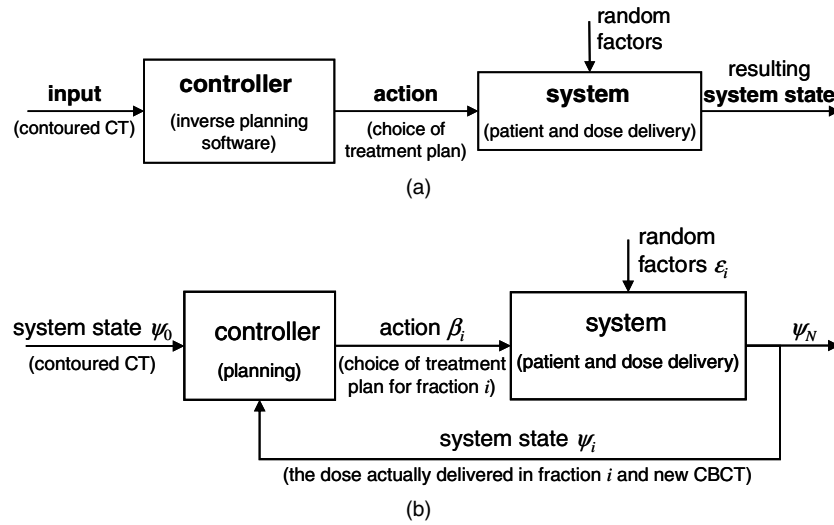


Figure 1. (a) Open-loop control (conventional treatment planning), (b) the closed-loop control framework.

by the controller to a simulation of the system: the input is fed to the algorithm which feeds an action to the system simulator which outputs the new state of the system.

Conventional inverse planning is a form of open-loop control. The *input* is a set of contoured CT images or the geometric model of the patient. The *controller* is the inverse planning software, and it outputs a treatment plan (the action). Finally, a simulator of the system (modeling the execution of the plan) uses the treatment plan and some random factors to output the resulting system state.

In closed-loop control, the controller is not run just once but repeatedly, each time receiving the current state of the system as its input (Widrow and Stearns 1985, Widrow and Walach 1995, Yan *et al* 1997). The term *feedback* is used because the effect of the current action (the new state of the system) is used by the controller to plan the next action. Suppose we have N time periods (treatment fractions), and we let β_i be the action in period i (treatment plan for fraction i) and ψ_i the resulting system state at the end of period i . Then a closed-loop controller is a function ξ_i producing an action for the current period from the state at the end of the last period $\beta_i = \xi_i(\psi_{i-1})$.

We propose to plan ART using closed-loop control, figure 1(b), where the controller is the inverse planning software and the state of the system after fraction i , $\psi_i = (\psi_i^{\text{geometry}}, \psi_i^{\text{cumdose}})$, has two parts: the patient's geometric model derived from contoured CBCT images, ψ_i^{geometry} , and the cumulative dose, ψ_i^{cumdose} , delivered in fractions 1 through i . (Clearly, $\psi_0^{\text{cumdose}} = 0$.) Unlike the open-loop control in conventional RT, in ART the controller outputs a plan (the action) for only the current treatment fraction. Table 1 summarizes the relationship between control theory and treatment planning.

Closed-loop control algorithms (which we propose for ART) are richer than static open-loop control algorithms (found for example in conventional RT). A static open-loop controller determines a treatment plan, $(\beta_1, \dots, \beta_N)$, and therefore the final state of the system ψ_N is based only on the initial state of the system ψ_0 , whereas dynamic closed-loop algorithms re-evaluate the state each time step and based on that, decide their current action. Because of that, the treatment plan, $(\beta_1, \dots, \beta_N)$, depends not only on the initial state ψ_0 , but also on the progress made during treatment, $(\psi_0, \dots, \psi_{N-1})$. Hence closed-loop control algorithms

Table 1. Relationship between control theory and treatment planning.

| Closed-loop control | Radiation therapy |
|--|--|
| Time period (i) | Fraction |
| Input (ψ_0) | Contoured CT images (ψ_0^{geometry}) |
| Action in period i , (β_i) | Treatment plan for fraction i |
| System state after period i , (ψ_i) | Contoured CBCT images (ψ_i^{geometry}), and cumulative delivered dose (ψ_i^{cumdose}) |
| Controller | RT inverse planning software |
| System | Patient geometry and treatment delivery |
| Random factors/noise in period i , (ε_i) | Setup errors, delivery errors, and organ deformations in fraction i |

are more likely to produce better results than static open-loop control algorithms for time-dependent systems.

Treatment planning is a critical step to realizing the potential of ART in clinical practice. Current IMRT inverse planning is designed for a conventional fractionated treatment scheme (Bortfeld 1999, Webb 2001, Censor 2003, Xing *et al* 2005, Xiao *et al* 2000, Yang and Xing 2004b) and is incapable of planning an adaptive treatment with consideration of the dose delivery history and updated patient geometry. However, anticipated improvements in imaging and verifying the dose and anatomy during dose delivery drive the need for sophisticated algorithms for ART planning. The purpose of this work is to establish a dynamic control framework for adaptive radiation therapy to explore the subtleties of possible planning algorithms. This planning framework allows us to individualize radiation therapy for each patient by taking advantage of newly available volumetric imaging information.

2.2. Formalization of ART plan optimization

Consider plans, $(\beta_1, \dots, \beta_N)$, for N fractions. Suppose at voxel v , the importance factor is $\alpha(v)$ (Xing *et al* 1999) and the prescribed dose is $D^{\text{prescribed}}(v)$. We define a dose delivery function $D(\cdot)$, such that under plan β and delivery error ε the cumulative dose after fraction i is $\psi_i^{\text{cumdose}}(v) = \psi_{i-1}^{\text{cumdose}}(v) + D(v; \beta, \varepsilon, \psi_{i-1}^{\text{geometry}})$ at voxel v . Without delivery errors (i.e., without any setup errors or deformations of organs), $\varepsilon = 0$. Let B be the set of deliverable (feasible) plans. Our ultimate goal is then to find the feasible plan $(\beta_1, \dots, \beta_N)$ optimizing

$$\begin{aligned} & \min_{\beta_1 \in B, \dots, \beta_N \in B} \sum_v \alpha(v) (D^{\text{prescribed}}(v) - \psi_N^{\text{cumdose}}(v))^2 \\ & = \min_{\beta_1 \in B, \dots, \beta_N \in B} \sum_v \alpha(v) \left(D^{\text{prescribed}}(v) - \sum_{i=1}^N D(v, \beta_i, \varepsilon_i, \psi_{i-1}^{\text{geometry}}) \right)^2. \end{aligned} \quad (1)$$

The objective function is the weighted quadratic deviation of the cumulative delivered dose from the prescribed dose. There is no way to find in advance the optimal solution to this problem because the actual delivery error in fraction i , ε_i , is unknown when we must decide on the plan β_i . In general, ART dose optimization is a degenerate problem as there exist numerous ways to take the delivery history and new patient geometry into consideration. Finding the optimal adaptation strategy is an important part of ART treatment planning. While the simplest idea is to optimize

$$\min_{\beta_1 \in B, \dots, \beta_N \in B} \sum_v \alpha(v) \left(D^{\text{prescribed}}(v) - \sum_{i=1}^N D(v, \beta_i, 0, \psi_0^{\text{geometry}}) \right)^2,$$

we suggest in the following sections some more sophisticated approximations to problem (1). These approaches provide valuable ART alternatives and, as demonstrated later, they often lead to improved accumulated dose distributions. For comparison purpose, the conventional planning strategy with population-based margin is also formalized and described.

2.3. Baseline algorithm 1: planning with population-based margins

A common approach adds large margins to structures to compensate for delivery errors and uses the same plan every fraction, $\beta_1 = \dots = \beta_N$ (i.e., a static plan). These margins lead to modified dose prescriptions, $\tilde{D}^{\text{prescribed}}(\cdot)$, and importance factors, $\tilde{\alpha}(\cdot)$. This yields the optimization problem:

$$\beta_1 = \dots = \beta_N = \arg \min_{\beta \in B} \sum_v \tilde{\alpha}(v) \left(\frac{\tilde{D}^{\text{prescribed}}(v)}{N} - D(v, \beta, 0, \psi_0^{\text{geometry}}) \right)^2.$$

While this approximation is simple, it exposes much healthy tissue to high dose.

2.4. Baseline algorithm 2: Adapting to Changing Geometry

Assume that a CBCT image is taken prior to every treatment fraction. Based on the current anatomy information, a new plan for the current fraction is formed. This means that we consider $\psi_{i-1}^{\text{geometry}}$ when determining β_i . In particular we choose for fraction i the feasible plan that minimizes the weighted quadratic deviation from the initial prescribed dose: for all fractions i ,

$$\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{D^{\text{prescribed}}(v)}{N} - D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2.$$

In other words, the algorithm tries to deliver the same daily prescribed dose, $D^{\text{prescribed}}/N$, in each fraction.

2.5. Adapting to Geometry and Delivered Dose

We can do better than in section 2.4, by taking into account not only the up-to-date CBCT imaging information $\psi_{i-1}^{\text{geometry}}$, but also other factors such as the cumulative delivered dose $\psi_{i-1}^{\text{cumdose}}$. Here we present two algorithms that take into account both $\psi_{i-1}^{\text{geometry}}$ and $\psi_{i-1}^{\text{cumdose}}$ when choosing a feasible plan for fraction i which minimizes the weighted quadratic deviation from the adaptive dose goal for this fraction. In general this adaptive dose goal may depend on the dose delivery history, $\psi_{i-1}^{\text{cumdose}}$, and predictions of future geometry. We assume that there is a mechanism that can calculate $\psi_{i-1}^{\text{cumdose}}$. The two algorithms we investigate differ in how they choose this adaptive dose goal. The first compensates every fraction for delivery errors in previous fractions, and the second incorporates predictions of the future patient geometry. (For example, the tumor shrinkage in response to the radiation therapy can be modeled by a tissue-feature based deformable registration (Chao *et al* 2007).)

2.5.1. Immediately Correcting Algorithm (ICA). This algorithm, figure 2, takes into account the dose delivery history as well as the anatomy model derived from daily CBCT images. No prediction is attempted about the patient model in the subsequent fractions. This algorithm adjusts the originally prescribed dose to completely compensate voxels which were overdosed (or underdosed) in previous fractions by decreasing (or increasing) the dose goal at those voxels. Specifically, it adds in fraction i the accumulated error from the previous fractions,




Figure 2. Block diagram for ICA.

$\frac{i-1}{N} D^{prescribed} - \psi_{i-1}^{cumdose}$, to the original daily prescription dose, $D^{prescribed}/N$, resulting in an adaptive dose goal of

$$D_i^{goal} = \frac{D^{prescribed}}{N} + \left(\frac{i-1}{N} D^{prescribed} - \psi_{i-1}^{cumdose} \right) = \frac{i}{N} D^{prescribed} - \psi_{i-1}^{cumdose}.$$

In the control theory literature, this is considered a form of proportional control. Note that the adaptive dose goal can be accurately calculated as long as the previous delivered doses, $\psi_{i-1}^{cumdose}$, are known. While this is currently not the case, the information is becoming increasingly available. This algorithm chooses the feasible plan $\beta_i \in B$ for fraction i that minimizes the weighted quadratic deviation of the planned dose, $D(\cdot, \beta_i, 0, \psi_{i-1}^{geometry})$ from the adaptive dose goal, D_i^{goal} :

$$\begin{aligned} \beta_i &= \arg \min_{\beta \in B} \sum_v \alpha(v) (D_i^{goal}(v) - D(v, \beta, 0, \psi_{i-1}^{geometry}))^2 \\ &= \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{i}{N} D^{prescribed}(v) - \psi_{i-1}^{cumdose}(v) - D(v, \beta, 0, \psi_{i-1}^{geometry}) \right)^2. \end{aligned} \quad (2)$$

2.5.2. Prudent Correcting Algorithm (PCA). Consider a treatment course of N fractions. Let i be the current fraction index and d be the number of fractions for which we forecast the patient's anatomy and position (e.g., $d = 2$ means predicting 2 days ahead of the current fraction). Note that d may be a function of the current fraction i . For instance, taking $d(i) = N - i$ means that we predict the anatomy and organ locations for the rest of the treatment course. In reality, we can interrupt a pre-designed re-planning schedule at any fraction and re-plan the subsequent fractions using a rolling horizon of d fractions: every time we re-plan we find the optimal plan for fractions $i, \dots, i + d$ (the current fraction and the subsequent d fractions) and then deliver the first part of the plan (i.e., until we re-plan again). In the control theory literature this is called finite horizon control. We then update the system when the next set of information becomes available (possibly before every fraction). Optimizing only for the current fraction ignores our (albeit imperfect) knowledge of the future and therefore misses the opportunity to compensate the dose in future fractions. To determine the plan chosen for fraction i we optimize

$$\begin{aligned} \min_{\beta'_1 \in B, \dots, \beta'_{i+d} \in B} \sum_v \alpha(v) & \left(\frac{i+1}{N} D^{prescribed}(v) - \psi_{i-1}^{cumdose}(v) \right. \\ & \left. - \sum_{j=1}^{j=i+d} D(v, \beta'_j, 0, \psi_j^{predicted-geometry}) \right)^2 \end{aligned} \quad (3)$$

where $\psi_i^{\text{predicted-geometry}}, \dots, \psi_{i+d}^{\text{predicted-geometry}}$ are the predicted anatomy locations for fractions i through $i + d$. The optimizers of problem (3) are a sequence of plans, $\beta_i, \dots, \beta_{i+d}$, the first of which is the plan, β_i , that we will implement in fraction i . Clinically, this scheme is useful when dealing with situations of tumor shrinkage after a re-planning CT/CBCT is done (see section 4 for more details).

We consider a simple prediction model in which the anatomy remains unchanged for the next d fractions to illustrate the approach (in practice one may use more sophisticated prediction models). That is, we assume $\psi_i^{\text{predicted-geometry}} = \psi_{i-1}^{\text{geometry}}$ for fractions i through $i + d$. For this prediction model, it can be shown that (3) reduces to

$$\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{i+d}{N} D_i^{\text{prescribed}}(v) - \psi_{i-1}^{\text{cumdose}}(v) - (1+d)D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2.$$

We can also give this model the form of problem (2) by defining

$$D_i^{\text{goal}} = \frac{D^{\text{prescribed}}}{N} + \frac{1}{d+1} \left(\frac{i-d}{N} D^{\text{prescribed}} - \psi_{i-1}^{\text{cumdose}} \right). \quad (4)$$

If $d = 0$, then this algorithm coincides with the previous ICA algorithm. In equation (4), we see that this algorithm differs from the ICA algorithm because the correction to the accumulated error is divided among the subsequent $d + 1$ fractions of therapy to achieve better uniformity and robustness of therapy. Alternative schedules for compensating the accumulated error that accommodate specific clinical considerations should be easily implementable.

2.6. Evaluation and case study

We developed an in-house inverse planning platform to evaluate this closed-loop control framework and the novel *Adapting to Geometry and Delivered Dose* algorithms. This platform implements various RT planning strategies by optimizing fluence maps using the commercial nonlinear optimization code SNOPT (Gill *et al* 2005). The platform also evaluates the effectiveness of these plans by simulating the dose delivery function $D(\cdot)$.

Using this platform we compare the following algorithms: *Adapting to Changing Geometry* (our baseline), *Immediately Correcting Algorithm* (ICA), *Prudent Correcting Algorithm* (PCA) and *Perfect Foresight*. Results of the population-based margins planning algorithm are not shown since the selection of margin size varies with institution and, in general, it performs much worse than our baseline. *Perfect Foresight* is an algorithm that plans the entire course of treatment using accurate predictions of the future errors and motions during the course of treatment (see table 2). Note that in the *Perfect Foresight* algorithm, the entire treatment course is planned as a whole (i.e., the fractions and their corresponding anatomies are taken into account all in one optimization problem). In contrast to PCA with $d = N$, the *Perfect Foresight* algorithm ‘knows’ *exactly* what the anatomy will be in each fraction. When the prediction model is perfect, then PCA with $d = N$ is the same as the *Perfect Foresight*. It is worth stressing the purpose of the *Perfect Foresight* algorithm: it is not an algorithm in the usual sense (because we assume that it knows the future) but only a bound on normal algorithms. It is also worth noting that not even *Perfect Foresight* can achieve complete conformity. We include *Perfect Foresight* in our comparisons to have a more realistic standard against which to compare our algorithms. Further, as our ability of predicting the future becomes better, the performance of our algorithms (and most other reasonable algorithms) should approach that of *Perfect Foresight*. Thus, even though the performance of this algorithm may not actually be attainable, it sets a valuable theoretical upper bound on the performance of any ART planning algorithm. It is critically important for

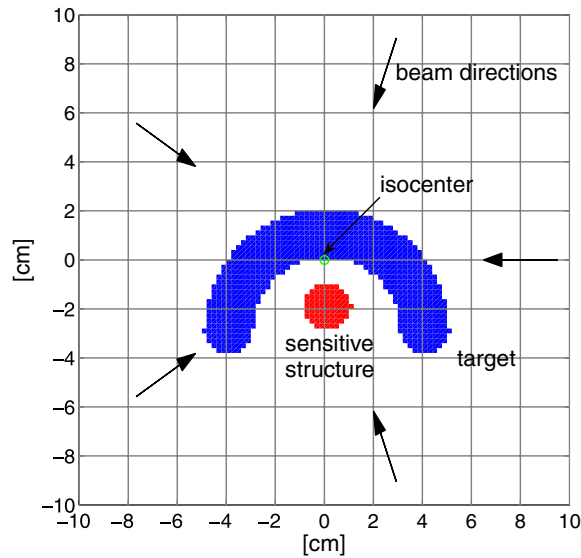


Figure 3. Phantom anatomy.

Table 2. Summary of various ART optimization schemes.

| Algorithm | Objective functions for different schemes of dose optimization |
|---|---|
| Perfect Foresight | $\min_{\beta_1, \dots, \beta_N \in B} \sum_v \alpha(v) \left(D^{\text{prescribed}}(v) - \sum_{i=1}^N D(v, \beta_i, \varepsilon_i, \psi_{i-1}^{\text{geometry}}) \right)^2$ <p>Minimize the difference between the prescription and the delivered dose accumulated over <i>all</i> treatment fractions with knowledge of the dose delivery errors ε_i.</p> |
| Baseline 1: population-based margins | $\beta_1 = \dots = \beta_N = \arg \min_{\beta \in B} \sum_v \tilde{\alpha}(v) \left(\frac{\bar{D}^{\text{prescribed}}(v)}{N} - D(v, \beta, 0, \psi_0^{\text{geometry}}) \right)^2$ <p>Add margins to the prescription and then minimize the difference between the prescription and the daily planned dose.</p> |
| Baseline 2: Adapting to Changing Geometry | $\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{D^{\text{prescribed}}(v)}{N} - D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2$ <p>Update the patient's geometric model every fraction using CBCT and plan a dose for that geometry that minimizes the difference to the daily prescription.</p> |
| Adapting to Geometry and Delivered Dose | |
| Immediately Correcting Algorithm (ICA) | $\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{i}{N} D^{\text{prescribed}}(v) - \psi_{i-1}^{\text{cumdose}}(v) - D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2$ <p>Update the patient's geometric model every fraction using CBCT and plan a dose for that geometry that minimizes the difference to the prescribed dose plus the accumulated error.</p> |
| Prudent Correcting Algorithm (PCA) | $\beta_i = \arg \min_{\beta \in B} \sum_v \alpha(v) \left(\frac{i+d}{N} D^{\text{prescribed}}(v) - \psi_{i-1}^{\text{cumdose}}(v) - (1+d)D(v, \beta, 0, \psi_{i-1}^{\text{geometry}}) \right)^2$ <p>Update the patient's geometric model every fraction using CBCT and using that geometry plan a dose for remaining $(d+1)$ fractions that minimizes the difference to the prescribed dose plus the accumulated error.</p> |

the algorithmic development of ART planning and provides an important measure helping the algorithm designer judge the potential for further improvement.

We compared the algorithms on a cubic phantom case and a prostate patient case. Phantom study is valuable for algorithm development, simply because the results are more intuitively

Table 3. Cumulative dose delivered to phantom for various algorithms relative to the baseline algorithm, *Adapting to Changing Geometry*.

| Algorithm | Tumor | | Sensitive structure | | Normal tissue | |
|-------------------|--------------|-------------|---------------------|-------------|---------------|-------------|
| | Avg dose (%) | Std dev (%) | Avg dose (%) | Std dev (%) | Avg dose (%) | Std dev (%) |
| ICA | +9 | +48 | -48 | -12 | +18 | +47 |
| PCA | +8 | -13 | -18 | -9 | +7 | +14 |
| Perfect foresight | +9 | -30 | -89 | -65 | -8 | +13 |

predictable/understandable and the ‘ground truth’ is readily attainable. A number of new tools/steps are involved in ART planning, such as deformable registration (Court and Dong 2003), CBCT-based dose calculation (Yang *et al* 2007) and dose optimization, each of them needs to be evaluated independently before they can be assembled together. Given the developmental status of these ‘enabling’ tools for ART, it seems logical to illustrate the methodology using some intuitive examples. The phantom, figure 3, is 20 cm in size and at its center has a C-shaped target enclosing a round sensitive structure. In the prostate case, the target volume was the prostate and the sensitive structures involved were rectum, femoral heads and bladder. In both cases, we asked the algorithms to generate an IMRT plan with 15 fractions and using five equally spaced beam directions. The tumor, sensitive structures and normal tissue were assigned importance factors of $\alpha = 10$, $\alpha = 3$ and $\alpha = 0.7$ respectively. The PCA parameter, d , was always set to the number of fractions left in the treatment $d = N - i$. We simulated setup errors by introducing, independently for every fraction, a random translation chosen uniformly at random from $[-1, 1]$ cm and a random rotation chosen uniformly at random from $[-2^\circ, 2^\circ]$. Other types of motion were not simulated in this study as a change in the motion model does not change the nature of the closed-loop control. Since the intentionally introduced motion has no systematic component, the prediction model for PCA assumes that the anatomy remains unchanged for the remainder of the treatment course (as described in section 2.5.2). In section 3, we compare the results visually in terms of the cumulative DVH and numerically for each organ in terms of the average dose and standard deviation of the dose.

3. Results

3.1. Phantom study

The results are shown graphically in figure 4 and statistically in table 3. Each row of table 3 compares an algorithm to our baseline, *Adapting to Changing Geometry*.

Figure 4(a) shows three benefits the *Adapting to Geometry and Delivered Dose* algorithms (and of course the *Perfect Foresight* algorithm) have over baseline 2, *Adapting to Changing Geometry*. For the first benefit, note how the 25% contour in our baseline is shifted slightly down from the sensitive structure. All the other algorithms do not show this shift. This shift is most likely due to some setup error (that is not corrected for in our baseline). The second benefit (also seen in the DVHs and table 3) is dose escalation to the tumor while keeping the sensitive structure close to 25% (unlike the baseline, these algorithms achieve tumor doses close to 100%). The third benefit is steeper gradients around the tumor (because dose correcting algorithms will not let errors accumulate).

Aside from these benefits, the dose distributions in figure 4(a) also show how much worse ICA is than PCA. As discussed above, ICA is able to escalate dose to the tumor while keeping the dose to the sensitive structure low, but along the way it overdoses normal tissue and decreases dose uniformity (see table 3). ICA may have fractions with higher doses than PCA

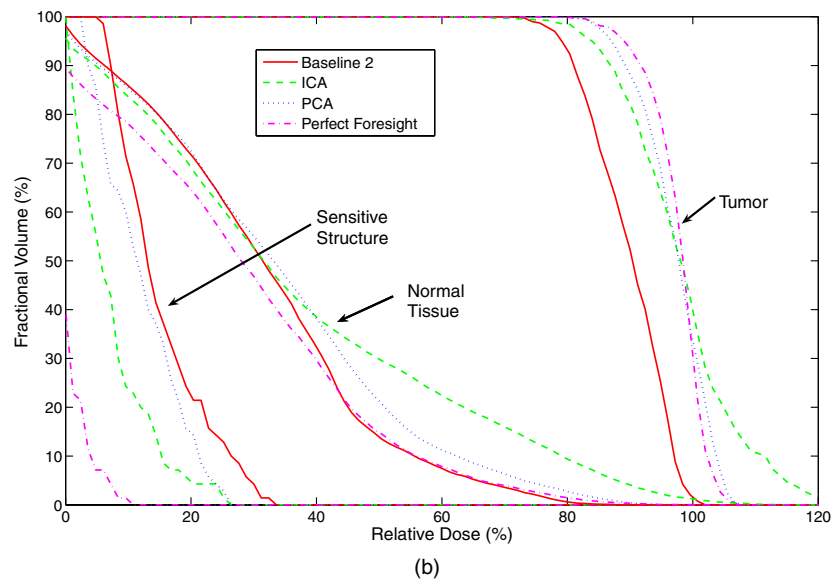
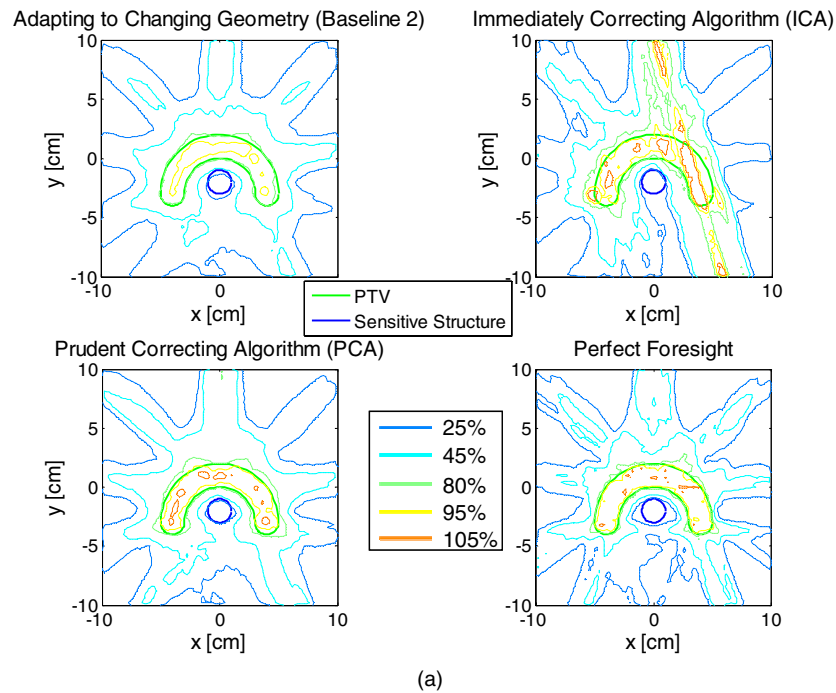


Figure 4. (a) Cumulative dose delivered to phantom as per cent of prescription. (b) DVHs of cumulative dose delivered to phantom as per cent of prescription.

since it tries to completely correct for dose delivery errors in the next fraction. We believe this makes ICA less robust than PCA because when combined with organ deformation or setup error these higher doses lead to bigger errors. The differing performance of the two *Adapting to Geometry and Delivered Dose* algorithms shows that the algorithmic details matter.

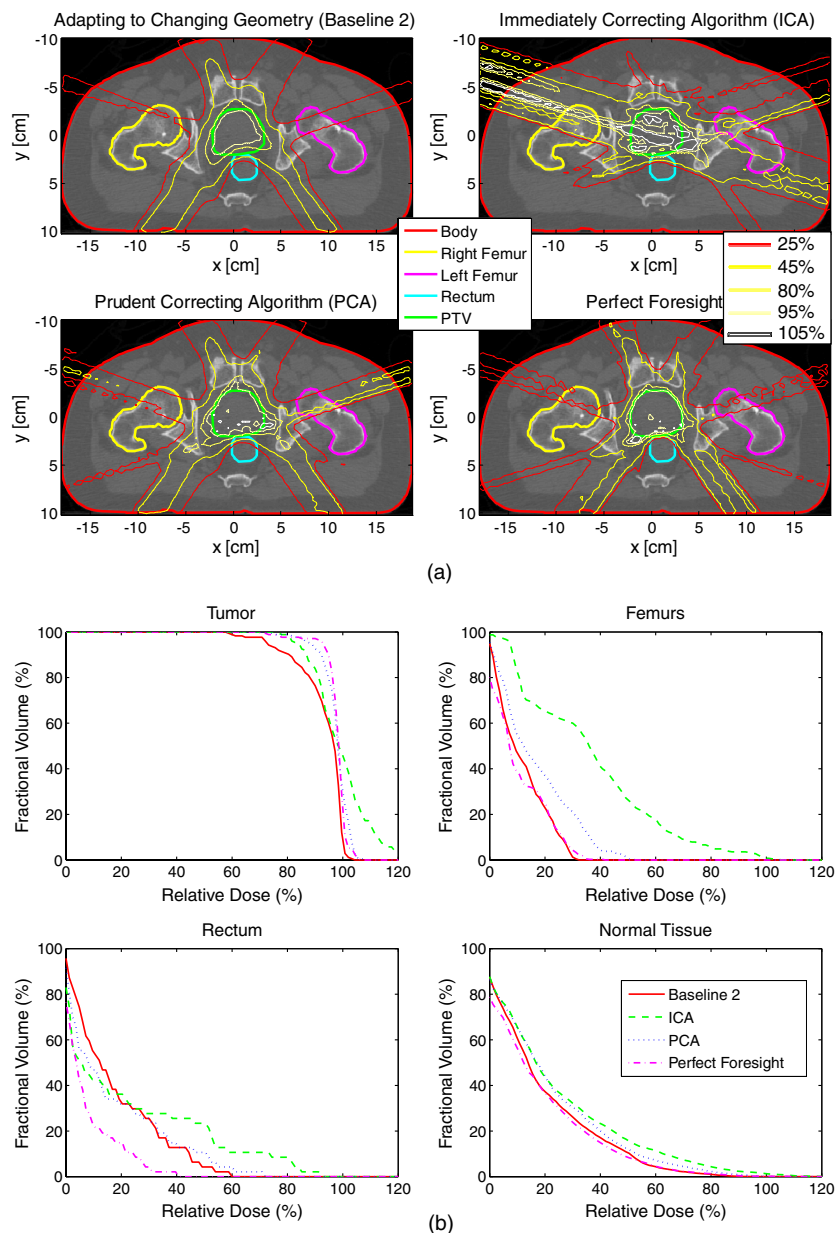


Figure 5. (a) Cumulative dose delivered in prostate study as per cent of prescription. (b) DVHs of cumulative dose delivered in prostate study as per cent of prescription.

Looking at the DVH in figure 4(b), we see that the *Perfect Foresight* algorithm achieves significantly better results than the adaptive algorithms (e.g., PCA). *Perfect Foresight* is superior not only in terms of tumor target coverage, but also in the sparing of sensitive structures. This is explained by the obvious fact that the *Perfect Foresight* algorithm has the fundamental advantage of accurately ‘knowing’ the future. The *Perfect Foresight* algorithm gives an upper limit on the potential of any closed-loop control algorithm. As treatment

Table 4. Cumulative dose delivered in a prostate study for various algorithms relative to the baseline algorithm, *Adapting to Changing Geometry*.

| Algorithm | Tumor | | Femurs | | Rectum | | Normal tissue | |
|-------------------|--------------|-------------|--------------|-------------|--------------|-------------|---------------|-------------|
| | Avg dose (%) | Std dev (%) | Avg dose (%) | Std dev (%) | Avg dose (%) | Std dev (%) | Avg dose (%) | Std dev (%) |
| ICA | +6 | +14 | +196 | +165 | +21 | +73 | +24 | +28 |
| PCA | +5 | -44 | +40 | +45 | -5 | +18 | +13 | +11 |
| Perfect Foresight | +5 | -49 | -8 | +12 | -54 | -40 | -7 | +2 |

imaging, planning and delivery techniques become more sophisticated, the performance of treatment will become closer and closer to the *Perfect Foresight* algorithm.

3.2. Prostate study

The resulting dose distributions and their DVHs are in figures 5(a) and (b). Each row of table 4 compares an algorithm to our baseline, *Adapting to Changing Geometry*.

As in the phantom case, the dose distributions (80% and 95% contour lines of figure 5(a)) show that PCA improves upon the baseline by being able to achieve dose escalation and increased uniformity to the tumor while keeping the dose to the rectum (the sensitive structure) low. Interestingly, ICA performs even worse than in the phantom case (it delivers more dose to the femurs and normal tissue). At the same time, even more than in the phantom case, the *Perfect Foresight* DVH shows how much the *Adapting to Geometry and Delivered Dose* algorithms can still improve. In particular, the dose could be lower to the femurs and to a lesser extent to normal tissue and the rectum. We believe that the difference between ICA and PCA arises from the fact that ICA achieves the adaptive goal by sequentially optimizing each fraction, whereas the PCA collectively optimizes a number of fractions with the goal of producing the best possible cumulative dose. ICA may not be optimal in producing the best possible cumulative dose because, when optimizing the treatment of a fraction, no knowledge about the subsequent fractions is taken into account. Thus it is not surprising that the performance of the PCA generally surpasses that of the ICA. Effectively, PCA benefits from the partial cancellation of random delivery errors caused by fractionation (Bortfeld *et al* 2002, Chui *et al* 2003).

4. Discussion

In current clinical practice, patient setup relies primarily on information from simulation and treatment planning (Balter 2003). During the whole course of treatment, usually the same treatment plan and setup DRRs (digitally reconstructed radiographs) are employed. Clinically, effort is focused on reproducing, with the aid of orthogonal planar images, the patient's geometry at the simulation stage using translations (and occasionally rotations). While this approach is justifiable for treatment of certain types of diseases such as brain tumors, it generally compromises the treatment because inter-fraction variations in volumes and shapes of the target and sensitive structures are not taken into account. Generally, these inter-fraction changes are multi-dimensional because organs can move relative to each other, and in an extreme situation, each voxel within a soft organ can move relative to other voxels in a complicated manner. Deformable image registration (Court and Dong 2003, Court *et al* 2005) helps by making these inter-fraction changes visible. But even then, compromised treatment is inevitable because the few degrees of freedom in patient setup (translation and

rotation) cannot completely correct for the multi-dimensional changes in the patient geometry: patient setup cannot simultaneously align all the involved structures.

ART solves the problems described in the previous paragraph by adjusting every fraction, not only the setup but also the treatment plan. We treat ART as a closed-loop control system where every fraction we re-optimize the radiation beams based on the latest information (e.g., coming from CBCT). The hope is that the many degrees of freedom available in selecting a plan allow us to compensate for the multi-dimensional changes in the patient geometry. In ART the plan is updated routinely, and the fraction-to-fraction variations of anatomy/physiology and dose delivery lead to modifications of the voxel-specific dose prescription.

ART can be implemented at different levels where the beams can be made to accommodate (i) the new patient setup and the deformed target shape; (ii) positional/anatomic/physiological changes of all involved organs; or (iii) deformable changes of organs and accumulated dose delivery errors. In conventional 3D conformal radiation therapy, it is not uncommon for a physician to modify a beam portal under the guidance of portal films while the patient is on the treatment couch. In a sense, this is an example of the first kind of ART listed above. However to adapt well to multi-dimensional organ deformations, the number of variables in the beam should be large. Thus the modality of choice for the treatments of type (ii) or (iii) is IMRT or alike. The information available in a particular treatment setup (e.g., whether information is captured about the delivered dose) restricts the types of ART possible in that setup. Our results show that correcting for dose delivery errors (i.e., type (iii) ART) incurred in previous treatment fractions is important. In conventional radiation therapy (and types (i) and (ii) ART) dose errors accumulate because they lack such a compensating mechanism. As technology improves, more sophisticated versions of ART with increased feedback can be implemented and our closed-loop control algorithms gain in relevance.

ART is a new strategy and its implementation entails the development of a number of 'enabling' tools, such as deformable image registration, automated contour mapping from planning CT to the CBCT just before treatment, and effective dose optimization algorithms. This paper tackles the dose optimization problem in ART by developing suitable dynamic control algorithms. These algorithms try to rapidly optimize the treatment plan each time a new set of input data is available. We propose two types of closed-loop control algorithms (*Adapting to Changing Geometry* and *Adapting to Geometry and Delivered Dose*, see table 2 for details) for different clinical applications. *Adapting to Changing Geometry* is useful when the accumulated dose is not known accurately. At this point, the deformable registration, which is essential for the calculation of accumulative dose, is still not robust enough for clinical use. *Adapting to Changing Geometry* is particularly helpful in this situation as it affords a currently implementable technique to cope with the nuisance caused by organ deformation. The 'manual' approach described by Mohan *et al* (2005) represents a special example of geometric adaptation. The algorithms *Adapting to Geometry and Delivered Dose* are designed to optimize the radiation treatment when both geometric and dosimetric updates are available from time to time. The two algorithms *Adapting to Geometry and Delivered Dose* (ICA and PCA) differ in how they use the update data to direct the treatment. First, ICA is proposed for the situation where the dose compensation needs to be performed right after each update of the system parameters. Clinically, ICA may be useful in dealing with unpredictable daily changes resulting from random and fraction-specific organ deformation such as rectum or bladder filling. Disease sites that are likely to benefit from the treatment include, but are not limited to, prostate, rectum and cervical cancers. In PCA, the task of dose compensation is accomplished by spreading the previous dosimetric errors over a number of subsequent fractions. For practical or clinical reason, the volumetric images just before treatment may not be available on a daily basis. In this case, PCA becomes a viable option for certain types of

diseases since it distributes the task of compensating a dosimetric error among a number of fractions instead of all-in-one. Indeed, there are clinical situations where the variation of the anatomy may not be notable from day to day but over a larger time span. Tumor shrinkage and weight loss in some head and neck cases represents a model example, where the CBCT and re-planning do not need to be performed on daily basis. Re-planning using PCA would allow us in this case to maximally benefit from state-of-the-art imaging information. In a way, PCA is similar to the off-line ‘dumped’ correction scheme proposed by Mackie *et al* (2003), Keller *et al* (2003), in which the correction is done at a certain point(s) of the treatment course and then applied to several subsequent fractions.

There is a point worth emphasizing here about the PCA algorithm. Even though the re-evaluation/re-planning are performed for a series of fractions, the specific plan used in PCA may vary from fraction to fraction as a result of PCA optimization. This may happen even if the anatomy/geometry does not change throughout the d fractions or the whole course of the treatment. In general, the deviation from the prescription consists of two factors: (i) the inherent difference between the prescription and the physically achievable dose distribution and (ii) the dose delivery error due to changes in the patient’s anatomy and/or actual dose delivery. Indeed, voxels within a target or a sensitive structure volume are generally not equivalent in achieving their dosimetric goals in IMRT planning (unless non-attenuating beams are used). Depending on the patient’s geometry, beam modality and field configuration, some regions may have a better chance to meet the prescription than others, and vice versa. Thus even when there is no geometric change at each fraction and the delivery can faithfully reproduce what is planned without any error, the dose goal will likely need to be modified to compensate for the inherent difference between the planned and prescribed doses in previous fraction(s). Dosimetrically, it is advantageous to optimally vary the plan for each fraction to accomplish the best possible cumulative dose distribution. A virtue of the PCA algorithm is that it is capable of incorporating our prediction about the patient’s anatomy. A simple prediction would be that the anatomy does not change in the following d fractions. A more sophisticated prediction would be, for example, a model predicting the tumor shrinkage for each of the following fractions (Chao *et al* 2007). In the special case where there is no geometric change in the whole course and no delivery error, the difference between ICA and PCA becomes how to better utilize the following fractional deliveries to make up the inherent dosimetric difference of previous fractions.

Although showing significant improvement, we note that the algorithms described in this paper are still far from the maximum theoretical performance limit. We encourage the development of more sophisticated closed-loop control algorithms that come closer to this limit. Better predictive schemes for future patient geometry should improve the PCA algorithm greatly. Our use of a weighted quadratic objective function for plan optimization (specifically the selection of beamlet weights) is probably suboptimal (Yang and Xing 2004a, 2004b, Wu *et al* 2002, Thieke *et al* 2003, Bedford and Webb 2006, Popple *et al* 2005, Oelfke and Bortfeld 1999, Choi and Deasy 2002). However, all the algorithms discussed in this paper would work with other objective functions. One could probably construct an optimal algorithm using the theory of optimal stochastic control. One hurdle to applying this theory is the fact that after delivering the treatment plan for some fraction, we must specify the probability of every possible system state. We are researching these ideas.

Finally, we emphasize that the proposed ART is purely dose based and does not consider any radiobiological effects. In principle treatment plan optimization should be based on biological models as they are clinically the most relevant. In practice, however, there is much controversy to this approach: the dose–response function linking the biological effect to the radiation dose is not sufficiently understood for various structures. Yang and Xing (2005)

have recently proposed a general time–dose–fractionation optimization strategy. This idea combined with our ART framework could potentially handle biologically adaptive radiation therapy. This is clearly a subject for future research.

5. Conclusion

The level of sophistication of adaptive therapy depends on the technology used to provide feedback data. In the last decade, the electronic portal imaging device (EPID) and the standalone CT scanner were the only technologies available for acquiring feedback information and the focus was on adaptively defining the margin for delineating the planning tumor target (PTV) (Yan *et al* 1997). In the era of onboard volumetric imaging, the meaning of ART is obviously different from previous studies. The feedback from onboard CBCT makes it possible for the first time to adaptively modify not only the margin but also the spatial dose distribution to best accommodate the geometric changes of the patient as well as the dosimetric deviations from the ideal prescription. The technical tools for implementing this new type of geometry- and dosimetry-based ART are still not available yet, and an effective treatment planning strategy capable of taking into account the dose delivery history and the patient's on-treatment geometric model must be in place in order for this new scheme of ART to reach its potential. We have formulated ART treatment planning into the framework of closed-loop control and developed three closed-loop ART algorithms: one *Adapting to Changing Geometry* and two *Adapting to Geometry and Delivered Dose* (ICA and PCA). The new formalism is useful in answering questions such as 'how to effectively utilize information about the current state and past delivery errors?', 'can we benefit from including the patient's future treatment(s) into today's adaptation strategy?', and 'what is the best achievable scenario of ART in an ideal situation?' All three algorithms are capable of incorporating volumetric imaging information acquired just before treatment and the latter two can also incorporate information on the accumulated dose. Application of the proposed algorithms to the phantom and clinical cases indicates that the algorithms utilizing the accumulated dose data (ICA and PCA) allow an escalation of the dose to the tumor target. For a patient positioning procedure with random setup errors, we found that the PCA performed best and significantly better than ICA even though the implemented algorithms differ only slightly. The over-correcting tendency of the ICA is not a desirable feature and should be avoided in the design of a practical ART planning algorithm. Generally, the performance of a closed-loop algorithm may depend on the motion model and other case-related issues. An important point here is that the details of closed-loop adaptation play a critical role in the success of ART. Clearly, much research remains to be done in developing more sophisticated closed-loop algorithms and searching for the best possible ART strategy to maximally utilize the geometric and dosimetric information that become or will soon become available. Given the complex interaction among errors, motion and the doses delivered to the patient, a sophisticated disease-site specific adaptation scheme may be useful.

Acknowledgments

We wish to thank Sachin Adlakha, Joseph Deasy, Yong Yang and Bernard Widrow for useful discussions. This work was supported in part by grants from the Department of Defense (PC040282), the National Cancer Institute (1R01 CA104205 and 5R01 CA98523), the Komen Breast Cancer Foundation (BCTR0504071) and a Graduate Research Fellowship from the

National Science Foundation. We also thank two anonymous referees for many helpful comments.

References

- AAPM IMRT Sub-committee 2003 Guidance document on delivery, treatment planning, and clinical implementation of IMRT: Report of the IMRT Sub-committee of the AAPM Radiation Therapy Committee *Med. Phys.* **30** 2089–115
- Balter J 2003 Target and critical structure definitions, dose prescription and reporting for IMRT *Intensity-Modulated Radiation Therapy: The State of the Art* ed J Palta and T Mackie (Colorado Springs, CO: Medical Physics Publishing) pp 183–98
- Balter J M, Sandler H M, Lam K, Bree R L, Lichter A S and ten Haken R K 1995 Measurement of prostate movement over the course of routine radiotherapy using implanted markers *Int. J. Radiat. Oncol. Biol. Phys.* **31** 113–8
- Bedford J L and Webb S 2006 Constrained segment shapes in direct-aperture optimization for step-and-shoot IMRT *Med. Phys.* **33** 944–58
- Bortfeld T 1999 Optimized planning using physical objectives and constraints *Semin. Radiat. Oncol.* **9** 20–34
- Bortfeld T, Jokivarsi K, Goitein M, Kung J and Jiang S 2002 Effects of intra-fraction motion on IMRT dose delivery: statistical analysis and simulation *Phys. Med. Biol.* **47** 2203–20
- Cao D, Earl M A, Luan S and Shepard D M 2006 Continuous intensity map optimization (CIMO): a novel approach to leaf sequencing in step and shoot IMRT *Med. Phys.* **33** 859–67
- Censor Y 2003 Mathematical optimization for the inverse problem of intensity-modulated radiation therapy *Intensity-Modulated Radiation Therapy: The State of the Art* ed J Palta and T Mackie (Colorado Springs, CO: Medical Physics Publishing) pp 25–50
- Chao M, Xie Y, Le Q and Xing L 2007 Modeling the volumetric change of head-and-neck tumor in response to radiation therapy 2007 *Annual Meeting of ASTRO (Los Angeles, CA)* (Abstract)
- Choi B and Deasy J O 2002 The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning *Phys. Med. Biol.* **47** 3579–89
- Court L E and Dong L 2003 Automatic registration of the prostate for computed-tomography-guided radiotherapy *Med. Phys.* **30** 2750–7
- Court L E, Dong L, Lee A K, Cheung R, Bonnen M D, O’Daniel J, Wang H, Mohan R and Kuban D 2005 An automatic CT-guided adaptive radiation therapy technique by online modification of multileaf collimator leaf positions for prostate cancer *Int. J. Radiat. Oncol. Biol. Phys.* **62** 154–63
- Chui C, Yorke E and Hong L 2003 The effects of intra-fraction organ motion on the delivery of intensity modulated field with a MLC *Med. Phys.* **30** 1736–46
- de la Zerda A, Armbruster B and Xing L 2006 Inverse planning for adaptive radiation therapy using dynamic programming *Annual Meeting of ASTRO (Philadelphia, PA)*
- Gill P E, Murray W and Saunders M A 2005 SNOPT: an SQP algorithm for large-scale constrained optimization *SIAM Rev.* **47** 99–131
- Keller H, Ritter M A and Mackie T R 2003 Optimal stochastic correction strategies for rigid-body target motion *Int. J. Radiat. Oncol. Biol. Phys.* **55** 261–70
- Langen K M, Meeks S L, Poole D O, Wagner T H, Willoughby T R, Kupelian P A, Ruchala K J, Haimeri J and Olivera G H 2005 The use of megavoltage CT (MVCT) images for dose recomputations *Phys. Med. Biol.* **50** 4259–76
- Mackie T R *et al* 2003 Image guidance for precise conformal radiotherapy *Int. J. Radiat. Oncol. Biol. Phys.* **56** 89–105
- Mohan R, Zhang X, Wang H, Kang Y, Wang X, Liu H, Ang K K, Kuban D and Dong L 2005 Use of deformed intensity distributions for on-line modification of image-guided IMRT to account for interfractional anatomic changes *Int. J. Radiat. Oncol. Biol. Phys.* **61** 1258–66
- Oelfke U and Bortfeld T 1999 Inverse planning for x-ray rotation therapy: a general solution of the inverse problem *Phys. Med. Biol.* **44** 1089–104
- Oldham M, Letourneau D, Watt L, Hugo G, Yan D, Lockman D, Kim L H, Chen P Y, Martinez A and Wong J W 2005 Cone-beam-CT guided radiation therapy: a model for on-line application *Radiother. Oncol.* **75** 271E1–8
- Olivera G H, Mackie T R, Ruchala K, Lu W and Kapatoes J 2006 Adaptive radiation therapy (art) strategies using helical tomotherapy *Image-Guided IMRT* ed T Bortfeld, R Schmidt-Ullrich, W De Deve and D E Wazer (Berlin: Springer) pp 235–246
- Popple R A, Prellow P B, Spencer S A, De Los Santos J F, Duan J, Fiveash J B and Brezovich I A 2005 Simultaneous optimization of sequential IMRT plans *Med. Phys.* **32** 3257–66
- Pouliot J *et al* 2005 Low-dose megavoltage cone-beam CT for radiation therapy *Int. J. Radiat. Oncol. Biol. Phys.* **61** 552–60

- Thieke C, Bortfeld T, Niemierko A and Nill S 2003 From physical dose constraints to equivalent uniform dose constraints in inverse radiotherapy planning *Med. Phys.* **30** 2332–9
- Trofimov A, Reitzel E, Lu H M, Martin B, Jiang S, Chen G and Bortfeld T 2005 Temporo-spatial IMRT optimization: concepts, implementation and initial results *Phys. Med. Biol.* **50** 2779–98
- van Herk M 2006 Errors and margins in radiotherapy *Semin. Radiat. Oncol.* **14** 52–64
- Webb S 2001 *Intensity-Modulated Radiation Therapy* (Bristol: Institute of Physics Publishing)
- Widrow B and Stearns S 1985 *Adaptive Signal Processing* (Upper Saddle River, NJ: Prentice-Hall)
- Widrow B and Walach E 1995 *Adaptive Inverse Control* (Upper Saddle River, NJ: Prentice-Hall)
- Wu Q, Liang J and Yan D 2006 Application of dose compensation in image-guided radiotherapy of prostate cancer *Phys. Med. Biol.* **51** 1405–19
- Wu Q, Mohan R, Niemierko A and Schmidt-Ullrich R 2002 Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose *Int. J. Radiat. Oncol. Biol. Phys.* **52** 224–35
- Xiao Y, Galvin J, Hossain M and Valicenti R 2000 An optimized forward-planning technique for intensity modulated radiation therapy *Med. Phys.* **27** 2093–9
- Xing L, Li J G, Donaldson S, Le Q T and Boyer A L 1999 Optimization of importance factors in inverse planning *Phys. Med. Biol.* **44** 2525–36
- Xing L, Thorndyke B, Schreiber E, Yang Y, Li T F, Kim G Y, Luxton G and Koong A 2006 Overview of image-guided radiation therapy *Med. Dosim.* **31** 91–112
- Xing L, Wu Q, Yang Y and Boyer A 2005 Physics of IMRT *Intensity-Modulated Radiation Therapy: A Clinical Perspective* ed A Mundt and J Roeske (Hamilton: BC Decker) pp 20–52
- Yan D, Vicini F, Wong J and Martinez A 1997 Adaptive radiation therapy *Phys. Med. Biol.* **42** 123–132
- Yang Y, Schreiber E, Li T and Xing L 2007 Dosimetric evaluation of kV cone-beam CT (CBCT) based dose calculation *Phys. Med. Biol.* **52** 685–705
- Yang Y and Xing L 2004a Clinical knowledge-based inverse treatment planning *Phys. Med. Biol.* **49** 5101–17
- Yang Y and Xing L 2004b Inverse treatment planning with adaptively evolving voxel-dependent penalty scheme *Med. Phys.* **31** 2839–44
- Yang Y and Xing L 2005 Optimization of radiation dose–time–fractionation scheme with consideration of tumor specific biology *Med. Phys.* **32** 3666–77