Proceedings of the 2016 Winter Simulation Conference
T. M. K. Roeder, P. I. Frazier, R. Szechtman, E. Zhou, T. Huschka, and S. E. Chick, eds.

A MODEL PREDICTIVE CONTROL APPROACH FOR DISCOVERING NONSTATIONARY FLUENCE-MAPS IN CANCER RADIOTHERAPY FRACTIONATION

Ali Ajdari
Archis Ghate

Department of Industrial and Systems Engineering
University of Washington, Box 352650
3900 Northeast Stevens Way
Seattle, WA 98195, USA

ABSTRACT

We consider an optimization problem in radiotherapy, where the goal is to maximize the biological effect on the tumor of radiation intensity profiles across multiple treatment sessions, while limiting their toxic effects on nearby healthy tissues. We utilize the standard linear-quadratic dose-response model, which yields a nonconvex quadratically constrained quadratic programming (QCQP) formulation. Since nonconvex QCQPs are in general computationally difficult, recent work on this problem has only considered *stationary* solutions. This restriction allows a convex reformulation, enabling efficient solution. All other generic convexification methods for nonconvex QCQPs also yield a stationary solution in our case. While stationary solutions could be sub-optimal, currently there is no efficient method for finding nonstationary solutions. We propose a model predictive control approach that can, in principle, efficiently discover nonstationary solutions. We demonstrate via numerical experiments on head-and-neck cancer that these nonstationary solutions could produce a larger biological effect on the tumor than stationary.

1 INTRODUCTION

In external beam photon radiotherapy for cancer, the goal is to maximize the so-called therapeutic ratio, that is, the differential between tumor-damage and toxic effects of radiation on nearby organs-at-risk (OAR). This is achieved via a two-pronged approach: spatial localization and temporal dispersion (also called fractionation) of radiation dose.

In the traditional view of spatial localization, a high radiation dose is prescribed to the tumor and upper bounds are forced on OAR-dose. An optimization problem is then solved to find a radiation intensity profile (also called a fluence-map) that delivers dose adhering to this protocol as closely as possible. The now prevalent Intensity Modulated Radiation Therapy (IMRT) technology and several well-established optimization models and algorithms are capable of finding highly conformal fluence-maps (Ehrgott et al. 2008; Romeijn et al. 2006; Romeijn and Dempsey 2008; Shepard et al. 1999; Webb 2010). In practice, the desired radiation dose (as found by spatial localization) is administered in a pre-determined number of multiple *equal-dosage* sessions. For example, for head-and-neck cancer, a dose of 70 Gy is administered in 35 sessions of 2 Gy each (Marks et al. 2010). Healthy cells often have better damage-repair capabilities than tumor cells. Temporal dispersion of dose therefore gives healthy cells some time to recover from radiation-damage between sessions, thus increasing the therapeutic ratio.

There has been a recent surge of interest in a *biological* view of spatial localization. Instead of finding a fluence-map that delivers a tumor-dose close to the prescription in a pre-determined number of equal-dosage sessions, the goal in the biological approach often is to find a fluence-map that maximizes the so-called biological effect (BE) on tumor over a pre-determined number of equal-dosage sessions. Similarly, instead of putting upper bounds on the total dose to OAR, upper bounds are forced on the so-called biologically

effective dose (BED) to OAR. The hope in this line of research is that such biological objectives and constraints might better-capture tumor- and OAR-response to radiation and hence increase survival rates.

The most well-known radiobiological model of dose-response is the linear-quadratic (LQ) model (Hall and Giaccia 2005). The tumor-BE and OAR-BED in biological optimization are commonly calculated using this LQ model. The reader is referred to Armpilia et al. (2004), Bertuzzi et al. (2013), Fowler (1990, 2001, 2007, 2008), Fowler and Ritter (1995), Jones et al. (1995), Keller et al. (2012), Saberian, Ghate, and Kim (2016, 2015a), Unkelbach et al. (2013, 2013), and references therein, for examples of such biological formulations, their analyses, and solution methods. In this LQ dose-response framework, the tumor-BE and the OAR-BED are both quadratic functions of dose. The aforementioned biological formulations thus call for maximizing a convex quadratic function subject to convex quadratic constraints. As such, they belong to the class of nonconvex quadratically constrained quadratic programs (QCQPs). Nonconvex QCQPs are in general NP-hard (Luo et al. 2010). Research in this area therefore has mainly progressed under two restrictive assumptions.

A large group of papers assumes that a fluence-map is available *a priori* via a traditional IMRT spatial treatment planning system that does not use biological objectives or constraints. According to the standard linear dose-deposition model, radiation dose is a linear function of fluence-map (Jeraj and Keall 1999; Siebers et al. 2001; Spirou and Chui 1998; Tian et al. 2013; Webb and Oldham 1996). Thus, a desired dose in one session can be administered by scaling the pre-determined fluence-map. Under the additional assumption of equal-dosage, i.e., *stationary* fractionation, this approach reduces the biological optimization formulation to a problem where the decision variable equals the identical dose to be administered in every session. This single-variable problem can be solved in closed-form. This class of stylized formulations is often called *separated* models, because they separate the spatial and the biological components of the problem. Three recent papers in this area have shown via a toy counterexample and/or supporting rigorous analyses that equal-dosage fractionation may not be optimal even when dose-response parameters do not change over time (Mizuta et al. 2012; Saberian, Ghate, and Kim 2015c, 2016). Another paper proved that unequal-dosage fractionation may be optimal when dose-response parameters do change over time (Unkelbach et al. 2013).

Two papers took an alternative *integrated* view whereby they attempted to directly optimize fluence-maps using biological formulations (Saberian, Ghate, and Kim 2015a; Unkelbach et al. 2013). The resulting formulations were computationally difficult. One of the papers (Saberian, Ghate, and Kim 2015a) therefore simplified the formulation by assuming that fluence-maps do not vary across sessions, i.e., a *stationary* solution. This significantly reduced the dimension of the problem, and in fact, allowed them to reformulate it as a convex program that enabled efficient solution. That paper numerically demonstrated that this integrated approach with stationary solutions can lead to higher tumor-BE compared to earlier stylized separated models. The other paper (Unkelbach et al. 2013) constructed a small toy example to show that a stationary solution need not be optimal. However, they did not provide a detailed algorithm, a rigorous analysis of its behavior, or computational results for finding nonstationary fluence-maps for nonconvex QCQPs. Two other papers numerically demonstrated that unequal-dosage fractionation could be optimal when dose-response parameters changed stochastically over time (Kim, Ghate, and Phillips 2012; Saberian, Ghate, and Kim 2015b).

This literature survey exposes the following two research questions: (1) could nonstationary fluence-maps be superior to stationary ones in integrated biological formulations of clinically realistic test-cases even when dose-response parameters do not vary over time? and (2) can we devise an efficient algorithm to discover such nonstationary fluence-maps? In this paper, we answer these questions in the affirmative.

The paper is organized as follows: a precise problem description and its mathematical formulation is presented in the next section; Section 3 describes an approximate solution algorithm for this formulation; numerical experiments in Section 4 demonstrate that it does indeed discover nonstationary solutions in some test-cases; we end with a few concluding remarks.

2 PROBLEM DESCRIPTION

The notation below is standard in the literature and is borrowed, for instance, from Saberian, Ghate, and Kim (2015a). Let n denote the number of tumor voxels. The radiation field is discretized into small segments called beamlets. Let k denote the number of beamlets. Let $u^t \in \Re^k_+$ denote the k-dimensional beamlet intensity vector (fluence-map) employed in treatment session t. The fixed number of treatment sessions, administered one-a-day, is denoted by N. Let A denote the $n \times k$ nonnegative tumor dose-deposition matrix; A_i denotes its ith row, which corresponds to the ith tumor voxel. According to the linear dose-deposition model, $A_i u^t$ is the dose delivered to the ith tumor voxel in session t. Let $\bar{A} = \sum_{i=1}^{n} A_i/n$; then, $\bar{A}u^t$ is the average dose delivered to the tumor in session t. Let S denote the matrix employed in writing smoothness constraints on u^t. Let α_0 and β_0 denote the dose-response parameters for the tumor's LQ model. A tumor proliferation term is also included in our model; it is defined by $\tau(N) \triangleq \frac{[(N-1)-T_{\text{lag}}]^+ \ln 2}{T_{\text{double}}}$, where $[(N-1)-T_{\text{lag}}]^+ = \max((N-1)-T_{\text{lag}}, 0)$. In this formula, T_{lag} represents the time (in days) after which the tumor starts proliferating following the start of treatment and T_{double} is the tumor doubling time (in days). The goal is to maximize the tumor-BE of average dose for the sequence of fluence-maps $(u^1; u^2; \ldots; u^t; \ldots; u^N)$. This tumor-BE objective is given by

$$\sum_{t=1}^{N} (\alpha_0 \bar{A}u^t + \beta_0 (\bar{A}u^t)^2) - \tau(N). \tag{1}$$

The set of OAR is denoted by $\mathscr{M} \triangleq \mathscr{M}_1 \cup \mathscr{M}_2$; here, $\mathscr{M}_1$ and $\mathscr{M}_2$ are mutually exclusive sets of OAR with maximum dose and mean dose constraints, respectively. OAR in $\mathscr{M}_1$ are often called serial, whereas those in $\mathscr{M}_2$ parallel. We use the subscript/superscript m to index quantities related to OAR $m \in \mathscr{M}$. The set of voxels in OAR m is denoted by $\mathscr{N}_m \triangleq \{1, 2, \ldots, n_m\}$. Let A^m denote the $n_m \times k$ nonnegative dose-deposition matrix for OAR m, with A^m_j being its jth row. Thus, the dose delivered to the jth voxel in $\mathscr{N}_m$ in the tth session is $A^m_j u^t$. Let $\rho_m \triangleq \beta_m/\alpha_m$ denote the inverse alpha-over-beta ratio of the α and β parameters of the LQ dose-response model for OAR m. Suppose for OAR $m \in \mathscr{M}_1$ that a total dose $D^m_{\max}$ is known to be tolerated by each voxel if administered in N^m_{conv} equal-dose fractions. Similarly, suppose for OAR $m \in \mathscr{M}_2$ that total mean dose D^m_{mean} is known to be tolerated if administered in N^m_{conv} equal-dose fractions. Let $\text{BED}^m_\square = D^m_\square + \rho_m (D^m_\square)^2 / N^m_{\text{conv}}$ be the BED of total dose $D^m_\square$ if administered in N^m_{conv} equal-dose fractions, where $\square$ represents either max or mean depending on the type of OAR.

Consider the following optimization problem:

$$\begin{aligned}
(P)\ F^* = \ & \max \ \sum_{t=1}^{N} (\alpha_0 \bar{A}u^t + \beta_0 (\bar{A}u^t)^2) - \tau(N), \\
\text{subject to } & \sum_{t=1}^{N} (A^m_j u^t) + \sum_{t=1}^{N} \rho_m (A^m_j u^t)^2 \le \text{BED}^m_{\max}, \ \forall j \in \mathscr{N}_m, \ m \in \mathscr{M}_1, \\
& \sum_{t=1}^{N} \sum_{j=1}^{n_m} (A^m_j u^t) + \sum_{t=1}^{N} \rho_m \sum_{j=1}^{n_m} (A^m_j u^t)^2 \le n_m \text{BED}^m_{\text{mean}}, \ m \in \mathscr{M}_2, \\
& Su^t \le 0, \ t = 1, 2, \ldots, N, \\
& u^t \ge 0, \ t = 1, 2, \ldots, N.
\end{aligned}$$

Since $\tau(N)$ is fixed and it does not affect optimal solutions of (P), we will ignore it in the rest of this paper. The first constraints enforce that the BED to each voxel in a serial OAR is no more than the conventional BED; these are called the maximum dose constraints. The second constraints ensure that the average BED of doses administered to different voxels of a parallel OAR is bounded above by the conventional

BED; these are the mean dose constraints. The third constraints ensure, via an appropriate smoothness matrix S, that the relative absolute difference between geometrically adjacent components of u^t is within a range attainable by IMRT. The feasible region of (P) is bounded and thus it has an optimal solution. All constraints in (P) are convex in the concatenated fluence-map vector $u \triangleq (u^1; u^2; \ldots; u^N) \in \Re^{k \times N}$. The objective is also convex in u, but since we wish to maximize this function, (P) is *not* a convex problem. Specifically, it is a nonconvex QCQP. If we restricted consideration to stationary fluence-maps by setting $u^1 = u^2 = \ldots = u^N \triangleq v$ as in Saberian, Ghate, and Kim (2015a), the objective function becomes monotone in $\bar{A}v$ and thus the problem can be reformulated as an equivalent convex one. In this paper, we compare our nonstationary solutions against an optimal (stationary) solution to that convex problem.

To get a sense of the large scale of (P), note that for head-and-neck cancer, the typical number of beamlets, k, is about 3000 and the number of sessions N is 35. Thus, the dimension of u is about a 100,000. The number of voxels in a serial OAR can be about a 1000. As such, the total number of constraints in (P) can be as high as several thousand. Thus, efficient exact solution of this nonconvex QCQP is computationally difficult in practice.

First note that the objective and constraint functions in (P) are symmetric with respect to permutations over t. Second, standard approaches for approximate solution of large-scale nonconvex QCQPs call for convexification (Luo et al. 2010). For (P), since the objective function is the only source of nonconvexity, these approximation methods would amount to replacing the objective with a concave function that is additively separable and symmetric over t. This would create a symmetric, convex optimization problem. Such problems are known to possess symmetric, that is, stationary optimal solutions (Waterhouse 1983). Thus, standard convexification methods for (P) would not produce nonstationary fluence-maps. We thus borrow an approach called model predictive control (MPC) from the literature on finite-horizon, discrete-time, deterministic, constrained, non-linear control theory (Bertsekas 2007).

3 MODEL PREDICTIVE CONTROL

The idea in MPC, for a problem with N sessions such as (P), is simple. Beginning with the initial "state" of the problem, we first solve an N-session problem assuming stationary fluence-maps. Suppose an optimal stationary sequence of fluence-maps for this problem is

$$(\underbrace{u^1_*; u^1_*; \ldots; u^1_*}_{N \text{ times}}).$$

Implement fluence-map u^1_* in the first session only. This transforms the state of the problem to a new state at the beginning of the second session. Now solve an $N-1$- session problem to obtain an optimal stationary sequence of fluence-maps

$$(\underbrace{u^2_*; u^2_*; \ldots; u^2_*}_{N-1 \text{ times}})$$

and implement the fluence-map u^2_* in the second session only. Repeat this process until the last session, where a single-period problem is solved and the resulting optimal fluence-map is implemented.

To implement MPC on (P), we first need to define the "state" of the problem at the beginning of sessions $t \geq 1$. Toward this end, for each serial OAR $m \in \mathscr{M}_1$ and for each voxel $j \in \mathscr{N}_m$ in this OAR, let $z^{t,m}_j$ denote the total BED administered in the first $t-1$ treatment sessions. Let $z^{t,m} \triangleq (z^{t,m}_1, z^{t,m}_1, \ldots, z^{t,m}_{n_m})$, and z^t be the vector formed by concatenating vectors $z^{t,m}$ for all $m \in \mathscr{M}_1$. Similarly, for each parallel OAR $m \in \mathscr{M}_2$, let $w^{t,m}$ denote the total average (over all voxels in $\mathscr{N}_m$) BED administered in the first $t-1$ sessions. Let w^t be the vector formed by concatenating vectors $w^{t,m}$ for all $m \in \mathscr{M}_2$. Then we define the

state as $[z^t, w^t]$. We will need the optimization problem

$$(P_t) \ \max \ \sum_{t=1}^{N} (\alpha_0 \bar{A} u^t + \beta_0 (\bar{A} u^t)^2)$$

$$\sum_{l=t}^{N} (A_j^m u^l + \rho_m (A_j^m u^l)^2) \leq \text{BED}_{\max}^m - z_j^{t,m}, \ \forall j \in \mathscr{N}_m, \ m \in \mathscr{M}_1,$$

$$\sum_{j=1}^{n_m} \sum_{l=t}^{N} ((A_j^m u^l) + \rho_m (A_j^m u^t)^2) \leq n_m \text{BED}_{\text{mean}}^m - n_m w^{t,m}, \ m \in \mathscr{M}_2,$$

$$Su^l \leq 0, \ l = t, t+1, \ldots, N,$$

$$u^l \geq 0, \ l = t, t+1, \ldots, N$$

in our precise listing of the MPC algorithm below.

The MPC Algorithm

1. Set $t = 1$, and begin with the initial state $[z^1, w^1] = [\vec{0}, \vec{0}]$.
2. **WHILE** $t \leq N$
   A. Let $u^t = u^{t+1} = \ldots = u^N$ in (P_t) and solve it. Let $(\underbrace{u_*^t; u_*^t; \ldots; u_*^t}_{N-t+1 \text{ times}})$ denote a stationary optimal solution of (P_t) obtained in this manner.
   B. Implement fluence-map u_*^t in session t and update the state as

$$z_j^{t+1,m} = z_j^t + A_j^m u_*^t + \rho_m (A_j^m u_*^t)^2, \ j = 1, 2, \ldots, n_m, \ m \in \mathscr{M}_1,$$

$$w^{t+1,m} = w^{t,m} + \frac{\sum_{j=1}^{n_m} \left[A_j^m u_*^t + \rho_m (A_j^m u_*^t)^2 \right]}{n_m}, \ m \in \mathscr{M}_2.$$

   C. set $t = t + 1$ and go back to Step 2.
3. **END WHILE**

This algorithm delivers the sequence of fluence-maps $(u_*^1, u_*^2, \ldots, u_*^N)$ that is feasible to problem (P). It seems plausible that this sequence could be *nonstationary*. The next theorem dashes this hope.

Theorem 1 Consider an optimization problem of the form

$$\max \ \sum_{t=1}^{N} (\alpha_0 \bar{A} u^t + \beta_0 (\bar{A} u^t)^2),$$

$$g_m(u^1, u^2, \ldots, u^N) \leq b_m, \ m = 1, 2, \ldots, M.$$

Suppose that functions g_m are convex and symmetric. When MPC is implemented on this problem, it *cannot* deliver a nonstationary solution that is strictly better than all stationary solutions. In particular, since the constraint functions in (P) are convex and symmetric, this stationarity property holds for our MPC implementation on problem (P).

Proof. We provide a proof by contradiction. Suppose that MPC delivers a nonstationary solution of the form

$$\mathbf{u_N} = \big(\underbrace{u^1, \ldots, u^1}_{K_1 \text{ times}}, \underbrace{u^2, \ldots, u^2}_{K_2 \text{ times}}, \ldots, \underbrace{u^N, \ldots, u^N}_{K_N \text{ times}} \big),$$

where $u^i \neq u^j$ for $i \neq j$. Here, $K_1 + K_2 + \ldots, K_N = N$. Then MPC must have returned

$$\mathbf{u_{N\text{-}1}} = (\underbrace{u^1, \ldots, u^1}_{K_1 \text{ times}}, \underbrace{u^2, \ldots, u^2}_{K_2 \text{ times}}, \ldots, \underbrace{u^{N-1}, \ldots, u^{N-1}}_{K_{N-1}+K_N \text{ times}})$$

when $t = 1 + K_1 + K_2 + \ldots + K_{N-2}$. Continuing backwards, MPC must have returned

$$\mathbf{u}_2 = (\underbrace{u^1, \ldots, u^1}_{K_1 \text{ times}}, \underbrace{u^2, \ldots, u^2}_{K_2+K_3 \ldots + K_N \text{ times}})$$

when $t = 1 + K_1$. Finally, MPC must have delivered

$$\mathbf{u}_1 = (\underbrace{u^1, \ldots, u^1}_{N \text{ times}})$$

when $t = 1$. We define the notation $F^t(u^1, u^2, \ldots, u^N) = \sum_{l=t}^{N} f(u^l)$, where $f(u^l) = \alpha_0 \bar{A} u^l + \beta_0 (\bar{A} u^l)^2$ for $l = 1, 2, \ldots, N$. Now consider the following three cases.

Case 1: $\bar{A}u^1 > \bar{A}u^2$. In this case, $f(u^1) > f(u^2)$. Then

$$F^{1+K_1}(\mathbf{u_1}) = (K_2 + \ldots + K_N) f(u^1) > (K_2 + K_3 \ldots + K_N) f(u^2) = F^{1+K_1}(\mathbf{u_2}).$$

Since

$$\mathbf{u_1} = (\underbrace{u^1, u^1, \ldots, u^1}_{N \text{ times}})$$

is a feasible solution to the N-session problem, this strict inequality contradicts the optimality of

$$(\underbrace{u^2, \ldots, u^2}_{K_2+K_3 \ldots + K_N \text{ times}})$$

in session $t = 1 + K_1$.

Case 2: $\bar{A}u^1 < \bar{A}u^2$. Let

$$v = \frac{K_1 u^1 + (K_2 + K_3 \ldots K_N) u^2}{N},$$

and consider the alternative solution

$$\mathbf{u_{alt}} = (\underbrace{v, v, \ldots, v}_{N \text{ times}}).$$

Now, since constraints g_m are convex and symmetric, Jensen's inequality can be employed to show that this alternative solution is feasible to the N-session problem. Moreover,

$$\bar{A}v = \frac{\bar{A}(K_1 u^1 + (K_2 + \ldots + K_N) u^2)}{N} = \frac{K_1 \bar{A} u^1 + (K_2 + \ldots + K_N) \bar{A} u^2}{N} > \frac{K_1 \bar{A} u^1 + (K_2 + \ldots + K_N) \bar{A} u^1}{N} = \bar{A} u^1.$$

Therefore, $f(v) > f(u^1)$. Consequently,

$$F(\mathbf{u_{alt}}) = \sum_{t=1}^{N} f(v) > \sum_{t=1}^{N} f(u^1) = F(\mathbf{u_1}).$$

This contradicts the optimality of $\mathbf{u_1}$ when $t = 1$.

Case 3: $\bar{A}u^1 = \bar{A}u^2$. In this case, $f(u^1) = f(u^2)$ and hence $F(\mathbf{u_1}) = F(\mathbf{u_2})$. Therefore, we can replace $\mathbf{u_2}$ with $\mathbf{u_1}$ when $t = 1 + K_1$.

Applying this argument recursively, we see that MPC cannot return a nonstationary solution that is strictly better than all stationary solutions. □

In view of this theorem, in order to force MPC to return a nonstationary sequence of fluence-maps with a high tumor-BE than any stationary fluence-map, we need to use a surrogate objective for which the conclusion of the theorem fails. We propose the total number of tumor cells remaining (TNTCR) objective from (Kim, Ghate, and Phillips 2012; Saberian, Ghate, and Kim 2015b) for this purpose. TNTCR is given by $\sum_{i=1}^{n} x_i^{N+1}$, where x_i^t denotes the total number of remaining tumor cells at the beginning of the tth session with dynamics $x_i^{t+1} = x_i^t \exp(-\alpha_0(A_i u^t) - \beta_0(A_i u^t)^2)$ for $i = 1, 2, \ldots, n$. In the next section, we study whether or not this surrogate objective induces MPC to return nonstationary fluence-maps with a higher tumor-BE than the best stationary fluence-map.

4 NUMERICAL RESULTS

We conducted numerical experiments on five different head-and-neck test cases from Saberian, Ghate, and Kim (2015b). These cases included four OAR: spinal cord (serial), brainstem (serial), left and right parotids (parallel; parotid glands are salivary glands located just in front of the two ears). Parameters ρ_m were fixed at $1/3$ for all OAR as is common in the clinical literature (Fowler 1990, 2001, 2007, 2008). All tolerance doses were set as in Saberian, Ghate, and Kim (2015a), which were in turn taken from a standard head-and-neck treatment protocol (Marks et al. 2010). The number of sessions was fixed at $N = 35$. We also included a maximum BED constraint (corresponding to a total dose of 90 Gy in 35 sessions) on the tumor to facilitate dose homogeneity. We fixed $\alpha_0 = 0.35$ Gy^{-1} and $\beta_0 = 0.035$ Gy^{-2} as is standard in the clinical literature (Fowler 1990, 2001, 2007, 2008). The initial cell density was assumed to be homogeneous over all tumor voxels. All computer simulations were performed on a 3.1 GHz iMac desktop with 16 GB RAM using MATLAB.

We used the TNTCR objective function to search for nonstationary solutions via MPC, and compared them with optimal solutions to the stationary version of problem (P). This comparison was based on the tumor-BE objective values in problem (P). The results showed that MPC was able to discover nonstationary solutions better than stationary solutions in three of the five cases. The percentage improvements attained by a nonstationary solution over the best stationary solution were 5.2%, 3.2%, and 2.1% for these three cases. In the other two cases, stationary solutions turned out to be superior than the nonstationary solutions returned by MPC by 3.2% and 0.92%. Figure 1 shows the tumor-BE for every voxel averaged over 35 sessions delivered by the nonstationary and best stationary solution for the first test-case. It illustrates that the nonstationary tumor-BE seems higher than the best stationary tumor-BE in this case. Figure 2 shows the nonstationary intensity profiles administered by MPC for ten sample tumor voxels for the same test-case. Figure 3 shows the nonstationary doses delivered to 500 sample tumor voxels by MPC. It also illustrates the average dose over all tumor voxels for this nonstationary solution (red line) as well as for the best stationary one (light blue line).

We also numerically optimized N for one of the cases, both for the stationary and MPC methods. This was achieved by fixing N at all its possible values $\{1, 2, \ldots, 50\}$, solving the problem for each fixed N, and comparing the tumor-BE. Parameters T_{lag} and T_{double} were fixed at 2 and 10 days, respectively, for both methods. The optimal number of treatment session were 28 and 31 for the stationary and MPC methods, respectively. Moreover, the nonstationary solution delivered 8.7% more tumor-BE in optimality than the stationary solution. A 5-10% increase in tumor-BE is comparable to other results reported in the existing literature on the stationary fractionation problem (Saberian, Ghate, and Kim 2015a). Another interesting result was that the tumor-BE using the MPC method behaved similar to the one previously observed with

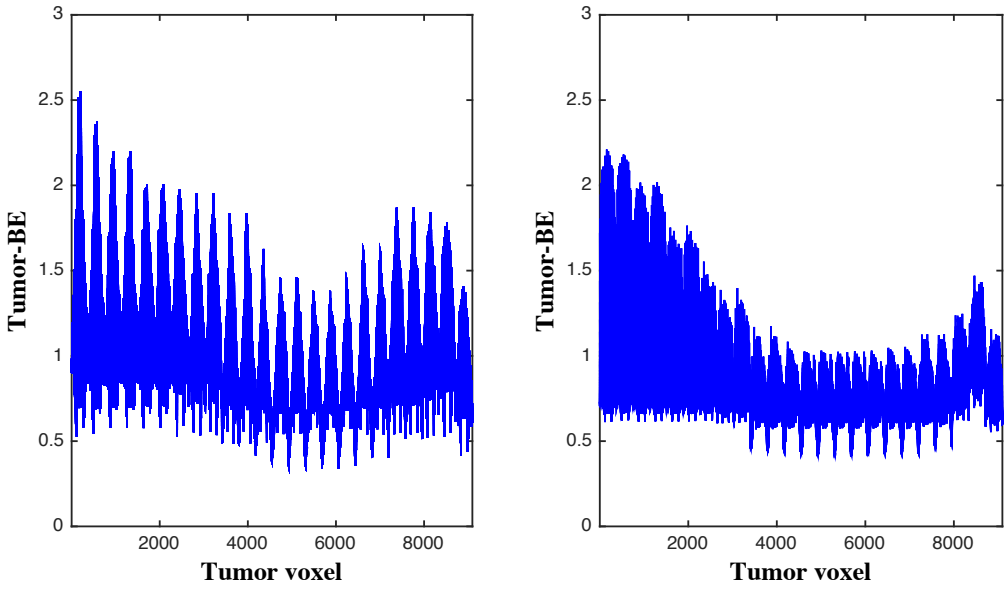

Figure 1: Tumor-BE averaged over 35 sessions administered by the nonstationary (left panel) and best stationary (right panel) solutions for test-case 1.

the stationary method (Saberian, Ghate, and Kim 2015a). That is, the tumor-BE increases with increasing N, reaches its maximum value, and then decreases.

In summary, we conclude that MPC may be able to discover nonstationary fluence-maps with a higher tumor-BE than stationary ones. An interesting direction for future research would be to optimize the number of treatment sessions N by using nonstationary fluence-maps.

ACKNOWLEDGMENT

This research was funded in part by the National Science Foundation through grant CMMI 1054026.

REFERENCES

Armpilia, C. I., R. G. Dale, and B. Jones. 2004. "Determination of the Optimum Dose Per Fraction in Fractionated Radiotherapy when there is Delayed Onset of Tumour Repopulation During Treatment". *The British Journal of Radiology* 77 (921): 765–767.

Bertsekas, D. P. 2007. *Dynamic Programming and Optimal Control*. third ed, Volume 1 and 2. Nashua, NH: Athena Scientific.

Bertuzzi, A., C. B. F. Papa, and C. Sinisgalli. 2013. "Optimal Solution for a Cancer Radiotherapy Problem". *Journal of Mathematical Biology* 66 (1-2): 311–349.

Ehrgott, M., C. Guler, H. W. Hamacher, and L. Shao. 2008. "Mathematical Optimization in Intensity Modulated Radiation Therapy". *4OR* 6 (3): 199–262.

Fowler, J. F. 1990. "How Worthwhile are Short Schedules in Radiotherapy?: A Series of Exploratory Calculations". *Radiotherapy and Oncology* 18 (2): 165–181.

Fowler, J. F. 2001. "Biological Factors Influencing Optimum Fractionation in Radiation Therapy". *Acta Oncologica* 40 (6): 712–717.

Fowler, J. F. 2007. "Is there an Optimal Overall Time for Head and Neck Radiotherapy? A Review with New Modeling". *Clinical Oncology* 19 (1): 8–27.

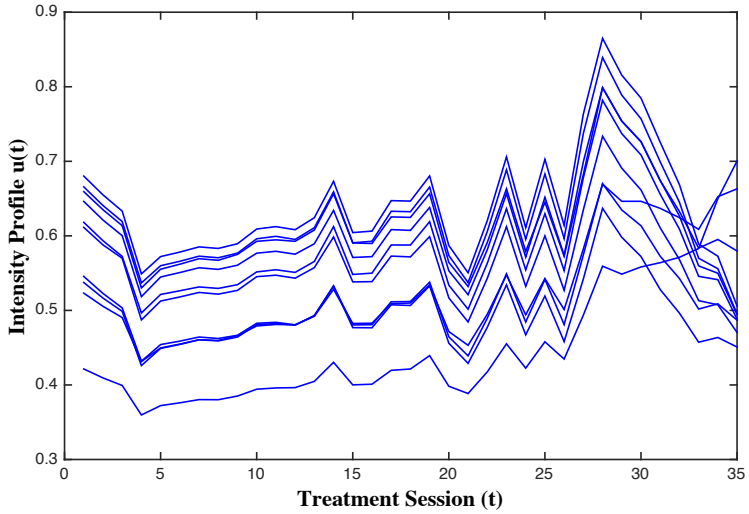

Figure 2: Nonstationary intensity profiles for ten sample tumor voxels.

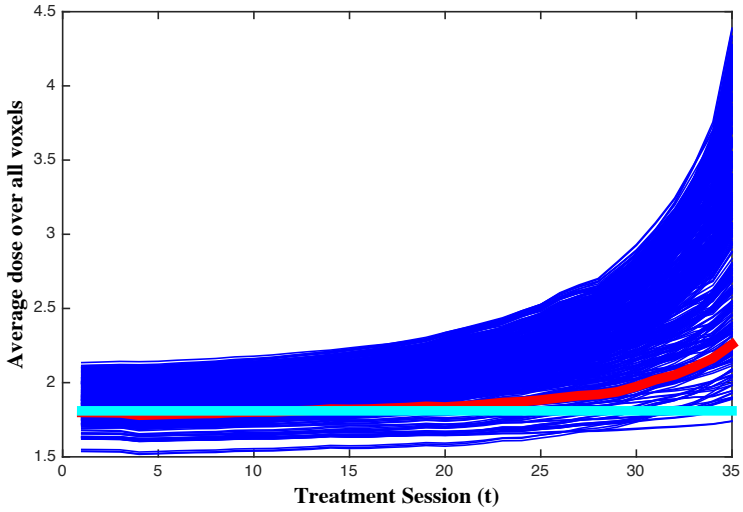

Figure 3: Average tumor doses delivered over 35 treatment sessions by the stationary (light blue line) and the nonstationary (red line) solutions. The dark blue lines show the doses administered by the nonostationary solution for 500 sample tumor voxels.

Fowler, J. F. 2008. "Optimum Overall Times II: Extended Modelling for Head and Neck Radiotherapy". *Clinical Oncology* 20 (2): 113–126.

Fowler, J. F., and M. A. Ritter. 1995. "A Rationale for Fractionation for Slowly Proliferating Tumors Such as Prostatic Adenocarcinoma". *International Journal of Radiation Oncology Biology Physics* 32 (2): 521–529.

Hall, E. J., and A. J. Giaccia. 2005. *Radiobiology for the Radiologist*. Philadelphia, Pennsylvania, USA: Lippincott Williams & Wilkins.

Jeraj, R., and P. Keall. 1999. "Monte Carlo-based Inverse Treatment Planning". *Physics in Medicine and Biology* 44 (8): 1885–1896.

Jones, B., L. T. Tan, and R. G. Dale. 1995. "Derivation of the Optimum Dose per Fraction from the Linear Quadratic Model". *The British Journal of Radiology* 68 (812): 894–902.

Keller, H., G. Meier, A. Hope, and M. Davison. 2012. "Fractionation Schedule Optimization for Lung Cancer Treatments using Radiobiological and Dose Distribution Characteristics". *Medical Physics* 39 (6): 3811–3811.

Kim, M., A. Ghate, and M. H. Phillips. 2012. "A Stochastic Control Formalism for Dynamic Biologically Conformal Radiation Therapy". *European Journal of Operational Research* 219 (3): 541 – 556.

Luo, Z.-Q., W.-K. Ma, A. M. C. So, and Y. Ye. 2010. "Semidefinite Relaxation of Quadratic Optimization Problems". *IEEE Signal Processing Magazine* 27 (3): 20–34.

Marks, L. B., E. D. Yorke, A. Jackson, R. K. T. Haken, L. S. Constine, A. Eisbruch, S. M. Bentzen, J. Nam, and J. O. Deasy. 2010. "Use of Normal Tissue Complication Probability Models in the Clinic". *International Journal of Radiation Oncology Biology Physics* 76 (3): S10–S19.

Mizuta, M., S. Takao, H. Date, N. Kishimoto, K. L. Sutherland, R. Onimaru, and H. Shirato. 2012. "A Mathematical Study to Select Fractionation Regimen Based on Physical Dose Distribution and the Linear-Quadratic Model". *International Journal of Radiation Oncology Biology Physics* 84 (3): 829 – 833.

Romeijn, H. E., R. K. Ahuja, J. F. Dempsey, and A. Kumar. 2006. "A New Linear Programming Approach to Radiation Therapy Treatment Planning Problems". *Operations Research* 54 (2): 201–216.

Romeijn, H. E., and J. F. Dempsey. 2008. "Intensity Modulated Radiation Therapy Treatment Plan Optimization". *TOP* 16 (2): 215–243.

Saberian, F., A. Ghate, and M. Kim. 2015a, April. "Spatiotemporally Integrated Fractionation in Radiotherapy". *under review at INFORMS Journal on Computing, preprint available at http://faculty.washington.edu/archis/upaper-apr-2015.pdf*.

Saberian, F., A. Ghate, and M. Kim. 2015b. "A Theoretical Stochastic Control Framework for Adapting Radiotherapy to Hypoxia". *under review at Physics in Medicine and Biology, preprint available at http://faculty.washington.edu/archis/hypoxia.pdf*.

Saberian, F., A. Ghate, and M. Kim. 2015c. "A Two-variable Linear Program Solves the Standard Linear-quadratic Formulation of the Fractionation Problem in Cancer Radiotherapy". *Operations Research Letters* 43 (3): 254–258.

Saberian, F., A. Ghate, and M. Kim. 2016. "Optimal Fractionation in Radiotherapy with Multiple Normal Tissues". *Mathematical Medicine and Biology* 33 (2): 211–252.

Shepard, D. M., M. C. Ferris, G. H. Olivera, and T. R. Mackie. 1999. "Optimizing the Delivery of Radiation Therapy to Cancer Patients". *SIAM Review* 41 (4): 721–744.

Siebers, J. V., S. D. Tong, M. Lauterbach, Q. W. Wu, and R. Mohan. 2001. "Acceleration of Dose Calculations for Intensity-modulated Radiotherapy". *Medical Physics* 28 (6): 903–910.

Spirou, S. V., and C. S. Chui. 1998. "A Gradient Inverse Planning Algorithm with Dose-volume Constraints". *Medical Physics* 25 (3): 321–333.

Tian, Z., M. Zarepisheh, X. Jia, and S. B. Jiang. 2013. "The Fixed-point Iteration Method for IMRT Optimization with Truncated Dose Deposition Coefficient Matrix". Technical report, arxiv.

Unkelbach, J., D. Craft, E. Saleri, J. Ramakrishnan, and T. Bortfeld. 2013. "The Dependence of Optimal Fractionation Schemes on the Spatial Dose Distribution". *Physics in Medicine and Biology* 58 (1): 159–167.

Unkelbach, J., C. Zeng, and M. Engelsman. 2013. "Simultaneous Optimization of Dose Distributions and Fractionation Schemes in Particle Radiotherapy". *Medical Physics* 40 (9): 091702.

Waterhouse, W. C. 1983. "Do Symmetric Problems have Symmetric Solutions?". *The American Mathematical Monthly* 90 (6): 378–387.

Webb, S. 2010. *Contemporary IMRT: Developing Physics and Clinical Implementation*. Bristol, UK: IOP Publishing.

Webb, S., and M. Oldham. 1996. "A Method to Study the Characteristics of 3D Dose Distributions Created by Superposition of Many Intensity-modulated Beams Delivered via a Slit Aperture with Multiple Absorbing Vanes". *Physics in Medicine and Biology* 41 (10): 2135–2153.

AUTHOR BIOGRAPHIES

ARCHIS GHATE is an Associate Professor in the Department of Industrial and Systems Engineering at the University of Washington, Seattle. He holds a PhD in Industrial and Operations Engineering from University of Michigan. His research interests lie in stochastic and dynamic optimization. His email address is archis@uw.edu.

ALI AJDARI is a PhD student in Industrial and Systems Engineering at the University of Washington, Seattle. He obtained his M.Sc. degree in Industrial and Systems Engineering from the Sharif University of Technology, Iran in 2012. His research interests are in applied operations research, simulation optimization, and mathematical programming. email address is ajdari@uw.edu.