Physics and Imaging in Radiation Oncology 39 (2026) 101003

Contents lists available at ScienceDirect

Physics and Imaging in Radiation Oncology

journal homepage: www.sciencedirect.com/journal/physics-and-imaging-in-radiation-oncology

Original Research Article

Optimal radiotherapy dose scheduling with variable fraction sizes and breaks via sequential mixed-integer convex programming

Anqi Fu [a],*, Zeno Gouw [b], Jeho Jeong [a], Joseph O. Deasy [a]

[a] Department of Medical Physics, Memorial Sloan-Kettering, 321 East 61st Street, 10065, NY, USA
[b] Department of Radiation Oncology, Netherlands Cancer Institute, Plesmanlaan 121, 1066 CX, Amsterdam, Netherlands

ARTICLE INFO

Keywords:
Radiation oncology
Treatment planning
Numerical optimization
Mixed-integer programming
Optimal control

ABSTRACT

Background and Purpose: Radiotherapy is typically delivered in consecutive equi-dose fractions, but research suggests a non-uniform dose schedule may produce a higher tumor control probability (TCP). We developed an optimization method that automatically constructs the best dose schedule with variable fraction sizes and treatment breaks based on a tumor dose–response model calibrated to head-and-neck squamous cell carcinoma, which captures the impact of cellular resource competition and hypoxia.

Materials and Methods: We formulated the dose scheduling problem as a finite-horizon optimal control problem. Fraction size was constrained by an upper bound on the biologically effective dose to normal tissue, along with a daily dose limit. This problem is nonconvex, so we employed a heuristic called the convex-concave procedure to solve a sequence of mixed-integer convex approximations that converges to a good estimate of the solution.

Results: The optimal schedules adhered to a pattern consisting of an initial "primer shot", followed by a 1 week break, and concluding with six small equi-dose fractions and a final large fraction. The primer shot killed proliferating cells, freeing up resources so hypoxic cells could reoxygenate during the treatment break. These reoxygenated cells are more radiosensitive, therefore the schedule waited until all cells have reoxygenated before delivering its largest dose. In computational experiments, our schedule achieved a 12% higher TCP than the standard equi-dose weekday schedule.

Conclusions: An optimization method was developed to construct non-uniform dose schedules based on a model of tumor dose–response in the presence of hypoxia, yielding significant improvements in TCP.

1. Introduction

In conventional radiotherapy, treatment is delivered in consecutive weekday fractions of uniform size. However, research suggests that alternate dose schedules can achieve better tumor control at lower risk of normal tissue complication [1]. One reason for this is hypoxia. Hypoxic cells are radioresistant, and hypoxia in solid tumors correlates with poor outcomes of radiotherapy [2,3]. On the other hand, clinical studies have shown that hypoxia often resolves within the first few weeks of radiation treatment [2,4,5]. These observations imply that the biological impact of a dose fraction varies over time, and thus, the optimal dose schedule must vary as well in response to the tumor.

To incorporate this observation into treatment planning, we need a mathematical model of tumor dose–response over time. Early frameworks like the linear-quadratic (LQ) model [6,7] and Cohen's kinetic models [8,9] do not take into account the role of hypoxia and reoxygenation in cellular response to radiation. More recent publications propose models that capture a wide range of radiobiological phenomena, including the effect of oxygen, along with algorithms such as simulated annealing [10], reinforcement learning [11], and numerical search heuristics [12,13] to solve the associated treatment planning problem. In this paper, we focus on a tumor dose–response model that captures the interplay between cellular competition for limited resources, hypoxia-induced radioresistance, and proliferation. The model has been validated on several types of tumors, including head-and-neck squamous cell carcinoma (HNSCC) [14] and non-small cell lung cancer (NSCLC) [15]. Using this model, Gouw et al. [16] simulated tumor response to radiotherapy for six common fractionation schemes, each with constant fraction size, while varying the length and location of interfractional breaks. Through a grid search, they found that the best schedule comprises an initial "primer shot", followed by a treatment break to allow tumor cells to reoxygenate, and concluding with the remaining fractions delivered in succession.

* Corresponding author.
E-mail addresses: anqif@alumni.stanford.edu (A. Fu), z.gouw@nki.nl (Z. Gouw), jeongj@mskcc.org (J. Jeong), deasyj@mskcc.org (J.O. Deasy).
URL: https://anqif.github.io (A. Fu).

https://doi.org/10.1016/j.phro.2026.101003
Received 28 January 2026; Received in revised form 12 May 2026; Accepted 19 May 2026
Available online 26 May 2026

Our paper seeks to build on the work of Gouw et al. with three key aims. First, we will formally define the dose scheduling problem as a mathematical optimization problem that incorporates both the tumor dose–response equations in Jeong et al. [14], as well as other clinical constraints. Second, we will expand the solution space by permitting fraction sizes to vary along with treatment breaks. Finally, we will develop and release an open-source Python package that implements our optimal dose scheduling method.

2. Methods and materials

2.1. Optimal dose scheduling problem formulation

In this section, we provide a sketch of the optimal dose scheduling problem and leave the details of its derivation to Supplementary Section A. We formulate the dose scheduling problem as a finite-horizon optimal control problem [17] using the tumor dose–response model with parameter values calibrated to HNSCC described in Jeong et al. [14] (See Supplementary Section B for an overview of their model and Supplementary Table S1 for a list of associated variables and parameters). Cells reside in one of three compartments: proliferative (P) cells with access to both oxygen and glucose, intermediately hypoxic (I) cells with access to glucose but not oxygen, and extremely hypoxic (H) cells without access to oxygen or glucose. During radiotherapy, a fraction of cells in each compartment are doomed, i.e., lethally damaged such that they die after failed mitosis, freeing up nutrients so viable cells can reoxygenate and recompartmentalize.

A course of treatment takes place over a fixed number of time steps T, with each dose fraction d_t delivered during time step t. Note that we can set $d_t = 0$ to indicate a break. We represent the number of doomed (resp. viable) cells in compartment $X \in \{P, I, H\}$ at the beginning of step t as $N_t^{X,d}$ (resp. $N_t^{X,v}$). The state variables of the optimal control problem are the viable/doomed cell counts in each compartment $N_t := (N_t^{P,v}, N_t^{P,d}, N_t^{I,v}, N_t^{I,d}, N_t^{H,v}, N_t^{H,d})$ and the control variables are the dose fractions d_t.

Our objective is to find a dose schedule $d = (d_1, \dots, d_T) \in \mathbf{R}_+^T$ that minimizes the tumor cell survival fraction (SF) at the end of treatment:

$$f_0(N_{T+1}) := (N_{T+1}^P + N_{T+1}^I + N_{T+1}^H)/N_1^{tot}, \tag{1}$$

where $N_1^{tot} = N_1^P + N_1^I + N_1^H$ is the total number of cells at the start of treatment. To mitigate damage to normal tissue, we enforce an upper bound on the biologically effective dose (BED_3) to a single organ-of-interest (OOI)

$$f_1(d) := \sum_{t=1}^{T} \left(d_t + \frac{\beta_N}{\alpha_N} d_t^2 \right) \le B^{\max}, \tag{2}$$

where $\alpha_N > 0$ and $\beta_N > 0$ are the normal tissue LQ parameters, which we select so that $\alpha_N/\beta_N = 3$, and $B^{\max}$ is a clinician-specified maximum BED_3.[1] Additional dose constraints, such as mean and maximum dose to an organ, can also be represented via a linear constraint of the form $Cd \le b$.

The cells in each compartment evolve according to the processes described in Supplementary Section B. To model these dynamics, we consider pre- and post-recompartmentalization cell distributions separately in our analysis. Let $\dot{N}_t = (\dot{N}_t^{H,v}, \dot{N}_t^{H,d})$ be the H compartment cell counts in time t after radiation cell-kill takes place, but before hypoxic cell loss or recompartmentalization, and let $\tilde{N}_t = (\tilde{N}_t^{P,v}, \tilde{N}_t^{P,d}, \tilde{N}_t^{I,v}, \tilde{N}_t^{I,d}, \tilde{N}_t^{H,v}, \tilde{N}_t^{H,d})$ be the cell counts for all compartments in time step t after radiation cell-kill, proliferation, and cell loss take place, but before recompartmentalization. By combining supplementary equations S4–S9, taking logs, and applying some algebraic manipulation, we can write the pre-recompartmentalization cell dynamics in the form

$$Av_t = 0, \quad t = 1, \dots, T, \tag{3}$$

$$g(v_t) + h(v_t) \le 0, \quad t = 1, \dots, T, \tag{4}$$

where $v_t := (d_t, N_t, \dot{N}_t, \tilde{N}_t)$, g is a convex function, h is a concave function, and A is a constant matrix that defines a set of simultaneous linear equations over v_t.

After recompartmentalization, cells must be distributed in one of three configurations, determined by the capacity of the compartments. Each of these configurations can be represented as a mutually exclusive set of constraints on our state variables:

1. All cells lie in the P compartment

$$\tilde{N}_t^{tot} \le N_1^P, \quad N_{t+1}^{P,v} + N_{t+1}^{P,d} = \tilde{N}_t^{tot}, \quad N_{t+1}^{I,v} + N_{t+1}^{I,d} = 0, \quad N_{t+1}^{H,v} + N_{t+1}^{H,d} = 0. \tag{5}$$

2. The P compartment is full, all excess cells lie in the I compartment, and the H compartment is empty

$$N_1^P \le \tilde{N}_t^{tot} \le N_1^P + N_1^I, \quad N_{t+1}^{P,v} + N_{t+1}^{P,d} = N_1^P, \quad N_{t+1}^{I,v} + N_{t+1}^{I,d} = \tilde{N}_t^{tot} - N_1^P, \quad N_{t+1}^{H,v} + N_{t+1}^{H,d} = 0. \tag{6}$$

3. The P and I compartments are full, and all excess cells lie in the H compartment

$$\tilde{N}_t^{tot} \ge N_1^P + N_1^I, \quad N_{t+1}^{P,v} + N_{t+1}^{P,d} = N_1^P, \quad N_{t+1}^{I,v} + N_{t+1}^{I,d} = N_1^I, \quad N_{t+1}^{H,v} + N_{t+1}^{H,d} = \tilde{N}_t^{tot} - N_1^P - N_1^I. \tag{7}$$

Here we define $\tilde{N}_t^{tot} := \tilde{N}_t^P + \tilde{N}_t^I + \tilde{N}_t^H$.

We do not know *a priori* which cell configuration corresponds to the optimal dose schedule, and consequently, we do not know which set of constraints to enforce in time step t. Thus, we incorporate the post-recompartmentalization configuration as a state in the optimal control problem. Define a new state variable $z_t \in \{0,1\}^3$, where

$$z_{t,s} = \begin{cases} 1 & \text{recompartmentalized cells in configuration } s \text{ at time } t \\ 0 & \text{otherwise} \end{cases}$$

is an indicator of configuration $s \in \{1,2,3\}$, and restrict $z_{t,1} + z_{t,2} + z_{t,3} = 1$ as only one set of constraints (5)–(7) can be active at a time. These indicator variables allow us to write the recompartmentalization dynamics as a set of binary mixed-integer linear inequality constraints

$$f_2(N_{t+1}, \tilde{N}_t, z_t) := FN_{t+1} + \tilde{F}\tilde{N}_t + Gz_t + c \le 0, \quad t = 1, \dots, T, \tag{8}$$

where the constants F, $\tilde{F}$, and G are matrices and c is a vector. (See Supplementary Section A for a more thorough discussion). Together, expressions (3), (4), and (8) fully characterize the cell dynamics in our optimal control problem.

We are now ready to define the optimal dose scheduling problem:

$$\begin{array}{ll} \text{minimize} & f_0(N_{T+1}) \\ \text{subject to} & f_1(d) \le B^{\max}, \quad Cd \le b \\ & Av_t = 0, \quad g(v_t) + h(v_t) \le 0, \quad t = 1, \dots, T \\ & f_2(N_{t+1}, \tilde{N}_t, z_t) \le 0, \quad t = 1, \dots, T \\ & d \ge 0, \quad N \ge 0, \quad \dot{N} \ge 0, \quad \tilde{N} \ge 0 \\ & z_t \in \{0,1\}^3, \quad z_{t,1} + z_{t,2} + z_{t,3} = 1 \end{array} \tag{9}$$

with variables $d = (d_1, \dots, d_T)$, $N = (N_2, \dots, N_{T+1})$, $\dot{N} = (\dot{N}_1, \dots, \dot{N}_T)$, $\tilde{N} = (\tilde{N}_1, \dots, \tilde{N}_T)$, and $z = (z_1, \dots, z_T)$. This problem is nonconvex because h is a concave function and z_t is restricted to the binary integer domain.

[1] This generic normal tissue BED_3 "budget" is a simplification intended to prevent unlimited dose escalation by the optimization algorithm, and not meant as a realistic representation of all OOIs.

2.2. Sequential mixed-integer convex optimization

Prob. (9) is a mixed-integer program (MIP) with nonconvex inequality constraints. It is in general computationally difficult to solve exactly. We will derive an estimate of its optima by constructing and solving a sequence of mixed-integer convex programs (MICPs), which provide a computationally-feasible approximation to the original problem.

A general MICP has the form [18]

$$\begin{array}{ll} \text{minimize} & f_0(x,z) \\ \text{subject to} & f_i(x,z) \le 0, \quad i=1,\dots,M \\ & z \in \mathbf{Z}^{n_z}, \end{array} \tag{10}$$

where $x \in \mathbf{R}^{n_x}$ is a continuous variable, $z \in \mathbf{Z}^{n_z}$ is a discrete variable, and $f_i : \mathbf{R}^{n_x} \times \mathbf{Z}^{n_z} \to \overline{\mathbf{R}}$ are convex functions for $i = 0,\dots,M$. This problem can be solved using branch-and-bound [19,20] or polyhedral outer approximation combined with linear/conic programming methods [21–23]. For details, see the review of MICP algorithms in Lubin et al. [24] and any textbook on convex optimization [25,26].

The objective of Prob. (9) is convex, and all constraints are convex or mixed-integer convex except for (4). This is because the term $h(v_t)$ in the constraint function is concave. To form an MICP approximation of Prob. (9), we will replace $h(v_t)$ with its linearization: given a point $v^k := (v_1^k,\dots,v_T^k)$, define the linearization of h around $v_t^k = (d_t^k, N_{t-1}^k, \tilde{N}_t^k)$ to be

$$\hat{h}(v_t; v_t^k) := h(v_t^k) + \nabla h(v_t^k)^T (v_t - v_t^k). \tag{11}$$

Then, the mixed-integer convex approximation of (11) is

$$g(v_t) + \hat{h}(v_t; v_t^k) \le \delta_t, \quad t = 1,\dots,T,$$

where $\delta_t \ge 0$ is a slack variable we have introduced to allow room for error in the estimate. (Since h is concave, $\hat{h}(\cdot; v_t^k) \ge h(\cdot)$ for all v_t^k). The slack is penalized in the objective, *e.g.*, with a penalty function

$$r(\delta) := \lambda \sum_{t=1}^{T} \mathbf{1}^T \delta_t$$

and a slack penalty parameter $\lambda > 0$, which is typically determined via trial-and-error.

With this modification, we can now define MICP k to be

$$\begin{array}{ll} \text{minimize} & f_0(N_{T+1}) + r(\delta) \\ \text{subject to} & f_1(d) \le B^{\max}, \quad Cd \le b \\ & Av_t = 0, \quad g(v_t) + \hat{h}(v_t; v_t^k) \le \delta_t, \quad t = 1,\dots,T \\ & f_2(N_{t+1}, \tilde{N}_t, z_t) \le 0, \quad t = 1,\dots,T \\ & d \ge 0, \quad N \ge 0, \quad \dot{N} \ge 0, \quad \tilde{N} \ge 0, \quad \delta \ge 0 \\ & z_t \in \{0,1\}^3, \quad z_{t,1} + z_{t,2} + z_{t,3} = 1 \end{array} \tag{12}$$

with respect to $d, z, N, \dot{N}, \tilde{N}$, and $\delta = (\delta_1,\dots,\delta_T)$. Starting from an initial point v^0, we iteratively solve MICP k with linearization point v^k, obtain an optimal point $v^\star$, and define the linearization point of MICP $k+1$ to be $v^{k+1} := v^\star$. This process continues until a stopping criterion is met, which is usually taken to be when the change in the objective falls below some user-specified threshold $\epsilon > 0$, *i.e.*,

$$p_{\text{opt}}^k - p_{\text{opt}}^{k+1} < \epsilon, \tag{13}$$

where p_{opt}^k is the optimal objective of Prob. (12) in iteration k. In our computational experiments, we found that a threshold of $\epsilon = 10^{-3}$ struck a good balance between runtime and treatment quality.

Algorithm 1 (*Sequential Mixed-Integer Convex Optimization*). *Given an initial point v^0 and a parameter $\lambda > 0$, repeat for iterations $k = 0, 1, \dots$*

1. *For $t = 1,\dots,T$, form the linearization of h around v_t^k according to Eq. (11).*
2. *Set v^{k+1} equal to an optimal point of problem (12).*

until stopping criterion (13) is satisfied.

Algorithm 1 falls into the broad category of majorization-minimization methods [27,28]. It is a mixed-integer variant of the convex-concave procedure (CCP) [29–32], a heuristic for finding local optima of nonconvex optimization problems. CCP has seen success in applications to adaptive radiotherapy [33, Section 4]. In the special case when no slack is allowed (*i.e.*, $\delta := 0$), it is guaranteed to converge provided v^0 is feasible for the original problem (9); see Supplementary Section C for a proof.

2.3. Computational implementation

We implemented the sequential mixed-integer convex optimization algorithm in Python 3.10 with CVXPY [34] and solved the MICPs using MOSEK [35,36]. To assess our algorithm's results, we carried out several computational experiments using tumor radiobiology parameter values typical of HNSCC [14] – in particular, we set an initial cell count of $N_1^{tot} = 6.4 \times 10^7$, growth fraction of 0.25, and cell loss factor of 0.92. Bounds on the maximum dose and normal tissue BED_3 were derived from the fractionation schedules studied in Gouw et al. [16] (The values of these bounds were chosen to facilitate comparison of our algorithm's results with the schedules in Gouw et al. and do not reflect actual clinical constraints for HNSCC). Total treatment duration varied between 10 and 19 days with weekend breaks, and the daily dose was restricted to a maximum of $D^{\max} = 18$ Gy. First, we fixed the number of fractions to 8 and compared the optimal dose schedule with/without a constant dose constraint, and with/without enforcing a 1 week treatment break after the initial dose. Then, we altered the value of $B^{\max}$ to examine the impact of normal tissue dose restrictions on the dose schedule. We also analyzed the optimal dose schedule under different treatment lengths T and numbers of fractions. Finally, we released a Python library that implements our dose scheduling method: https://github.com/anqif/opt_frac. For details on the implementation, including algorithm initialization and refinements, we refer the reader to Supplementary Section D along with Supplementary Tables S1 and S2.

3. Results

Fig. 1 depicts the daily dose and tumor SF of the optimal 8-fraction weekday schedule when doses were constrained to be (a) constant, (b) variable, and (c) variable with a treatment break after the initial fraction. In the first two cases, reoxygenation of the H compartment was complete by day 5, but the I compartment remained non-empty, so the cell-kill of subsequent doses was reduced due to the radioresistance of the intermediately hypoxic cells. However, when we allowed fraction sizes to vary *and* enforced a week-long treatment break (Fig. 1(c)), we achieved a much higher final TCP. The break provided enough time for all cells to reoxygenate, and the variability allowed us to deliver a much larger dose on the final day, so that we fully exploited the radiosensitivity of the remaining proliferating cells. We dub the initial dose the "primer shot", as it crucially "primes" the tumor cells to begin reoxygenation and recompartmentalization during the treatment break. (The decision to explicitly enforce a treatment break was made to eliminate small numerical artifacts, which do not contribute significantly to the TCP. See Supplementary Section E.1 for details).

We next examined how the optimal primer shot schedule is impacted by changes in the upper bound on the normal tissue BED_3 ($B^{\max}$). As illustrated in Fig. 2, increasing $B^{\max}$ allowed fraction sizes to increase across the board, with the greatest gains occurring in the primer shot and final dose. The larger fractions induced cells to reoxygenate and recompartmentalize faster, thus producing a higher TCP. However, across our computational experiments, we found there were diminishing returns to increasing $B^{\max}$ beyond a certain point due to hard limits on the daily dose and required time for reoxygenation, which reflect clinical experience.

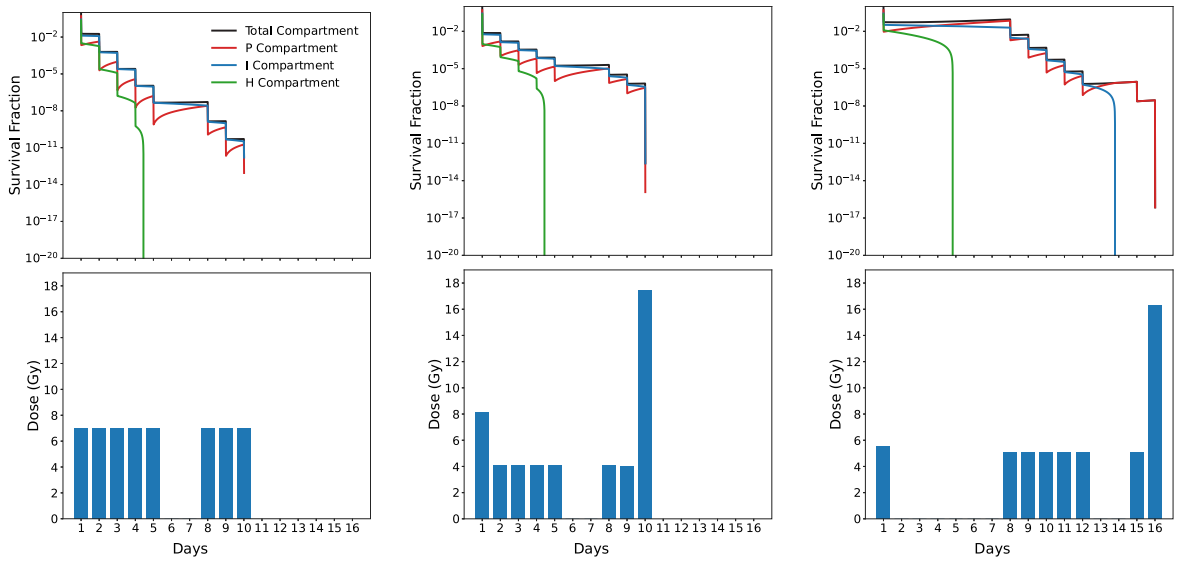

(a) Constant Fractions (TCP = 0.83) (b) Variable Fractions (TCP = 0.87) (c) Variable Fracs + Break (TCP = 0.93)

Fig. 1. Comparison of tumor SF for optimal 8-fraction weekday schedules with (a) constant fraction size, (b) variable fraction size, and (c) variable fraction size with a 1 week break after the initial primer shot dose, using a normal tissue BED_3 bound of $B^{max} = 210$ Gy. The primer shot schedule achieved the highest TCP of 0.93.

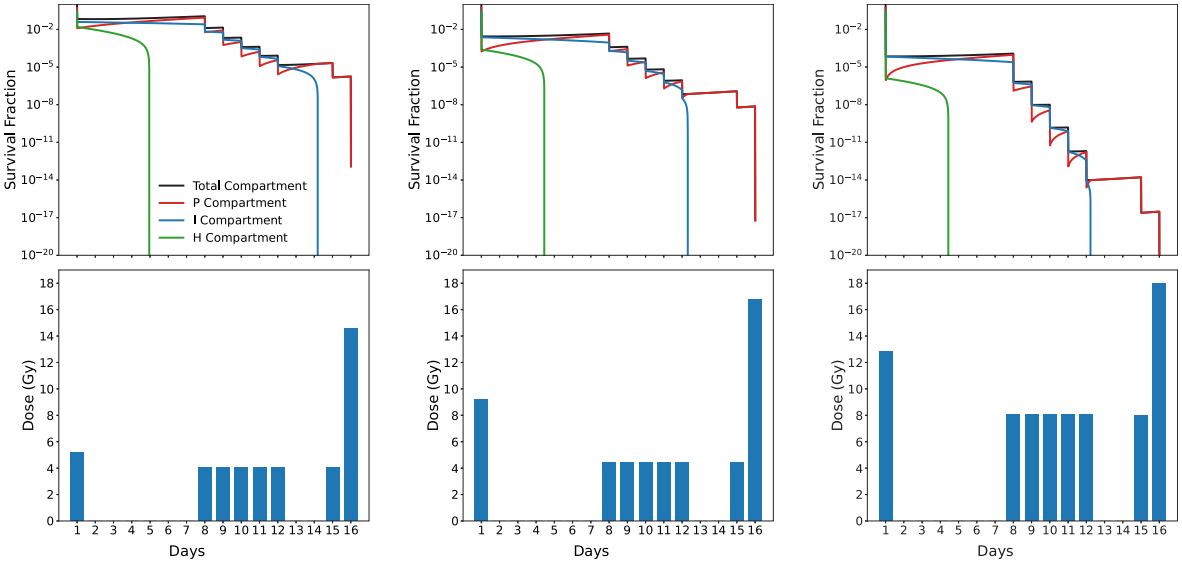

(a) $B^{max} = 160$ (TCP = 0.88) (b) $B^{max} = 216$ (TCP = 0.94) (c) $B^{max} = 378$ (TCP = 0.95)

Fig. 2. Comparison of tumor SF for the 8-fraction/16-day primer shot schedule optimized with a normal tissue BED_3 bound of (a) $B^{max} = 160$, (b) $B^{max} = 216$, and (c) $B^{max} = 378$ Gy.

In Fig. 3, we show how the primer shot schedule varied with the number of fractions and treatment length (T) when B^{max} was fixed at 210 Gy. Generally, the final TCP increased with treatment length, rising to a maximum at $T = 19$ days, while the average fraction size, except on the first and last day of treatment, tended to decrease with T as the total dose was spread out over more fractions. The overall dose pattern nevertheless remained the same, with a primer shot that kicked off reoxygenation and a large final dose once recompartmentalization was complete.

4. Discussion

In this paper, we formulated the dose scheduling problem as an optimal control problem with cell dynamics derived from a model of tumor dose–response that captures the impact of hypoxia on radioresistance. We developed a solution algorithm using sequential mixed-integer convex programming and demonstrated its effectiveness in computational experiments with parameter values calibrated to HNSCC. Our results showed that the optimal dose schedule consists of an initial primer shot, a break to allow for reoxygenation, followed by a series of doses that target the remaining proliferating cells.

Prior research on dose scheduling either relied on simple dose–response models, which did not take into account the role of hypoxia [6,9], or employed numerical search methods to solve the associated nonconvex optimization problem [10,12,13]. In contrast, we formally defined the dose scheduling problem as an optimal control problem, explicitly modeling the impact of cellular dynamics on radioresistance, and leveraged the tools of convex programming to solve

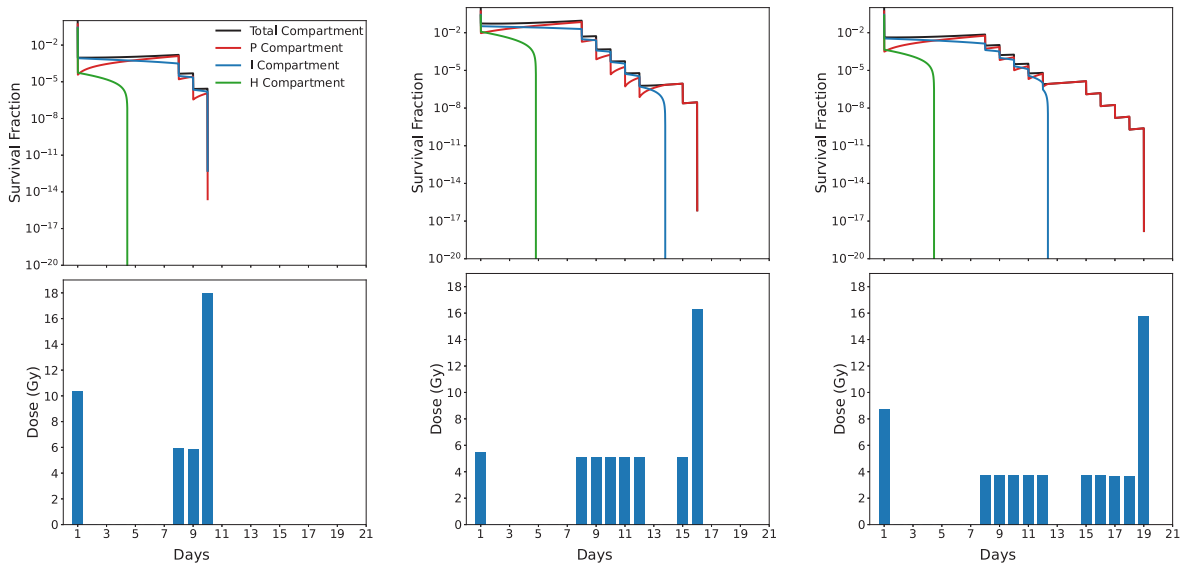

(a) $T = 10$, 4 Fractions (TCP = 0.85) (b) $T = 16$, 8 Fractions (TCP = 0.93) (c) $T = 19$, 11 Fractions (TCP = 0.94)

Fig. 3. Comparison of tumor SF for optimal primer shot schedules with (a) 4 fractions, (b) 8 fractions, and (c) 11 fractions, using a normal tissue BED_3 bound of $B^{\max} = 210$ Gy.

the problem. Our method differs from recent papers that adopted a deep learning approach [37–40] because it does not require training data, making it easy to adapt to other tumor environments with a recalibration of the parameters. Compared to the results in Gouw et al. [16], our typical dose schedule diverged in two important respects. First, small doses were distributed across the reoxygenation period rather than concentrated near the end of treatment. This is likely due to the presence of multiple solutions to the optimization problem, from which our algorithm chose a solution that satisfied the normal tissue BED_3 bound by the widest margin. (The quadratic term in the BED_3 constraint function $f_1(d)$ is lower for several small doses than one large dose). Second, the size of the final dose fraction in all our schedules was significantly larger than the size of all earlier fractions. This is because we allowed individual fraction sizes to vary in our optimization problem. A dose delivered at the end of treatment, when tumor cells are most radiosensitive, will yield the highest cell-kill rate. Thus, it makes sense that the majority of the dose "budget" was concentrated in the final day.

One concern in clinical implementation of our dose schedule is the potential for accelerated repopulation of tumor cells during the treatment break [41]. In our model, the growth factor and initial proliferation fraction (f^P_{pro}) are fixed, but these parameter values may be updated over the course of treatment and the schedule reoptimized via a procedure called model predictive control (MPC) [42,43]. With MPC, we solve for the optimal dose schedule, but only deliver the dose assigned to the current treatment day. Then, on the patient's next visit, we reassess the tumor's status, update the parameters to reflect any changes, and repeat the process. In this way, MPC allows us to quickly respond to adverse developments like accelerated repopulation without needing to explicitly model the effect in our problem.

Algorithm 1 solves an *approximation* of the optimal dose scheduling problem (9). Due to the problem's nonconvexity, other dose schedules may exist that perform as well or better than the solution output by our algorithm. We examined this possibility in Supplementary Section E.2 and discovered that it is possible for a constant fractionation schedule to achieve a TCP close to an algorithmically generated schedule while maintaining the same normal tissue BED_3. However, to find these schedules, we needed to perform a brute force search of all potential dose schedules – a feat that is computationally intractable beyond very short treatment lengths. Moreover, it is difficult for such a brute force approach to take into account complex clinical constraints, such as patient dose limits on specific days. Algorithm 1, on the other hand, can handle such constraints flexibly and automatically using numerical optimization.

The numerical results presented in Section 3 are quantitatively specific to the HNSCC parameter values used in our optimization problem. To apply our algorithm to other disease sites, recalibration of the parameters is required, followed by site-specific validation. Our computational experiments were limited to 1–3 week long treatment schedules due to the size and complexity of the optimization problem. Specifically, the number of time steps T coupled with the binary representation of recompartmentalization resulted in a large MIP, which had to be solved in each iteration of Algorithm 1. To scale Algorithm 1 up, there are two possible approaches. The first is to improve the computational performance of the solver using parallelization, GPU acceleration, and other hardware-based techniques. The second is to restrict the search space of Prob. (12). For example, we could assume that in the optimal dose schedule, reoxygenation/recompartmentalization takes place over a single time interval and optimize with respect to its start/end points. This would result in fewer binary variables, and thus a smaller MICP.

When constructing our mathematical model, our goal was to move beyond the conventional LQ description of dose–response, while ensuring the optimization problem was computationally tractable and solvable. Thus, we limited ourselves to a spatially averaged microenvironment with constant nutrient influx and considered only a single OOI with fixed $\alpha_N/\beta_N = 3$. Other models exist that capture interdependencies in a spatial tumor microenvironment. For example, Harriss-Phillips et al. [44] proposed a stochastic, spatially-resolved model based on the therapeutic ratio and found that early high-dose fractions followed by lower doses were optimal. In contrast, our model has a different structure — it is built on radiobiological principles, but with parameters fitted to clinical data [15], allowing us to bridge the gap between *in silico* predictions and actual patient cohorts. In the future, we plan to extend our model to the voxel level via baseline imaging, which will allow us to explicitly incorporate spatial heterogeneity and vascular dynamics, and add support for joint TCP/normal tissue complication probability-based optimization, as well as multiple OOIs with different α/β ratios and dose-volume effects.

In summary, we have developed an optimization-based method that efficiently generates the best radiotherapy dose schedule with support for variable fraction sizes and treatment breaks. Our typical schedule achieved a higher TCP than comparable equi-dose schedules by delivering an initial primer shot, then concentrating dose after reoxygenation had completed, when tumor cells were most radiosensitive.

CRediT authorship contribution statement

Anqi Fu: Writing – review & editing, Writing – original draft, Visualization, Software, Methodology, Investigation, Formal analysis, Conceptualization. **Zeno Gouw:** Writing – review & editing, Validation, Methodology, Investigation, Formal analysis. **Jeho Jeong:** Writing – review & editing, Validation, Methodology, Investigation, Formal analysis. **Joseph O. Deasy:** Writing – review & editing, Supervision, Investigation, Funding acquisition, Conceptualization.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

This research was funded in part by the the US National Institutes of Health (NIH)/National Cancer Institute (NCI) support grant P30 CA008748.

Appendix A. Supplementary data

Supplementary material related to this article can be found online at https://doi.org/10.1016/j.phro.2026.101003.

References

[1] Ahmed KA, Correa CR, Dilling TJ, Rao NG, Shridhar R, Trotti AM, Wilder RB, Caudell JJ. Altered fractionation schedules in radiation treatment: A review. Semin Oncol 2014;41:730–50. http://dx.doi.org/10.1053/j.seminoncol.2014.09.012.

[2] Löck S, Perrin R, Seidlitz A, Bandurska-Luque A, Zschaeck S, Zöphel K, Krause M, Steinbach J, Kotzerke J, Zips D, Troost EGC, Baumann M. Residual tumour hypoxia in head-and-neck cancer patients undergoing primary radiochemotherapy, final results of a prospective trial on repeat FMISO-PET imaging. Radiother Oncol 2017;124:533–40. http://dx.doi.org/10.1016/j.radonc.2017.08.010.

[3] Beckers C, Pruschy M, Vetrugno I. Tumor hypoxia and radiotherapy: A major driver of resistance even for novel radiotherapy modalities. Semin Cancer Biol 2024;98:19–30. http://dx.doi.org/10.1016/j.semcancer.2023.11.006.

[4] Suzuki Y, Nakano T, Ohno T, Niibe Y, Morita S, Tsujii H. Oxygenated and reoxygenated tumors show better local control in radiation therapy for cervical cancer. Int J Gynecol Cancer 2006;16:306–11. http://dx.doi.org/10.1111/j.1525-1438.2006.00341.x.

[5] Zips D, Zöphel K, Abolmaali N, Perrin R, Abramyuk A, Haase R, Appold S, Steinbach J, Kotzerke J, Baumann M. Exploratory prospective trial of hypoxia-specific PET imaging during radiochemotherapy in patients with locally advanced head-and-neck cancer. Radiother Oncol 2012;105:21–8. http://dx.doi.org/10.1016/j.radonc.2012.08.019.

[6] Fowler JF. The linear-quadratic formula and progress in fractionated radiotherapy. Brit. J Radiol 1989;62:679–94. http://dx.doi.org/10.1259/0007-1285-62-740-679.

[7] McMahon SJ. The linear quadratic model: Usage, interpretation and challenges. Phys Med Biol 2019;64:01TR01. http://dx.doi.org/10.1088/1361-6560/aaf26a.

[8] Cohen L. Cell population kinetics in radiation therapy: Optimization of tumor dosage. Cancer 1973;32:236–44. http://dx.doi.org/10.1002/1097-0142(197307)32:1<236::AID-CNCR2820320135>3.0.CO;2-V.

[9] Cohen L. Optimization of dose-time factors for a tumor and multiple associated normal tissues. Int J Radiat Oncol Biol Phys 1987;13:251–8. http://dx.doi.org/10.1016/0360-3016(87)90135-0.

[10] Leder K, Pitter K, LaPlant Q, Hambardzumyan D, Ross BD, Chan TA, Holland EC, Michor F. Mathematical modeling of PDGF-driven glioblastoma reveals optimized radiation dosing schedules. Cell 2014;156:603–16. http://dx.doi.org/10.1016/j.cell.2013.12.029.

[11] Jalalimanesh A, Haghighi HS, Ahmadi A, Soltani M. Simulation-based optimization of radiotherapy: Agent-based modeling and reinforcement learning. Math Comput Simulation 2017;133:235–48. http://dx.doi.org/10.1016/j.matcom.2016.05.008.

[12] Kuznetsov M, Kolobov A. Investigation of the effects of angiogenesis on tumor growth using a mathematical model. Biophysics 2015;60:449–56. http://dx.doi.org/10.1134/S0006350915030082.

[13] Kuznetsov M, Kolobov A. Optimization of dose fractionation for radiotherapy of a solid tumor with account of oxygen effect and proliferative heterogeneity. Mathematics 2020;8:1204. http://dx.doi.org/10.3390/math8081204.

[14] Jeong J, Shoghi KI, Deasy JO. Modelling the interplay between hypoxia and proliferation in radiotherapy tumour response. Phys Med Biol 2013;58:4897. http://dx.doi.org/10.1088/0031-9155/58/14/4897.

[15] Jeong J, Oh JH, Sonke J-J, Belderbos J, Bradley JD, Fontanella AN, Rao SS, Deasy JO. Modeling the cellular response of lung cancer to radiation therapy for a broad range of fractionation schedules. Clin Cancer Res 2017;23:5469—-79. http://dx.doi.org/10.1158/1078-0432.ccr-16-3277.

[16] Gouw ZAR, Jeong J, Rimner A, Lee NY, Jackson A, Fu A, Sonke J-J, Deasy JO. "Primer shot" fractionation with an early treatment break is theoretically superior to consecutive weekday fractionation schemes for early-stage non-small cell lung cancer. Radiother Oncol 2024;190:110006. http://dx.doi.org/10.1016/j.radonc.2023.110006.

[17] Kirk DE. Optimal Control Theory: An Introduction. Dover Publications; 2004.

[18] Cauligi A, Culbertson P, Stellato B, Bertsimas D, Schwager M, Pavone M. Learning mixed-integer convex optimization strategies for robot planning and control. Proc IEEE Conf Decis Control 2020;59:1698–705. http://dx.doi.org/10.1109/CDC42340.2020.9304043.

[19] Gupta OK, Ravindran A. Branch and bound experiments in convex nonlinear integer programming. Manag Sci 1985;31:1533–46. http://dx.doi.org/10.1287/mnsc.31.12.1533.

[20] Bonami P, Lee J, Leyffer S, Wächter A. On branching rules for convex mixed-integer nonlinear optimization. ACM J Exp Algorithmics 2013;18:2.1–31. http://dx.doi.org/10.1145/2532568.

[21] Duran M, Grossmann I. An outer-approximation algorithm for a class of mixed-integer nonlinear programs. Math Program 1986;36:307–39. http://dx.doi.org/10.1007/BF02592064.

[22] Hijazi H, Bonami P, Ouorou A. An outer-inner approximation for separable mixed-integer nonlinear programs. INFORMS J Comput 2014;26:31–44. http://dx.doi.org/10.1287/ijoc.1120.0545.

[23] Vielma JP, Dunning I, Huchette J, Lubin M. Extended formulations in mixed integer conic quadratic programming. Math Program Comput 2017;9:369–418. http://dx.doi.org/10.1007/s12532-016-0113-y.

[24] Lubin M, Yamangil E, Bent R, Vielma JP. Polyhedral approximation in mixed-integer convex optimization. Math Program 2018;172:139–68. http://dx.doi.org/10.1007/s10107-017-1191-y.

[25] Boyd S, Vandenberghe L. Convex Optimization. Cambridge University Press; 2004.

[26] Ben-Tal A, Nemirovski A. Lectures on Modern Convex Optimization: Analysis, Algorithms, and Engineering Applications. Society for Industrial and Applied Mathematics; 2001.

[27] Hunter DR, Lange K. A tutorial on MM algorithms. Am Stat 2004;58:30–7. http://dx.doi.org/10.1198/0003130042836.

[28] Sun Y, Babu P, Palomar DP. Majorization-minimization algorithms in signal processing, communications, and machine learning. IEEE Trans Signal Process 2017;65:794–816. http://dx.doi.org/10.1109/TSP.2016.2601299.

[29] Yuille AL, Rangarajan A. The concave-convex procedure (CCCP). Neural Comput 2003;15:915–36, https://dl.acm.org/doi/10.5555/2980539.2980672.

[30] Sriperumbudur BK, Lanckriet GR. On the convergence of the concave-convex procedure. NeurIPS 2009;23:1759–67, https://dl.acm.org/doi/10.5555/2984093.2984290.

[31] Lipp T, Boyd S. Variations and extension of the convex-concave procedure. Optim Eng 2016;17:263–87. http://dx.doi.org/10.1007/s11081-015-9294-x.

[32] Shen X, Diamond S, Gu Y, Boyd S. Disciplined convex-concave programming. Proc IEEE Conf Decis Control 2016;55:1009–14. http://dx.doi.org/10.1109/CDC.2016.7798400.

[33] Fu A, Xing L, Boyd S. Operator splitting for adaptive radiation therapy with nonlinear health dynamics. Optim Methods Softw 2022;27:2300–23. http://dx.doi.org/10.1080/10556788.2022.2078824.

[34] Diamond S, Boyd S. CVXPY: A Python-embedded modeling language for convex optimization. J Mach Learn Res 2016;17:1–5, https://dl.acm.org/doi/10.5555/2946645.3007036.

[35] Andersen ED, Andersen KD. The MOSEK interior point optimizer for linear programming: An implementation of the homogeneous algorithm. In: High performance optimization. Springer-Verlag; 2000, p. 197–232.

[36] ApS M. The MOSEK optimizer API for python manual. Version 10.0.. 2023, Available from: http://docs.mosek.com/10.0/pythonapi/index.html.

[37] Archambault Y, Boylan C, Bullock D, Morgas T, Peltola J, Ruokokoski E, Genghi A, Haas B, Suhonen P, Thompson S. Making on-line adaptive radiotherapy possible using artificial intelligence and machine learning for efficient daily re-planning. Med Phys Int 2020;8:77–86, http://www.mpijournal.org/pdf/2020-02/MPI-2020-02-p077.pdf.

[38] Lee D, Hua Y-C, Kuo L, Alam S, Yorke E, Li A, Rimner A, Zhang P. Deep learning driven predictive treatment planning for adaptive radiotherapy of lung cancer. Radiother Oncol 2022;169:57–63. http://dx.doi.org/10.1016/j.radonc.2022.02.013.

[39] Hooshangnejad H, Chen Q, Feng X, Zhang R, Farjam R, Voong KR, Hales RK, Du Y, Jia X, Ding K. DAART: A deep learning platform for deeply accelerated adaptive radiation therapy for lung cancer. Front Oncol. 2023;13:1201679. http://dx.doi.org/10.3389/fonc.2023.1201679.

[40] Maniscalco A, Liang X, Lin M-H, Jiang S, Nguyen D. Single patient learning for adaptive radiotherapy dose prediction. Med Phys 2023;50:7324–37. http://dx.doi.org/10.1002/mp.16799.

[41] Bourhis J, Overgaard J, Audry H, Ang KK, Saunders M, Bernier J, Horiot J-C, Maître AL, Pajak TF, Poulsen MG, O'Sullivan B, Dobrowsky W, Hliniak A, Skladowski K, Hay JH, Pinto LHJ, Fallai C, Fu KK, Sylvester R, Pignon J-P. Hyperfractionated or accelerated radiotherapy in head and neck cancer: A meta-analysis. Lancet 2006;368:843–54. http://dx.doi.org/10.1016/S0140-6736(06)69121-6.

[42] Camacho EF, Alba CB. Model Predictive Control. Springer; 2004.

[43] Rawlings JB, Mayne DQ, Diehl MM. Model Predictive Control: Theory, Computation, and Design. Nob Hill Publishing; 2022.

[44] Harriss-Phillips WM, Bezak E, Potter A. Stochastic predictions of cell kill during stereotactic ablative radiation therapy: Do hypoxia and reoxygenation really matter? Int J Radiat Oncol Biol Phys 2016;95:1290–7. http://dx.doi.org/10.1016/j.ijrobp.2016.03.014.