

©Copyright 2015
Fateme Saberian

Convex and Dynamic Optimization with Learning for Adaptive
Biologically Conformal Radiotherapy

Fatemeh Saberian

A dissertation submitted in partial fulfillment
of the requirements for the degree of

Doctor of Philosophy

University of Washington

2015

Program Authorized to Offer Degree: Industrial and System Engineering

University of Washington

Abstract

Convex and Dynamic Optimization with Learning for Adaptive Biologically
Conformal Radiotherapy

Fatemeh Saberian

In external beam radiotherapy for cancer, ionizing radiation kills tumor cells but also damages nearby normal tissue. The goal therefore is to maximize tumor-control while keeping toxic effects of radiation on healthy anatomies within tolerable limits. This is attempted by following a two-pronged approach that includes a spatial and a temporal component. On the spatial side, the goal is to modulate the intensity profile of the radiation fields so as to deliver high radiation-dose to the tumor while limiting the dose to nearby healthy tissue. This is called fluence-map optimization. On the temporal side, the planned dose is split into multiple treatment sessions, called fractions, which are often administered daily over several weeks. Such fractionation gives the healthy tissue sufficient time to recover between sessions, as healthy cells are believed to possess better damage-repair capabilities than tumor cells. In addition, emerging advances in functional imaging are beginning to offer new opportunities to dynamically incorporate tumor's observed response into the treatment planning process; some experts believe that this could start a new era of spatiotemporally integrated treatment planning.

The research objective of this dissertation is to apply convex and dynamic optimization methods to establish a rigorous mathematical framework called Adaptive Biologically Conformal Radiotherapy (ABCRT) for spatiotemporally integrated radiotherapy planning. The standard log-linear-quadratic survival model is employed to model the tumor's and the normal tissue's dose-response throughout this dissertation.

As a first step toward the above objective, a spatiotemporally separated optimal fractionation problem is studied in Chapter 2. In this problem, a nominal fluence-map is assumed to

be available and the goal is to find the number of treatment sessions and the corresponding tolerable sequence of doses (obtained by scaling the given fluence-map) that maximize the biological effect on the tumor. This formulation incorporates any number and any combination of maximum dose, mean dose, and dose-volume type constraints for serial and parallel normal tissues. Furthermore, it does not a priori assume that the doses are invariant across fractions. Sufficient conditions for optimality of equal-dosage and single-dosage fractionation are provided. These conditions require the treatment planner to order the values of the tumor's and the normal tissue's dose-response parameters, and they yield a closed-form formula for the optimal dose when the number of fractions is fixed. This formula is then employed to prove that the biological effect on the tumor is a quasiconcave function of the number of fractions. This leads to a simple procedure for finding an optimal number of fractions when the aforementioned sufficient conditions are met. Extensive computer simulations and sensitivity analyses on ten head-and-neck and prostate cancer test cases are performed to gain clinically relevant insights.

Chapter 3 further investigates the spatiotemporally separated optimal fractionation problem for a fixed number of sessions, but this time, without restricting to any specific ordering of the dose-response parameters. This yields a nonconvex quadratically constrained quadratic programming (QCQP) formulation. Nonconvex QCQPs are typically difficult to solve, and in general, belong to the class NP-hard. We prove that an optimal solution to the nonconvex QCQP formulation of the fractionation problem can be derived in closed-form by instead solving a two-variable linear program with a few constraints.

In Chapter 4, a spatiotemporally integrated formulation of the optimal fractionation problem, where the fluence-map and the number of fractions are optimized simultaneously, is considered. An efficient convex programming method is proposed to approximately solve the resulting large-scale model. Through extensive computer simulations on ten head-and-neck and prostate cancer test cases with a broad range of radiobiological parameters, we compare the biological effect on tumor obtained by this integrated approach relative to that from two other models. The first is a traditional IMRT fluence-map optimization model that

does not optimize the number of fractions. The second (from Chapter 2) assumes that a fluence-map is available a priori from a traditional fluence-map optimization model and then optimizes the number of fractions, thus separating the spatial and temporal components. The results of these simulations suggest that integrated optimization of the fluence-map and the number of fractions could improve tumor-control.

In Chapter 5, we start to focus on developing fluence-maps that adapt to the uncertain, spatiotemporal evolution of various key determinants of treatment efficacy. Specifically, we propose a theoretical stochastic control framework for adapting radiotherapy to hypoxia. Hypoxia, that is, an inadequate supply of oxygen in living tissue, has long been known to adversely affect the outcome of radiotherapy in solid tumors. The aim of Chapter 5 therefore is to estimate any potential benefits of such adaptive planning. This requires the solution of computationally challenging, high-dimensional stochastic control problems. We propose an approximate control scheme that involves solving a sequence of large-scale, convex, non-linear problems. We present a simulation-based comparison of this method against static treatment plans on five head-and-neck test cases. These simulations suggest that adaptive planning could improve tumor-control. Through these simulations, we gain insights into when and why adaptive planning is likely to yield the largest benefits.

In Chapter 6, we present a framework based on Bayesian dynamic programming to learn a patient-specific distribution of tumor-response uncertainty over the treatment course while simultaneously optimizing fluence-maps. Specifically, we assume that a particular tumor-response parameter has a Categorical distribution and that the decision maker uses a Dirichlet prior on this Categorical distribution. We then exploit the conjugate property of the Categorical-Dirichlet pair to formulate a Bayesian dynamic programming formulation of our optimal learning problem. We then extend the approximate control scheme from Chapter 5 to tackle this problem. We show via computer simulations on one head-and-neck test case that we are able to efficiently learn the tumor-response parameter distribution over the treatment course. Opportunities for future research are outlined in Chapter 7.

TABLE OF CONTENTS

	Page
List of Figures	iv
List of Tables	vi
Chapter 1: Introduction	1
1.1 Research contributions and thesis organization	4
Chapter 2: Optimal fractionation in radiotherapy with multiple normal tissues	8
2.1 Introduction	8
2.2 Limitations of existing work	9
2.3 Our contributions	12
2.4 Problem formulation	16
2.4.1 Expression for the objective function	16
2.4.2 Normal tissue constraints	18
Serial normal tissue	19
Mean dose constraints for parallel normal tissue	20
Dose-volume constraints	21
A formulation of the optimal fractionation problem	22
2.5 Sufficient conditions for optimality of equal-dosage fractionation	24
2.6 Optimal equal-dosage fractionation	35
2.6.1 Characterization of the limiting dose and the optimal tumor-BE as functions of the number of fractions	37
2.7 The special case of a single normal tissue	44
2.8 Numerical results	47
2.8.1 Description of test cases	48
Head-and-neck cancer cases	48
Prostate cancer cases	50
2.8.2 Switching of the limiting normal tissues and nonmonotonic behavior of the tumor-BE without proliferation	52

2.8.3	Comparison with results obtained by separately considering each normal tissue one-by-one	53
2.8.4	Sensitivity to tumor doubling time	56
2.8.5	Sensitivity to time-lag before proliferation begins	58
2.9	Discussion	60
Chapter 3:	A two-variable linear program to solve an optimal fractionation problem	66
Chapter 4:	Spatiotemporally optimal fractionation	78
4.1	Introduction	78
4.2	Problem formulation	81
4.2.1	Expression for the tumor objective function	81
4.2.2	Normal tissue tolerance and fluence-map smoothness constraints	84
	Maximum BED constraints for serial normal tissues	84
	Mean BED constraints for parallel normal tissues	85
	Dose-volume constraints for parallel normal tissues	86
	Fluence-map smoothness constraints	87
4.2.3	Complete optimization model	87
4.3	An efficient solution method	89
4.4	Results	94
4.4.1	Description of test cases	95
	Head-and-neck cancer cases	95
	Prostate cancer cases	97
4.4.2	Sensitivity to tumor doubling time and to time-lag before proliferation	98
4.4.3	Improvement in tumor-BE	103
	Comparison with conventional IMRT	103
	Comparison with a spatiotemporally separated model	107
4.5	Discussion	108
Chapter 5:	A theoretical stochastic control framework for adapting radiotherapy to hypoxia	111
5.1	Introduction	111
5.2	A theoretical stochastic control formalism	112
5.2.1	Tumor state dynamics	113
5.2.2	Control variables and normal tissue dynamics	115
	Maximum BED constraints for serial normal tissues	115

Mean BED constraints for parallel normal tissues	116
Dose-volume constraints for parallel normal tissues	117
Fluence-map smoothness constraints	117
Set of feasible fluence-map policies	118
5.2.3 Bellman’s equations of dynamic programming	118
5.3 Solution method: Certainty Equivalent Control	121
5.3.1 Efficient solution of problem (P_t)	124
Constraint generation method	127
Log-barrier interior point algorithm for solving (\hat{Q}_t) and (\bar{Q}_t)	128
5.3.2 A streamlined version of CEC	131
5.4 Results	133
5.5 Conclusions and discussion	141
Chapter 6: Optimal learning of tumor response	148
6.1 Introduction	148
6.2 Numerical results	151
6.3 Conclusions	152
Chapter 7: Future work	155
Appendix A: PhanC: A Matlab software for creating phantom test cases on a computer	157
A.1 Input data	158
A.2 Method	160
A.3 Validation	163

LIST OF FIGURES

Figure Number	Page
2.1 A schematic illustration of the structure of the limiting dose $d^*(N)$	38
2.2 Switching of the limiting normal tissues and nonmonotonic behavior of the tumor-BE without proliferation for head-and-neck case 1	54
2.3 Switching of the limiting normal tissues and nonmonotonic behavior of the tumor-BE without proliferation for prostate case 2	55
2.4 Tumor-BE with proliferation obtained by considering different normal tissues one by one for head-and-neck case 1	56
2.5 Sensitivity to tumor doubling time	57
2.6 Sensitivity of the optimal number of fractions to T_{double} as a function of the relative difference between the α/β ratios for the tumor, and for the parotids	59
3.1 A geometric proof of the equivalence between two variable LP and problem (3.37)-(3.42)	73
4.1 Sensitivity of the optimal number of fractions and the optimal tumor-BE to tumor doubling time T_{double} (days) when $T_{lag} = 7$ days. (a) Head-and-neck case 1; tumor α/β was 10 Gy, and the α/β ratio for all normal tissues was 3 Gy. (b) Prostate case 1; tumor α/β was 6 Gy, and the α/β ratio for all normal tissues was 3 Gy.	100
4.2 The average percentage improvement achieved by our spatiotemporally integrated approach over: (a) IMRT for head-and-neck, (b) IMRT for prostate, (c) spatiotemporally separated model in Chapter 2 for head-and-neck, and (d) spatiotemporally separated model in Chapter 2. These numbers are for $T_{lag} = 7$ days and averaged over all combinations of tumor and normal tissue α/β ratios.	109
4.3 The average percentage improvement achieved by our spatiotemporally integrated approach over IMRT for: (a) head-and-neck case 1, (b) prostate case 1. These numbers are for $T_{lag} = 7$ days with all normal tissue α/β ratios fixed at 3 for different tumor α/β ratios as shown in the legends.	110
5.1 Improvement in TNTCR obtained by CEC using hypoxia images over the static approach for: (a) exponential covariance function, (b) rational-quadratic covariance function.	137

5.2	Improvement in TNTCR obtained by CEC with cell density images over the static approach for: (a) exponential covariance function, (b) rational-quadratic covariance function.	138
5.3	The average (over all tumor voxels) dose (Gy) per session delivered by CEC with hypoxia images in each one of the 30 simulations over 35 treatment sessions in case 1: (a) exponential covariance function (b) rational-quadratic covariance function. The solid flat line shows the constant dose (Gy) delivered by the static method in every session.	140
5.4	The average (over all tumor voxels) dose (Gy) per session delivered by CEC with cell density images in each one of the 30 simulations over 35 treatment sessions in case 1: (a) exponential covariance function (b) rational-quadratic covariance function. The solid flat line shows the constant dose (Gy) delivered by the static method in every session.	141
5.5	The dose (Gy) delivered by CEC with hypoxia images and by the static method in various treatment sessions for case 3 with the rational-quadratic covariance function.	143
5.6	The dose (Gy) delivered by CEC with cell density images and by the static method in various treatment sessions for case 2 with the exponential covariance function.	144
5.7	A histogram of oxygen partial pressure in different tumor voxels for various treatment sessions for case 2 with the exponential covariance function.	145
5.8	A histogram of tumor alpha values (Gy^{-1}) over different tumor voxels for various treatment sessions for case 2 with the exponential covariance function.	146
5.9	A histogram of tumor beta values (Gy^{-2}) over different tumor voxels for various treatment sessions for case 2 with the exponential covariance function.	147
6.1	Three different probability vectors for tumor dose response parameter	152
6.2	Learning the probability vector over 35 treatment sessions for normal and beta distributions	153
6.3	Percentage difference between the TNTCRs for 10 simulations for normal and beta distributions	154
A.1	PhanC Graphic User Interface (GUI)	159
A.2	Geometry to calculate the dose deposition coefficients, A_{ij} , in matrix A	160
A.3	Definition of (a) interior and exterior beamlets (b) primary beamlets P_i for voxel i (c) x-axis in profiles	162
A.4	Measured penumbra from 10cm x 10cm field at various depths	163

LIST OF TABLES

Table Number	Page
2.1 Description of the geometry used in head-and-neck cancer cases.	49
2.2 Tolerance doses for normal tissues in head-and-neck test cases	49
2.3 Description of the geometry used in prostate cancer cases.	51
2.4 Dose-volume constraints for various normal tissues in prostate cancer cases .	52
2.5 Comparison with results obtained by separately considering each normal tissue one- by-one	63
2.6 Sensitivity to tumor doubling time	64
2.7 Sensitivity to time-lag before proliferation begins	65
4.1 A summary of some optimal fractionation models that use the LQ framework.	80
4.2 Description of the geometry used in head-and-neck cancer cases. This data is identical to Chapter 2.	96
4.3 Tolerance doses for various normal tissues in our head-and-neck test cases where the dose is administered in $N_{\text{conv}} = 35$ equal-dose fractions. Recall that for dose-volume type constraints no more than a volume fraction ϕ of the normal tissue can receive dose more than D_{dv} . This treatment protocol is similar, for example, to [1].	96
4.4 Values of various radiobiological parameters used in sensitivity analyses for head-and-neck cancer.	97
4.5 Description of the geometry used in prostate cancer cases. This data is identical to Chapter 2.	97
4.6 Dose-volume constraints for various normal tissues when dose is administered in $N_{\text{conv}} = 45$ equal-dose fractions — no more than a volume fraction ϕ of the normal tissue can receive dose more than D_{dv} . This treatment protocol is similar, for example, to [1].	98
4.7 Values of various radiobiological parameters used in sensitivity analyses for prostate cancer.	99
4.8 Range of N_{99}^* across five head-and-neck cases with $T_{\text{lag}} = 7$ days and all combinations of tumor and normal tissue α/β ratios.	101
4.9 Range of N_{99}^* for head-and-neck case 1 with $T_{\text{double}} = 10$ days, $T_{\text{lag}} = 7$ days, and different combinations of tumor and parotid α/β ratios with α/β for unspecified normal tissue, spinal cord, and brainstem fixed at 3 Gy.	102

4.10	Range of N_{99}^* for five prostate cases with $T_{lag} = 7$ days and $\alpha/\beta = 4, 6$ for tumor and all combinations of normal tissue α/β ratios.	102
4.11	Range of N_{99}^* for prostate case 1 with $T_{double} = 30$ days, $T_{lag} = 7$ days, and different combinations of tumor, rectum and bladder α/β ratios with α/β for unspecified normal tissue and femurs fixed at 3 Gy.	103
4.12	Range of N_{99}^* over five head-and-neck cases over all combinations of tumor and normal tissue α/β ratios.	104
4.13	Range of N_{99}^* over five prostate cases for $\alpha/\beta = 4, 6$ for tumor over all combinations of normal tissue α/β ratios.	104
5.1	Description of the geometry used in head-and-neck cancer cases.	134
5.2	Tolerance doses for various normal tissues in our head-and-neck test cases where the dose is administered in $N_{conv} = 35$ equal-dose fractions. Recall that for dose-volume type constraints no more than a volume fraction ϕ of the normal tissue can receive dose more than D_{dv}	135
5.3	Improvement in average TNTCR obtained by CEC over the static approach for exponential and rational-quadratic covariance functions.	137
A.1	OF is the ratio of the dose from an arbitrary field size to the dose from 10cm x 10cm field size at d_{max}	165

ACKNOWLEDGMENT

I am heartily thankful to my adviser, professor Archis Ghate, whose encouragement, supervision and support from the preliminary to the concluding level helped me in developing and understanding of this subject. I would also like to thank professor Minsun Kim, a member of my reading committee, for her generous support throughout my Ph.D. studies. Last but not the least, I want to thank my parents, family and friends for their never ending moral support.

DEDICATION

To my family and friends
who helped to make this dream come true

Chapter 1

INTRODUCTION

The 2014 World Cancer Report [2] estimated that by 2030 there could be 17 million cancer-deaths annually worldwide. There are different types of cancer treatment such as surgery, chemotherapy, and radiotherapy. Surgery is used to remove tumors, or parts of tumors, and surrounding normal tissue during an operation. Surgery is often followed by other therapies to make sure all cancerous cells have been removed. Chemotherapy uses chemical substances to damage cancer cells. It uses one or more anti-cancer drugs to target all rapidly dividing cells which include tumor cells. Radiation therapy uses radiation to kill cancer cells by damaging their DNA permanently. Radiation targeted toward the cancerous regions however also passes through other nearby healthy anatomies, thus damaging both tumor and healthy cells. The goal therefore is to maximize tumor-damage while limiting toxic effects of radiation on nearby healthy cells.

According to the American Cancer Society, at least 50% of cancer patients undergo some type of radiation therapy as a part of their treatment [3]. Two common types of radiotherapy are external beam radiation therapy and brachytherapy. The main difference between these two types of radiotherapy is the location of radiation source. In brachytherapy, small radioactive seeds are placed into the patient's body close to cancerous region. Therefore, the delivered radiation is localized as much as possible. In external beam radiotherapy, a linear accelerator is used to pass ionizing radiation through the affected part of the patient's body in order to kill tumor cells. In this thesis, we focus on external beam radiotherapy.

Intensity modulated radiation therapy (IMRT) technology is used to deliver ex-

ternal beams of radiation. One popular way to deliver IMRT is to use a dynamic multi-leaf collimator (MLC), that is fitted in the aperture of the linear accelerator gantry. The MLC is comprised of dozens of leaf pairs that can stay stationary or move during the treatment. Individual leaves in the MLC work together to dynamically change the shape of the exposed beam. The position and speed of movements of MLC leaves can be controlled to create intensity modulation of the beam. Radiation fields from multiple directions are often used to avoid tumor cold-spots. These are regions where insufficient dose would otherwise be delivered in an attempt to spare the normal tissue. Typically, radiotherapy is administered daily over several weeks to give healthy cells time to recover between sessions, which is called fractionation. That is, the total amount of radiation that is to be delivered is split into smaller portions that are administered over the entire treatment horizon. It is partly motivated by the empirical observation that healthy cells have better damage repair ability than tumors. Therefore, by spreading out the treatment over many days, healthy cells have a chance to recover between treatment sessions.

The starting point in radiation therapy planning is an image of the internal anatomies including the cancerous targets, and nearby unspecified normal tissue and organ(s)-at-risk (OAR) (collectively, “normal tissue”). OARs are healthy anatomies, such as the spinal cord or heart, that are close to the cancerous regions and need to be spared as much as possible. The necessary geometrical information is obtained using an anatomical image such as computed tomography (CT) or a magnetic resonance imaging (MRI) scan. Once a CT or MRI scan is taken, a physician can contour the relevant tumor volume and delineate the critical organs close to the cancerous regions. The physician also specifies the set of constraints that need to be met such as a limit on the maximum or mean dose delivered to each OAR. The distribution of the radiation dose over the patient’s anatomy is then tuned to optimize some quantitative measure of treatment efficacy.

Optimization plays an important role in achieving an effective radiation treatment plan. There are different decision variables that can be optimized in order to deliver

a suitable treatment plan. The main variables, that are often considered, are the intensity of the radiation beams, and the angles from which to deliver the radiation. Several approaches have been developed for dealing with the beam angle optimization problem. Due to the nonconvexity of this problem [4], many of these approaches are based on heuristic methods [5]. After determining angles of radiation fields, the intensities of the beams are optimized. In this thesis, as is common in the literature, we will focus on the optimization of beam intensities, assuming that the beam directions are given.

The topic of beam intensity optimization in radiation therapy has received a lot of attention from the optimization community [6–8]. The references [9, 10] provide comprehensive reviews of the existing literature. Conventional radiation treatment plans decouple the spatial and temporal aspects of radiotherapy, and mostly ignore uncertainty in dynamic biological processes, and in particular, cannot adapt to the patient’s actual response over the treatment course. Repair of sublethal damage in normal tissue, and redistribution, repopulation, and reoxygenation of tumor cells, called the four Rs of radiobiology, determine efficacy of fractionated radiotherapy [11]. Mathematical modeling of such cell activity has traditionally been difficult owing to incomplete understanding of the underlying biochemical processes, variations in cells within a tumor, differences between types of tumors, and between individual patients. Therefore, most existing optimization approaches ignore this crucial temporal-biological, uncertain component of fractionated radiotherapy.

Recent technological advances in biological imaging are enabling researchers to track tumor-response over time. For example, positron emission tomography (PET) radio-labeled with fluorine-18 (^{18}F) and fluorodeoxyglucose (FDG) tracer, or magnetic resonance spectroscopic imaging (MRSI) of choline/citrate ratio can provide spatiotemporal information about tumor clonogen density [12–15]. PET with iodine-124-iododeoxyuridine can measure tumor proliferation rates. ^{18}F -MISO PET can map spatiotemporal evolution of hypoxia (lack of oxygen reduces radiosensitivity). At the same time, mathematical models of tumor-response are also becoming more

sophisticated [16–20]. Potential benefits of adapting treatment based on information obtained from biological images taken over several weeks of treatment, and also of using spatiotemporally integrated mathematical models of tumor-response, have recently been documented [14, 15, 21].

1.1 Research contributions and thesis organization

Our research objective in this thesis is to apply convex and dynamic optimization methods to establish a rigorous mathematical framework called Adaptive Biologically Conformal Radiotherapy (ABCRT) for spatiotemporally integrated radiation treatment planning. An initial foundation for this type of treatment planning was recently laid in [22–24]. We further build upon this foundation by making the following contributions.

In Chapter 2, we present a formulation of the optimal fractionation problem (the problem of determining the optimal number of fractions and the corresponding radiation dose per fraction) that includes multiple normal tissues. This formulation uses the standard log-linear-quadratic survival model. Our model can tackle any combination of maximum dose, mean dose, and dose-volume type constraints for serial and parallel normal tissues as this is characteristic of most treatment protocols. We also allow for a spatially heterogeneous dose distribution within each normal tissue. Furthermore, we do not a priori assume that the doses are invariant across fractions. We establish sufficient conditions under which equal-dosage fractionation is optimal. Finally, we perform extensive numerical experiments using ten head-and-neck and prostate test cases to uncover several clinically relevant insights.

In Chapter 3, we develop a two-variable linear program that can solve the nonconvex QCQP fractionation problem from Chapter 2 to optimality. Nonconvex QCQPs are typically difficult to solve, and in general, they belong to the class NP-hard. Existing mathematical literature on the fractionation problem therefore uses heuristic methods without any analyses about solution quality, or derives exact solutions in stylized special case. However, we show that an optimal solution to this nonconvex

QCQP can in fact be derived in closed-form. A particularly appealing feature of this two-variable LP model is that the number of constraints is equal to the number of OARs which is typically a small number.

In Chapter 4, we present a spatiotemporally integrated formulation of the optimal fractionation problem. Specifically, we simultaneously optimize the fluence-map and the number of fractions. Our objective is to maximize the total biological effect of tumor dose averaged over all voxels subject to maximum dose, mean dose, and dose-volume constraints for various normal tissues. We propose an efficient convex programming method to approximately solve this model. This method yields an exact solution when the treatment protocol does not include dose-volume constraints. Through extensive computer simulations on ten head-and-neck and prostate cancer test cases with a broad range of radiobiological parameters, we compare the biological effect on tumor obtained by our integrated approach relative to that from two other models. The first is a traditional IMRT fluence-map optimization model that does not optimize the number of fractions. The second (from Chapter 2) assumes that a fluence-map is available a priori from a traditional IMRT optimization model and then optimizes the number of fractions, thus separating the spatial and temporal components. The improvements in tumor biological effect over IMRT were 9%-52% with average 22%, and 53%-108% with average 69%, for head-and-neck and prostate, respectively. The improvements in tumor biological effect over the spatiotemporally separated model were 15%-45% with average 27%, and 17%-23% with average 21%, for head-and-neck and prostate, respectively. This suggests that integrated optimization of the fluence-map and the number of fractions could improve treatment efficacy as measured within the linear-quadratic framework.

In Chapter 5, we start to focus on developing fluence-maps that adapt to the uncertain, spatiotemporal evolution of various key determinants of treatment efficacy. Specifically, we propose a theoretical stochastic control framework for adapting radiotherapy to hypoxia. Hypoxia, that is, an inadequate supply of oxygen in living tissue, has long been known to adversely affect the outcome of radiotherapy in solid tumors

[15, 25, 26]. Moreover, oxygen partial pressure within tumors varies both spatially and temporally over the treatment course and there is a considerable inter-patient variation in intra-tumor oxygen partial pressure. The aim of Chapter 5 therefore is to estimate any potential benefits of such adaptive planning. This requires the solution of computationally challenging, high-dimensional stochastic control problems. We propose an approximate control scheme that involves solving a sequence of large-scale, convex, non-linear problems. We perform numerical experiments on five test cases by using a first-order vector autoregressive process with exponential and rational-quadratic covariance functions from the spatiotemporal statistics literature to simulate hypoxia evolution. Our results suggest that dynamic planning could lead to 13%-66% improvement in the number of tumor cells remaining at the end of the treatment course.

In Chapter 6, we present a framework based on Bayesian dynamic programming to learn a patient-specific distribution of tumor-response uncertainty over the treatment course while simultaneously optimizing fluence-maps. This is in contrast to Chapter 5 and the work in [22–24], where the distribution of tumor-response was assumed to be known a priori. Specifically, we assume that a particular tumor-response parameter has a Categorical distribution (note that this is the most general discrete distribution with finite support) and that the decision maker uses a Dirichlet prior on this Categorical distribution. We then exploit the conjugate property of the Categorical-Dirichlet pair to formulate a Bayesian dynamic programming formulation of our optimal learning problem. We then extend the approximate control scheme from Chapter 5 to efficiently tackle this problem. We show via computer simulations on one head-and-neck test case that we are able to efficiently learn the tumor-response parameter distribution over the treatment course. Our simulations suggest that the TNTCR attained by our learning approach is close to that obtained by a hypothetical, idealized, clairvoyant algorithm.

In Chapter 7, we outline directions for future research. These include a non-stationary formulation of the optimal fractionation problem from Chapter 4; a robust

formulation of the optimal fractionation problem from Chapter 3 and 4; and an optimal stopping extension of the optimal fractionation problem from Chapters 2, 3, and 4.

The test cases used for computer simulations reported in this thesis were generated using an in-house MATLAB software. This software is described in the Appendix.

Chapter 2

**OPTIMAL FRACTIONATION IN RADIOTHERAPY WITH
MULTIPLE NORMAL TISSUES****2.1 Introduction**

In external beam radiation therapy for cancer, ionizing radiation targeted toward the cancerous region also passes through nearby normal tissue and thus damages both. The goal therefore is to maximize tumor damage while ensuring that the healthy anatomies can safely tolerate the resulting toxic effects of radiation. In intensity modulated radiation therapy (IMRT) [27], this is attempted by following a two-pronged approach that includes a spatial and a temporal component.

The first, spatial component of this approach in essence is to prescribe a uniform or non-uniform high dose to the cancerous region and put upper limits on the dose delivered to the healthy anatomies. The intensity profile of the radiation field, which is also called the fluence-map, is then optimized so that the delivered dose is as close as possible to this tumor-conforming prescription. This is called fluence-map optimization. The fluence-map optimization problem has been studied extensively over the last two decades. Increasingly sophisticated dose prescriptions have been proposed and efficient optimization algorithms to find the corresponding intensity profiles have been designed, tested clinically, and incorporated into commercial treatment planning systems [6, 9, 28–31].

The second, temporal, component of the two-pronged approach involves breaking the total planned dose into several *equal-dose* fractions that are administered in multiple, well-separated treatment sessions (called fractions) spread over many weeks. This is called fractionation and it is designed mainly to give healthy anatomies sufficient time to recover between consecutive treatment sessions as healthy cells have better damage-repair capabilities than tumor cells [11, 32]. The key challenge on this

temporal side of the equal-dosage fractionation problem is to find the optimal number of fractions and the corresponding dose in different fractions. Roughly speaking, using several sessions with a small dose in each session reduces the damage done to the healthy anatomies. This may enable the treatment planner to administer a much larger total dose to the tumor as compared to using only a few sessions with a higher dose per session. However, especially in fast proliferating tumors, long treatment courses may adversely affect the likelihood of adequate tumor control. Several alternative number of fractions have been clinically studied over more than a hundred years [33–43]. Many of these studies were conducted before the IMRT technology became widely available in the late 1990s. It is generally believed that the optimal number of fractions depends on the relative difference between the radiation dose-response of the tumor and that of the healthy anatomies, and also on the anatomy of the tumor and the healthy tissues relative to the radiation field. Currently, the most widely accepted mathematical framework for radiation dose-response is the linear-quadratic (LQ) model [11]. Using this model, some researchers have mathematically explored the effect of model parameters on the optimal number of fractions. We review these studies next.

2.2 *Limitations of existing work*

To the best of our knowledge, the use of the LQ model to determine the optimal number of fractions started in the 1990s with Fowler’s work [44, 45] and the investigation in Jones et al. [46]. Specifically, Fowler showed in a series of papers using insightful calculations that some treatment courses are too short and some treatment courses are too long; there is a range in between these two extremes that should include the optimal number of treatment sessions [44, 45, 47–49]. These papers consider a tumor and a single normal tissue both receiving dose d in each of the N fractions. From an IMRT perspective, the assumption that the tumor and the normal tissue receive a homogeneous and identical dose essentially ignores the spatial side of the problem. The key idea that repeatedly occurs in all papers in this area is to choose the number

of fractions N and the corresponding dose d per fraction that is delivered to the tumor to maximize the biological effect (BE) on the tumor for a fixed biologically effective dose (BED) delivered to the normal tissue [11]. The papers by Fowler and coauthors were arithmetic in nature in that they calculated and plotted the BE on tumor as a function of N to reach their conclusions (without rigorously proving the structure of the tumor-BE as a function of N).

Jones et al. [46] also considered a single normal tissue and proposed a model that was conceptually similar to Fowler's. They employed calculus to maximize the tumor-BE by setting its derivative with respect to N to equal zero. This led to a quadratic equation for the optimal dose per fraction. The positive solution of this equation was then used in turn to derive a formula for the optimal number of fractions. Thus, to the best of our knowledge, Jones et al. obtained the first analytical results in this area. However, Jones et al. did not include a time-lag before tumor proliferation begins after treatment initiation (see [49] for a discussion of the importance of including this time-lag). Jones and co-authors later improved their earlier model by including such a time-lag [50]. However, these two papers by Jones and co-authors ignored the constraint that N must be an integer and did not rigorously prove that setting the derivative with respect to N to zero was a sufficient condition for global optimality in their model (that is, they did not investigate concavity properties of tumor-BE by computing its second derivative).

Yang and Xing [51] also studied a model similar to Fowler's but they utilized an extension of the basic LQ model to incorporate redistribution and reoxygenation of tumor cells [52] and included two normal tissues. They applied the simulated annealing heuristic to optimize the number of fractions without any guarantee of optimality. Their broad conclusions were similar to Fowler's. Bertuzzi et al. [53] recently considered a model with two normal tissues and used Karush-Kuhn-Tucker (KKT) conditions to characterize the structure of optimal solutions, but the number of fractions was fixed in their model.

Mizuta et al. [54] recently considered a slight generalization of Fowler's calcula-

tions, where the tumor and the single normal tissue received doses d and sd in each fraction, respectively, but they did not consider tumor proliferation. Here, s is called the sparing factor and it attempts to model the realistic possibility that the tumor and the normal tissue do not receive identical doses. They reached the intuitive conclusion that the optimal number of fractions also depends on the sparing factor s in addition to the α/β ratios of the tumor and the normal tissue in the LQ model. In fact, a similar observation was made much earlier by Jones et al. [46].

Unkelbach et al. [55] recently generalized the result in Mizuta et al. They also considered a single normal tissue but explicitly allowed for a spatially heterogeneous normal tissue dose distribution. Consequently, the analysis in Unkelbach et al. is a notable exception to all previous efforts toward optimizing the number of fractions. In particular, they used voxel-dependent sparing factors for each normal tissue voxel. The use of such sparing factors is justified because optimization models in radiotherapy routinely model dose as a linear transformation of the fluence-map. See [56–60] and Section 2.4 for more details about how this transformation is defined using a dose-deposition coefficient matrix. Unkelbach et al. derived an effective sparing factor, $s_{effective}$, which was a function only of the voxel-dependent sparing factors. The specific functional form of this effective sparing factor depends on whether the normal tissue constraint is a maximum dose constraint or a mean dose constraint. Their key conclusion was that a single fraction is optimal if $(\alpha/\beta)_{tumor} \leq \frac{(\alpha/\beta)_{normaltissue}}{s_{effective}}$ and an infinite number of fractions is optimal otherwise. A similar conclusion was also reached independently by Keller et al. in [61]. However, and crucially, Unkelbach et al. and Keller et al. did not incorporate tumor proliferation and this was the reason for the optimality of an infinite number of fractions.

Bortfeld et al. [62] extended the model in Unkelbach et al. by including tumor proliferation, but they did not include a time-delay before proliferation starts. Specifically, their focus was on modeling accelerated tumor re-population whereby the tumor exhibits faster growth toward the end of the treatment course. They used dynamic programming to demonstrate how dose per fraction increases over the treat-

ment course under accelerated re-population; this was the main contribution of their paper. As an aside, they rigorously proved the correctness of a previously known closed-form formula (as for example in Jones et al. [46] and Armpilia et al. [50]) for the optimal number of fractions under constant re-population. However, as in Unkelbach et al., Bortfeld et al. also only considered a single normal tissue.

In fact, the inability to simultaneously include two or more normal tissues is a major practical limitation of all work in this area that derives a provably optimal number of fractions. Cancer sites always include at least two normal tissues — unspecified normal tissue and an organ at risk such as the spinal cord. In fact, in most cases, several serial and parallel normal tissues play a role in limiting the tumor-dose. Finally, the existing works do not explicitly describe precisely how their optimal fractionation models could be implemented in practice. Specifically, it is not clear exactly how sparing factors are to be calculated and how the resulting temporal dosing schedule is to be combined with a spatially optimized fluence-map. In short, these works do not integrate their models with state-of-the-art IMRT methods. In this thesis, we are able to simultaneously overcome these limitations, and yet rigorously derive optimal dosing schedules as summarized next.

2.3 Our contributions

We present a mathematical formulation of the optimal fractionation problem based on the LQ model using sparing factors. This formulation allows for doses to vary across fractions. It incorporates (i) any combination of several maximum dose, mean dose and dose-volume type constraints on multiple normal tissues that receive spatially heterogeneous doses; and (ii) tumor proliferation with a time-lag. Finally, since the starting point for our model is a spatially optimized fluence-map (say from any treatment planning system) it is implementable after appropriate clinical validation in the future without requiring any additional hardware or technological developments. The generality of our model introduces several challenges. In this section, we highlight our contributions in the context of these challenges.

- **Optimality of equal-dosage fractionation:** Although equal-dosage fractionation is ubiquitous in practice and in the existing literature, it is not known whether this strategy is in fact optimal in the case of multiple normal tissues. Consequently, our formulation *does not a priori* assume equal-dosage fractionation. As a result, our initial formulation is a mixed-integer, *nonconvex*, quadratically constrained quadratic program (QCQP); such problems belong to the class NP-hard [63]. Fortunately, we are able to prove two simple sufficient conditions under which equal-dosage fractionation *is* optimal (see Propositions 2.5.1 and 2.5.4). Based on the prevalent estimates of the parameters of the LQ model, we expect these conditions to hold in many commonly studied tumors such as head-and-neck and prostate. We thus devise an equal-dosage reformulation of our problem and study it in detail in all subsequent sections.
- **Switching of the dose-limiting normal tissue:** When the problem formulation includes a single normal tissue, that normal tissue, by definition, is what limits tumor-dose. This facilitates the derivation of a simple closed-form formula for the optimal number of fractions and the corresponding dose per fraction. However, in the case of multiple normal tissues, we show that the dose-limiting normal tissue may depend on the number of fractions (see, for example, our Figure 2.2(a), which demonstrates that the limiting normal tissue switches from the unspecified normal tissue to the right parotid). Furthermore, it is not evident if and when such a switch occurs as this is a function of many model parameters especially when the formulation includes several maximum dose, mean dose and dose-volume type constraints to mimic realistic anatomies and treatment protocols, and allows for heterogeneous dose distributions.
- **Nonmonotonic behavior of tumor-BE:** The tumor-BE *without proliferation* is monotonic in the number of fractions in optimal fractionation problems that include a single normal tissue. We show that this is no longer true with multiple normal tissues; see, for example, our Figure 2.2(b), which illustrates that

the tumor-BE without proliferation can first increase and then decrease as the number of fractions is increased. In particular, *even without proliferation*, when $(\alpha/\beta)_{tumor} > \frac{(\alpha/\beta)_{normaltissue}}{s_{effective}}$, a finite number of fractions may be optimal in the case of multiple normal tissues (contrast this with the aforementioned conclusion in Unkelbach et al. and Keller et al. that with a single normal tissue, an infinite number of fractions is optimal). More generally, when a proliferation term is included, the tumor-BE is either decreasing or strictly concave in the number of fractions in the case of a single normal tissue (see Section 2.7). Again, this is at the heart of the derivation of a simple closed-form formula for the optimal number of fractions and the corresponding dose per fraction. In the case of multiple normal tissues, the behavior of tumor-BE (with or without proliferation) as a function of the number of fractions is complicated owing to the possible switching of the limiting normal tissue and this has not been investigated in the existing literature.

- Piecewise convexity of the limiting dose per fraction and quasiconcavity of tumor-BE: We perform a rigorous mathematical analysis of our model to completely characterize how the limiting dose and tumor-BE behave as functions of the number of fractions (see Lemmas 2.6.1 and Theorem 2.6.3). Specifically, in Equation (2.94), we obtain a closed-form formula for the limiting dose, and show that it is a strictly decreasing, piecewise strictly convex function of the number of fractions (see Lemma 2.6.1 and Figure 2.1). This leads to a closed-form formula for the corresponding tumor-BE. Furthermore, we establish in Theorem 2.6.3 that this tumor-BE eventually decreases and that it is a quasiconcave [64] function of the number of fractions.
- Simple calculation procedure for problem solution: We propose a simple calculation procedure rooted in the above quasiconcavity property to obtain an optimal number of fractions and the corresponding dose per fraction. Specifically, we calculate the tumor-BE using our closed-form formula for dose per

fraction sequentially for $N = 1, 2, \dots$ until the tumor-BE “turns around” to correctly identify the optimal number of fractions (see Theorem 2.6.3 and Corollary 2.6.4). We note here that in the case of multiple normal tissues, it is not possible to find the optimal number of fractions in closed-form by setting a derivative to equal zero unlike the single normal tissue case as in Jones et al., Bortfeld et al., and our Section 2.7.

- Uncovering clinically relevant insights through numerical experiments: In Section 5.4, we perform extensive numerical experiments and sensitivity analyses using five head-and-neck cancer cases and five prostate cancer cases to gain further insight into the optimal fractionation problem. Some examples of such insights are listed here.
 - Mismatch between predictions of our model with multiple normal tissues and existing models with a single normal tissue: We show in Section 2.8.3 that the optimal number of fractions in our model with multiple normal tissues cannot in general be found simply by choosing from the set of optimal number of fractions obtained from separately solving single normal tissue problems one by one. This point is especially important as all existing analytical literature considers only a single normal tissue, and furthermore, the difference between the optimal number of fractions from our model and that obtained from a single normal tissue model can be rather large.
 - Sensitivity to tumor doubling time: We demonstrate in Section 2.8.4 that the optimal number of fractions increases as the tumor doubling time increases — that is, a longer treatment schedule is optimal for slower growing tumors. We demonstrate that the sensitivity of the optimal number of fractions to the tumor doubling time is higher when the alpha over beta ratio of the tumor is larger as compared to that of a normal tissue.
 - Effect of time-lag before proliferation starts: We show that when the suffi-

cient condition in Proposition 2.5.1 is met, it is not optimal to administer less than $1 + T_{lag}$ fractions where T_{lag} is the lag in days before tumor proliferation begins (see Remark 2.5.2). When the tumor doubling time is smaller than a threshold determined by T_{lag} and other model parameters, it is optimal to administer exactly $1 + T_{lag}$ fractions. When the doubling time is larger than this threshold, the optimal number of fractions is larger than $1 + T_{lag}$ but (otherwise) no longer sensitive to T_{lag} . Our numerical results in Section 2.8.5 further illustrate this behavior for our head-and-neck cases.

The rest of this chapter is organized as follows. In the next section, we describe our problem formulation in detail. In Section 2.5, we prove sufficient conditions for equal-dosage fractionation. Then in Section 2.6, we present an equal-dosage reformulation of our original model. We then present its mathematical analysis, and propose a simple solution procedure. A previously known closed-form formula for the optimal number of fractions in the case of a single normal tissue is derived in Section 2.7 for the sake of completeness and notational consistency as we will need this formula in Section 2.8.3 when its results are compared with our multiple normal tissue model. We present numerical results in Section 5.4. Finally, in Section 5.5, we conclude by discussing opportunities for extensions of our work here.

2.4 Problem formulation

Our overall methodology is to choose the number of fractions to maximize the BE of the average tumor-dose subject to constraints on normal tissue doses. All components of this methodology are described next.

2.4.1 Expression for the objective function

Let N be the number of fractions. This is the key decision variable in our optimization model. We assume that a single fraction is administered every day. As we will describe in Section 5.5, our model can be easily extended to tackle other fractionation

schedules such as the one that administers treatment on weekdays only or the one that administers treatment twice daily on weekdays. Let T_{double} be the doubling time for the tumor [44–46] and let T_{lag} be the time-lag after which tumor proliferation starts after treatment initiation [49]; both these are measured in days. Finally, let α_0 and β_0 denote the parameters of the LQ model for the tumor.

In optimization models in IMRT, various anatomical regions of interest in the patient’s body are discretized into equal-volume cubes called voxels. Let n denote the number of such voxels in the tumor. We index these voxels by $i = 1, 2, \dots, n$. Similarly, the radiation field is also discretized into small segments called beamlets. Let k be the number of such beamlets. We use $u \in \mathfrak{R}_+^k$ to denote a generic k -dimensional beamlet intensity vector (fluence-map). Let A be the $n \times k$, nonnegative, tumor-dose deposition coefficient matrix. Let A_i denote its i th row. This row corresponds to the i th tumor-voxel. Then, according to the linear dose deposition model, the dose delivered to the i th voxel in the tumor equals $A_i u$ if fluence-map u is used.

We assume that a nominal fluence-map is available to the treatment planner. For example, this could be obtained by solving any of the standard fluence-map optimization models as is routinely done in commercial treatment planning systems available in hospitals. We denote this fluence-map by $u_{nominal}$ and it serves as an input and a starting point for our model. An important benefit of using such a fluence-map as input is that we are assured that the spatial properties of the treatment plan such as its dose-volume histograms are clinically acceptable. Let $d_{nominal}$ denote the average tumor-dose over all tumor-voxels delivered by this nominal fluence-map. Consider any dose $d \geq 0$. According to the linear dose deposition model, a fluence-map that delivers an average tumor-dose of d can be constructed by scaling $u_{nominal}$ to $u_{nominal} d / d_{nominal}$.

Now suppose that average tumor-doses $\vec{d} = (d_1, d_2, \dots, d_N)$ are delivered over the N fractions; \vec{d} is the second decision variable in our problem. According to the LQ

model, the BE of delivering average tumor-doses \vec{d} over N fractions equals

$$\alpha_0 \sum_{t=1}^N d_t + \beta_0 \sum_{t=1}^N d_t^2 - \tau(N). \quad (2.1)$$

Here, we have used the notation

$$\tau(N) = \frac{[(N-1) - T_{lag}]^+ \ln 2}{T_{double}}. \quad (2.2)$$

The expression $[(N-1) - T_{lag}]^+$, which is defined as $\max((N-1) - T_{lag}, 0)$, models the possibility that tumor proliferation does not start until T_{lag} days after treatment begins [49]. The appearance of $N-1$ instead of N in the above expression may seem surprising at first; however, note that the time that *elapses* between the first and the N th fractions is actually $N-1$ days and hence this is the time over which the tumor could proliferate. Note that $\tau(N) = 0$ for all N such that $N \leq 1 + T_{lag}$, and that $\tau(N)$ is strictly increasing and linear for all N such that $N \geq 1 + T_{lag}$. In particular, $\tau(N)$ is nondecreasing for all $N \geq 1$; it is also convex for all $N \geq 1$. These properties of $\tau(N)$ will be used in our mathematical analysis. We next describe the constraints in our model.

2.4.2 Normal tissue constraints

Let O_1, O_2, \dots, O_M denote the M different normal tissues under consideration. Let $\mathcal{M} = \{1, \dots, M\}$ be the set of indices of these normal tissues. For $m \in \mathcal{M}$, let n_m denote the number of voxels in O_m . These voxels are indexed by $j = 1, 2, \dots, n_m$. Let \mathcal{N}_m denote the set $\{1, 2, \dots, n_m\}$ of these voxels. All normal tissue voxels are assumed to have equal volume as is standard practice in IMRT. Let A^m be the $n_m \times k$, nonnegative, dose deposition coefficient matrix for O_m . Let A_j^m be the j th row of this matrix; this is the row that corresponds to the j th voxel in O_m . That is, according to the linear dose deposition model, $A_j^m u$ is the dose delivered to the j th voxel in O_m if fluence-map u is used. For each $m \in \mathcal{M}$ and $j \in \mathcal{N}_m$, let s_j^m denote the sparing factor for the j th voxel in normal tissue O_m . This sparing factor is given by

$$s_j^m = \frac{A_j^m u_{nominal}}{d_{nominal}}. \quad (2.3)$$

Then, according to the linear dose deposition model, the scaled fluence-map $u_{nominal}d_t/d_{nominal}$ delivers a dose of

$$\frac{A_j^m u_{nominal} d_t}{d_{nominal}} = s_j^m d_t \quad (2.4)$$

to voxel j in normal tissue O_m . Let α_m and β_m denote the parameters of the LQ model for normal tissue O_m , and let $\rho_m = 1/(\alpha_m/\beta_m)$ denote the inverse of the alpha over beta ratio for this normal tissue. Then, the BED corresponding to the doses $s_j^m \vec{d}$ that are delivered to voxel j in O_m over N fractions is given by

$$\sum_{t=1}^N (s_j^m d_t) + \rho_m \sum_{t=1}^N (s_j^m d_t)^2. \quad (2.5)$$

The above BED formula will be used in our formulation to incorporate the three most common types of constraints on normal tissue that are used in practice — maximum dose constraints for serial normal tissue, and mean dose or dose-volume type constraints for parallel normal tissue. We describe each of these in detail in the next three sections.

Serial normal tissue

Let $\mathcal{M}_1 \subseteq \mathcal{M}$ be the set of indices of serial normal tissues; these are the normal tissues whose function is hampered even when a small portion suffers damage [1]. Examples include spinal cord and brainstem. Suppose for any $m \in \mathcal{M}_1$ that a total dose D_{max}^m is known to be tolerated by each voxel in O_m if administered in N_{conv}^m equal-dose fractions. The BED corresponding to this schedule equals

$$BED_{max}^m = D_{max}^m (1 + \rho_m (D_{max}^m / N_{conv}^m)). \quad (2.6)$$

Then, doses $s_j^m \vec{d}$ over N fractions can be tolerated by the j th voxel in normal tissue O_m if

$$\sum_{t=1}^N (s_j^m d_t) + \rho_m \sum_{t=1}^N (s_j^m d_t)^2 \leq BED_{max}^m, \quad \forall j \in \mathcal{N}_m. \quad (2.7)$$

Thus, for each $m \in \mathcal{M}_1$, our problem formulation will include constraints (2.7).

Mean dose constraints for parallel normal tissue

Let $\mathcal{M}_2 \subseteq \mathcal{M}$ be the set of indices of parallel normal tissues with mean dose constraints; these are the normal tissues where a sufficiently small portion can be damaged without affecting the organ function [1]. Examples include parotid glands, heart, liver, kidneys, and lungs. Suppose for any $m \in \mathcal{M}_2$ that mean dose D_{mean}^m is known to be tolerated by O_m if administered in N_{conv}^m equal-dose fractions. The BED corresponding to this mean dose is given by

$$BED_{mean}^m = D_{mean}^m (1 + \rho_m (D_{mean}^m / N_{conv}^m)). \quad (2.8)$$

Then, doses $s_j^m \vec{d}$ in over N fractions can be tolerated by O_m if

$$\sum_{t=1}^N \sum_{j=1}^{n_m} (s_j^m d_t + \rho_m (s_j^m d_t)^2) \leq n_m BED_{mean}^m. \quad (2.9)$$

As in Unkelbach et al. [55], the quadratic constraint (2.9) can be simplified as

$$p_m \sum_{t=1}^N d_t + \rho_m q_m \sum_{t=1}^N d_t^2 \leq n_m BED_{mean}^m, \quad (2.10)$$

where we have used the notation

$$p_m \triangleq \sum_{j=1}^{n_m} s_j^m, \text{ and } q_m \triangleq \sum_{j=1}^{n_m} (s_j^m)^2. \quad (2.11)$$

We also define the effective sparing factor

$$s_{mean}^m \triangleq q_m / p_m, \quad (2.12)$$

and let

$$B_{mean}^m \triangleq q_m n_m BED_{mean}^m / (p_m)^2. \quad (2.13)$$

Then, multiplying both sides by the nonnegative number $q_m / (p_m)^2$, inequality (2.10) can be rewritten as

$$\sum_{t=1}^N s_{mean}^m d_t + \rho_m \sum_{t=1}^N (s_{mean}^m d_t)^2 \leq B_{mean}^m. \quad (2.14)$$

Thus, for each $m \in \mathcal{M}_2$, our optimization model will include constraint (2.14).

Dose-volume constraints

Let $\mathcal{M}_3 \subseteq \mathcal{M}$ be the set of indices of normal tissues with dose-volume constraints. Examples include rectum, bladder, and lung. In particular, suppose for any $m \in \mathcal{M}_3$ that no more than a volume fraction ϕ_m of normal tissue O_m can receive a dose more than D_{dv}^m if administered in N_{conv}^m fractions. The BED of total dose D_{dv}^m administered in N_{conv}^m equal-dose fractions is given by

$$BED_{dv}^m = D_{dv}^m(1 + \rho_m(D_{dv}^m/N_{conv}^m)). \quad (2.15)$$

Since all voxels in O_m are assumed to have equal volume, the volume fraction is the same as the voxel fraction. For each $m \in \mathcal{M}_3$ and for each voxel $j \in \mathcal{N}_m$, we thus define binary-valued functions $f_j(\vec{d}, N)$ such that

$$f_j(\vec{d}, N) = \begin{cases} 1 & \text{if } \sum_{t=1}^N (s_j^m d_t) + \rho_m \sum_{t=1}^N (s_j^m d_t)^2 > BED_{dv}^m, \\ 0 & \text{if } \sum_{t=1}^N (s_j^m d_t) + \rho_m \sum_{t=1}^N (s_j^m d_t)^2 \leq BED_{dv}^m. \end{cases} \quad (2.16)$$

We use the integer K_m to denote $\lfloor n_m \phi_m \rfloor$. Then the dose-volume constraints are written in our formulation as

$$\sum_{j=1}^{n_m} f_j(\vec{d}, N) \leq K_m, \text{ for } m \in \mathcal{M}_3. \quad (2.17)$$

Here, for simplicity of notation, we have assumed that there is at most one dose-volume constraint for each normal tissue. This simplifying assumption can be easily relaxed to allow for multiple dose-volume constraints on one or more of the normal tissues. In fact, we do this in our computational experiments for prostate cancer in Section 5.4 (see Table 4.6). For each $m \in \mathcal{M}_3$, our formulation will include constraint (2.17). We emphasize here that, to the best of our knowledge, despite the prevalence of dose-volume constraints in practice, they have thus far not been incorporated into existing optimal fractionation models.

A formulation of the optimal fractionation problem

Putting together the various components discussed in the above sections, we formulate our optimal fractionation problem as

$$(OPTFRAC1) \max_{N, \vec{d}} \alpha_0 \sum_{t=1}^N d_t + \beta_0 \sum_{t=1}^N d_t^2 - \tau(N), \quad (2.18)$$

$$\text{subject to } \sum_{t=1}^N (s_j^m d_t) + \rho_m \sum_{t=1}^N (s_j^m d_t)^2 \leq BED_{max}^m, \quad j \in \mathcal{N}_m, \quad m \in \mathcal{M}_1, \quad (2.19)$$

$$\sum_{t=1}^N s_{mean}^m d_t + \rho_m \sum_{t=1}^N (s_{mean}^m d_t)^2 \leq B_{mean}^m, \quad m \in \mathcal{M}_2, \quad (2.20)$$

$$\sum_{j=1}^{n_m} f_j(\vec{d}, N) \leq K_m, \quad \text{for } m \in \mathcal{M}_3, \quad (2.21)$$

$$\vec{d} \geq 0, \quad (2.22)$$

$$N \geq 1, \quad \text{integer.} \quad (2.23)$$

Note that for each $m \in \mathcal{M}_1$, constraint (2.19) is most restrictive for a voxel $j \in \mathcal{N}_m$ that has the largest sparing factor. Therefore, as in Unkelbach et al. [55], we define the effective sparing factor for each $m \in \mathcal{M}_1$ as

$$s_{max}^m \triangleq \max_{j \in \mathcal{N}_m} s_j^m. \quad (2.24)$$

Then, for each $m \in \mathcal{M}_1$, the group of constraints (2.19) is equivalent to the single constraint

$$\sum_{t=1}^N (s_{max}^m d_t) + \rho_m \sum_{t=1}^N (s_{max}^m d_t)^2 \leq BED_{max}^m. \quad (2.25)$$

Now consider constraint (2.21) for any $m \in \mathcal{M}_3$. This dose-volume constraint requires that $\sum_{t=1}^N (s_j^m d_t) + \rho_m \sum_{t=1}^N (s_j^m d_t)^2 \leq BED_{dv}^m$ for at least $n_m - K_m$ voxels; this gives us the flexibility to pick and choose *which* voxels satisfy this upper bound. Since the objective function in (OPTFRAC1) is increasing in each d_t , for $t = 1, 2, \dots, N$, the largest objective function value is obtained by enforcing this upper bound on $n_m - K_m$

voxels with the smallest sparing factors. That is, we should enforce this upper bound only on the voxels with the smallest, the second smallest, the third smallest, ..., and the $(n_m - K_m)$ th smallest sparing factors. But this is equivalent to enforcing the upper bound on a single voxel with the $(n_m - K_m)$ th smallest sparing factor. Using standard order statistic subscript notation, we thus define the effective sparing factor for normal tissue $m \in \mathcal{M}_3$ as

$$s_{dv}^m \triangleq s_{(n_m - K_m)}^m. \quad (2.26)$$

Thus, constraint (2.21) reduces to the much simpler constraint

$$\sum_{t=1}^N s_{dv}^m d_t + \rho_m \sum_{t=1}^N (s_{dv}^m d_t)^2 \leq BED_{dv}^m. \quad (2.27)$$

Consequently, we rewrite problem (OPTFRAC1) equivalently as

$$(OPTFRAC2) \max_{N, \vec{d}} \alpha_0 \sum_{t=1}^N d_t + \beta_0 \sum_{t=1}^N d_t^2 - \tau(N), \quad (2.28)$$

$$\text{subject to } \sum_{t=1}^N (s_{max}^m d_t) + \rho_m \sum_{t=1}^N (s_{max}^m d_t)^2 \leq BED_{max}^m, \quad m \in \mathcal{M}_1, \quad (2.29)$$

$$\sum_{t=1}^N s_{mean}^m d_t + \rho_m \sum_{t=1}^N (s_{mean}^m d_t)^2 \leq B_{mean}^m, \quad m \in \mathcal{M}_2, \quad (2.30)$$

$$\sum_{t=1}^N s_{dv}^m d_t + \rho_m \sum_{t=1}^N (s_{dv}^m d_t)^2 \leq BED_{dv}^m, \quad \text{for } m \in \mathcal{M}_3, \quad (2.31)$$

$$\vec{d} \geq 0, \quad (2.32)$$

$$N \geq 1, \text{ integer}. \quad (2.33)$$

Observe that (OPTFRAC2) is a nonconvex QCQP in the decision vector \vec{d} and also includes the integer valued decision variable N . As a result, exact solution of this optimization problem appears computationally challenging. Fortunately, we are able to derive two sufficient conditions under which an equal-dosage schedule is optimal to this problem. Furthermore, we expect these sufficient conditions to hold in many commonly studied tumors such as head-and-neck and prostate. This motivates the

equal-dosage reformulation of (OPTFRAC2) that we study in Sections 2.6 and 5.4 of this Chapter.

2.5 Sufficient conditions for optimality of equal-dosage fractionation

Recall that the set of all normal tissues is denoted by $\mathcal{M} = \{1, \dots, M\}$ and assume for simplicity of exposition that $\mathcal{M}_1 \cap \mathcal{M}_2 = \mathcal{M}_2 \cap \mathcal{M}_3 = \mathcal{M}_1 \cap \mathcal{M}_3 = \emptyset$. Then we define C_m to equal the right hand side in the quadratic inequality constraint in (OPTFRAC2) for normal tissue $m \in \mathcal{M}$; similarly, we define σ_m as the effective sparing factor for normal tissue $m \in \mathcal{M}$. Then (OPTFRAC2) can be compactly written as

$$\text{OPTFRAC2} \max_{N, \vec{d}} \alpha_0 \sum_{t=1}^N d_t + \beta_0 \sum_{t=1}^N d_t^2 - \tau(N) \quad (2.34)$$

$$\sigma_m \sum_{t=1}^N d_t + \sigma_m^2 \rho_m \sum_{t=1}^N d_t^2 \leq C_m, \quad m \in \mathcal{M}, \quad (2.35)$$

$$\vec{d} \geq 0, \quad (2.36)$$

$$N \geq 1, \text{ integer}. \quad (2.37)$$

Moreover, for any fixed integer $N \geq 1$, we formulate

$$(\text{OPTFRAC2}(N)) \max_{\vec{d}} \alpha_0 \sum_{t=1}^N d_t + \beta_0 \sum_{t=1}^N d_t^2 - \tau(N) \quad (2.38)$$

$$\sigma_m \sum_{t=1}^N d_t + \sigma_m^2 \rho_m \sum_{t=1}^N d_t^2 \leq C_m, \quad m \in \mathcal{M}, \quad (2.39)$$

$$\vec{d} \geq 0. \quad (2.40)$$

First note that the objective function and the constraint functions in this problem are continuous and the feasible region is bounded. Thus, by the extreme value theorem of Weierstrass, the problem has an optimal solution.

We next establish two conditions under which equal-dosage fractionation is optimal in (OPTFRAC2(N)).

Proposition 2.5.1. *Suppose $\frac{\alpha_0}{\beta_0} \geq \max_{m \in \mathcal{M}} \frac{\alpha_m/\beta_m}{\sigma_m}$. Then it is optimal in (OPTFRAC2(N)) to deliver dose*

$$d^* = \min_{m \in \mathcal{M}} \left(\frac{-1 + \sqrt{1 + 4\rho_m C_m/N}}{2\sigma_m \rho_m} \right) \quad (2.41)$$

in each one of the N fractions.

We provide two different proofs of this result, one is algebraic and the other uses KKT conditions. We expect the condition in this proposition to hold in tumors similar to head-and-neck cancer where the tumor α/β ratio is believed to be much bigger than that of surrounding normal tissues [48, 49, 65]. We also emphasize that the key result of Proposition 2.5.1 is the conclusion that it is optimal to use an identical dose in every session. Once this is established, it is easy to derive the actual value of the dose per session by solving a quadratic equation as has been done in most of the existing literature with a single normal tissue. Finally, in the special case where there is a single normal tissue, our condition correctly reduces to the condition that $(\alpha/\beta)_{tumor} \geq \frac{(\alpha/\beta)_{normaltissue}}{\sigma_{effective}}$ as in the existing literature.

Remark 2.5.2. *We note that if $\alpha_0/\beta_0 > \max_{m \in \mathcal{M}} \frac{\alpha_m/\beta_m}{\sigma_m}$ as in Proposition 2.5.1 then it is not optimal to administer less than $1 + T_{lag}$ fractions. This is because the tumor-BE $g^*(N) = \min_{m \in \mathcal{M}} g_m(N)$ as defined later in (2.112) is increasing over the range $1 \leq N < 1 + T_{lag}$. This holds because each $g_m(N)$ has this same property as shown in our proof of Lemma 2.7.1 in the case of a single normal tissue. A further discussion of this issue is included in Section 2.8.5.*

An algebraic proof of Proposition 2.5.1: Suppose there exists an optimal solution $\vec{d}^* = (d_1^*, d_2^*, \dots, d_N^*)$ where $d_i^* \neq d_j^*$ for some distinct i and j chosen from the set $\{1, 2, \dots, N\}$. We will show that it is possible to construct an equal-dosage solution with the same tumor-BE as \vec{d}^* . Since N is fixed, $\tau(N)$ has no effect on optimal doses; we thus ignore $\tau(N)$ in this proof.

Since \vec{d}^* is optimal, at least one of the constraints in (2.39) must be active. Without

loss of generality, we assume that it is the constraint with $m = 1$. Thus, we have,

$$\sum_{t=1}^N (d_t^*)^2 = \frac{C_1 - \sigma_1 \sum_{t=1}^N d_t^*}{\sigma_1^2 \rho_1}. \quad (2.42)$$

Substituting this in the objective function, we see that the optimal objective function value is given by

$$BE^* \triangleq \frac{C_1 \beta_0}{\sigma_1^2 \rho_1} + \left(\alpha_0 - \frac{\beta_0}{\rho_1 \sigma_1} \right) \sum_{t=1}^N d_t^*. \quad (2.43)$$

To obtain a contradiction to the optimality of \vec{d}^* , we will now construct an alternative feasible solution with a larger objective value than BE^* . Toward this end, we first define dose

$$\delta = \frac{\sum_{t=1}^N d_t^*}{N}, \quad (2.44)$$

and the dose vector $\vec{\delta} = (\delta, \delta, \dots, \delta)$. We claim that this dose vector is *strictly* feasible to all constraints in OPTFRAC2(N). To see this, we first define the notation $\|\vec{d}\|_1$, for any nonnegative dose vector \vec{d} , to denote the l_1 norm of \vec{d} . Similarly, $\|\vec{d}\|_2$ denotes the l_2 norm of \vec{d} . We now evaluate the LHS in (2.39) for some any arbitrary $m \in \mathcal{M}$ by substituting the dose vector $\vec{\delta}$. We get,

$$\sigma_m \sum_{t=1}^N \delta + \sigma_m^2 \rho_m \sum_{t=1}^N \delta^2 = N \sigma_m \delta + \sigma_m^2 \rho_m N \delta^2 = \sigma_m \sum_{t=1}^N d_t^* + \sigma_m^2 \rho_m N \frac{\left(\sum_{t=1}^N d_t^* \right)^2}{N^2} \quad (2.45)$$

$$= \sigma_m \sum_{t=1}^N d_t^* + \sigma_m^2 \rho_m \frac{(\|\vec{d}^*\|_1)^2}{N} < \sigma_m \sum_{t=1}^N d_t^* + \sigma_m^2 \rho_m (\|\vec{d}^*\|_2)^2 \quad (2.46)$$

$$= \sigma_m \sum_{t=1}^N d_t^* + \sigma_m^2 \rho_m \sum_{t=1}^N (d_t^*)^2 \leq C_m. \quad (2.47)$$

Here, the strict inequality follows from the well-known norm-inequality $\|\vec{d}\|_1 \leq \sqrt{N} \|\vec{d}\|_2$ (with equality holding only when all components of \vec{d} are identical; note that at least two components of our vector \vec{d}^* are distinct by assumption). The second inequality

holds because \vec{d}^* is feasible to OPTFRAC2(N). Owing to the *strict* feasibility of $\vec{\delta}$ to OPTFRAC2(N), there is some slack in every constraint. In particular, there exists an $\epsilon > 0$ such that delivering dose $\delta + \epsilon$ in each one of the N fractions is feasible. In fact, we will choose ϵ to be large enough such that at least one of the constraints becomes active. Specifically, we let $\epsilon = \min_{m \in \mathcal{M}} \epsilon_m$, where $\epsilon_m > 0$ for $m \in \mathcal{M}$ is uniquely defined such that

$$N\sigma_m(\delta + \epsilon_m) + N\sigma_m^2\rho_m(\delta + \epsilon_m)^2 = C_m. \quad (2.48)$$

Suppose without loss of generality that $\epsilon = \epsilon_2$. That is, the dose delivered in each fraction now equals $\delta + \epsilon_2$ and

$$N\sigma_2(\delta + \epsilon_2) + N\sigma_2^2\rho_2(\delta + \epsilon_2)^2 = C_2, \text{ and hence,} \quad (2.49)$$

$$N(\delta + \epsilon_2)^2 = \frac{C_2 - N\sigma_2(\delta + \epsilon_2)}{\sigma_2^2\rho_2}. \quad (2.50)$$

Now we evaluate the objective function value BE for this new dosing schedule. We have,

$$\text{BE} = \alpha_0 N(\delta + \epsilon_2) + \beta_0 N(\delta + \epsilon_2)^2 = \frac{C_2\beta_0}{\sigma_2^2\rho_2} + \left(\alpha_0 - \frac{\beta_0}{\sigma_2\rho_2}\right)N(\delta + \epsilon_2) \quad (2.51)$$

$$= \frac{C_2\beta_0}{\sigma_2^2\rho_2} + \left(\alpha_0 - \frac{\beta_0}{\sigma_2\rho_2}\right) \sum_{t=1}^N d_t^* + \left(\alpha_0 - \frac{\beta_0}{\sigma_2\rho_2}\right)N\epsilon_2. \quad (2.52)$$

Now we show that $\text{BE} = \text{BE}^*$. Since \vec{d}^* is feasible for constraint (2.39) with $m = 2$, we have,

$$\sigma_2 \sum_{t=1}^N d_t^* + \sigma_2^2\rho_2 \sum_{t=1}^N d_t^{*2} \leq C_2. \quad (2.53)$$

Now, by substituting for $\sum_{t=1}^N d_t^{*2}$ from (2.42), the above inequality simplifies as

$$\sigma_2 \sum_{t=1}^N d_t^* + \sigma_2^2\rho_2 \left(\frac{C_1 - \sigma_1 \sum_{t=1}^N d_t^*}{\sigma_1^2\rho_1} \right) = \frac{\sigma_2^2\rho_2 C_1}{\sigma_1^2\rho_1} + \left(\sigma_2 - \frac{\sigma_2^2\rho_2}{\sigma_1\rho_1} \right) \sum_{t=1}^N d_t^* \leq C_2. \quad (2.54)$$

Now, by multiplying both sides of this inequality by the positive term $\frac{\beta_0}{\rho_2\sigma_2^2}$, we get,

$$\frac{\beta_0 C_1}{\sigma_1^2\rho_1} + \frac{\beta_0}{\sigma_2^2\rho_2} \left(\sigma_2 - \frac{\sigma_2^2\rho_2}{\sigma_1\rho_1} \right) \sum_{t=1}^N d_t^* \leq \frac{\beta_0 C_2}{\sigma_2^2\rho_2}. \quad (2.55)$$

Now by adding $(\alpha_0 - \beta_0/\sigma_2\rho_2) \sum_{t=1}^N d_t^*$ to both sides of this inequality, we obtain,

$$\text{BE}^* = \frac{C_1\beta_0}{\sigma_1^2\rho_1} + \left(\alpha_0 - \frac{\beta_0}{\sigma_1\rho_1}\right) \sum_{t=1}^N d_t^* \leq \frac{C_2\beta_0}{\sigma_2^2\rho_2} + \left(\alpha_0 - \frac{\beta_0}{\sigma_2\rho_2}\right) \sum_{t=1}^N d_t^* \leq \text{BE}. \quad (2.56)$$

Here, the last inequality follows from the definition of BE in Equation (2.52) because $\epsilon_2 > 0$ and $\left(\alpha_0 - \frac{\beta_0}{\sigma_2\rho_2}\right) = \left(\alpha_0 - \frac{\beta_0(\alpha_2/\beta_2)}{\sigma_2}\right) \geq 0$ by our hypothesis in the statement of the proposition. But since BE* is the optimal tumor-BE, the above inequality implies that $\text{BE} = \text{BE}^*$. This shows that there is no loss of optimality in only considering equal-dosage fractionation schedules in (OPTFRAC2(N)) under the hypothesis of this proposition. The fact that the dose in an optimal equal-dosage schedule is given by (2.41) follows by finding the smallest from the M positive roots of M quadratic equations defined by the constraints (2.39).

Remark 2.5.3. *The idea of using the average of the components of a vector optimal solution as in (2.44), has often been used to prove optimality of symmetric solutions to convex problems (see, for example, [66] for an interesting historical account of this topic). In fact, we used this idea in the context of a convex problem in radiotherapy in Proposition 4.2.1 in Kim's doctoral thesis [22]. Since OPTFRAC2(N) is not a convex problem, our proof of Proposition 2.5.1 relies on a sufficient condition under which a modification of this "average of components" method works. Below we also provide a different proof based on KKT conditions.*

Proof of Proposition 2.5.1 using KKT conditions: Problem OPTFRAC2(N) is not convex. However, it is easy to construct a strictly feasible dose vector. Furthermore, the feasible region is convex. In particular, this problem satisfies Slater's constraint qualification [64]. Consequently, KKT conditions are necessary for optimality. To write these, we associate Lagrange multipliers $\lambda_1, \lambda_2, \dots, \lambda_M$ with the inequality constraints in OPTFRAC2(N) and Lagrange multipliers $\mu_1, \mu_2, \dots, \mu_N$ with the nonnegativity constraints on doses. We let \vec{e}_t , for $t = 1, 2, \dots, N$, denote the N -dimensional unit vector with a one as the t th element and zeros everywhere else.

Then, the KKT conditions are given by

$$-\begin{bmatrix} \alpha_0 + 2\beta_0 d_1 \\ \alpha_0 + 2\beta_0 d_2 \\ \vdots \\ \alpha_0 + 2\beta_0 d_N \end{bmatrix} + \sum_{m=1}^M \lambda_m \sigma_m \begin{bmatrix} 1 + 2\rho_m \sigma_m d_1 \\ 1 + 2\rho_m \sigma_m d_2 \\ \vdots \\ 1 + 2\rho_m \sigma_m d_N \end{bmatrix} - \sum_{t=1}^N \mu_t \vec{e}_t = \vec{0}, \quad (2.57)$$

$$\vec{d} \geq 0, \quad (2.58)$$

$$\sigma_m \sum_{t=1}^N d_t + \sigma_m^2 \rho_m \sum_{t=1}^N (d_t)^2 \leq C_m, \quad m \in \mathcal{M}, \quad (2.59)$$

$$(\sigma_m \sum_{t=1}^N d_t + \sigma_m^2 \rho_m \sum_{t=1}^N (d_t)^2 - C_m) \lambda_m = 0, \quad m \in \mathcal{M}, \quad (2.60)$$

$$\mu_t d_t = 0, \quad t = 1, 2, \dots, N, \quad (2.61)$$

$$\lambda_m \geq 0, \quad m \in \mathcal{M}, \quad (2.62)$$

$$\mu_t \geq 0, \quad t = 1, 2, \dots, N. \quad (2.63)$$

Now, for each $k = 1, 2, \dots, N$, let $\vec{d}(k) \triangleq (d_1, d_2, \dots, d_k, 0, 0, \dots, 0)$, where $d_1 > 0$, $d_2 > 0$, \dots , $d_k > 0$. That is, $\vec{d}(k)$ is a dose vector where the doses in the first k fractions are positive but the doses in the other fractions are zero. We use the above KKT conditions to characterize these dose vectors. Note that it suffices to consider dose vectors where the leading elements are positive and the tail elements are zero; that is, we need not separately consider vectors such as $(0, d_2, 0, 0, \dots, 0)$. This is because the objective function and the constraints are symmetric with respect to permutations of elements of dose vectors.

For $\vec{d}(k)$, KKT conditions (2.61) imply that $\mu_1 = \mu_2 = \dots = \mu_k = 0$. Thus,

Equations (2.57) reduce to

$$-(\alpha_0 + 2\beta_0 d_t) + \sum_{m=1}^M \lambda_m \sigma_m (1 + 2\rho_m \sigma_m d_t) = 0, \quad t = 1, 2, \dots, k, \quad (2.64)$$

$$\mu_t = \sum_{m=1}^M \lambda_m \sigma_m - \alpha_0, \quad t = k+1, \dots, N. \quad (2.65)$$

We rewrite the first group of equations above as

$$2d_t(-\beta_0 + \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2) = \alpha_0 - \sum_{m=1}^M \lambda_m \sigma_m, \quad t = 1, 2, \dots, k. \quad (2.66)$$

Since the above right hand side does not depend on t , the left hand sides must also be equal for different t . Therefore, we have,

$$d_s(-\beta_0 + \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2) = d_t(-\beta_0 + \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2), \quad s, t = 1, 2, \dots, k. \quad (2.67)$$

We now consider two cases: (i) $(\alpha_0/\beta_0) > \max_{m \in \mathcal{M}} (\alpha_m/\beta_m)/\sigma_m$, and (ii) $(\alpha_0/\beta_0) = \max_{m \in \mathcal{M}} (\alpha_m/\beta_m)/\sigma_m$.

Case 1: $(\alpha_0/\beta_0) > \max_{m \in \mathcal{M}} (\alpha_m/\beta_m)/\sigma_m$. We claim that $\beta_0 \neq \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2$. This would imply that $d_1 = d_2 = \dots = d_k \triangleq d^*(k)$ for some $d^*(k)$ from the above equation. We proceed by contradiction. So suppose that $\beta_0 = \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2$. Substituting this into

(2.66), we get, $\alpha_0 = \sum_{m=1}^M \lambda_m \sigma_m$. Thus, $(\beta_0/\alpha_0) = (\sum_{m=1}^M \lambda_m \rho_m \sigma_m^2) / (\sum_{m=1}^M \lambda_m \sigma_m)$. But by (2.62), $\lambda_m \geq 0$ for $m \in \mathcal{M}$. Thus, (β_0/α_0) is a convex combination of $\sigma_m \beta_m / \alpha_m$.

But this is a contradiction since $(\alpha_0/\beta_0) > (\alpha_m/\beta_m)/\sigma_m$ for all m by assumption.

Now, since the right hand side in (2.65) does not depend on t , we must have that $\mu_{k+1} = \mu_{k+2} = \dots = \mu_N = \mu$ for some μ . Then, (2.65) implies that $\sum_{m=1}^M \lambda_m \sigma_m = \alpha_0 + \mu$.

Substituting this in any one of the group (2.64), we get,

$$\mu = 2d^*(k) \left(\beta_0 - \sum_{m=1}^M \lambda_m \sigma_m^2 \rho_m \right). \quad (2.68)$$

From this, we claim that if $\mu \geq 0$ then $d^*(k) < 0$ and hence the only possible KKT point is where $\mu_t = 0$ for $t = 1, 2, \dots, N$. Toward this end, suppose $\mu \geq 0$. Then,

from (2.65), we have that $\sum_{m=1}^M \lambda_m \sigma_m \geq \alpha_0$. Without loss of generality, suppose that $\min_{m \in \mathcal{M}} \sigma_m \rho_m = \sigma_1 \rho_1$. We have, $\sum_{m=2}^M \lambda_m \sigma_m \geq \alpha_0 - \lambda_1 \sigma_1$. That is, $-\sum_{m=2}^M \lambda_m \sigma_m \leq \lambda_1 \sigma_1 - \alpha_0$. Thus,

$$\beta_0 - \sum_{m=1}^M \lambda_m \sigma_m^2 \rho_m \leq \beta_0 - \sigma_1 \rho_1 \sum_{m=1}^M \lambda_m \sigma_m = \beta_0 - \lambda_1 \sigma_1^2 \rho_1 - \sigma_1 \rho_1 \sum_{m=2}^M \lambda_m \sigma_m \leq \beta_0 - \alpha_0 \sigma_1 \rho_1 < 0. \quad (2.69)$$

Here, the last strict inequality follows because $\alpha_0/\beta_0 > 1/(\rho_m \sigma_m)$ for all m , and in particular, $\alpha_0/\beta_0 > 1/(\rho_1 \sigma_1)$, by assumption. Substituting this in (2.68) implies that $d^*(k) < 0$ as claimed. In particular, $\vec{d}^*(N) \triangleq (d^*(N), d^*(N), \dots, d^*(N))$ is the only possible KKT point and hence an equal-dosage solution must be optimal. The actual value of $d^*(N)$ can be computed by solving quadratic equations as explained in the algebraic proof above.

Case 2: $(\alpha_0/\beta_0) = \max_{m \in \mathcal{M}} (\alpha_m/\beta_m)/\sigma_m$. We assume without loss of generality that there are no ties in this maximum. In particular, assume without loss of generality that $(\alpha_0/\beta_0) = (\alpha_1/\beta_1)/\sigma_1 = \max_{m \in \mathcal{M}} (\alpha_m/\beta_m)/\sigma_m$. Now suppose as above that $\mu \geq 0$ and by similar algebra we see that $\beta_0 = \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2 \leq 0$. Then, we consider two further subcases: (a) $\beta_0 = \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2$, and (b) $\beta_0 < \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2$.

Case 2a: In this case, by using the above convex combination argument, we see that $\lambda_1 > 0$ and $\lambda_m = 0$ for all $m \in \mathcal{M}$, $m \neq 1$. Then (2.66) implies that $\lambda_1 = \alpha_0/\sigma_1 > 0$, and in particular, (2.60) implies that $\sigma_1 \sum_{t=1}^k d_t + \sigma_1^2 \rho_1 \sum_{t=1}^k (d_t^2) = C_1$. That is, the corresponding inequality constraint is active and we can eliminate $\sum_{t=1}^k (d_t^2)$ from the objective function by substituting

$$\sum_{t=1}^k (d_t^2) = \frac{C_1 - \sigma_1 \sum_{t=1}^k d_t}{\sigma_1^2 \rho_1}. \quad (2.70)$$

But since $(\alpha_0/\beta_0) = 1/(\rho_1 \sigma_1)$, this reduces the objective function to the constant $\beta_0 C_1 / (\sigma_1^2 \rho_1)$. Similarly, we can eliminate $\sum_{t=1}^k (d_t^2)$ from all other constraints and this

reduces OPTFRAC2(N) to a problem with a single nonnegative variable $z = \sum_{t=1}^k d_t$. Then, an optimal solution to this simplified problem can be found such that $d_1 = d_2 = \dots = d_k$. Furthermore, the objective function value itself is independent of the value of $k = 1, 2, \dots, N$. That is, an equal-dosage solution is optimal. **Case 2b:** In this case, we can show that $d_i < 0$ for $i = 1, 2, \dots, k$ by algebra similar to above and hence this case cannot lead to a KKT solution.

Proposition 2.5.4. *Suppose $\frac{\alpha_0}{\beta_0} \leq \min_{m \in \mathcal{M}} \frac{\alpha_m/\beta_m}{\sigma_m}$. Then it is optimal in OPTFRAC2(N) to deliver dose*

$$d^* = \min_{m \in \mathcal{M}} \left(\frac{-1 + \sqrt{1 + 4\rho_m C_m}}{2\sigma_m \rho_m} \right) \quad (2.71)$$

in any one of the N fractions and set the doses in the other $N - 1$ fractions to zero. Since this dosing schedule is optimal for every $N \geq 1$, it is optimal to use a single fraction in (OPTFRAC2).

We provide two different proofs of this result, one is algebraic and the other uses KKT conditions. We expect this condition to hold in tumors similar to prostate cancer where the tumor α/β ratio is believed to be no larger than that of surrounding normal tissues [67–69]. The key contribution of Proposition 2.5.4 is the conclusion that it is optimal to deliver a single dose. Once this is established, it is again easy to derive the value of this dose by solving a quadratic equation. Finally, in the special case where there is a single normal tissue, our condition correctly reduces to the condition that $(\alpha/\beta)_{tumor} \leq \frac{(\alpha/\beta)_{normaltissue}}{\sigma_{effective}}$ as in the existing literature.

An algebraic proof of Proposition 2.5.4: Suppose that dose vector $\vec{d}^* = (d_1^*, d_2^*, \dots, d_N^*)$ is optimal, and that there exist distinct i and j chosen from the set $\{1, 2, \dots, N\}$ such that $d_i^* > 0$ and $d_j^* > 0$. Without loss of generality, we set $i = 1$ and $j = 2$ in the rest of this proof. We will construct a dosing schedule where the dose in the first fraction is positive but the dose in the second fraction is zero; we will show that this schedule has the same tumor-BE as that of \vec{d}^* . Then, the claim in this proposition follows by applying this procedure recursively. Since N is fixed, $\tau(N)$ has no effect on optimal doses; we thus ignore $\tau(N)$ in this proof.

Since \vec{d}^* is optimal, at least one of the constraints in (2.39) must be active. Without loss of generality, we assume that it is the constraint with $m = 1$. Thus, we have,

$$\sigma_1(d_1^* + d_2^* + d_3^* + \dots + d_N^*) + \sigma_1^2 \rho_1 ((d_1^*)^2 + (d_2^*)^2 + (d_3^*)^2 + \dots + (d_N^*)^2) = C_1. \quad (2.72)$$

Thus, as in the proof of Proposition 2.5.4, we again define the optimal objective function value as

$$\text{BE}^* \triangleq \frac{C_1 \beta_0}{\sigma_1^2 \rho_1} + \left(\alpha_0 - \frac{\beta_0}{\rho_1 \sigma_1} \right) \sum_{t=1}^N d_t^*. \quad (2.73)$$

Now consider an alternative feasible dose vector $\vec{\delta} \triangleq (d_1^* + \epsilon, 0, d_3^*, \dots, d_N^*)$, where $\epsilon = \min_{m \in \mathcal{M}} \epsilon_m$, and $\epsilon_m > 0$ for all $m \in \mathcal{M}$ are chosen such that

$$\sigma_m(d_1^* + \epsilon_m + d_3^* + \dots + d_N^*) + \sigma_m^2 \rho_m ((d_1^* + \epsilon_m)^2 + (d_3^*)^2 + \dots + (d_N^*)^2) = C_m. \quad (2.74)$$

We assume without loss of generality that $\epsilon = \epsilon_2$. Therefore,

$$(d_1^* + \epsilon_2)^2 + (d_3^*)^2 + \dots + (d_N^*)^2 = \frac{C_2 - \sigma_2(d_1^* + \epsilon_2 + d_3^* + \dots + d_N^*)}{\sigma_2^2 \rho_2}. \quad (2.75)$$

Note that because $\vec{\delta}$ is feasible to constraint (2.39) for $m = 1$ and that this constraint is assumed to be active at \vec{d}^* (recall Equation (2.72)), we have,

$$\sigma_1(d_1^* + \epsilon_2 + d_3^* + \dots + d_N^*) + \sigma_1^2 \rho_1 ((d_1^* + \epsilon_2)^2 + (d_3^*)^2 + \dots + (d_N^*)^2) \leq \quad (2.76)$$

$$\sigma_1(d_1^* + d_2^* + d_3^* + \dots + d_N^*) + \sigma_1^2 \rho_1 ((d_1^*)^2 + (d_2^*)^2 + (d_3^*)^2 + \dots + (d_N^*)^2). \quad (2.77)$$

After simplifying the above inequality, we get

$$\epsilon_2 + \sigma_1 \rho_1 (\epsilon_2^2 + 2\epsilon_2 d_1^*) \leq d_2^* + \sigma_1 \rho_1 (d_2^*)^2, \quad (2.78)$$

which implies that

$$\epsilon_2 + \sigma_1 \rho_1 \epsilon_2^2 < d_2^* + \sigma_1 \rho_1 (d_2^*)^2 \quad (2.79)$$

because $2\epsilon_2 d_1^* > 0$. This shows that

$$\epsilon_2 < d_2^*. \quad (2.80)$$

The objective function value at $\vec{\delta}$ is given by

$$\alpha_0(d_1^* + \epsilon_2 + d_3^* + \dots + d_N^*) + \beta_0((d_1^* + \epsilon_2)^2 + (d_3^*)^2 + \dots + (d_N^*)^2) \quad (2.81)$$

$$= \alpha_0(d_1^* + \epsilon_2 + d_3^* + \dots + d_N^*) + \beta_0 \frac{C_2 - \sigma_2(d_1^* + \epsilon_2 + d_3^* + \dots + d_N^*)}{\sigma_2^2 \rho_2} \quad (2.82)$$

$$= \frac{\beta_0 C_2}{\sigma_2^2 \rho_2} + \left(\alpha_0 - \frac{\beta_0}{\sigma_2 \rho_2} \right) (d_1^* + \epsilon_2 + d_3^* + \dots + d_N^*) \quad (2.83)$$

$$\geq \frac{\beta_0 C_2}{\sigma_2^2 \rho_2} + \left(\alpha_0 - \frac{\beta_0}{\sigma_2 \rho_2} \right) (d_1^* + d_2^* + d_3^* + \dots + d_N^*) \geq \text{BE}^*. \quad (2.84)$$

Here, the first inequality follows because $(\alpha_0 - \frac{\beta_0}{\sigma_2 \rho_2}) \leq 0$ by our hypothesis in the statement of the proposition and $\epsilon_2 < d_2^*$ from (2.80). The proof to show that the second inequality holds is the same as the corresponding part in the proof for Proposition 2.5.1. But since BE^* is the optimal tumor-BE, the above inequality implies that $\text{BE} = \text{BE}^*$. This shows that there is no loss of optimality in administering a single fraction with positive dose in (OPTFRAC2(N)) under the hypothesis of this proposition. The value of this single dose is easily obtained by finding the smallest among the M positive roots of the M quadratic equations implied by constraints (2.39) just as in the proof of Proposition 2.5.1.

Proof of Proposition 2.5.4 using KKT conditions: We only provide a brief sketch of this proof to avoid repetition. Using a convex combination argument identical to the KKT proof of Proposition 2.5.1, we can again show that $\beta_0 \neq \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2$ if $(\alpha_0/\beta_0) < \min_{m \in \mathcal{M}} (\alpha_m/\beta_m)/\rho_m$. In this case, using algebra similar to the corresponding proof of Proposition 2.5.1, we can show that $\vec{d}^*(k) \triangleq (\underbrace{d^*(k), \dots, d^*(k)}_{k \text{ times}}, 0, \dots, 0)$, for some $d^*(k) > 0$, and for $k = 1, 2, \dots, N$, are the only possible KKT points. In fact, the best value of $d^*(k)$ is defined precisely as in Equation (2.94) by

$$d^*(k) = \min_{m \in \mathcal{M}} b_m(k), \quad (2.85)$$

where $b_m(k)$ is defined as in (2.93). That is,

$$b_m(k) \triangleq \frac{-1 + \sqrt{1 + 4\rho_m C_m/k}}{2\sigma_m \rho_m}, \quad m \in \mathcal{M}. \quad (2.86)$$

Among the N different dosing schedules $\vec{d}^*(k)$, the one that maximizes the tumor-BE must be optimal. The fact that $\vec{d}^*(1)$ is this tumor-BE maximizing schedule, can be seen briefly as follows. We define the tumor-BE without proliferation of $\vec{d}^*(k)$ as

$$f^*(k) \triangleq \alpha_0 k (d^*(k)) + \beta_0 k (d^*(k))^2 = \alpha_0 k (\min_{m \in \mathcal{M}} b_m(k)) + \beta_0 k (\min_{m \in \mathcal{M}} b_m(k))^2 \quad (2.87)$$

$$= \min_{m \in \mathcal{M}} f_m(k), \quad (2.88)$$

where $f_m(k) \triangleq \alpha_0 k (b_m(k)) + \beta_0 k (b_m(k))^2$. Now, following our calculus based approach for the single normal tissue case in Section 2.7, we can see that $f_m(k)$ is a decreasing function of k for each $m \in \mathcal{M}$. Consequently, $f^*(k)$ is also a decreasing function of k . This shows that $\vec{d}^*(1)$ maximizes tumor-BE. That is, a single-dosage solution is optimal.

In the other case where $(\alpha_0/\beta_0) = \min_{m \in \mathcal{M}} (\alpha_m/\beta_m)/\sigma_m$, we again let this minimum equal $(\alpha_1/\beta_1)/\sigma_1$ without loss of generality. Then there are two further subcases: (a) $\beta_0 = \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2$, and (b) $\beta_0 > \sum_{m=1}^M \lambda_m \rho_m \sigma_m^2$. The proof of subcase (a) is similar to that of Case 2(a) in the proof of Proposition 2.5.1 and the proof of subcase (b) is similar to the proof of the first case above. In both scenarios, we can conclude that a single-dosage solution is optimal.

2.6 Optimal equal-dosage fractionation

We now focus on an equal-dosage variation of (OPTFRAC2). When the condition in Proposition 2.5.1 holds, this equal-dosage variation yields an optimal solution to (OPTFRAC2). When the condition in Proposition 2.5.4 holds, no further calculations are necessary as we know that it is optimal to administer a single fraction with dose given by (2.71). The equal-dosage variation is obtained by setting $d_1 = d_2 = \dots = d_N$

and calling this dose d . Thus, (OPTFRAC2) now reduces to the two-variable problem

$$(P) \max_{N,d} \alpha_0 N d + \beta_0 N d^2 - \tau(N) \quad (2.89)$$

$$N(\sigma_m d + \rho_m(\sigma_m^2 d^2)) \leq C_m, \quad m \in \mathcal{M}, \quad (2.90)$$

$$d \geq 0, \quad (2.91)$$

$$N \geq 1, \text{ integer}. \quad (2.92)$$

Because $\sigma_m d \geq 0$, the quadratic constraints (2.90) can be equivalently written as the linear constraints

$$d \leq b_m(N) \triangleq \frac{-1 + \sqrt{1 + 4\rho_m C_m / N}}{2\sigma_m \rho_m}, \quad m \in \mathcal{M}. \quad (2.93)$$

Moreover, since the objective function in (P) is increasing in d for each fixed $N \geq 1$, we can eliminate d from (P) by simply substituting the largest possible value of d that is allowed by these linear constraints. We denote this largest possible value of d by $d^*(N)$ and note that it is given by

$$d^*(N) \triangleq \min_{m \in \mathcal{M}} b_m(N). \quad (2.94)$$

For each $N \geq 1$, we call dose $d^*(N)$ and any corresponding normal tissue (that achieves the minimum in (2.94)) “limiting”. Now, with $d^*(N)$ as defined in (2.94) at our disposal, problem (P) simplifies to

$$(P) \max g^*(N) \triangleq \alpha_0 N(d^*(N)) + \beta_0 N(d^*(N))^2 - \tau(N) \quad (2.95)$$

$$N \geq 1, \text{ integer}. \quad (2.96)$$

In practice, there is an implicit upper bound N_{max} on the number of fractions that can be administered based on logistical considerations [45]. Consequently, the optimal number of fractions can be found by evaluating $g^*(N)$ defined in (2.95) for every integer $1 \leq N \leq N_{max}$ and then picking the best. In Theorem 2.6.3 in the next section, we go beyond this brute-force yet practical approach and perform a careful mathematical analysis of the structure of the dose $d^*(N)$ in (2.94) and the corresponding biological effect $g^*(N)$ in (2.95). This provides additional insights into problem (P) and leads to a more streamlined procedure for solving (2.95)-(2.96).

2.6.1 *Characterization of the limiting dose and the optimal tumor-BE as functions of the number of fractions*

Lemma 2.6.1. *The limiting dose $d^*(N)$ defined in (2.94) is strictly decreasing in N . When viewed as a function of all real numbers $N \geq 1$, $d^*(N)$ is piecewise strictly convex. See Figure 2.1.*

Proof. Observe that dose $b_m(N)$ as defined in (2.93) is a strictly decreasing functions of N for each $m \in \mathcal{M}$. Since $d^*(N)$ is the smallest among these doses, it is also strictly decreasing. Now we view $b_m(N)$, for $m \in \mathcal{M}$, as functions of all real numbers $N \geq 1$ and show that they are strictly convex. Thus, it suffices to show that $a_m(N) \triangleq \sqrt{1 + 4\rho_m C_m/N}$ is strictly convex for each $m \in \mathcal{M}$. After some algebraic simplification, we obtain,

$$\frac{da_m(N)}{dN} = \frac{-2\rho_m C_m}{\sqrt{N^4 + 4\rho_m C_m N^3}}. \quad (2.97)$$

As a result, we get,

$$\frac{d^2 a_m(N)}{dN^2} = \frac{\rho_m C_m (4N^3 + 12\rho_m C_m N^2)}{(N^4 + 4\rho_m C_m N^3)^{3/2}}. \quad (2.98)$$

The above second derivative is positive, and hence, $a_m(N)$ is strictly convex. If the dose-limiting normal tissue is invariant over all real numbers $N \geq 1$, then $d^*(N)$ is strictly convex (and hence trivially piecewise strictly convex). So we suppose that this is not the case. Then, owing to the continuity of doses $b_m(N)$ in N , there exist real numbers $1 < i_1 < i_2 < i_3, \dots$ such that the dose-limiting normal tissues do not change within the intervals $[1, i_1), [i_1, i_2), [i_2, i_3), \dots$. Furthermore, in each of these intervals, the limiting dose $d^*(N)$ equals one of the strictly convex functions $b_m(N)$. Therefore, $d^*(N)$ is piecewise strictly convex. \square

The rate at which $d^*(N)$ decreases affects the structure of $g^*(N)$. Our next lemma puts an N -dependent upper bound on $d^*(N)$.

Lemma 2.6.2. *There exists a positive bounding constant \bar{B} such that $d^*(N) \leq \bar{B}/N$ for all $N \geq 1$.*




Figure 2.1: A schematic illustration of the structure of the limiting dose $d^*(N)$:(a) for a single normal tissue, the limiting dose is strictly decreasing and strictly convex; (b) with multiple normal tissues, it is strictly decreasing and piecewise strictly convex; the kinks at fractions $N1$ and $N2$ represent switching of the limiting normal tissues. Also see Figure 2.2(a).

Proof. We show that

$$\bar{B} = \max_{m \in \mathcal{M}} (C_m / \sigma_m). \quad (2.99)$$

works. To see this, first consider any $m \in \mathcal{M}$. We have,

$$1 + 4\rho_m C_m / N \leq 1 + 4\sigma_m \rho_m \bar{B} / N \leq 1 + 4\sigma_m \rho_m \bar{B} / N + 4\sigma_m^2 \rho_m^2 \bar{B}^2 / N^2. \quad (2.100)$$

Here the first inequality follows from the definition of \bar{B} in (2.99) and the second inequality follows because $4\sigma_m^2 \rho_m^2 \bar{B}^2 / N^2 > 0$. Then, by taking square roots of both sides, we obtain

$$\sqrt{1 + 4\rho_m C_m / N} \leq 1 + 2\sigma_m \rho_m \bar{B} / N, \quad (2.101)$$

or equivalently,

$$\frac{-1 + \sqrt{1 + 4\rho_m C_m / N}}{2\sigma_m \rho_m} = b_m(N) \leq \frac{\bar{B}}{N}. \quad (2.102)$$

This shows from Equation (2.94) that $d^*(N) \leq \bar{B} / N$ for all N . \square

The next theorem is a key result of this Chapter as it characterizes the structure of the optimal tumor-BE $g^*(N)$.

Theorem 2.6.3. *The biological effect $g^*(N)$ in the objective function (2.95) of (P) has the following properties:*

Property 1. *there exists an $\hat{N} \geq 2$ such that $g^*(\hat{N} - 1) > g^*(\hat{N})$;*

Property 2. *$g^*(N)$ is either (i) nonincreasing in N for all integers $N \geq 1$ or (ii) nondecreasing in N up to a positive integer and then nonincreasing after this positive integer; and*

Property 3. *suppose $\hat{N} \geq 2$ is the smallest fraction N with the property that $g^*(N - 1) > g^*(N)$ (note that such an \hat{N} exists by Property 1), then $\hat{N} - 1$ is the optimal number of fractions.*

Proof of Property 1. To establish the existence of an \hat{N} with the said property, it suffices to focus on $N \geq 1 + T_{lag}$.

We consider two cases: the first case is where there exists an $\hat{N} \geq 1 + T_{lag}$ such that $g^*(\hat{N}) + \tau(\hat{N}) \leq g^*(\hat{N} - 1) + \tau(\hat{N} - 1)$; the second case is where such an \hat{N} does not exist.

In the first case, we have,

$$g^*(\hat{N}) + \tau(\hat{N}) \leq g^*(\hat{N} - 1) + \tau(\hat{N} - 1), \quad (2.103)$$

$$\tau(\hat{N} - 1) < \tau(\hat{N}), \quad (2.104)$$

where the second inequality follows because $\tau(N)$ is strictly increasing for $N \geq 1 + T_{lag}$. Adding inequalities (2.103) and (2.104) yields that $g^*(\hat{N}) < g^*(\hat{N} - 1)$ as required.

The second case is more difficult and we resort to an asymptotic analysis over N to establish our claim. We have that $g^*(N) + \tau(N) > g^*(N - 1) + \tau(N - 1)$ for all $N \geq 1 + T_{lag}$. We use $z^*(N)$ to denote $g^*(N) + \tau(N)$; thus, $z^*(N)$ is a strictly increasing sequence for all $N \geq 1 + T_{lag}$.

From Lemma 2.6.2, we have that

$$z^*(N) = g^*(N) + \tau(N) = \alpha_0 N d^*(N) + \beta_0 N (d^*(N))^2 \leq \alpha_0 \bar{B} + \beta_0 \frac{\bar{B}^2}{N} \leq \alpha_0 \bar{B} + \beta_0 \bar{B}^2, \quad (2.105)$$

for all $N \geq 1$. That is, $z^*(N)$ is a bounded sequence. Since $z^*(N)$ is strictly increasing, it must then converge. As a result, there exists an \bar{N} large enough such that

$$z^*(N) - z^*(N-1) < \ln 2 / T_{double}, \quad \forall N \geq \bar{N}. \quad (2.106)$$

For all $N \geq 2 + T_{lag}$, the definition of $\tau(N)$ in (2.2) implies that

$$\tau(N-1) = \frac{((N-2) - T_{lag}) \ln 2}{T_{double}}, \quad (2.107)$$

$$\tau(N) = \frac{((N-1) - T_{lag}) \ln 2}{T_{double}}, \quad \text{and}, \quad (2.108)$$

$$\tau(N) - \tau(N-1) = \frac{\ln 2}{T_{double}}. \quad (2.109)$$

Let $\bar{\bar{N}} = \max\{\bar{N}, 2 + T_{lag}\}$. Then, (2.106) and (2.109) imply that $z^*(N) - z^*(N-1) < \tau(N) - \tau(N-1)$ for all $N \geq \bar{\bar{N}}$. Then using the fact that $z^*(N) = g^*(N) + \tau(N)$ and $z^*(N-1) = g^*(N-1) + \tau(N-1)$, we get,

$$g^*(N) < g^*(N-1), \quad \forall N \geq \bar{\bar{N}}. \quad (2.110)$$

Then setting $\hat{N} = \bar{\bar{N}}$, we see that $g^*(\hat{N}) < g^*(\hat{N}-1)$ as required.

Proof of Property 2. We define the tumor-BE corresponding to dose $b_m(N)$ per fraction over N fractions as

$$g_m(N) \triangleq \alpha_0 N b_m(N) + \beta_0 N b_m^2(N) - \tau(N), \quad m \in \mathcal{M}. \quad (2.111)$$

In this proof of Property 2, we view $d^*(N)$ in (2.94), the above tumor-BE, as well as $g^*(N)$ in (2.95) as functions of all real numbers $N \geq 1$. Based on the definition of $d^*(N)$ in (2.94) and also of $g^*(N)$ in (2.95), we observe that

$$g^*(N) = \min_{m \in \mathcal{M}} g_m(N). \quad (2.112)$$

We will show that $g_m(N)$ is quasiconcave in N over $N \geq 1$ for each fixed $m \in \mathcal{M}$. This would imply that $g^*(N)$ is also quasiconcave because it is the minimum of a collection of quasiconcave functions (see Section 3.4.4 in [64]). According to Section 3.4.4 in [64] there are only three possibilities for the structure of a quasiconcave function such as $g^*(N)$: it is nondecreasing; or it is nonincreasing; or it is nondecreasing first and then nonincreasing. The first possibility is not feasible in view of Property 1. This would prove Property 2.

Consider any fixed $m \in \mathcal{M}$. To show that $g_m(N)$ is quasiconcave, we consider the function

$$f(N) \triangleq \alpha_0 N b_m(N) + \beta_0 N b_m^2(N). \quad (2.113)$$

For algebraic simplicity, we define $D(N) \triangleq N b_m(N)$ as the total tumor-dose delivered in N fractions when the tumor-dose per fraction equals $b_m(N)$. For all $N > 0$, we note that

$$\frac{df(N)}{dN} = \alpha_0 \frac{dD(N)}{dN} + \beta_0 \frac{2ND(N) \frac{dD(N)}{dN} - D^2(N)}{N^2} \quad (2.114)$$

$$= \alpha_0 \frac{dD(N)}{dN} + 2\beta_0 b_m(N) \frac{dD(N)}{dN} - \beta_0 b_m^2(N). \quad (2.115)$$

We have,

$$\frac{dD(N)}{dN} = \frac{1}{2\sigma_m \rho_m} \left(-1 + \frac{2N + 4\rho_m C_m}{2\sqrt{N^2 + 4\rho_m C_m N}} \right) \quad (2.116)$$

$$= \frac{1}{2\sigma_m \rho_m} \frac{(\sqrt{1 + 4\rho_m C_m/N} - 1)^2}{2\sqrt{1 + 4\rho_m \rho_m/N}} \quad (2.117)$$

for all $N > 0$. We use the notation $r_m \triangleq 1/(2\sigma_m \rho_m)$ and also $\chi(N) \triangleq \sqrt{1 + 4\rho_m C_m/N} - 1$ for brevity. With this notation, we have $b_m(N) = r_m \chi(N)$ and $dD(N)/dN = r_m \chi^2(N)/(2(\chi(N) + 1))$. Then substituting these two into Equation (2.114) we obtain,

$$\frac{df(N)}{dN} = \alpha_0 \frac{r_m \chi^2(N)}{2(\chi(N) + 1)} + 2\beta_0 \frac{r_m^2 \chi^3(N)}{2(\chi(N) + 1)} - \beta_0 r_m^2 \chi^2(N) = \frac{(\alpha_0 - 2\beta_0 r_m) r_m \chi^2(N)}{2(\chi(N) + 1)} \quad (2.118)$$

$$= \frac{(\alpha_0 - \alpha_m \beta_0 / \sigma_m \beta_m) r_m \chi^2(N)}{2(\chi(N) + 1)}. \quad (2.119)$$

We now consider two cases. The first case is where $\sigma_m \alpha_0 / \beta_0 \leq \alpha_m / \beta_m$ and the second case is when $\sigma_m \alpha_0 / \beta_0 > \alpha_m / \beta_m$.

Case 1: $\sigma_m \alpha_0 / \beta_0 \leq \alpha_m / \beta_m$. In this case, (2.119) implies that $df/dN \leq 0$ for all $N > 0$. That is, $f(N)$ is nonincreasing for all $N > 0$, and in particular, for all $N \geq 1$. Since $\tau(N)$ is nondecreasing, this implies that the objective function $g_m(N) = f(N) - \tau(N)$ is nonincreasing over $N \geq 1$.

Case 2: $\sigma_m \alpha_0 / \beta_0 > \alpha_m / \beta_m$. Now, from (2.119), for all $N > 0$, we have,

$$\frac{d^2 f(N)}{dN^2} = \frac{d}{dN} \left(\frac{df(N)}{dN} \right) = \frac{d}{d\chi} \left(\frac{df(N)}{dN} \right) \frac{d\chi(N)}{dN} \quad (2.120)$$

$$= \frac{(\alpha_0 - \alpha_m \beta_0 / \sigma_m \beta_m) r_m}{2} \frac{2\chi(N)(\chi(N) + 1) - \chi^2(N)}{(\chi(N) + 1)^2} \frac{d\chi(N)}{dN} \quad (2.121)$$

$$= \frac{(\alpha_0 - \alpha_m \beta_0 / \sigma_m \beta_m) r_m}{2} \frac{\chi^2(N) + 2\chi(N)}{(\chi(N) + 1)^2} \frac{d\chi(N)}{dN}. \quad (2.122)$$

Now recall that $\chi(N) = \sqrt{1 + 4C_m \rho_m / N} - 1$ and hence $\chi(N) > 0$ for all $N > 0$. Moreover, $d\chi(N)/dN < 0$ for all $N > 0$. Therefore, $d^2 f(N)/dN^2 < 0$ and hence $f(N)$ is strictly concave for all $N > 0$ (a proof of strict concavity based on a second derivative was also developed independently by Bortfeld et al. [62] in their work on a single normal tissue; we nevertheless decided to present our derivatives using our notation here for the sake of completeness). We now have,

$$\tau(N) = \frac{[(N - 1) - T_{lag}]^+ \ln 2}{T_{double}} = \frac{\max((N - 1) - T_{lag}, 0) \ln 2}{T_{double}}, \quad (2.123)$$

and hence $\tau(N)$ is a convex function. As a result, $g_m(N) = f(N) - \tau(N)$ is a concave function over $N \geq 1$.

The above two cases show that $g_m(N)$ is either nonincreasing or concave and hence it is quasiconcave as claimed. This completes the proof of Property 2.

Proof of Property 3. This property is an immediate consequence of Property 2.

The Theorem 2.6.3 motivates the following streamlined procedure for solving problem (P).

A calculation procedure for solving (P)

Step 1. Set $N = 1$.

Step 2. Compute the optimal dose per fraction $d^*(N)$ using Equation (2.94). Use this $d^*(N)$ to compute the biological effect $g^*(N)$ as in (2.95).

Step 3. If $N = 1$, then set $N = 2$ and go to Step 2 above.

Step 4. If $N > 1$ then STOP if $g^*(N - 1) > g^*(N)$, report $N - 1$ as the optimal number of fractions, $d^*(N - 1)$ as the optimal average tumor-dose per fraction, and finally, $u_{nominal}d^*(N - 1)/d_{nominal}$ as the fluence-map as per the linear dose deposition model; otherwise, set $N = N + 1$ and go to Step 2 above.

The following corollary of Theorem 2.6.3 establishes the correctness of this procedure.

Corollary 2.6.4. *The stopping condition in Step 4 of the above procedure is eventually met; that is, the procedure terminates finitely. Moreover, it correctly identifies the optimal number of fractions in (P).*

Proof. The first claim follows from Property 1 in Theorem 2.6.3. The second claim follows from Property 3 in Theorem 2.6.3. □

We acknowledge the possibility that neither the condition in Proposition 2.5.1 nor that in Proposition 2.5.4 holds; that is, $\min_{m \in \mathcal{M}} \frac{\alpha_m/\beta_m}{\sigma_m} < \alpha_0/\beta_0 < \max_{m \in \mathcal{M}} \frac{\alpha_m/\beta_m}{\sigma_m}$. We provide an example here.

$$\max d_1 + 0.2d_1^2 + d_2 + 0.2d_2^2 \tag{2.124}$$

$$d_1 + 0.1666d_1^2 + d_2 + 0.1666d_2^2 \leq 44.8762 \tag{2.125}$$

$$d_1 + 0.3571d_1^2 + d_2 + 0.3571d_2^2 \leq 79.5918 \tag{2.126}$$

$$d_1, d_2 \geq 0. \tag{2.127}$$

It turns out that this problem has exactly two optimal solutions: the first where $d_1 \approx 1.04, d_2 \approx 13.46$ and the second where $d_1 \approx 13.46, d_2 \approx 1.04$. Thus, there is

no equal-dosage solution that is optimal. Also, it is not optimal to set one dose to zero and the other to a positive value. Note that this scenario cannot arise when there is a single normal tissue as in the existing literature. In these situations, our simple procedure for obtaining an optimal solution to the equal-dosage formulation (P) yields an approximate solution to (OPTFRAC2). In the next Chapter we propose a two-variable linear program which can solve this fractionation problem, even when both conditions are not satisfied, to optimality.

2.7 The special case of a single normal tissue

In this Section, we study a special case of our general model in the main body of this chapter; it includes a single normal tissue. In this special case, we are able to derive a closed-form formula for the optimal number of fractions. The formula is valid for a normal tissue with maximum dose, mean dose, or dose-volume type constraints. The formula is consistent with other similar closed-form formulas in the existing literature (reviewed in Section 2.2 of the main body of this Chapter), after accounting for minor idiosyncratic differences between some of the existing single normal tissue models and our single normal tissue model.

Consider a special case of problem (P) in the main text of this Chapter with a single normal tissue. Thus, \mathcal{M} now equals the singleton set $\{1\}$ and hence we denote various normal tissue parameters with a subscript of “1”. Then, problem (P) reduces to

$$(P_{\text{single}}) \max \alpha_0 N b_1(N) + \beta_0 N b_1^2(N) - \tau(N) \quad (2.128)$$

$$N \geq 1, \text{ integer}, \quad (2.129)$$

where $b_1(N)$ is now the limiting dose defined as $(-1 + \sqrt{1 + 4\rho_1 C_1/N})/(2\sigma_1\rho_1)$ from (2.93). We reiterate that problem (P_{single}) simultaneously, although in some cases only slightly, generalizes the models studied by Fowler and his co-authors, Jones et al., Unkelbach et al., Keller et al., and Bortfeld et al. We next state and prove one lemma that provides an optimal fractionation schedule for problem (P_{single}). This result is

not among the main contributions of this Chapter. We include it here mainly for the sake of completeness, notational consistency, and to ensure that our slightly general result correctly recovers a previously known formula in the literature. Moreover, our formula here is used in Section 2.8.3 of the main body of this Chapter while comparing the optimal number of fractions obtained by separately considering different normal tissues one by one with those from our multiple normal tissue model. Before we state our formula, we emphasize here that if $\alpha_0/\beta_0 \leq \frac{\alpha_1/\beta_1}{\sigma_1}$, then Proposition 2.5.4 implies that it is optimal to administer a single fraction. We therefore focus here on the only other possibility for these parameters of our single normal tissue model, namely, wherein $\alpha_0/\beta_0 > \frac{\alpha_1/\beta_1}{\sigma_1}$.

Lemma 2.7.1. *Suppose that $\alpha_0/\beta_0 > \frac{\alpha_1/\beta_1}{\sigma_1}$. Then it is not optimal to administer less than $1 + T_{lag}$ fractions. Let $r_1 \triangleq 1/(2s_{effective}\rho_1)$ and $\eta \triangleq \ln 2/T_{double}$. We also define*

$$N^* \triangleq \frac{4\rho_1 C_1}{\left(\frac{\eta + \sqrt{\eta^2 + 2\eta r_1(\alpha_0 - \alpha_1\beta_0/\sigma_1\beta_1)}}{r_1(\alpha_0 - \alpha_1\beta_0/\sigma_1\beta_1)} + 1\right)^2 - 1}. \quad (2.130)$$

If $N^ < 1 + T_{lag}$, then it is optimal to administer $1 + T_{lag}$ fractions; otherwise it is optimal to administer either $\lfloor N^* \rfloor$ or $\lceil N^* \rceil$ fractions depending on which one of these two integers maximizes the objective function in (P_{single}) .*

Proof. Since $\tau(N) = 0$ for all $1 \leq N \leq 1 + T_{lag}$, the objective function in problem (P_{single}) equals $f(N) \triangleq \alpha_0 N b_1(N) + \beta_0 N b_1^2(N)$ over this range of N . Then with algebra identical to the proof of Property 2 in Theorem 2.6.3, we obtain,

$$\frac{df(N)}{dN} = \frac{(\alpha_0 - \alpha_1\beta_0/\sigma_1\beta_1)r_1\chi^2(N)}{2(\chi(N) + 1)}, \quad (2.131)$$

where $\chi(N) = -1 + \sqrt{1 + 4\rho_1 C_1/N}$. Since $\sigma_1\alpha_0/\beta_0 > \alpha_1/\beta_1$, and $\chi(N) + 1 = \sqrt{1 + 4\rho_1 C_1/N} > 0$, the above derivative is positive. This implies that $f(N)$ is strictly increasing over $1 \leq N < 1 + T_{lag}$. Thus, it is not optimal to administer less than $1 + T_{lag}$ fractions. We investigate the objective function in (P_{single}) for $N \geq 1 + T_{lag}$. Let $\delta = \frac{(1+T_{lag})\ln 2}{T_{double}}$. Then, the objective function is given by

$$g^*(N) \triangleq f(N) - \eta N + \delta, \text{ for } N \geq 1 + T_{lag}, \quad (2.132)$$

where $f(N) = \alpha_0 N b_1(N) + \beta_0 N b_1^2(N)$. Then, similar to (2.131), we get

$$\frac{dg^*(N)}{dN} = \frac{(\alpha_0 - \alpha_1 \beta_0 / \sigma_1 \beta_1) r_1 \chi^2(N)}{2(\chi(N) + 1)} - \eta. \quad (2.133)$$

Also, similar to (2.122), we get,

$$\frac{d^2 g^*(N)}{dN^2} = \frac{(\alpha_0 - \alpha_1 \beta_0 / \sigma_1 \beta_1) r_1}{2} \frac{\chi^2(N) + 2\chi(N)}{(\chi(N) + 1)^2} \frac{d\chi(N)}{dN}. \quad (2.134)$$

Now recall that $\chi(N) = \sqrt{1 + 4\rho_1 C_1 / N} - 1$ and hence $\chi(N) > 0$. Moreover, $d\chi(N)/dN < 0$. Therefore, $d^2 g^*(N)/dN^2 < 0$ and hence $g^*(N)$ is concave. Moreover, by Property 1 in Theorem 2.6.3, we know that $g^*(N)$ has a stationary point. If $g^*(N)$ has a stationary point that is at least $1 + T_{lag}$, then this stationary point maximizes the objective function in our problem over all real numbers $N \geq 1 + T_{lag}$. If, on the other hand, $g^*(N)$ does not have such a stationary point (that is, the stationary point is smaller than $1 + T_{lag}$), then $N = 1 + T_{lag}$ maximizes the objective function in our problem over all real numbers $N \geq 1 + T_{lag}$. Thus, we need to first find a stationary point of $g^*(N)$. In particular, (2.133) implies that we need to solve the quadratic equation

$$r_1(\alpha_0 - \alpha_1 \beta_0 / \sigma_1 \beta_1) \chi^2(N) - 2\eta \chi(N) - 2\eta = 0 \quad (2.135)$$

for $\chi(N)$. Since $\chi(N)$ must be positive, we investigate whether the above equation has a positive solution. Solutions of this equation are given by

$$\chi(N) = \frac{\eta \pm \sqrt{\eta^2 + 2\eta r_1(\alpha_0 - \alpha_1 \beta_0 / \sigma_1 \beta_1)}}{r_1(\alpha_0 - \alpha_1 \beta_0 / \sigma_1 \beta_1)}, \quad (2.136)$$

and since $\sigma_1 \alpha_0 / \beta_0 > \alpha_1 / \beta_1$,

$$\chi^*(N) = \frac{\eta + \sqrt{\eta^2 + 2\eta r_1(\alpha_0 - \alpha_1 \beta_0 / \sigma_1 \beta_1)}}{r_1(\alpha_0 - \alpha_1 \beta_0 / \sigma_1 \beta_1)} \quad (2.137)$$

is the unique positive solution. Moreover, since $\chi^*(N) = \sqrt{1 + 4\rho_1 C_1 / N^*} - 1$, where N^* is the unique stationary point of $g^*(N)$, Equation (2.137) implies that

$$N^* = \frac{4\rho_1 C_1}{\left(\frac{\eta + \sqrt{\eta^2 + 2\eta r_1(\alpha_0 - \alpha_1 \beta_0 / \sigma_1 \beta_1)}}{r_1(\alpha_0 - \alpha_1 \beta_0 / \sigma_1 \beta_1)} + 1 \right)^2 - 1}. \quad (2.138)$$

Finally, because $g^*(N)$ is concave, and the optimal number of fractions must be an integer, we substitute $N = \lfloor N^* \rfloor$ and $N = \lceil N^* \rceil$ into $g^*(N)$ and find which one of these two yields the higher biological effect. That is the optimal number of fractions. \square

Recall from Section 6.1 that Jones et al. did not prove that setting the derivative of biological effect to zero was sufficient for global optimality in their model. Moreover, they ignored the integer restriction on N . Lemma 2.7.1 overcomes these limitations. Also recall that the optimality of an infinite number of fractions when $\sigma_1\alpha_0/\beta_0 > \alpha_1/\beta_1$ was proven in Unkelbach et al. and Keller et al. However, those papers did not model tumor proliferation. In a special case of our model where there is no tumor proliferation, we have, $\tau(N) = 0$ for all N ; thus, the objective function in problem (P_{single}) is given simply by $f(N)$ and the above proof shows that it is strictly increasing over N for all N when $\sigma_1\alpha_0/\beta_0 > \alpha_1/\beta_1$. That is, an infinite number of fractions is optimal and hence Lemma 2.7.1 recovers the result in Unkelbach et al. and Keller et al. Bortfeld et al. also independently proved strict concavity of tumor-BE using a second derivative similar to ours. One very minor difference between our result here and that of Bortfeld et al. is that we allow for a time-lag before proliferation begins; setting $T_{lag} = 0$ thus correctly recovers their result.

2.8 Numerical results

We are now ready to present numerical experiments and sensitivity analyses. We first describe in detail the test cases that we used in our experiments. Then, in Section 2.8.2, we demonstrate that the limiting normal tissue can switch as the number of fractions is varied and that, consequently, tumor-BE without proliferation may not be monotonic. In Section 2.8.3, we illustrate that the optimal number of fractions cannot be found by solving separate optimization problems for different normal tissues one by one. Then in Section 2.8.4, we investigate the sensitivity of the optimal number of fractions and the optimal tumor-BE to T_{double} . We find that the optimal number of fractions is smaller for smaller values of T_{double} , which represent faster growing

tumors. Also, the optimal tumor-BE itself is smaller for these faster growing tumors. We demonstrate that the optimal number of fractions can be more sensitive to T_{double} when the tumor alpha over beta ratio is larger relative to that of the normal tissues. We also find that for each fixed value of T_{double} , and especially when this value is small, the range for optimal number of fractions over all our head-and-neck cases and for a wide range of other problem parameters is quite narrow. In Section 2.8.5, we study the sensitivity of the optimal number of fractions to T_{lag} . We find that it is insensitive to T_{lag} other than the fact that $1 + T_{lag}$ provides a lower bound on the optimal number of fractions as in Remark 2.5.2 under the sufficient condition in Proposition 2.5.1. We comment that to avoid overcrowding in our figures and in the interest of brevity, we have included below plots for a single representative case; we believe that this is sufficient because the qualitative trends in all cases were similar.

2.8.1 Description of test cases

In this section, we first describe the ten test cases that were used in our numerical experiments. Five of these were head-and-neck cases and the other five were prostate. All test cases were generated using our in-house software Phantom Creator (PhanC) written in MATLAB (see Appendix A). These test cases were three-dimensional and were carefully developed to be representative of clinical scenarios in terms of geometry and size. All cases used equally spaced coplanar beams and the beamlet resolution was $5 \times 5 \text{ mm}^2$. All voxels were $3 \times 3 \times 3 \text{ mm}^3$. All computer simulations were carried out on a 3.1 GHz iMac desktop with 16 GB RAM using MATLAB.

Head-and-neck cancer cases

All cases used seven beams and included spinal cord, brainstem, left and right parotids and unspecified normal tissue between these critical organs. The total number of voxels in the head-and-neck target and in the normal tissues, and the total number of beamlets is shown in Table 5.1 below.

case #	# of beamlets (k)	# of tumor voxels (n)	# of normal tissues voxels
1	3910	27576	67386
2	3888	31930	67270
3	4128	36320	76160
4	3003	22372	53176
5	3256	28638	64713

Table 2.1: Description of the geometry used in head-and-neck cancer cases.

The conventional fractionation schedule was assumed to include $N_{conv} = 35$ fractions. While formulating problem (P), we included maximum dose constrains for spinal cord, brainstem and unspecified normal tissue. A dose-volume constraint for unspecified normal tissue was also added. Mean dose constrains were used for left and right parotids. The tolerance dose values for various normal tissues were similar to [1, 70–72] and are listed in Table 5.2 below.

Normal tissue	D_{max} (Gy)	D_{mean} (Gy)	D_{dv} (Gy), ϕ
Spinal cord	45	N/A	N/A
Brainstem	50	N/A	N/A
Left and right parotids	N/A	28	N/A
Unspecified normal tissue	77	N/A	70, 0.05

Table 2.2: Tolerance doses for various normal tissues in our head-and-neck test cases where the dose is administered in $N_{conv} = 35$ equal-dose fractions. Recall that for dose-volume type constraints no more than a volume fraction ϕ of the normal tissue can receive dose more than D_{dv} .

We needed a nominal fluence-map $u_{nominal}$ as a starting point to calculate the sparing factors for different normal tissues as explained in Section 4.2.2. In practice, a

spatially optimized fluence-map that conforms to an individual hospital’s treatment protocol would be readily available from a commercial treatment planning system and could be used for this purpose. We obtained $u_{nominal}$ by solving a standard fluence-map optimization problem where the objective was to minimize the total squared-deviation of tumor voxel doses from a prescription-dose of 70 Gy subject to maximum dose constraints for spinal cord and brainstem and mean dose constraints for parotids with dose tolerance levels as in Table 5.2 (see, for example, [9]). To ensure that the intensity profile would be deliverable in practice using a multi-leaf collimator, we put linear smoothness constraints on each radiation field. These smoothness constraints ensured that the absolute percentage difference between intensity values of two adjacent beamlets was smaller than a threshold. The resulting optimization problem was convex and it was solved using the MATLAB toolbox CVX [73].

Our parameter values were based on [38, 44, 47–49, 51, 65, 74, 75]. Head-and-neck tumor was assumed to have $\alpha = 0.35 \text{ Gy}^{-1}$. For unspecified normal tissue, α/β value was fixed at 3 Gy. Our model had four remaining groups of parameters: α/β ratio for tumor; α/β ratios for spinal cord, brainstem, left and right parotids; and T_{double} and T_{lag} for tumor. We performed sensitivity analyses with several values for these parameters. The α/β ratio for head-and-neck tumor was selected from the set $\{8, 10, 12\}$. The α/β ratios for left and right parotids were assumed to be equal and the value of this ratio was selected from the set $\{3, 4, 5, 6\}$. The α/β ratios for spinal cord and brainstem were chosen from the set $\{2, 3, 4, 5, 6\}$. The doubling time T_{double} for head and neck was selected from the set $\{2, 3, 5, 8, 10, 20, 40, 50\}$; and T_{lag} was selected from the set $\{7, 14, 21, 28, 35\}$.

Prostate cancer cases

All cases used five beams and included rectum, bladder, left and right femurs and unspecified normal tissue between these critical organs. The total number of voxels in the prostate target and in the normal tissues, and the total number of beamlets is

case #	# of beamlets (k)	# of tumor voxels (n)	Number of normal tissues voxels
1	938	6180	145545
2	847	7225	326703
3	935	4628	367656
4	930	4956	314544
5	870	4840	269450

Table 2.3: Description of the geometry used in prostate cancer cases.

shown in Table 4.5 below.

The conventional fractionation schedule was assumed to include $N_{conv} = 45$ fractions. In our formulation of problem (P), we included a maximum dose constraint for unspecified normal tissue, and included dose-volume constraints for all normal tissues. There were no mean dose constraints. The dose-volume constraints for all normal tissues were similar to [1, 70–72] and are listed in Table 4.6.

We needed a nominal fluence-map $u_{nominal}$ to calculate sparing factors for different normal tissues. As in the head-and-neck cases, $u_{nominal}$ was obtained by solving a standard fluence-map optimization problem where the goal was to minimize the total squared-deviation of tumor voxel doses from a prescription-dose of 81 Gy with maximum dose constraints of 85, 89 and 65 Gy for rectum, bladder, and femurs, respectively. Smoothness constraints were also included as explained in the head-and-neck case. Again, the resulting optimization problem was convex and it was solved using the MATLAB toolbox CVX [73].

Here, our parameter values were based on [41, 45, 68, 69, 71, 76–79]. Prostate tumor was assumed to have $\alpha = 0.15 \text{ Gy}^{-1}$. For unspecified normal tissue, α/β value was fixed at 3 Gy. After fixing these values, our model had four remaining groups of parameters: α/β ratio for tumor; α/β ratios for rectum, bladder, and femurs; and T_{double} and T_{lag} for tumor. We performed sensitivity analyses with several values for

Normal tissue	D_{dv} (Gy), ϕ
Rectum	50, 0.5
	85, 0.05
Bladder	50, 0.5
	70, 0.35
	75, 0.25
	80, 0.15
	89, 0.05
Femurs	65, 0.05
Unspecified normal tissue	81, 0.05

Table 2.4: Dose-volume constraints for various normal tissues when dose is administered in $N_{conv} = 45$ equal-dose fractions — no more than a volume fraction ϕ of the normal tissue can receive dose more than D_{dv} .

these parameters. The α/β ratio for prostate was selected from the set $\{2, 3, 4, 6\}$. The α/β ratios for left and right femurs were assumed to be equal and the value of this ratio was selected from the set $\{3, 4, 5, 6\}$. The α/β ratios for rectum and bladder were chosen to be equal and from the set $\{3, 4, 5, 6\}$. The doubling time T_{double} for prostate was selected from the set $\{5, 20, 40, 60, 80\}$; and T_{lag} was selected from the set $\{7, 14, 21, 28, 35\}$.

2.8.2 *Switching of the limiting normal tissues and nonmonotonic behavior of the tumor-BE without proliferation*

As stated in Section 6.1, the dose-limiting normal tissue of course does not change with the number of fractions when the optimization model includes only one normal tissue to begin with. We demonstrate by example here that this is no longer the case when the optimization model includes multiple normal tissues. In particular, Figure 2.2(a) plots the limiting doses corresponding to individual normal tissues as functions

of N . The overall limiting dose itself is the minimum of such individual limiting doses. The figure shows that the limiting normal tissue switches from the unspecified normal tissue to the right parotid as the number of fractions is varied from 35 to 45 in this case. The qualitative trend in this figure is similar to our schematic in Figure 2.1(b) and is consistent with Lemma 2.6.1. This type of switching between limiting normal tissues is essentially why it is not possible to develop a closed-form formula for the optimal number of fractions with multiple normal tissues. A related point is that, as stated in Section 6.1, the tumor-BE without proliferation is either increasing or decreasing in N when the problem formulation includes only one normal tissue. We demonstrate in Figure 2.2(b) that this is no longer the case when the optimization model includes multiple normal tissues. In fact, the figure illustrates that whether or not the tumor-BE without proliferation is monotonic depends on problem parameters. Moreover, since some believe [68, 69] that the α/β ratio for prostate tumors could be quite low, and in particular, it could be smaller than that of nearby normal tissues, there are even more possibilities for the qualitative trend in prostate tumor-BE as a function of the number of fractions. These are illustrated in Figures 2.3(a), (b).

2.8.3 Comparison with results obtained by separately considering each normal tissue one-by-one

Since a closed-form formula for the optimal number of fractions is available for the single normal tissue case (see for example Jones et al. [46], Armpilia et al. [50], Bortfeld et al. [62] and other references therein, as well as our Section 2.7), the question arises as to whether the optimal number of fractions in the multiple normal tissue case can be obtained somehow by separately applying the closed-form formula to each normal tissue one by one. We show in Table 2.5 that this is not the case. In particular, the columns corresponding to different values of T_{double} in Table 2.5 with parotid $\alpha/\beta = 3$ Gy show that there is no single normal tissue that can always correctly predict the optimal number of fractions obtained from our model. For




Figure 2.2: This figure was generated for our head-and-neck case 1 with tumor $\alpha/\beta = 8$ Gy. The α/β ratios for all normal tissues were fixed at 3 Gy, except for the parotids. (a) Switching of the limiting normal tissue with varying number of fractions; the α/β ratio for the parotids was fixed at 6 Gy; (b) Monotonic and nonmonotonic behavior of the tumor BE without proliferation with varying number of fractions for different values of the α/β ratio for the parotids. Also note that the kink in the tumor-BE without proliferation occurs exactly where there is a switch in the limiting normal tissue.

instance, when $T_{double} = 2, 3, 5$ days, unspecified normal tissue correctly predicts the number of fractions from our model; when $T_{double} = 8$ days, brainstem best-predicts the number of fractions from our model; when $T_{double} = 10$ days, spinal cord is the best predictor; and finally, when $T_{double} = 20, 40, 50$ days, parotids are the best predictors. Moreover, in some cases, the optimal number of fractions from our model does not equal any of the numbers obtained from the single normal tissue model; for instance, see the column of $T_{double} = 20$ with parotid $\alpha/\beta = 4$ Gy, and the columns of $T_{double} = 20, 40, 50$ with parotid $\alpha/\beta = 5$ Gy.

The intuition behind why models with a single normal tissue cannot predict the

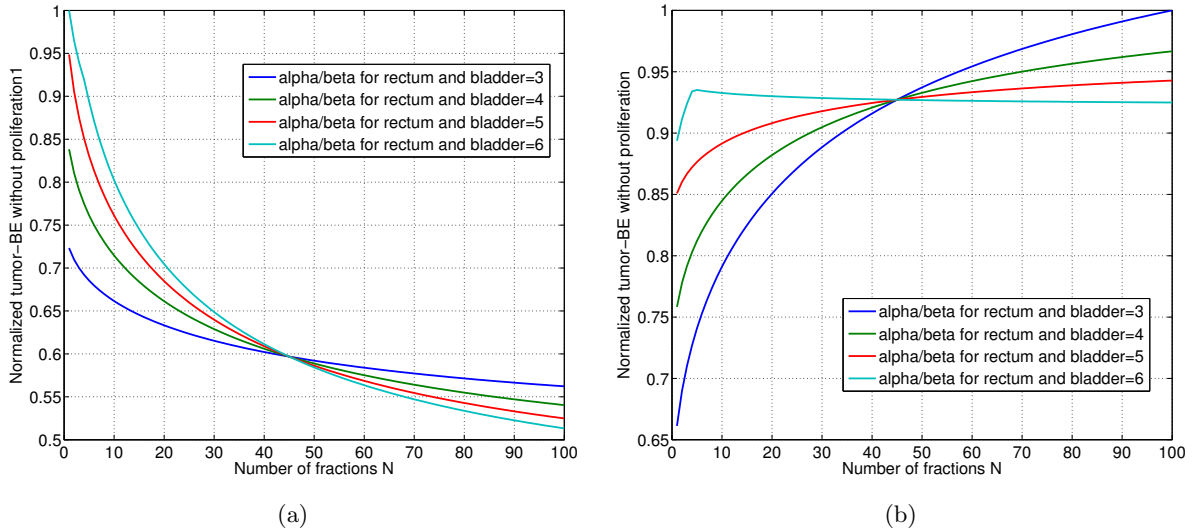


Figure 2.3: This figure was generated for our prostate case 2. The α/β ratios for all normal tissues were fixed at 3 Gy, except for rectum and bladder for which the ratios were fixed as shown in the figure. (a) The α/β ratio for the prostate tumor was fixed at 2 Gy; (b) The α/β ratio for the prostate tumor was fixed at 6 Gy. When the α/β for the rectum and the bladder was fixed at 6 Gy, the limiting normal tissue switched at $N = 4$ as suggested by the kink in the corresponding tumor-BE plot; there was no switching for the other values of this α/β ratio.

optimal number of fractions from our model is illustrated in Figure 2.4. In particular, the actual tumor-BE as a function of N in our model is the “lower-envelop” of the tumor-BEs from the group of models that include a single normal tissue one by one as formalized in Equation (2.112). As a result, the maximizer of this lower-envelop can be far from the maximizers of the individual tumor-BE curves for different normal tissues.




Figure 2.4: Tumor-BE with proliferation obtained by considering different normal tissues one by one. The actual tumor-BE with proliferation is given by the lower-envelop of all curves. These figures were generated using our head-and-neck case 1. Tumor α/β ratio was fixed at 10 Gy and normal tissue α/β ratios were fixed at 3 Gy; T_{lag} was fixed at 7 days. (a) T_{double} was 2 days. In this case, the optimal number of fractions obtained from our model happened to equal the number that would have resulted if only unspecified normal tissue with a maximum dose constraint were included in a single normal tissue model. More generally speaking, the maximizer of the lower-envelop occurs at the same location as that of one of the normal tissues. (b) T_{double} was 8 days. Here, the maximizer of the lower-envelop occurs far from the maximizers of any of the individual normal tissue curves.

2.8.4 Sensitivity to tumor doubling time

It is known from previous work on optimization models with a single normal tissue that shorter treatment courses would be better when the tumor grows fast as characterized by smaller values of the tumor doubling time T_{double} (see, for example, Figure 1 in [44]). We observed this trend with multiple normal tissues as well. As an example, the tumor-BE with proliferation is plotted against N for various values of T_{double} in

Figures 4.1(a), (b) for our head-and-neck case 1 and prostate case 1. The figures show that the optimal number of fractions is smaller for faster growing tumors. Also, the optimal tumor-BE itself is smaller for faster growing tumors. The figures also show




Figure 2.5: Sensitivity of the optimal number of fractions and the optimal tumor-BE to tumor doubling time T_{double} (days) when $T_{lag} = 7$ days. (a) Head-and-neck case 1; tumor α/β was 8 Gy, and the α/β ratio for all normal tissues was 3 Gy. (b) Prostate case 1; tumor α/β was 6 Gy, and the α/β ratio for all normal tissues was 3 Gy. Some believe that the prostate α/β is much smaller than 6, and in particular, even smaller than 3; in most of those cases, as suggested by Proposition 2.5.1, it is optimal to administer a single fraction. We used this higher value of the α/β ratio here mainly to generate a broader range of curves for completeness.

that the tumor-BE can sometimes grow very slowly with N before it starts decreasing. When this is the case, the benefit of continuing treatment beyond a certain point, all the way up to the optimal number of fractions, is negligible especially when compared to other factors such as financial cost of treatment and logistical inconvenience to the patient. For each T_{double} value, we therefore tracked the smallest number of fractions

at which the tumor-BE reached 99% of its optimal value. We denote this number of fractions by N_{99}^* . The range of this number over all combinations of tumor α/β ratios and normal tissue α/β ratios for our five head-and-neck cases is reported in Table 4.8. The table shows that for each T_{double} value, the range is quite narrow over all cases and over all combinations of α/β ratios. This observation may be helpful in practice given the uncertainty in and difficulty in estimating α/β values.

The table also shows that the effect of increasing T_{double} values on N_{99}^* seems to be saturating beyond $T_{double} \approx 10$ (there are only a very few combinations for which the right-end of the range is a relatively high number such as 56 or 63).

All optimal fractionation numbers in this table are at least 8, that is, $1 + T_{lag}$ is a lower bound on the optimal number of fractions (see Remark 2.5.2).

Finally, we did not include a similar table for our prostate cases here because it was optimal to administer a single fraction in most parameter combinations (see Proposition 2.5.4).

We also investigated how the sensitivity of the optimal number of fractions to T_{double} depended on the relative values of the α/β ratios for the head-and-neck tumor and the normal tissues. To achieve this, we fixed the α/β ratios of the spinal cord and the brainstem at 3 Gy, and varied the α/β ratios for the parotids over the set $\{3, 4, 5, 6\}$. We also varied the head-and-neck tumor α/β ratio over the set $\{8, 10, 12\}$. The optimal number of fractions was plotted against T_{double} in Figure 2.6. This figure shows that the sensitivity of the optimal number of fractions to T_{double} increases with increasing α/β ratio for the tumor. The qualitative trend in N_{99}^* was similar although it was less sensitive to T_{double} .

2.8.5 Sensitivity to time-lag before proliferation begins

We now investigate sensitivity of the optimal number of fractions to T_{lag} . Our Remark 2.5.2 shows that the optimal number of fractions is not sensitive to T_{lag} other than the fact that $1 + T_{lag}$ is a lower bound on the optimal number of fractions when the

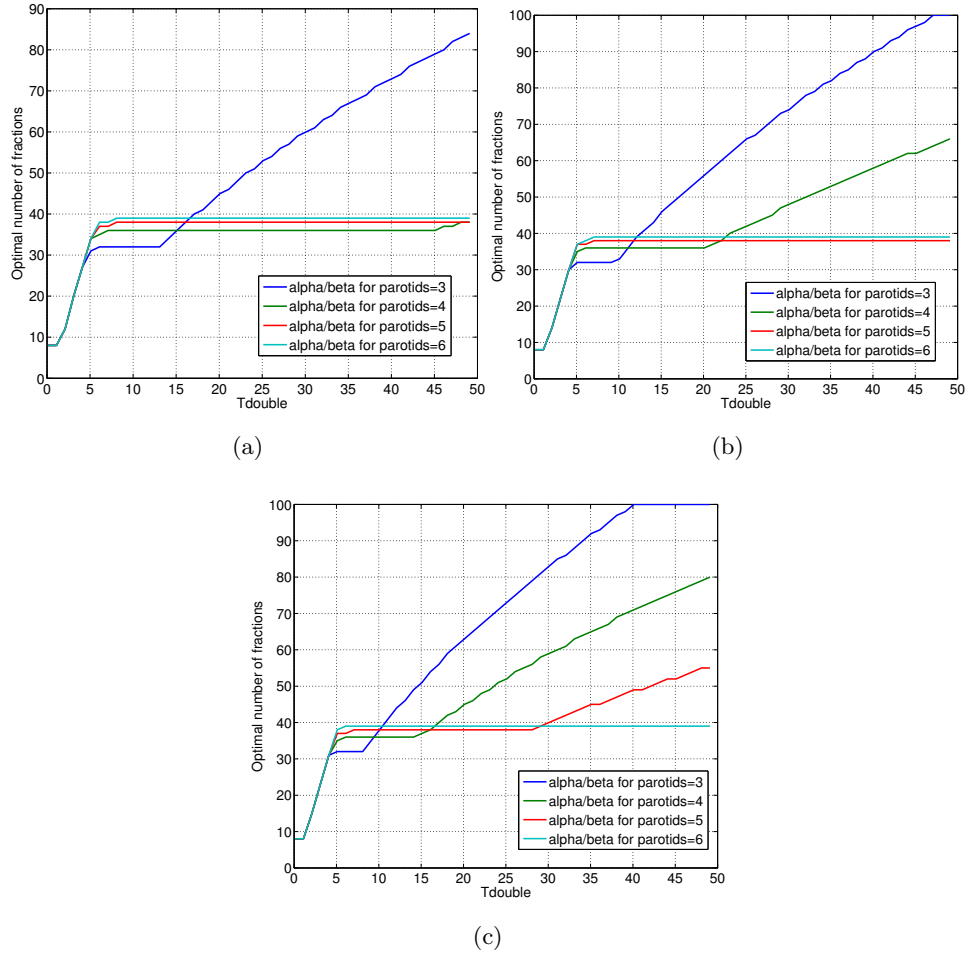


Figure 2.6: Sensitivity of the optimal number of fractions to T_{double} as a function of the relative difference between the α/β ratios for the tumor, and for the parotids. This figure was generated for our head-and-neck case 2. The α/β ratios for the spinal cord and the brainstem were fixed at 3 Gy. T_{lag} was 7 days. (a) The α/β ratio for the tumor was 8 Gy; (b) The α/β ratio for the tumor was 10 Gy; (c) The α/β ratio for the tumor was 12 Gy.

sufficient condition in Proposition 2.5.1 holds. A closer look at our closed-form formula (2.130) for the optimal number of fractions in the single normal tissue case and the statement of Lemma 2.7.1 shows that this lower bound is relevant mainly when T_{double} is small. Although we do not have a closed-form formula for the optimal number of

fractions in the case of multiple normal tissues, we intuitively expect this logic to hold. To obtain further insight into this matter, we performed sensitivity analyses for head-and-neck cases as summarized in Table 4.12. Our analyses confirmed that for small values of T_{double} relative to T_{lag} , the optimal number of fractions is equal to $1 + T_{lag}$; for larger value of T_{double} , the optimal number of fractions is not sensitive to T_{lag} . Again, we did not include a similar table for our prostate cases because the optimal number of fractions was one in most instances and thus it was not sensitive to T_{lag} .

2.9 Discussion

As stated in Section 6.1, several alternative fractionation schedules have been tested clinically. For example, one popular schedule involves administering once daily treatment only over weekdays; another schedule involves twice daily treatment only over weekdays. Our formulation (P) can be easily modified to accommodate such fractionation schemes. To see how this can be achieved, suppose that the N fractions are well-separated and are administered in $T(N)$ days. Here, the mathematical meaning of the phrase well-separated is that there exists a time-interval of Δ days such that $T(N + 1) - T(N) \geq \Delta$ for all $N \geq 1$. Thus, $T(N)$ is a strictly increasing function of N . Note also that $T(1) = 0$. For example, if treatment is administered once daily as in Sections 2.4 and 5.4, then $T(N) = N - 1$ for all $N \geq 1$. Similarly, if treatment is administered once daily but only over weekdays, then $T(N) = N - 1$ for $N = 1, 2, \dots, 5$; $T(N) = N + 1$ for $N = 6, 7, \dots, 10$; and so on. We emphasize here that this notation is general enough to encompass essentially any plausible fractionation schedule. For instance, consider a hypothetical schedule that administers two fractions on the first day with a twelve hour gap in between and then skips the next day and administers a third fraction forty eight hours after the second fraction; in this case, we would have $T(1) = 0$, $T(2) = 1/2$ a day, and $T(3) = 2.5$ days. On the other hand, if the gap between the first two fractions is forty eight hours and the

third fraction is administered twelve hours after the second fraction, then we would have $T(1) = 0$, $T(2) = 2$ days and $T(3) = 2.5$ days. To accommodate such schedules in our formulation (P), we only need to generalize the definition of $\tau(N)$ in (2.2) as

$$\tau(N) = \frac{[T(N) - T_{lag}]^+ \ln 2}{T_{double}}. \quad (2.139)$$

Lemmas 2.6.1 and 2.6.2, and Property 1 in Theorem 2.6.3 continue to hold in this general case without change. However, we have not been able to prove that Properties 2 and hence Property 3 in Theorem 2.6.3 hold. The reason is that our proof of Property 2 relies on convexity of $\tau(N)$ as defined in (2.2) but the generalization in (2.139) may not be convex. Thus, our algorithm in Section 2.6.1 will become a heuristic. However, this is not a limitation in practice. As stated earlier, the optimal number of fractions can be found by a brute-force method that utilizes an upper bound N_{max} on the number of fractions.

In this Chapter, as in all existing literature on determining the optimal number of fractions using the LQ model with tumor proliferation, we treated the spatial and the temporal components of the problem separately. That is, a fluence-map $u_{nominal}$ was assumed to be available and the optimal number of fractions was determined by scaling this fluence-map. Such separation of the two components is suboptimal. Unkelbach et al. [80] and Kim et al. [23] have recently proposed optimization models where the fluence-map is optimized also using the BED concept; our problem (OPTFRAC2(N)) can be seen as a special case of their general models. However, their general models were not solved to optimality owing to computational difficulties rooted in nonconvexity and the large size of their problems. Those two papers also did not explicitly include tumor proliferation and hence did not optimize the number of fractions. In Chapter 4, we will extend our work here to simultaneously optimize the fluence-map and the number of fractions. However, that formulation will be a large-scale (with several thousand variables and tens of thousands of constraints), nonconvex, mixed integer optimization problem and hence it will be computationally challenging to solve. As a result, it will not be possible to rigorously guarantee optimality. So the research there

will focus on the development of an efficient algorithm for approximate solution.

parotid $\alpha/\beta = 3$ Gy		T_{double} (days)							
		2	3	5	8	10	20	40	50
	Our model	12	20	35	39	39	44	71	82
One normal tissue	Spinal cord	9	14	23	34	41	70	112	130
	Brainstem	9	15	25	38	46	79	126	146
	Left parotid	8	8	14	22	26	44	71	82
	Right parotid	8	8	14	22	26	44	71	81
	Unspec. normal tissue (max)	12	20	35	55	67	117	192	223
	Unspec. normal tissue (dv)	12	20	34	52	63	109	178	206
parotid $\alpha/\beta = 4$ Gy		T_{double} (days)							
		2	3	5	8	10	20	40	50
	Our model	12	20	35	39	39	40	40	45
One normal tissue	Spinal cord	9	14	23	34	41	70	112	130
	Brainstem	9	15	25	38	46	79	126	146
	Left parotid	8	8	8	11	13	23	39	45
	Right parotid	8	8	8	11	13	23	39	45
	Unspec. normal tissue (max)	12	20	35	55	67	117	192	223
	Unspec. normal tissue (dv)	12	20	34	52	63	109	178	206
parotid $\alpha/\beta = 5$ Gy		T_{double} (days)							
		2	3	5	8	10	20	40	50
	Our model	12	20	35	39	40	40	40	40
One normal tissue	Spinal cord	9	14	23	34	41	70	112	130
	Brainstem	9	15	25	38	46	79	126	146
	Left parotid	1	1	1	1	1	1	8	8
	Right parotid	1	1	1	1	1	1	1	8
	Unspec. normal tissue (max)	12	20	35	55	67	117	192	223
	Unspec. normal tissue (dv)	12	20	34	52	63	109	178	206

Table 2.5: A comparison of the optimal number of fractions obtained from our model with the optimal number of fractions that would have resulted if we had used a model with a single normal tissue and then had applied our results in Section 2.7. The table was generated for our head-and-neck case 1. The α/β ratio for the tumor was fixed at 10 Gy; the α/β ratios for spinal cord and brainstem were fixed at 3 Gy. The α/β ratio for the parotids was varied as shown in the table. T_{lag} was set at 7 days.

Head-and-neck case #	T_{double} (days)							
	2	3	5	8	10	20	40	50
1	8-10	13-16	23-28	33-35	35-36	37-38	37-50	37-56
2	8-10	14-17	24-29	28-35	29-36	31-42	34-63	34-70
3	8-10	13-16	24-28	37-41	41-44	44-47	45-57	45-65
4	8-10	13-16	24-28	36-39	38-41	41-43	42-59	42-66
5	8-10	14-17	24-29	34-37	36-39	38-40	39-56	39-63

Table 2.6: The range of values of N_{99}^* for five head-and-neck cases with $T_{lag} = 7$ days over all combinations of tumor and normal tissue α/β ratios. Here, in case N_{99}^* was less than $1 + T_{lag}$ we reset it to $1 + T_{lag}$ as we know that this is a lower bound on the optimal number of fractions as in Remark 2.5.2.

$T_{lag} = 14$ (days)	T_{double} (days)							
Head-and-neck case #	2	3	5	8	10	20	40	50
1	15	15-16	23-28	33-35	35-36	37-38	37-50	37-56
2	15	15-17	23-29	28-35	29-36	31-41	34-63	34-70
3	15	15-16	23-28	37-41	41-44	44-47	45-57	45-65
4	15	15-16	23-28	36-39	38-41	41-43	42-59	42-66
5	15	15-17	24-29	34-37	36-39	38-40	39-56	39-63

$T_{lag} = 21$ (days)	T_{double} (days)							
Head-and-neck case #	2	3	5	8	10	20	40	50
1	22	22	23-28	33-35	35-36	37-38	37-50	37-56
2	22	22	23-29	28-35	29-36	31-41	34-63	34-70
3	22	22	23-28	37-41	41-44	44-47	45-57	45-65
4	22	22	23-28	36-39	38-41	41-43	42-59	42-66
5	22	22	24-29	34-37	36-39	38-40	39-56	39-63

$T_{lag} = 28$ (days)	T_{double} (days)							
Head-and-neck case #	2	3	5	8	10	20	40	50
1	29	29	29	33-35	35-36	37-38	37-50	37-56
2	29	29	29	29-35	29-36	31-41	34-63	34-70
3	29	29	29	36-41	40-44	44-47	45-57	45-65
4	29	29	29	35-39	38-41	41-43	42-59	42-66
5	29	29	29	34-37	36-38	38-40	39-56	39-63

$T_{lag} = 35$ (days)	T_{double} (days)							
Head-and-neck case #	2	3	5	8	10	20	40	50
1	36	36	36	36	36	36-38	37-50	37-56
2	36	36	36	36	36	36-41	36-63	36-70
3	36	36	36	36-41	40-44	44-47	45-57	45-65
4	36	36	36	36-39	38-41	41-43	42-59	42-66
5	36	36	36	36-37	36-38	38-40	39-56	39-63

Table 2.7: The range of values of N_{99}^* for five head-and-neck cases with different combinations of T_{double} and T_{lag} over all combinations of tumor and normal tissue α/β ratios.

Chapter 3

A TWO-VARIABLE LINEAR PROGRAM TO SOLVE AN OPTIMAL FRACTIONATION PROBLEM

The optimal fractionation problem involves finding a sequence $\vec{d} = (d_1, d_2, \dots, d_N)$ of tumor doses for $N > 1$ treatment sessions. Within the well-known linear-quadratic (LQ) dose-response framework [11], one common approach to this problem is to maximize the biological effect of \vec{d} on the tumor while ensuring that the corresponding biologically effective dose (BED) given to various nearby serial and parallel organs-at-risk (OAR) can be tolerated.

A majority of research on the optimal fractionation problem has considered a single OAR and in that case an optimal solution is known in closed-form (see, for example, [44–46, 54, 55, 62] and references therein). Over the last few years, researchers have started studying more general and realistic formulations with the recent focus being on the multiple OAR case (see, for example, Chapter 2, [80] and literature reviews therein). To the best of our knowledge, the most recent formulation was discussed in Chapter 2, and it has the following form.

$$\text{(OPTFRAC)} \quad \max_{\vec{d}} \quad \alpha_0 \sum_{t=1}^N d_t + \beta_0 \sum_{t=1}^N d_t^2, \quad (3.1)$$

$$\text{subject to} \quad \sum_{t=1}^N d_t + s_{\max}^m \rho_m \sum_{t=1}^N (d_t)^2 \leq \text{BED}_{\max}^m / s_{\max}^m, \quad m \in \mathcal{M}_1, \quad (3.2)$$

$$\sum_{t=1}^N d_t + s_{\text{mean}}^m \rho_m \sum_{t=1}^N (d_t)^2 \leq B_{\text{mean}}^m / s_{\text{mean}}^m, \quad m \in \mathcal{M}_2, \quad (3.3)$$

$$\sum_{t=1}^N d_t + s_{\text{dv}}^m \rho_m \sum_{t=1}^N (d_t)^2 \leq \text{BED}_{\text{dv}}^m / s_{\text{dv}}^m, \quad \text{for } m \in \mathcal{M}_3, \quad (3.4)$$

$$\vec{d} \geq 0. \quad (3.5)$$

Here, α_0 and β_0 are the LQ dose-response parameters for the tumor. Set \mathcal{M}_1

is the set of serial OAR with maximum dose constraints; \mathcal{M}_2 is the set of parallel OAR with mean dose constraints; and \mathcal{M}_3 is the set of parallel OAR with dose-volume constraints. We assume for simplicity that an OAR has either a maximum dose constraint or a mean dose constraint or a dose-volume constraint; that is, $\mathcal{M}_1 \cap \mathcal{M}_2 = \mathcal{M}_2 \cap \mathcal{M}_3 = \mathcal{M}_1 \cap \mathcal{M}_3 = \emptyset$. Thus, the set of all OAR is denoted by $\mathcal{M} = \mathcal{M}_1 \cup \mathcal{M}_2 \cup \mathcal{M}_3$. We also assume for simplicity that an OAR has at most one dose-volume constraint. Both these assumptions can be removed by using more complicated notation. For OAR $m \in \mathcal{M}$, $\rho_m = \beta_m/\alpha_m$ is the ratio of its LQ dose-response parameters α_m and β_m . The upper limit parameters BED_{\max}^m , $\text{BED}_{\text{mean}}^m$, and BED_{dv}^m are BED values that the various OAR are known to tolerate, as for example, those derived from standard treatment guidelines available in [1]. Examples of precise formulas for BED_{\max}^m , $\text{BED}_{\text{mean}}^m$, and BED_{dv}^m are included in Chapter 2. Finally, parameters s_{\max}^m , s_{mean}^m , s_{dv}^m are the so-called effective sparing factors for the OAR with maximum dose, mean dose and dose-volume constraints, respectively. Again, precise formulas for these are available in Chapter 2 and, for other examples, in [55, 61, 62].

As all constraints in (OPTFRAC) have the same structure, we henceforth use C_m to denote the right hand side upper bound in the constraint for OAR $m \in \mathcal{M}$ and use s^m to denote the effective sparing factor for OAR $m \in \mathcal{M}$. As a result, (OPTFRAC) can now be written more compactly as

$$\text{(OPTFRAC)} \quad \max_{\vec{d}} \quad \alpha_0 \sum_{t=1}^N d_t + \beta_0 \sum_{t=1}^N d_t^2, \quad (3.6)$$

$$\text{subject to} \quad \sum_{t=1}^N d_t + s^m \rho_m \sum_{t=1}^N (d_t)^2 \leq C_m, \quad m \in \mathcal{M}, \quad (3.7)$$

$$\vec{d} \geq 0. \quad (3.8)$$

Formulation (3.6)-(3.8) is a *nonconvex* quadratically constrained quadratic program (QCQP) — although the constraints and hence the feasible region is convex in \vec{d} , the objective is to *maximize* a convex function (this latter being the source of nonconvexity). Such problems are typically computationally difficult to solve, and in

general belong to the class NP-hard [63]. A recurrent theme in the aforementioned literature on optimal fractionation therefore is either to contemplate algorithms for approximate solution or to derive sufficient conditions on problem parameters under which single-dosage or equal-dosage solutions* are optimal. However, and somewhat unsatisfactorily, such sufficient conditions are not necessary; for instance, an equal-dosage solution may be optimal even when the corresponding sufficient condition is violated. The special interest in single-dosage and equal-dosage solutions stems partly from the belief that, based on currently available parameter values, they are likely to be optimal for prostate and head-and-neck tumors, respectively. Moreover, optimal single-dosage and equal-dosage solutions can be obtained in closed-form even in the multiple OAR case (see Chapter 2). However, in the multiple OAR case, such single-dosage or equal-dosage solutions may not be optimal. In that scenario, an optimal unequal multiple-dosage solution is desired[†]. Unfortunately, there has been no known method that is guaranteed to find an optimal unequal multiple-dosage solution (in closed-form or otherwise). Moreover, the existing literature does not provide sufficient or necessary conditions for the optimality of unequal multiple-dosage solutions. Neither does it provide sufficient or necessary conditions for precisely when no single-dosage and no equal-dosage solution is optimal and thus one must seek unequal multiple-dosage solutions. As such, problem (OPTFRAC) has thus remained unsolved.

We show in Theorem 3.0.1 below that an optimal solution to (OPTFRAC) can be derived in closed-form from the solution of a two-variable linear program (LP). This theorem and its Corollaries 3.0.2 and 3.0.3 also provide all the aforementioned, and previously elusive, sufficient and necessary conditions.

Before we present our two-variable LP, we introduce additional notation. We

*Equal-dosage solutions use $d_1 = d_2 = \dots = d_N$; single-dosage solutions set all doses except one to zero.

[†]By this we mean any solution that is neither single-dosage nor equal-dosage.

define, as in Chapter 2 and elsewhere,

$$b_{\max}^m(N) \triangleq \frac{-1 + \sqrt{1 + 4\rho_m \text{BED}_{\max}^m/N}}{2s_{\max}^m \rho_m}, \quad m \in \mathcal{M}_1, \quad (3.9)$$

$$b_{\text{mean}}^m(N) \triangleq \frac{-1 + \sqrt{1 + 4\rho_m \text{BED}_{\max}^m/N}}{2s_{\text{mean}}^m \rho_m}, \quad m \in \mathcal{M}_2, \quad (3.10)$$

$$b_{\text{dv}}^m(N) \triangleq \frac{-1 + \sqrt{1 + 4\rho_m \text{BED}_{\text{dv}}^m/N}}{2s_{\text{dv}}^m \rho_m}, \quad m \in \mathcal{M}_3. \quad (3.11)$$

We also use $b_{\max}^m(1)$, $b_{\text{mean}}^m(1)$, and $b_{\text{dv}}^m(1)$ to denote the quantities obtained by substituting $N = 1$ into formulas (3.9)-(3.11). Moreover, we define, as in Chapter 2,

$$\gamma^* \triangleq \min \left(\min_{m \in \mathcal{M}_1} b_{\max}^m(1), \min_{m \in \mathcal{M}_2} b_{\text{mean}}^m(1), \min_{m \in \mathcal{M}_3} b_{\text{dv}}^m(1) \right), \quad \text{and} \quad (3.12)$$

$$c^* \triangleq \min \left(\min_{m \in \mathcal{M}_1} b_{\max}^m(N), \min_{m \in \mathcal{M}_2} b_{\text{mean}}^m(N), \min_{m \in \mathcal{M}_3} b_{\text{dv}}^m(N) \right). \quad (3.13)$$

Here, γ^* is the largest possible (and hence optimal) dose in a single-dosage solution, whereas c^* is the largest possible (and hence optimal) dose per session in an equal-dosage solution. Now consider the two-variable LP

$$(2\text{VARLP}) \max_{x,y} \alpha_0 x + \beta_0 y, \quad (3.14)$$

$$\text{subject to } x + s^m \rho_m y \leq C_m, \quad m \in \mathcal{M}, \quad (3.15)$$

$$y \leq \gamma^* x, \quad (3.16)$$

$$c^* x \leq y, \quad (3.17)$$

$$x \geq 0, \quad (3.18)$$

$$y \geq 0. \quad (3.19)$$

Theorem 3.0.1. *Let x^* , y^* be[‡] an optimal solution to (2VARLP). Then exactly one of the following three situations must hold.*

1. $x^* = \sqrt{y^*}$: in this case, it is optimal to set $d_t = \gamma^*$ in exactly one session t and set the other $N - 1$ doses d_s , for $s \neq t$, to zero; that is, a single-dosage solution is optimal.

[‡](2VARLP) does indeed have an optimal solution because its feasible region is bounded.

2. $x^* = \sqrt{Ny^*}$: in this case, it is optimal to set $d_t = c^*$, for $t = 1, 2, \dots, N$; that is, an equal-dosage solution is optimal.
3. $\sqrt{y^*} < x^* < \sqrt{Ny^*}$: in this case, we have an uncountable number of unequal multiple-dosage optimal solutions; as a concrete example of one such solution, the two-dose solution where $d_3 = d_4 = \dots = d_N = 0$, and

$$d_1 = \frac{x^* + \sqrt{2y^* - (x^*)^2}}{2}, \quad d_2 = x^* - d_1, \quad (3.20)$$

is optimal when $x^* < \sqrt{2y^*}$.

Proof. First, we use the variable transformations $x = \sum_{t=1}^N d_t$ and $y = \sum_{t=1}^N d_t^2$ and equivalently reformulate (OPTFRAC) as

$$\max_{\vec{d}, x, y} \alpha_0 x + \beta_0 y, \quad (3.21)$$

$$\text{subject to } x + s^m \rho_m y \leq C_m, \quad m \in \mathcal{M}, \quad (3.22)$$

$$x = \sum_{t=1}^N d_t, \quad (3.23)$$

$$y = \sum_{t=1}^N d_t^2, \quad (3.24)$$

$$\vec{d} \geq 0, \quad (3.25)$$

$$x \geq 0, \quad (3.26)$$

$$y \geq 0. \quad (3.27)$$

Since $\vec{d} \geq 0$, x and \sqrt{y} can be seen as the l_1 and l_2 norms of \vec{d} , respectively. Consequently, every x, y, \vec{d} combination that is feasible to constraints (3.23)-(3.25) also satisfies the two inequalities $\sqrt{y} \leq x \leq \sqrt{Ny}$ (this is a well-known relationship between l_1 and l_2 norms). Thus, we first add these two inequalities to the above problem

without altering its feasible region. This yields

$$\max_{\vec{d}, x, y} \alpha_0 x + \beta_0 y, \quad (3.28)$$

$$\text{subject to } x + s^m \rho_m y \leq C_m, \quad m \in \mathcal{M}, \quad (3.29)$$

$$x = \sum_{t=1}^N d_t, \quad (3.30)$$

$$y = \sum_{t=1}^N d_t^2, \quad (3.31)$$

$$\vec{d} \geq 0, \quad (3.32)$$

$$\sqrt{y} \leq x, \quad (3.33)$$

$$x \leq \sqrt{N y}, \quad (3.34)$$

$$x \geq 0, \quad (3.35)$$

$$y \geq 0. \quad (3.36)$$

We claim that variable \vec{d} and constraints (3.30)-(3.32) can now be dropped to equivalently rewrite the above problem as

$$\max_{x, y} \alpha_0 x + \beta_0 y, \quad (3.37)$$

$$\text{subject to } x + s^m \rho_m y \leq C_m, \quad m \in \mathcal{M} \quad (3.38)$$

$$\sqrt{y} \leq x, \quad (3.39)$$

$$x \leq \sqrt{N y}, \quad (3.40)$$

$$x \geq 0, \quad (3.41)$$

$$y \geq 0. \quad (3.42)$$

In particular, an optimal sequence of doses that satisfies (3.30)-(3.32) can be recovered in closed-form from any optimal solution of (3.37)-(3.42). To prove this intermediate claim, we let \bar{x} , \bar{y} denote an optimal solution to (3.37)-(3.42). If $\sqrt{\bar{y}} = \bar{x}$, then the single-dosage solution $d_1 = \bar{x}$, $d_t = 0$ for $t = 2, 3, \dots, N$, \bar{x} , \bar{y} is optimal to (3.28)-(3.36). This follows from the property that the l_2 norm of a nonnegative vector

equals its l_1 norm if and only if exactly one of the elements of the vector is positive and the others are zero. If $\sqrt{N\bar{y}} = \bar{x}$, then the equal-dosage solution $d_t = \bar{x}/N$ for $t = 1, 2, \dots, N$, \bar{x}, \bar{y} is optimal to (3.28)-(3.36). This again follows from the property that the l_2 norm of an N -dimensional nonnegative vector is \sqrt{N} times its l_1 norm if and only if all elements of the vector are identical. Finally, if $\sqrt{\bar{y}} < \bar{x} < \sqrt{N\bar{y}}$, then any nonnegative \vec{d} that satisfies the equations $\bar{x} = \sum_{t=1}^N d_t$ and $\bar{y} = \sum_{t=1}^N (d_t)^2$ is optimal to (3.28)-(3.36) and the set of such \vec{d} is uncountable. Again, from a property of l_1 and l_2 norms, no single-dosage or equal-dosage solution can satisfy this system of equations. In this case, for example, setting $d_3 = d_4 = \dots = d_N = 0$, $d_1 = \frac{\bar{x} + \sqrt{2\bar{y} - \bar{x}^2}}{2}$, $d_2 = \bar{x} - d_1$ works. These values of d_1 and d_2 were derived by solving the two equations $d_1 + d_2 = \bar{x}$ and $d_1^2 + d_2^2 = \bar{y}$ using the standard quadratic equation formula. Also note that these values of d_1 and d_2 are indeed nonnegative but in fact are both positive as required for an unequal multiple-dosage solution. This holds because, again, from the l_1, l_2 properties, we have, $\sqrt{\bar{y}} < \bar{x} \leq \sqrt{2\bar{y}}$, which implies that $\bar{y} < \bar{x}^2 \leq 2\bar{y}$. These last two inequalities imply that $0 \leq 2\bar{y} - \bar{x}^2 < \bar{x}^2$, which means that $\bar{x}/2 \leq d_1 < \bar{x}$ and $0 < d_2 \leq \bar{x}/2$.

Although problem (3.37)-(3.42) only includes two variables, it is nonconvex because of constraint (3.33). We surmount this difficulty by showing that constraints (3.33)-(3.34) can be replaced with constraints (3.16)-(3.17) without loss of optimality, which yields (2VARLP). A geometric proof of this fact is shown in Figure 3.1. The idea is that replacing the constraints this way does not exclude any optimal solution(s) of problem (3.37)-(3.42) and does not introduce any new optimal solutions. That is, (2VARLP) cannot have any optimal solutions in the newly added region (shown with vertical hatch lines) between the dotted line segment OA and the thick black curve $y = x^2$ also connecting O and A. Note in the figure that that the function $y = x^2$ is convex in x and hence line segment OA is above the curve $y = x^2$. Similarly, (3.37)-(3.42) does not have any optimal solutions in the chopped off feasible region (shown with horizontal hatch lines) between the dotted line segment OD and the thick black

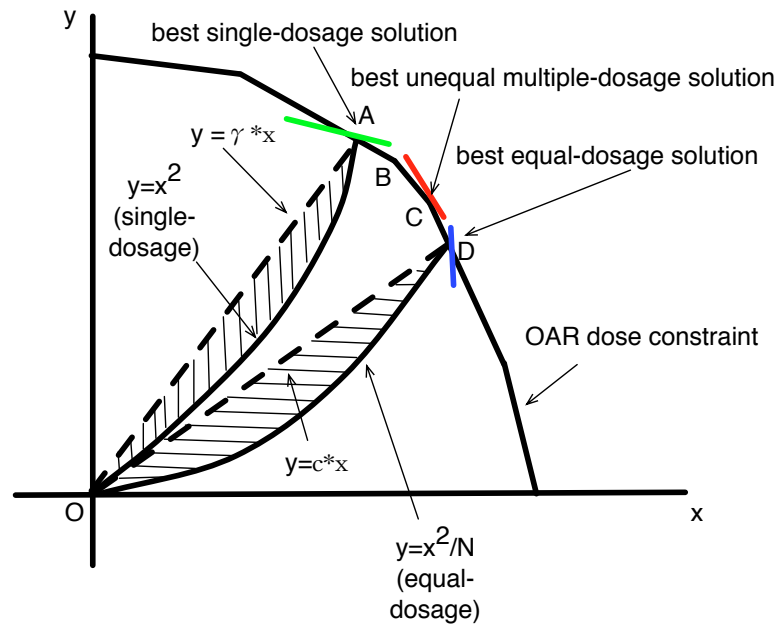


Figure 3.1: A geometric proof of the equivalence between problem (3.37)-(3.42) and problem (2VARLP). Based on properties of l_1 and l_2 norms discussed in the text, the curve $y = x^2$ corresponds to the set of single-dosage solutions whereas the curve $y = x^2/N$ corresponds to the set of equal-dosage solutions. Problem (2VARLP) is created by (i) adding the vertically hatched area to the feasible region of (3.37)-(3.42) and (ii) by removing the horizontally hatched area from the feasible region of (3.37)-(3.42). Nevertheless, optimal solutions to both problems occur at points such as A, B, C, or D depending on the slope of the objective function (see the green, blue, and red objective function lines for example).

curve $y = x^2/N$ also connecting O and D. Note again that the function $y = x^2/N$ is convex in x and hence the line segment OD is above the curve $y = x^2/N$. These facts about the precise locations of optimal solutions (3.37)-(3.42) and (2VARLP) are rooted in the structure of problem (3.37)-(3.42) in that (i) all of its linear constraints as well as its objective function have negative slopes; (ii) the only location where an optimal solution to (3.37)-(3.42) can occur on the curve $y = x^2$ is at its corner point (point A in Figure 3.1) and this property also holds for the segment OA given

by $y = \gamma^*x$ in (2VARLP); and (iii) the only location where an optimal solution to (3.37)-(3.42) can occur on the curve $y = x^2/N$ is at its corner point (point B in Figure 3.1) and this property also holds for the segment OD given by $y = c^*x$ in (2VARLP). In summary, (2VARLP) is equivalent to (OPTFRAC) and every optimal solution of one problem can be transformed in closed-form into an optimal solution of the other problem.

The three possibilities stated in the theorem then follow from the properties of l_1 and l_2 norms discussed above and then deriving the best single-dosage, the best equal-dosage, and any feasible unequal-dosage solution that satisfies the system $x^* = \sum_{t=1}^N d_t$ and $y^* = \sum_{t=1}^N (d_t)^2$, respectively. In particular, the best single-dosage solution delivers dose γ^* in a single session; the best equal-dosage solution delivers dose c^* in each one of the N sessions; and finally, the stated two-dose solution is one optimal unequal multiple-dosage solution that satisfies $x^* = d_1 + d_2$ and $y^* = d_1^2 + d_2^2$ as explained above. \square

More generally, given x^* and y^* , an optimal dosing schedule can be obtained by setting the dose in one session to d_1 , the doses in the other $N - 1$ sessions to d_2 and solving the two equations

$$d_1 + (N - 1)d_2 = x^* \quad (3.43)$$

$$d_1^2 + (N - 1)d_2^2 = y^* \quad (3.44)$$

for d_1 and d_2 . This yields

$$d_2 = \frac{x^*}{N} \left[1 - \sqrt{1 - \left(1 - \frac{y^*}{(x^*)^2}\right) \left(\frac{N}{N-1}\right)} \right] \quad (3.45)$$

$$d_1 = x^* - (N - 1)d_2. \quad (3.46)$$

We emphasize that the sufficient conditions for the optimality of single-dosage and equal-dosage solutions in Theorem 3.0.1 are equivalent to the other known sufficient conditions in Chapter 2. For instance, referring back to Figure 3.1, we note that the

only point in the feasible region of (2VARLP) where the pair x^*, y^* satisfies $x^* = \sqrt{y^*}$ is at corner point A . An optimal solution to (2VARLP) occurs at corner point A if the magnitude of the slope of the objective function line is less than or equal to the magnitudes of the slope of all OAR constraint lines. That is, if

$$\frac{\alpha_0}{\beta_0} \leq \min_{m \in \mathcal{M}} \left\{ \frac{\alpha_m / \beta_m}{s^m} \right\}. \quad (3.47)$$

This is precisely the sufficient condition reported in Chapter 2 for the optimality of a single-dosage solution. Similarly, referring back to Figure 3.1 again, we note that the only point in the feasible region of (2VARLP) where the pair x^*, y^* satisfies $x^* = \sqrt{Ny^*}$ is at corner point D . An optimal solution to (2VARLP) occurs at corner point D if the magnitude of the slope of the objective function line is at least as big as the magnitude of the slope of all OAR constraint lines. That is, if

$$\frac{\alpha_0}{\beta_0} \geq \max_{m \in \mathcal{M}} \left\{ \frac{\alpha_m / \beta_m}{s^m} \right\}. \quad (3.48)$$

This is also precisely the sufficient condition reported in Chapter 2 for the optimality of an equal-dosage solution. However, these previously known sufficient conditions (3.47) and (3.48) for the optimality of single-dosage and equal-dosage solutions are somewhat unsatisfactory because they are not necessary. For instance, it is possible to construct numerical examples where condition (3.47) does not hold but a single-dosage solution is optimal. Similarly for (3.48) regarding equal-dosage solutions. We show next that (2VARLP) does not suffer from this limitation.

Corollary 3.0.2. *The three conditions stated in Theorem 3.0.1 are necessary. That is,*

1. *Suppose a single-dosage solution is optimal. Then there exists a pair (x^*, y^*) that is optimal to (2VARLP) such that $x^* = \sqrt{y^*}$.*
2. *Suppose an equal-dosage solution is optimal. Then there exists a pair (x^*, y^*) that is optimal to (2VARLP) such that $x^* = \sqrt{Ny^*}$.*

3. Suppose an unequal multiple-dosage solution is optimal. Then there exists a pair (x^*, y^*) that is optimal to (2VARLP) such that $\sqrt{y^*} < x^* < \sqrt{Ny^*}$.

Proof. Follows from the correspondence between optimal solutions to (OPTFRAC) and to (2VARLP). \square

Finally, the existing literature does not provide any insights into the uniqueness of optimal solutions to (OPTFRAC). Since single-dosage and equal-dosage solutions are easier to implement in practice, it would be helpful to know precisely when they are the only optimal solutions and also precisely when they are not optimal. The next corollary of Theorem 3.0.1 provides a complete answer to this issue.

Corollary 3.0.3. *We have the following three sufficient and necessary conditions for the uniqueness and lack of optimality of single-dosage and equal-dosage solutions.*

1. *There are no optimal solutions other than the N different single-dosage optimal solutions described in Theorem 3.0.1 if and only if there is a unique optimal solution (x^*, y^*) to (2VARLP) and it satisfies $x^* = \sqrt{y^*}$.*
2. *The equal-dosage solution described in Theorem 3.0.1 is the unique optimal solution if and only if there is unique optimal solution (x^*, y^*) to (2VARLP) and it satisfies $x^* = \sqrt{Ny^*}$.*
3. *No single-dosage solution is optimal, no equal-dosage solution is optimal, and an uncountable number of unequal multiple-dosage solutions (including the two-dosage solution defined in Theorem 3.0.1) are optimal if and only if every pair (x^*, y^*) that is optimal to (2VARLP) satisfies $\sqrt{y^*} < x^* < \sqrt{Ny^*}$.*

Proof. Follows from the correspondence between optimal solutions to (OPTFRAC) and to (2VARLP) and from the properties of l_1 , l_2 norms discussed in the proof of Theorem 3.0.1. \square

We conclude by discussing one extension and one limitation of our work.

Theorem 3.0.1 and its corollaries also hold without change if a tumor proliferation term is included in calculating the biological effect on tumor as in Chapter 2. Moreover, by including tumor proliferation this way, one can also optimize the number of treatment sessions N . This can be done by plotting the optimal biological effect versus N after solving (2VARLP) for different values of N within a clinically viable range and then choosing an N that maximizes this optimal biological effect as in Chapter 2 and elsewhere.

Although Theorem 3.0.1 and its corollaries completely resolve a particular form of the optimal fractionation problem that has recently been studied in the radiotherapy literature, their proofs and conclusions appear fragile in that they are highly dependent on the structure of (OPTFRAC). Consequently, similar results are unlikely to hold for other variations of the optimal fractionation problem that may arise in the future.

Chapter 4

SPATIOTEMPORALLY OPTIMAL FRACTIONATION

4.1 Introduction

In external beam radiotherapy, radiation damages both the cancer cells and the normal tissue. Thus, the goal is to maximize damage to the tumor while limiting toxic effects on nearby normal tissue. This is attempted by spatial localization and temporal dispersion of radiation dose.

Spatial localization is achieved by prescribing a high dose to the cancerous region and putting upper limits on the dose to healthy anatomies. The well-developed Intensity Modulated Raditation Therapy (IMRT) technology, the associated spatial optimization models and solution algorithms are employed to optimize the radiation intensity profile (also called the fluence-map). One standard formulation of this problem minimizes the total squared deviation from the tumor prescription dose of the doses delivered to all voxels in the tumor. This spatial side of radiation therapy has been studied extensively [6, 9, 27–31].

On the temporal side, the prevalent strategy is to break the total planned dose into multiple, well-separated treatment sessions, called fractions, that are administered over several weeks. An identical dose is planned for each fraction; this is called equal-dosage fractionation. Since normal cells typically have better damage-repair capabilities than tumor cells, such temporal dispersion gives the normal tissue some time to recover between sessions. For many tumors, using a large number of fractions with a small dose in each fraction may allow the treatment planner to administer a larger total tumor dose as compared to that using a small number of fractions with a large dose in each fraction. Thus, it would seem that for these tumors, the longer the treatment course the better. However, tumors may proliferate over the treatment

course, and then shorter treatment courses are believed to work better as they kill tumor cells quickly before significant proliferation. Moreover, there is an increasing interest in shorter treatment courses as they are logistically more convenient. Such tradeoffs in determining an optimal number of fractions have been clinically studied over the last several decades [33–43]. This question of choosing the number of fractions and the corresponding dose per fraction is generally referred to as the *optimal fractionation problem*.

The optimal number of fractions depends on the relative difference between the tumor’s and the normal tissue’s response to radiation as well as on the anatomy of the cancerous region with respect to the positioning of the radiation fields. The linear-quadratic (LQ) model is the prevalent framework for quantifying the tumor’s and the normal tissue’s response to radiation [11]. Competing fractionation schedules are thus compared using the LQ model. However, existing research on optimal fractionation that utilizes the LQ framework is limited to stylized formulations that often lead to a closed-form formula for the optimal number of fractions. This formula is derived using single-variable calculus. The idea is to maximize the biological effect (BE) of radiation dose on the tumor subject to an upper bound constraint on the biologically effective dose (BED) delivered to the normal tissue [11]. This type of literature includes [44–50, 54, 55, 61, 62]. Table 4.1 below summarizes the contributions of these models.

One limitation of most existing stylized models is that they only consider a single normal tissue. This not only may lead to an incorrect prediction of the number of fractions but also may yield a dose that cannot be tolerated by other nearby normal tissues that were excluded from the formulation. This important concern stems from the fact that essentially all anatomical regions of interest include multiple normal tissues. Another limitation of the stylized models is that they do not explicitly model intensity modulation and thus essentially ignore the spatial side of the problem even though IMRT technology is now ubiquitous. Such limitations curtail the practical

References	≥ 2 normal tissues	closed-form for dose	spatiotemporally optimal
[44, 45, 47–49]	No	No	No
[46, 50]	No	Yes	No
[51]	Yes	No	No
[54]	No	Yes	No
[55, 61, 62]	No	Yes	No
Chapter 2	Yes	Yes	No

Table 4.1: A summary of some optimal fractionation models that use the LQ framework.

applicability of these stylized models and thus they were addressed in Chapter 2.

Although, to the best of our knowledge, our model in Chapter 2 is currently the most comprehensive formulation of the optimal fractionation problem based on the LQ framework, it has an important limitation — it separated the spatial and temporal components of the problem. That is, a spatially optimized fluence-map was assumed to be available *a priori* and the number of fractions was then optimized with respect to this map using the concept of sparing factors. This approach simplified the optimal fractionation problem considerably — we were able to characterize the BE on tumor as a quasiconcave function of the number of fractions, which led to a simple procedure for optimizing this number. Unfortunately, this is suboptimal especially because the spatial optimization problem that is solved *a priori* does not directly depend on any biological dose-response parameters of the tumor or the normal tissue. Therefore, in this paper, we build a spatiotemporally integrated optimal fractionation model where the fluence-map and the number of fractions are both optimization variables. This formulation is computationally difficult to solve. We thus propose an efficient algorithm rooted in convex programming for its approximate solution.

Through computer simulations on head-and-neck and prostate cancer test cases, we compare the tumor-BE achieved by this spatiotemporally integrated model with

that attained by (i) a traditional IMRT fluence-map optimization model that does not optimize the number of fractions, and (ii) our spatiotemporally separated model in Chapter 2. In our simulations, the spatiotemporally integrated approach shows 22% and 69% improvements in tumor-BE over the traditional IMRT model on an average in our head-and-neck and prostate cases, respectively. Similarly, it achieves 27% and 21% improvements in tumor-BE over the spatiotemporally separated model on an average in our head-and-neck and prostate cases, respectively. This suggests that, within the LQ framework, even approximate solution of our computationally difficult, spatiotemporally integrated model may offer some benefit over solving existing stylized models.

4.2 Problem formulation

Our overall methodology is to choose a fluence-map and a number of fractions so as to maximize the total-BE of average dose over all tumor voxels subject to constraints on normal tissue dose and on the smoothness of the fluence-map. The mathematical notation and terminology here is standard in the literature and is borrowed from Chapter 2.

4.2.1 Expression for the tumor objective function

Let n denote the number of tumor voxels, indexed by $i = 1, 2, \dots, n$. The radiation field is discretized into small segments called beamlets. Let k be the number of beamlets and let $u \in \mathfrak{R}_+^k$ denote the k -dimensional beamlet intensity vector used in each treatment session. Let A be the $n \times k$, nonnegative tumor dose deposition matrix and let A_i denote its i th row, which corresponds to the i th tumor voxel. That is, according to the standard linear dose deposition model [56–60], $A_i u$ is the dose delivered to the i th tumor voxel and $\bar{A}u \triangleq \frac{\sum_{i=1}^n A_i u}{n}$ is the average dose over all tumor voxels in each session.

Now consider a treatment course with N once-daily fractions. Let T_{double} (days)

denote the doubling time for the tumor and T_{lag} (days) denote the time-lag after which tumor proliferation starts after treatment initiation. The expression $[(N-1) - T_{\text{lag}}]^+$, which is defined as $\max((N-1) - T_{\text{lag}}, 0)$, is the time over which the tumor proliferates.

We define

$$\tau(N) \triangleq \frac{[(N-1) - T_{\text{lag}}]^+ \ln 2}{T_{\text{double}}}, \quad (4.1)$$

and let α_0 and β_0 denote the parameters of the LQ model for the tumor. Then, according to the LQ model, the total N -session BE of the average tumor dose is given by

$$N\alpha_0(\bar{A}u) + N\beta_0(\bar{A}u)^2 - \tau(N). \quad (4.2)$$

We wish to choose N within a clinically viable range $1 \leq N \leq N_{\text{max}}$ and a fluence-map u to maximize (4.2). This objective is nonconvex in u because we are maximizing a convex function. Fortunately, we are able to show later in Section 5.3 that this objective can be easily re-written in a convex (in fact, linear) form. This is achieved by observing that, when N is fixed, maximizing (4.2) is equivalent to maximizing $\bar{A}u$, a linear function.

Note here that, as is common in IMRT, there are other possible choices for the objective function in this formulation. One option is to maximize the BE of the minimum tumor dose over all voxels. It turns out that, similar to our objective function, this alternative objective function can also be converted into a convex (in fact, linear) form. This is achieved by observing that, when N is fixed, maximizing the BE of the minimum tumor dose is equivalent to maximizing the minimum tumor-dose itself; the minimum tumor-dose can be maximized by instead maximizing a new variable t and adding a linear constraint for each tumor voxel enforcing that the dose to this voxel is at least t .

Both the average tumor-dose and the minimum tumor-dose are special cases of the well-known concept of generalized equivalent uniform dose (gEUD) [81]. This gEUD

is parameterized by a single parameter p , and is given by

$$\text{gEUD} \triangleq \left(\sum_{i=1}^n (A_i u)^p / n \right)^{1/p}. \quad (4.3)$$

The gEUD can be seen as a generalized average and it is well-known that it reduces to the average tumor-dose when $p = 1$ and to the minimum tumor-dose as $p \rightarrow -\infty$ (see [81]). Thus, both our approach of maximizing the BE of the average tumor-dose and the aforementioned alternative approach of maximizing the BE of the minimum tumor-dose can be seen as maximizing the BE of a particular tumor-gEUD. Consequently, both these choices of objective functions are consistent with the original motivation for the concept of EUD (see [82]). We decided not to use the alternative maximin objective function because, as one would expect in a formulation that maximizes the worst-case BE, it led to fluence-maps that were too conservative in our preliminary numerical experiments.

Another possibility for the objective function is to maximize the average BE of the doses delivered to all tumor-voxels. This alternative, nonconvex objective is perhaps biologically more meaningful because it adds the BE over individual voxels to quantify the combined effect. Interestingly, the standard approach for solving this nonconvex problem reduces to using our objective in expression (4.2). The reasoning for this is as follows. The alternative nonconvex objective function cannot be equivalently written in a convex form. The resulting optimization problem is thus computationally intractable. The standard approach for approximate solution of such problems is to instead solve a relaxation of the objective function (see, for example, the literature review in [63]). The standard relaxation of the alternative objective is based on the fact that the sum of squares is no bigger than the square of the sum; and this relaxation of the alternative objective function yields our objective function in (4.2).

We next describe our constraints in detail.

4.2.2 Normal tissue tolerance and fluence-map smoothness constraints

Let O_1, O_2, \dots, O_M denote the M different normal tissues under consideration. Let $\mathcal{M} = \{1, \dots, M\}$ be the set of indices of these normal tissues. For $m \in \mathcal{M}$, let n_m denote the number of voxels in O_m . These voxels are indexed by $j = 1, 2, \dots, n_m$. Let \mathcal{N}_m denote the set $\{1, 2, \dots, n_m\}$ of these voxels. All normal tissue voxels are assumed to have equal volume. Let A^m be the $n_m \times k$, nonnegative dose deposition matrix for O_m . Let A_j^m be the j th row of this matrix; this is the row that corresponds to the j th voxel in O_m . That is, $A_j^m u$ is the dose delivered to the j th voxel in O_m in each session. Let α_m and β_m be the parameters of the LQ model for normal tissue O_m and we define $\rho_m \triangleq 1/(\alpha_m/\beta_m)$ for brevity. Then, the BED of the dose delivered to the j th voxel in O_m over N fractions is given by

$$N(A_j^m u) + \rho_m N(A_j^m u)^2. \quad (4.4)$$

Our model includes the three most common types of constraints on normal tissues: maximum dose constraints, mean dose constraints, and dose-volume constraints.

Maximum BED constraints for serial normal tissues

Let $\mathcal{M}_1 \subseteq \mathcal{M}$ be the set of indices of serial normal tissues for which we wish to include maximum dose constraints. These are the normal tissues whose function is hampered even when a small region is damaged by radiation. Suppose, for any $m \in \mathcal{M}_1$, that a total dose D_{\max}^m is known to be tolerated by each voxel in O_m if administered in N_{conv}^m equal-dose fractions. The BED of this schedule equals

$$\text{BED}_{\max}^m = D_{\max}^m (1 + \rho_m (D_{\max}^m / N_{\text{conv}}^m)). \quad (4.5)$$

We use the standard approach of comparing normal tissue BED. Thus, a dose of $N(A_j^m u)$ over N fractions can be tolerated by the j th voxel in normal tissue O_m if

$$N(A_j^m u) + N \rho_m (A_j^m u)^2 \leq \text{BED}_{\max}^m, \quad \forall j \in \mathcal{N}_m. \quad (4.6)$$

Thus, for each $m \in \mathcal{M}_1$, our problem formulation will include constraints (5.7).

Mean BED constraints for parallel normal tissues

Let $\mathcal{M}_2 \subseteq \mathcal{M}$ be the set of indices of parallel normal tissues for which we wish to include mean dose constraints. These are the normal tissues where a sufficiently small portion can be damaged without affecting the organ function. Suppose, for any $m \in \mathcal{M}_2$, that mean dose D_{mean}^m is known to be tolerated by O_m if administered in N_{conv}^m equal-dose fractions. The BED of this mean dose is given by

$$\text{BED}_{\text{mean}}^m = D_{\text{mean}}^m (1 + \rho_m (D_{\text{mean}}^m / N_{\text{conv}}^m)). \quad (4.7)$$

Then, for normal tissue O_m , we write the mean BED constraint as

$$\frac{N \sum_{j=1}^{n_m} (A_j^m u) + N \rho_m \sum_{j=1}^{n_m} (A_j^m u)^2}{n_m} \leq \text{BED}_{\text{mean}}^m. \quad (4.8)$$

There is a subtle difference between the left hand side of this inequality and its right hand side. The left hand side quantifies the average BED of doses delivered to different voxels; the right hand side quantifies the BED of the average dose delivered to different voxels. An alternative way to express the left hand side is to use the BED of average dose. That is, to write

$$N(\bar{A}^m u) + N \rho_m (\bar{A}^m u)^2 \leq \text{BED}_{\text{mean}}^m, \quad (4.9)$$

where $\bar{A}^m u \triangleq \sum_{j=1}^{n_m} A_j^m u / n_m$ is the average dose. However, the left hand side of our quadratic constraint (4.8) is biologically more meaningful than the alternative expression in (4.9) because it adds the BED for individual voxels to construct the combined effect. The left hand side in (4.8) also leads to a more conservative fluence-map because this left hand side is an upper bound on the left hand side in (4.9) owing to the aforementioned property about sum of squares. On the other hand, there are two benefits to using the alternative expression in (4.9). The first is that it makes the left hand side consistent with the right hand side. Secondly, this alternative constraint can be equivalently expressed in terms of a linear constraint on the average dose (in Section 5.3, we show how to do this in the context of inequality (5.7), and

the conversion would be identical for inequality (4.9)). This latter will be a computational advantage as a single linear constraint will be easier to handle than the convex quadratic constraint in (4.8). Nevertheless, we decided to use the constraint in (4.8) for its aforementioned clinically sound foundation despite it being computationally more challenging than the linearized equivalent of constraint (4.9).

Dose-volume constraints for parallel normal tissues

Let $\mathcal{M}_3 \subseteq \mathcal{M}$ be the set of indices of normal tissues with dose-volume constraints. Suppose, for any $m \in \mathcal{M}_3$, that no more than a volume fraction ϕ_m of normal tissue O_m can receive a dose more than D_{dv}^m if administered in N_{conv}^m fractions. The BED of total dose D_{dv}^m administered in N_{conv}^m equal-dosage fractions is given by

$$\text{BED}_{\text{dv}}^m = D_{\text{dv}}^m (1 + \rho_m (D_{\text{dv}}^m / N_{\text{conv}}^m)). \quad (4.10)$$

Since all voxels in O_m have equal volume, the volume fraction is the same as the voxel fraction. For each $m \in \mathcal{M}_3$ and for $j = 1, 2, \dots, n_m$, we thus define binary-valued functions $f_j^m(N, u)$ as

$$f_j^m(N, u) = \begin{cases} 1 & \text{if } N(A_j^m u) + N\rho_m(A_j^m u)^2 > \text{BED}_{\text{dv}}^m, \\ 0 & \text{if } N(A_j^m u) + N\rho_m(A_j^m u)^2 \leq \text{BED}_{\text{dv}}^m. \end{cases} \quad (4.11)$$

In words, $f_j^m(N, u)$ is one if the BED of dose delivered by fluence-map u to voxel j in N sessions exceeds the tolerance BED_{dv}^m ; $f_j^m(N, u)$ is zero otherwise. We use the integer K_m to denote $\lfloor n_m \phi_m \rfloor$, that is, the largest integer that is at most $n_m \phi_m$. Then, the dose-volume constraints are written in our optimization model as

$$\sum_{j=1}^{n_m} f_j^m(N, u) \leq K_m, \quad m \in \mathcal{M}_3. \quad (4.12)$$

These constraint ensure that there are at most K_m voxels for which $N(A_j^m u) + N\rho_m(A_j^m u)^2 > \text{BED}_{\text{dv}}^m$; in other words, there are at least $n_m - K_m$ voxels for which $N(A_j^m u) + N\rho_m(A_j^m u)^2 \leq \text{BED}_{\text{dv}}^m$. Here, for simplicity, we have assumed that there

is at most one dose-volume constraint for each normal tissue. This assumption is not needed anywhere in our algorithm and hence can be removed. In fact, we do this in our computational experiments for prostate cancer in Section 5.4.

Fluence-map smoothness constraints

To ensure that the intensity profile is deliverable in practice using a multi-leaf collimator, we put a smoothness constraint on each radiation field [83, 84]. In particular, for each radiation field, we bound the absolute relative difference between intensities of each pair of nearest neighbor beamlets by a fraction ϵ . Then the smoothness constraints can be written compactly in matrix format as $Su \leq 0$, where S is a block diagonal matrix with entries $-(1 + \epsilon)$, $(1 - \epsilon)$, -1 , 0 , $+1$ at appropriate locations. We are now ready to provide our complete optimization model.

4.2.3 Complete optimization model

Based on the above discussion, we formulate the optimal fractionation problem as

$$(P) F^* = \max_{N,u} N\alpha_0(\bar{A}u) + N\beta_0(\bar{A}u)^2 - \tau(N), \quad (4.13)$$

$$N(A_j^m u) + N\rho_m(A_j^m u)^2 \leq \text{BED}_{\max}^m, \quad \forall j \in \mathcal{N}_m, m \in \mathcal{M}_1, \quad (4.14)$$

$$N \sum_{j=1}^{n_m} (A_j^m u) + N\rho_m \sum_{j=1}^{n_m} (A_j^m u)^2 \leq n_m \text{BED}_{\text{mean}}^m, \quad m \in \mathcal{M}_2, \quad (4.15)$$

$$\sum_{j=1}^{n_m} f_j^m(N, u) \leq K_m, \quad m \in \mathcal{M}_3, \quad (4.16)$$

$$Su \leq 0, \quad (4.17)$$

$$u \geq 0, \quad (4.18)$$

$$1 \leq N \leq N_{\max}, \text{ integer.} \quad (4.19)$$

This formulation includes all constraints that are essential for capturing the trade-off between tumor-BE and the BED for serial and parallel normal tissue. Additional constraints can also be added to this formulation if required by the treatment protocol.

For instance, a minimum dose constraint on tumor could be easily added to avoid cold-spots; similarly, a maximum dose constraint on tumor could be added to increase dose-uniformity. These constraints, being linear, pose no additional computational hurdles. Moreover, these additional constraints might make our formulation more relevant for clinical practice. However, we decided not to include these constraints in the formulation for two reasons: (i) in our preliminary numerical experiments, we observed that the qualitative trends discovered in our sensitivity analyses continue to hold even when these constraints are included; and (ii) our smoothness constraints at least to some extent attempt to ensure dose-uniformity.

Observe that for each fixed N , the functions $f_j^m(N, u)$ are discontinuous in u . In fact, it is well-known in the IMRT literature that dose-volume constraints are difficult to handle [6]. In realistic instances of (P) , the number of beamlets k is likely to equal a few thousand. The number of constraints can also be in the tens or hundreds of thousand depending on the total number of normal tissue voxels. As a result, exact solution of (P) is computationally intractable. We next develop an algorithm for efficient, approximate solution of (P) . This method uses a simple constraint generation approach that is rooted in the fact that when N is fixed, if we drop the dose-volume constraints, then the resulting problem, although seemingly nonconvex, can be equivalently re-written as a linear program with convex, quadratic constraints. We emphasize here that if the treatment protocol does not include any dose-volume constraints, then our method produces an exact solution to (P) .

4.3 An efficient solution method

We first define a sequence of problems $P(N)$, obtained by fixing N at $1, 2, \dots, N_{\max}$ in (P) . We have,

$$P(N) \quad F^*(N) = \max_u \quad N\alpha_0(\bar{A}u) + N\beta_0(\bar{A}u)^2, \quad (4.20)$$

$$N(A_j^m u) + N\rho_m(A_j^m u)^2 \leq \text{BED}_{\max}^m, \quad \forall j \in \mathcal{N}_m, \quad m \in \mathcal{M}_1, \quad (4.21)$$

$$N \sum_{j=1}^{n_m} (A_j^m u) + \rho_m N \sum_{j=1}^{n_m} (A_j^m u)^2 \leq n_m \text{BED}_{\text{mean}}^m, \quad m \in \mathcal{M}_2, \quad (4.22)$$

$$\sum_{j=1}^{n_m} f_j^m(N, \vec{u}(N)) \leq K_m, \quad m \in \mathcal{M}_3, \quad (4.23)$$

$$Su \leq 0, \quad (4.24)$$

$$u \geq 0. \quad (4.25)$$

Note that

$$F^* = \max_{N \in \{1, 2, \dots, N_{\max}\}} F^*(N) - \tau(N). \quad (4.26)$$

Thus, problem (P) can be solved by first solving the sequence of problems $P(N)$, for $N = 1, 2, \dots, N_{\max}$, and then choosing an N that yields the best tumor-BE and using the corresponding optimal fluence-map.

Without loss of generality, we assume that for every normal tissue $m \in \mathcal{M}$, every row of the dose deposition matrix A^m has at least one strictly positive entry. For if not, then the normal tissue voxel corresponding to a zero row in A^m can be removed from further consideration as radiation does not reach that voxel. Also without loss of generality, we assume that there exists either a serial normal tissue $m \in \mathcal{M}_1$ or a parallel normal tissue $m \in \mathcal{M}_2$ with the property that every column of its dose deposition matrix A^m includes at least one strictly positive entry. For if there exists a column c without any strictly positive entries, then it turns out that the c th component of u , denoted u_c , can be increased arbitrarily without damaging any normal tissue. This assumption is met in practice for example when maximum dose constraints are included on the unspecified normal tissue because each radiation beamlet must pass

through at least some unspecified normal tissue and hence the corresponding row of A^m will have a strictly positive entry.

Lemma 4.3.1. *Problem $P(N)$ has an optimal solution for each $N \geq 1$; this implies that (P) has an optimal solution as well.*

Proof. For every $m \in \mathcal{M}_3$, dose-volume constraints (4.23) imply that, the BED for at least $L_m \triangleq n_m - K_m$ and at most n_m voxels in O_m should be less than BED_{dv}^m . As a result, there are $W_m = \sum_{l=L_m}^{n_m} \binom{n_m}{l}$ ways to express the dose-volume constraint (4.23). Each of these enforce that some combination of $L_m \leq l \leq n_m$ out of the n_m voxels in O_m satisfy the appropriate BED limit. We index these distinct ways by $w_m = 1, 2, \dots, W_m$ and let $\mathcal{J}_{w_m}(N) \subseteq \mathcal{N}_m$ be the set of voxels in O_m for which the dose-limit is enforced in the w_m th way of expressing the dose-volume constraint for O_m . Thus, there are a total of $W = \prod_{m \in \mathcal{M}_3} W_m$ ways to express our dose-volume constraints. Let $\vec{w} \triangleq (w_1, w_2, \dots, w_{|\mathcal{M}_3|})$. We create W subproblems from $P(N)$, each representing one of these W ways of expression.

$$P(\vec{w}; N) \max_{\vec{u}(N)} N\alpha_0(\bar{A}u) + N\beta_0(\bar{A}u)^2, \quad (4.27)$$

$$N(A_j^m u) + N\rho_m(A_j^m u)^2 \leq \text{BED}_{\text{max}}^m, \quad \forall j \in \mathcal{N}_m, m \in \mathcal{M}_1, \quad (4.28)$$

$$N \sum_{j=1}^{n_m} (A_j^m u) + N\rho_m \sum_{j=1}^{n_m} (A_j^m u)^2 \leq n_m \text{BED}_{\text{mean}}^m, \quad m \in \mathcal{M}_2, \quad (4.29)$$

$$N(A_j^m u) + N\rho_m(A_j^m u)^2 \leq \text{BED}_{\text{dv}}^m, \quad j \in \mathcal{J}_{w_m}(N), m \in \mathcal{M}_3, \quad (4.30)$$

$$Su \leq 0, \quad (4.31)$$

$$u \geq 0. \quad (4.32)$$

An optimal solution to $P(N)$ can be recovered by choosing the best among optimal solutions to these subproblems. Thus, we prove that each of the above W problems has an optimal solution. Let $U(\vec{w}; N) \subset \mathfrak{R}^k$ denote the set of feasible fluence-map vectors u . This set is nonempty because the trivial fluence-map $u = \vec{0}$ is feasible to $P(\vec{w}; N)$. Suppose a serial normal tissue $m \in \mathcal{M}_1$ has a strictly positive entry in each

column of its dose deposition matrix A^m (if instead a parallel normal tissue possesses this property, then the proof can be modified easily). For any $l = 1, 2, \dots, k$, let \mathcal{I}_l denote the set of rows of A^m in which the entry in the l th column is strictly positive. For all $i \in \mathcal{I}_l$, we denote the corresponding positive entries of A^m by $A_{i,l}^m$. Since all entries in A^m are nonnegative, the maximum dose constraint on this normal tissue implies that $A_{i,l}^m u_l \leq \text{BED}_{\max}^m$ for $t = 1, 2, \dots, N$. In other words, $u_l \leq \frac{\text{BED}_{\max}^m}{N \min_{i \in \mathcal{I}_l} A_{i,l}^m}$. Thus, the feasible region is bounded. It is also closed as all constraint functions are continuous. Moreover, the objective function in $P(\vec{w}; N)$ is also continuous. Thus, $P(\vec{w}; N)$ has an optimal solution as claimed. \square

Now observe that because the objective in $P(N)$ is increasing in $\bar{A}u$, it is equivalent to maximizing $\bar{A}u$. Thus, we rewrite $P(N)$ as

$$P(N) \max_u \bar{A}u, \quad (4.33)$$

$$N(A_j^m u) + \rho_m N(A_j^m u)^2 \leq \text{BED}_{\max}^m, \quad \forall j \in \mathcal{N}_m, \quad m \in \mathcal{M}_1, \quad (4.34)$$

$$N \sum_{j=1}^{n_m} (A_j^m u) + \rho_m N \sum_{j=1}^{n_m} (A_j^m u)^2 \leq n_m \text{BED}_{\text{mean}}^m, \quad m \in \mathcal{M}_2, \quad (4.35)$$

$$\sum_{j=1}^{n_m} f_j^m(N, u) \leq K_m, \quad m \in \mathcal{M}_3, \quad (4.36)$$

$$Su \leq 0, \quad (4.37)$$

$$u \geq 0. \quad (4.38)$$

Moreover, since $A_j^m u \geq 0$, constraints (4.34) can be equivalently rewritten as linear constraints wherein the right hand side is obtained by solving a quadratic equation.

This yields the equivalent problem

$$P(N) \max_u \bar{A}u, \quad (4.39)$$

$$A_j^m u \leq \frac{-1 + \sqrt{1 + 4\rho_m \text{BED}_{\max}^m / N}}{2\rho_m}, \quad \forall j \in \mathcal{N}_m, m \in \mathcal{M}_1, \quad (4.40)$$

$$N \sum_{j=1}^{n_m} (A_j^m u) + \rho_m N \sum_{j=1}^{n_m} (A_j^m u)^2 \leq n_m \text{BED}_{\text{mean}}^m, \quad m \in \mathcal{M}_2, \quad (4.41)$$

$$\sum_{j=1}^{n_m} f_j^m(N, u) \leq K_m, \quad m \in \mathcal{M}_3, \quad (4.42)$$

$$Su \leq 0, \quad (4.43)$$

$$u \geq 0. \quad (4.44)$$

The objective function in this problem is linear; constraints (4.40) and (4.43) are linear; constraints (4.41) are convex, quadratic. Thus, the only remaining computational challenge is posed by the dose-volume constraints (4.42). We therefore propose a simple and natural constraint generation approach to surmount this difficulty. In particular, we first solve $P(N)$ *without* constraints (4.42). Suppose \hat{u} is an optimal solution to this problem. Then, for each $m \in \mathcal{M}_3$, we find $n_m - K_m$ voxels that receive the smallest doses among the n_m voxels in O_m under fluence-map \hat{u} . Let subset $\mathcal{N}_m(\hat{u}, N) \subseteq \mathcal{N}_m$ denote this group of voxels. We then re-solve $P(N)$ but this time by replacing constraints (4.42) with tolerance limits on all voxels in the set $\mathcal{N}_m(\hat{u}, N)$. Finally, we note that these tolerance limits, which appeared in definition (5.11) of functions $f_j(u, N)$, can be equivalently re-written as linear constraints whose right

hand side is obtained by solving a quadratic equation. This yields the problem

$$Q(N) \ G^*(N) \triangleq \max_u \bar{A}u, \quad (4.45)$$

$$A_j^m u \leq \frac{-1 + \sqrt{1 + 4\rho_m \text{BED}_{\max}^m / N}}{2\rho_m}, \quad \forall j \in \mathcal{N}_m, \ m \in \mathcal{M}_1, \quad (4.46)$$

$$N \sum_{j=1}^{n_m} (A_j^m u) + N\rho_m \sum_{j=1}^{n_m} (A_j^m u)^2 \leq n_m \text{BED}_{\text{mean}}^m, \quad m \in \mathcal{M}_2, \quad (4.47)$$

$$A_j^m u \leq \frac{-1 + \sqrt{1 + 4\rho_m \text{BED}_{\text{dv}}^m / N}}{2\rho_m}, \quad \forall j \in \mathcal{N}_m(\hat{u}, N), \ m \in \mathcal{M}_3, \quad (4.48)$$

$$Su \leq 0, \quad (4.49)$$

$$u \geq 0. \quad (4.50)$$

Problem $Q(N)$ has a linear objective and only includes linear or convex quadratic constraints and hence it can be solved efficiently. We conclude this section by summarizing our overall algorithm for approximate solution of (P) .

Algorithm for solving (P)

1. For $N = 1, 2, \dots, N_{\max}$,
 - (a) solve problem $P(N)$ as in (4.39)-(4.44) but without dose-volume constraints (4.42); let \hat{u} denote a fluence-map that is optimal to this problem, and let $\mathcal{N}_m(\hat{u}, N)$, for each $m \in \mathcal{M}_3$, be the set of voxels that receive the smallest doses among the n_m voxels in O_m under fluence-map \hat{u} ;
 - (b) solve problem $Q(N)$ to denote its optimal solution by $u^*(N)$ and use its optimal value $G^*(N)$ to approximate the optimal value $F^*(N)$ of problem $P(N)$ formulated in (4.20)-(4.25) as $F^*(N) \approx \bar{F}(N) \triangleq N\alpha_0 G^*(N) + N\beta_0(G^*(N))^2$;

end the loop over N .
2. Use $\bar{F}(N)$ as an approximation to $F^*(N)$ in (4.26) to obtain an optimal number of fractions N^* and use the corresponding optimal fluence-map $u^*(N^*)$ in each fraction.

In the next section, we apply this algorithm to ten test cases in head-and-neck and prostate cancer. We perform sensitivity analyses and also quantify the potential benefit of our spatiotemporally integrated approach.

4.4 Results

In this section, our computational experiments are designed to illustrate two key points — (i) the qualitative trends in the effect of various problem parameters on the optimal number of fractions and the corresponding optimal tumor-BE are identical to those in our spatiotemporally separated model in Chapter 2 and also to those in some of the earlier stylized models [44–50, 54, 55, 61, 62]; and, crucially, (ii) the spatiotemporally integrated approach achieves a higher tumor-BE than both a model

that uses a spatially optimized IMRT fluence-map without optimizing the number of fractions, and our spatiotemporally separated model in Chapter 2. The first of these two points is studied in Section 4.4.2, and it serves as an indirect validation of our model. The second point, studied in Section 4.4.3, hints at the potential increase in treatment efficacy, within the LQ framework, that could be obtained by solving our integrated model instead of solving existing stylized models in the literature.

4.4.1 Description of test cases

We first describe the ten test cases that were used in our experiments. Five of these were head-and-neck cases and the other five were prostate. All test cases were generated using our in-house phantom creator software PhanC written in MATLAB (see Appendix A), and they were identical to the ones we used in Chapter 2. These test cases were three-dimensional and were carefully developed to be representative of clinical scenarios in terms of geometry and size. Specifically, our test cases were similar to those discussed in [6, 85–88]. All voxels were $3 \times 3 \times 3 \text{ mm}^3$. The beamlet resolution was $5 \times 5 \text{ mm}^2$ in all cases. All cases used equally spaced coplanar beams (seven for head-and-neck and five for prostate). Computer simulations were performed on a 3.1 GHz iMac desktop with 16 GB RAM using the MATLAB convex optimization toolbox CVX [73].

Head-and-neck cancer cases

All cases used seven beams and included spinal cord, brainstem, left and right parotids and unspecified normal tissue between these critical organs. The total number of voxels in the head-and-neck target and in the normal tissues, and the total number of beamlets is shown in Table 5.1 below.

The conventional fractionation schedule was assumed to include $N_{\text{conv}} = 35$ fractions. While formulating problem (P), we included maximum dose constraints for spinal cord, brainstem and unspecified normal tissue. A dose-volume constraint for

case #	# of beamlets (k)	# of tumor voxels (n)	# of normal tissues voxels
1	3910	27576	67386
2	3888	31930	67270
3	4128	36320	76160
4	3003	22372	53176
5	3256	28638	64713

Table 4.2: Description of the geometry used in head-and-neck cancer cases. This data is identical to Chapter 2.

unspecified normal tissue was also added. Mean dose constrains were used for left and right parotids. The tolerance dose values for various normal tissues were similar to [1, 70–72] and identical to Chapter 2. These are listed in Table 5.2 below. All radiobiological parameter values used are listed in Table 4.4 along with their source.

Normal tissue	D_{\max} (Gy)	D_{mean} (Gy)	D_{dv} (Gy), ϕ
Spinal cord	45	N/A	N/A
Brainstem	50	N/A	N/A
Left and right parotids	N/A	28	N/A
Unspecified normal tissue	77	N/A	70, 0.05

Table 4.3: Tolerance doses for various normal tissues in our head-and-neck test cases where the dose is administered in $N_{\text{conv}} = 35$ equal-dose fractions. Recall that for dose-volume type constraints no more than a volume fraction ϕ of the normal tissue can receive dose more than D_{dv} . This treatment protocol is similar, for example, to [1].

Structure	parameter	value(s)	source
Head-and-neck tumor	α	0.35 Gy ⁻¹	[44, 47–49]
	α/β	{8, 10, 12} Gy	[48, 65, 75]
	T_{double}	{2, 3, 5, 8, 10, 20, 40, 50} days	[44, 48, 51, 75]
	T_{lag}	{7, 14, 21, 28, 35} days	[48]
Unspecified normal tissue	α/β	3 Gy	[89–91]
Parotids	α/β	{3,4,5,6} Gy	[44, 48, 51]
Spinal cord and brainstem	α/β	{2,3,4,5,6} Gy	[44, 48, 51]

Table 4.4: Values of various radiobiological parameters used in sensitivity analyses for head-and-neck cancer.

Prostate cancer cases

All cases used five beams and included rectum, bladder, left and right femurs and unspecified normal tissue between these critical organs. The total number of voxels in the prostate target and in the normal tissues, and the total number of beamlets is shown in Table 4.5 below.

case #	# of beamlets (k)	# of tumor voxels (n)	Number of normal tissues voxels
1	938	6180	145545
2	847	7225	326703
3	935	4628	367656
4	930	4956	314544
5	870	4840	269450

Table 4.5: Description of the geometry used in prostate cancer cases. This data is identical to Chapter 2.

The conventional fractionation schedule was assumed to include $N_{\text{conv}} = 45$ fractions. In our formulation of problem (P), we included a maximum dose constraint for unspecified normal tissue, and included dose-volume constraints for all normal tissues. There were no mean dose constraints. The dose-volume constraints for all normal tissues were similar to [1, 70–72] and identical to Chapter 2. These are listed in Table 4.6. All radiobiological parameter values used are listed in Table 4.7 along with their source.

Normal tissue	D_{dv} (Gy), ϕ
Rectum	50, 0.5
	85, 0.05
Bladder	50, 0.5
	70, 0.35
	75, 0.25
	80, 0.15
	89, 0.05
Femurs	65, 0.05
Unspecified normal tissue	81, 0.05

Table 4.6: Dose-volume constraints for various normal tissues when dose is administered in $N_{\text{conv}} = 45$ equal-dose fractions — no more than a volume fraction ϕ of the normal tissue can receive dose more than D_{dv} . This treatment protocol is similar, for example, to [1].

4.4.2 Sensitivity to tumor doubling time and to time-lag before proliferation

As in previous work in this area, we found that the optimal number of fractions is smaller for faster growing tumors as characterized by smaller values of the tumor doubling time T_{double} . As an example, the tumor-BE is plotted against N for various

Structure	parameter	value(s)	source
Prostate tumor	α	0.15 Gy ⁻¹	[44, 47–49]
	α/β	{2, 3, 4, 6} Gy	[45, 68, 69, 76]
	T_{double}	{5, 20, 40, 60, 80} days	[45, 68, 78, 79]
	T_{lag}	{7, 14, 21, 28, 35} days	[78, 79]
Unspecified normal tissue	α/β	3 Gy	[89–91]
Femurs	α/β	{3,4,5,6} Gy	[45, 77, 90]
Rectum and bladder	α/β	{3,4,5,6} Gy	[41, 45, 77, 90]

Table 4.7: Values of various radiobiological parameters used in sensitivity analyses for prostate cancer.

values of T_{double} in Figures 4.1(a), (b) for our head-and-neck case 1 and prostate case 1. Also, the optimal tumor-BE itself is smaller for faster growing tumors.

Our simulations showed that the tumor-BE can sometimes grow very slowly with N before it starts decreasing. We therefore tracked the smallest number of fractions at which the tumor-BE reached 99% of its optimal value. We denote this number of fractions by N_{99}^* . The range* of this number over all combinations of tumor α/β ratios and normal tissue α/β ratios for our five head-and-neck cases is reported in Table 4.8. The table shows that the ranges did not change significantly across different cases. The table also shows that for some fixed values of T_{double} , the range of N_{99}^* is somewhat broad; this is especially true for slower proliferating tumors as characterized by larger T_{double} values. These broad ranges for N_{99}^* resulted from our choice of a somewhat broad range for α/β values. To further illustrate this point, we present a more detailed set of results for head-and-neck case 1 as an example in Table 4.9. These results were obtained for $T_{\text{double}} = 10$ days and $T_{\text{lag}} = 7$ days by fixing the α/β ratios for the

Based on our analysis in Chapter 2, $1 + T_{\text{lag}}$ is likely to be a lower bound on the optimal number of fractions; thus, when N_{99}^ was smaller than $1 + T_{\text{lag}}$, it was reset to $1 + T_{\text{lag}}$.




Figure 4.1: Sensitivity of the optimal number of fractions and the optimal tumor-BE to tumor doubling time T_{double} (days) when $T_{\text{lag}} = 7$ days. (a) Head-and-neck case 1; tumor α/β was 10 Gy, and the α/β ratio for all normal tissues was 3 Gy. (b) Prostate case 1; tumor α/β was 6 Gy, and the α/β ratio for all normal tissues was 3 Gy.

unspecified normal tissue, spinal cord and brainstem at 3 Gy and then varying the α/β ratios for the two parotids over the set $\{3, 4, 5, 6\}$ Gy and the α/β ratios for the tumor over the set $\{8, 10, 12\}$ Gy. The table shows that for a fixed value of parotid α/β ratio, N_{99}^* increases with increasing values of tumor α/β . Similarly, for a fixed value of tumor α/β ratio, N_{99}^* decreases with increasing values of parotid α/β . In summary, N_{99}^* becomes larger as the relative difference between tumor and parotid α/β grows.

Head-and-neck case #	T_{double} (days)							
	2	3	5	8	10	20	40	50
1	8	8-11	8-19	10-28	12-33	23-55	38-73	44-75
2	8	8-13	9-22	15-33	19-39	34-63	52-76	58-78
3	8	8-13	9-22	15-32	18-38	32-63	52-76	58-78
4	8	8-12	8-20	13-30	16-36	29-59	49-75	55-77
5	8-9	8-14	10-22	16-34	20-40	36-64	55-77	60-78

Table 4.8: Range of N_{99}^* across five head-and-neck cases with $T_{\text{lag}} = 7$ days and all combinations of tumor and normal tissue α/β ratios.

In all prostate cases where the tumor α/β ratio was 2 or 3, it was optimal to administer a single fraction for all combinations of α/β ratios for normal tissues as has been noted in the existing literature on stylized models. The range of N_{99}^* for all combinations of α/β ratios for normal tissues and $\alpha/\beta = 4, 6$ for tumor is shown in Table 4.10 for our five test cases. The table shows that the ranges did not change significantly across different cases. To shed further light on the somewhat broad range of N_{99}^* values for some fixed values of T_{double} , we present more detailed results for prostate case 1 as an example in Table 4.11. These results were obtained by fixing $T_{\text{double}} = 30$ days and $T_{\text{lag}} = 7$ days, fixing the unspecified normal tissue and femur

	tumor α/β		
parotid α/β	8	10	12
3	22	29	33
4	18	25	30
5	15	22	27
6	12	20	25

Table 4.9: Range of N_{99}^* for head-and-neck case 1 with $T_{\text{double}} = 10$ days, $T_{\text{lag}} = 7$ days, and different combinations of tumor and parotid α/β ratios with α/β for unspecified normal tissue, spinal cord, and brainstem fixed at 3 Gy.

α/β at 3 Gy, varying the bladder and rectum α/β ratios over the set $\{3, 4, 5, 6\}$ Gy, and varying the tumor α/β ratio over the set $\{4, 6\}$ Gy. As in head-and-neck, N_{99}^* increases with increasing relative difference between tumor and normal tissue α/β .

	T_{double} (days)							
Prostate case #	5	8	10	20	30	40	60	80
1	8	8-12	8-16	8-31	8-44	8-56	8-68	8-72
2	8	8-12	8-16	8-32	8-46	8-58	8-69	8-73
3	8	8-11	8-15	8-30	8-43	8-54	8-67	8-71
4	8	8-12	8-16	8-32	8-46	8-57	8-68	8-72
5	8	8-12	8-16	8-32	8-46	8-58	8-69	8-73

Table 4.10: Range of N_{99}^* for five prostate cases with $T_{\text{lag}} = 7$ days and $\alpha/\beta = 4, 6$ for tumor and all combinations of normal tissue α/β ratios.

We also investigated the sensitivity of the optimal number of fractions to T_{lag} .

	tumor α/β	
rectum and bladder α/β	4	6
3	8	44
4	8	26
5	8	20
6	8	18

Table 4.11: Range of N_{99}^* for prostate case 1 with $T_{\text{double}} = 30$ days, $T_{\text{lag}} = 7$ days, and different combinations of tumor, rectum and bladder α/β ratios with α/β for unspecified normal tissue and femurs fixed at 3 Gy.

Our experiments confirmed (see Tables 4.12 and 4.13) that for small values of T_{double} relative to T_{lag} , the optimal number of fractions is equal to $1 + T_{\text{lag}}$; for larger value of T_{double} , the optimal number of fractions is not sensitive to T_{lag} . As in Tables 4.8 and 4.10, the range of N_{99}^* did not vary across different test cases; therefore, in Tables 4.12 and 4.13, we simply listed the range of N_{99}^* over all head-and-neck and prostate cases, respectively, rather than separating them case-by-case. For each T_{double} , T_{lag} combination, the somewhat broad range of N_{99}^* again resulted from our broad range of α/β values as previously demonstrated in Tables 4.9 and 4.11.

4.4.3 Improvement in tumor-BE

Comparison with conventional IMRT

First, we compare the tumor-BE obtained by our spatiotemporally integrated approach with that from a conventional IMRT optimization model that does not optimize the number of fractions. The number of fractions and the prescription dose were fixed at $N_{\text{conv}}^{\text{hn}} = 35$, $D_{\text{presc}}^{\text{hn}} = 70$ Gy for head-and-neck cancer, and $N_{\text{conv}}^{\text{pr}} = 45$,

	T_{double} (days)							
T_{lag} (days)	2	3	5	8	10	20	40	50
14	15	15	15-22	15-33	15-40	22-64	38-76	44-78
21	22	22	22	22-33	22-40	22-64	38-76	44-78
28	29	29	29	29	29-33	29-55	38-73	44-75
35	36	36	36	36	36	36-55	38-73	44-75

Table 4.12: Range of N_{99}^* over five head-and-neck cases over all combinations of tumor and normal tissue α/β ratios.

	T_{double} (days)							
T_{lag} (days)	5	8	10	20	30	40	60	80
14	8-15	10-15	8-16	8-32	10-46	8-58	10-69	8-73
21	8-22	8-22	8-22	8-32	10-46	10-58	10-69	8-73
28	8-29	8-29	8-29	8-32	8-46	8-58	8-69	8-73
35	8-36	8-36	8-36	8-36	8-45	8-57	8-69	8-73

Table 4.13: Range of N_{99}^* over five prostate cases for $\alpha/\beta = 4, 6$ for tumor over all combinations of normal tissue α/β ratios.

$D_{\text{presc}}^{\text{pr}} = 81$ Gy for prostate cancer similar to [1]. The IMRT optimization model minimized the total squared-deviation of tumor voxel doses from the prescription doses. Maximum dose constraints for spinal cord and brainstem and mean dose constraints for parotids with dose tolerance levels tabulated in Section 4.4.1 were included for head-and-neck cases. Maximum dose constraints of 85, 89 and 65 Gy for rectum, bladder, and femurs, respectively were included for the prostate cases similar to [1]. Smoothness constraints were also included in all ten cases.

The resulting model for head-and-neck was as follows:

$$\begin{aligned}
(HN - CONV) \quad & \min \sum_{i=1}^n |A_i^{\text{hn}} u - D_{\text{presc}}^{\text{hn}} / N_{\text{conv}}^{\text{hn}}|^2 \\
& N_{\text{conv}}^{\text{hn}} A_j^{\text{spinal}} u \leq D_{\text{max}}^{\text{spinal}}, \text{ for all voxels } j \text{ in spinal cord,} \\
& N_{\text{conv}}^{\text{hn}} A_j^{\text{stem}} u \leq D_{\text{max}}^{\text{stem}}, \text{ for all voxels } j \text{ in brainstem,} \\
N_{\text{conv}}^{\text{hn}} \sum_{j=1}^{n_{\text{left parotid}}} \frac{A_j^{\text{left parotid}} u}{n_{\text{left parotid}}} & \leq D_{\text{mean}}^{\text{left parotid}}, \\
N_{\text{conv}}^{\text{hn}} \sum_{j=1}^{n_{\text{right parotid}}} \frac{A_j^{\text{right parotid}} u}{n_{\text{right parotid}}} & \leq D_{\text{mean}}^{\text{right parotid}}, \\
S_{\text{hn}} u & \leq 0, \\
u & \geq 0.
\end{aligned}$$

Suppose u^{hn} is an optimal solution to this problem. Then, its tumor-BE is computed using formula (4.2) as

$$N_{\text{conv}}^{\text{hn}} \alpha_0^{\text{hn}} (\bar{A}^{\text{hn}} u^{\text{hn}}) + N_{\text{conv}}^{\text{hn}} \beta_0^{\text{hn}} (\bar{A}^{\text{hn}} u^{\text{hn}})^2 - \tau(N_{\text{conv}}^{\text{hn}}).$$

The model for prostate was as follows:

$$\begin{aligned}
(PR - CONV) \quad & \min \sum_{i=1}^n |A_i^{\text{pr}} u - D_{\text{presc}}^{\text{pr}} / N_{\text{conv}}^{\text{pr}}|^2 \\
& N_{\text{conv}}^{\text{pr}} A_j^{\text{rectum}} u \leq D_{\text{max}}^{\text{rectum}}, \text{ for all voxels } j \text{ in rectum,} \\
& N_{\text{conv}}^{\text{pr}} A_j^{\text{bladder}} u \leq D_{\text{max}}^{\text{bladder}}, \text{ for all voxels } j \text{ in bladder,} \\
& N_{\text{conv}}^{\text{pr}} A_j^{\text{left femur}} u \leq D_{\text{max}}^{\text{left femur}}, \text{ for all voxels } j \text{ in left femur,} \\
& N_{\text{conv}}^{\text{pr}} A_j^{\text{right femur}} u \leq D_{\text{max}}^{\text{right femur}}, \text{ for all voxels } j \text{ in right femur,} \\
S_{\text{pr}} u & \leq 0, \\
u & \geq 0.
\end{aligned}$$

Suppose u^{pr} is an optimal solution to this problem. Then, its tumor-BE is computed using formula (4.2) as

$$N_{\text{conv}}^{\text{pr}} \alpha_0^{\text{pr}} (\bar{A}^{\text{pr}} u^{\text{pr}}) + N_{\text{conv}}^{\text{pr}} \beta_0^{\text{pr}} (\bar{A}^{\text{pr}} u^{\text{pr}})^2 - \tau(N_{\text{conv}}^{\text{pr}}).$$

The improvements in tumor-BE obtained by our spatiotemporally integrated approach over conventional IMRT are listed in Figures 4.2(a) and (b) for head-and-neck and prostate, respectively, for $T_{\text{lag}} = 7$ days as an example since the trends in improvements were not sensitive to T_{lag} .

Figure 4.2(a) for head-and-neck cases shows that the improvement in tumor-BE achieved by our spatiotemporally integrated approach over conventional IMRT first decreases and then increases with increasing values of T_{double} . Specifically, the potential benefit of our integrated approach is likely to be the largest for fast and slowly proliferating head-and-neck tumors. In particular, the most significant benefits are likely to be achieved for fast proliferating tumors. This overall trend is consistent with the fact that the optimal number of fractions obtained by our integrated approach is closest to the conventional value of 35 for moderately proliferating tumors as seen in Table 4.8.

Figure 4.2(b) for prostate cases shows a different trend. For each case, the improvement decreases with increasing values of T_{double} . The magnitude of the improvement itself seems to be bigger than that in head-and-neck. This is because, as seen in Table 4.10, the optimal number of fractions is significantly different from the conventional value of 45 for most parameter combinations.

A more detailed analysis of the trends in Figures 4.2(a), (b) is provided in Figures 4.3(a), (b) for head-and-neck case 1 and prostate case 1, respectively. For head-and-neck, Figure 4.3(a) shows that, for most values of T_{double} , the spatiotemporally integrated model achieves larger improvements over conventional IMRT for larger values of tumor α/β when all normal tissue α/β ratios are fixed at 3 Gy. This is because the difference between the optimal number of fractions derived from our integrated model and the conventional value of 35 increases for larger values of tumor α/β . For prostate, Figure 4.3(b) shows that the largest improvement is achieved when tumor α/β is 2 or 3 Gy when all normal tissue α/β ratios are fixed at 3 Gy. This is because, as stated earlier, a single fraction is optimal in those cases and this value is significantly different from the conventional value of 45.

Comparison with a spatiotemporally separated model

Second, we compare the tumor-BE achieved by our spatiotemporally integrated model with that attained by our spatiotemporally separated model in Chapter 2. In that model, the objective was to maximize the biological effect of average tumor dose subject to the same dose constraints on normal tissues as in our formulation (P). However, there, a spatially optimized IMRT fluence-map as in Section 4.4.3 was assumed to be available as input and fixed *a priori*; the number of fractions was then optimized with respect to this fluence-map using a sparing factors approach. This led to a closed-form formula for the average tumor-dose per fraction as a function of the number of fractions. The tumor-BE of this average tumor-dose was then characterized as a quasiconcave function of the number of fractions using formula (4.2). This resulted in a simple procedure for finding an optimal number of fractions.

The improvements in tumor-BE obtained by our spatiotemporally integrated approach over the spatiotemporally separated model are plotted in Figures 4.2(c) and (d) for head-and-neck and prostate, respectively, for $T_{\text{lag}} = 7$ days as an example. The qualitative trend in the improvements shown in Figure 4.2(c) for our head-and-neck cases is similar to that in Figure 4.2 (a). For our prostate cases plotted in Figure 4.2(d), the improvement is much less sensitive to T_{double} since the optimal number of fractions obtained by our integrated approach and the separated model is similar in most cases as this number is often 1 or $1 + T_{\text{lag}}$.

We conclude this section by elaborating on one observation that might seem surprising at first. A comparison of Figures 4.2(a) and (c) shows that the improvement achieved by our spatiotemporally integrated approach over the spatiotemporally separated approach is sometimes larger than that over conventional IMRT. This seems counterintuitive, because the spatiotemporally separated model uses a conventional, spatially optimized IMRT fluence-map for a conventional number of fractions as input and improves upon it by optimizing the number of fractions. Specifically, the improvement over IMRT should be at least as large as the improvement over the spa-

tiotemporally separated model. This intuition is correct. The apparent contradiction is rooted in the fact that the constraints in our spatiotemporally separated model in Chapter 2 were more conservative than in the conventional IMRT model used here. This is because, for reasons explained in Section 5.2.2, the spatiotemporally separated model in Chapter 2 also puts an upper bound on the average BED over all normal tissue voxels as in this paper, whereas conventional IMRT uses an upper bound on the average dose over all voxels.

4.5 Discussion

We built a spatiotemporally integrated model for the optimal fractionation problem using the LQ framework. Decision variables in this formulation correspond to the number of fractions as well as the fluence-map. We proposed an efficient, convex programming algorithm for approximate solution of this problem. Computer simulations on test cases suggest that this approach could potentially increase the tumor-BE within the LQ framework as compared to a conventional IMRT optimization model that does not optimize the number of fractions and also as compared to a recently proposed model that separates the spatial and the temporal components of the problem. Our sensitivity analyses provided insights into the effect of various model parameters on the resulting solutions.

Our algorithm tackles dose-volume constraints in a simple, natural, and efficient manner while preserving convexity. It should also be possible to use other existing methods to handle dose-volume constraints although we did not pursue this direction here as it is not the main focus of this paper (see, for example, the literature review and methods in [6]). We believe that our qualitative observations in this paper will hold for such alternative formulations.

Many alternative fractionation schedules have been tested clinically. Examples include once daily treatment but only over weekdays and twice daily treatment only over weekdays. Our formulation (P) can be easily modified to accommodate such

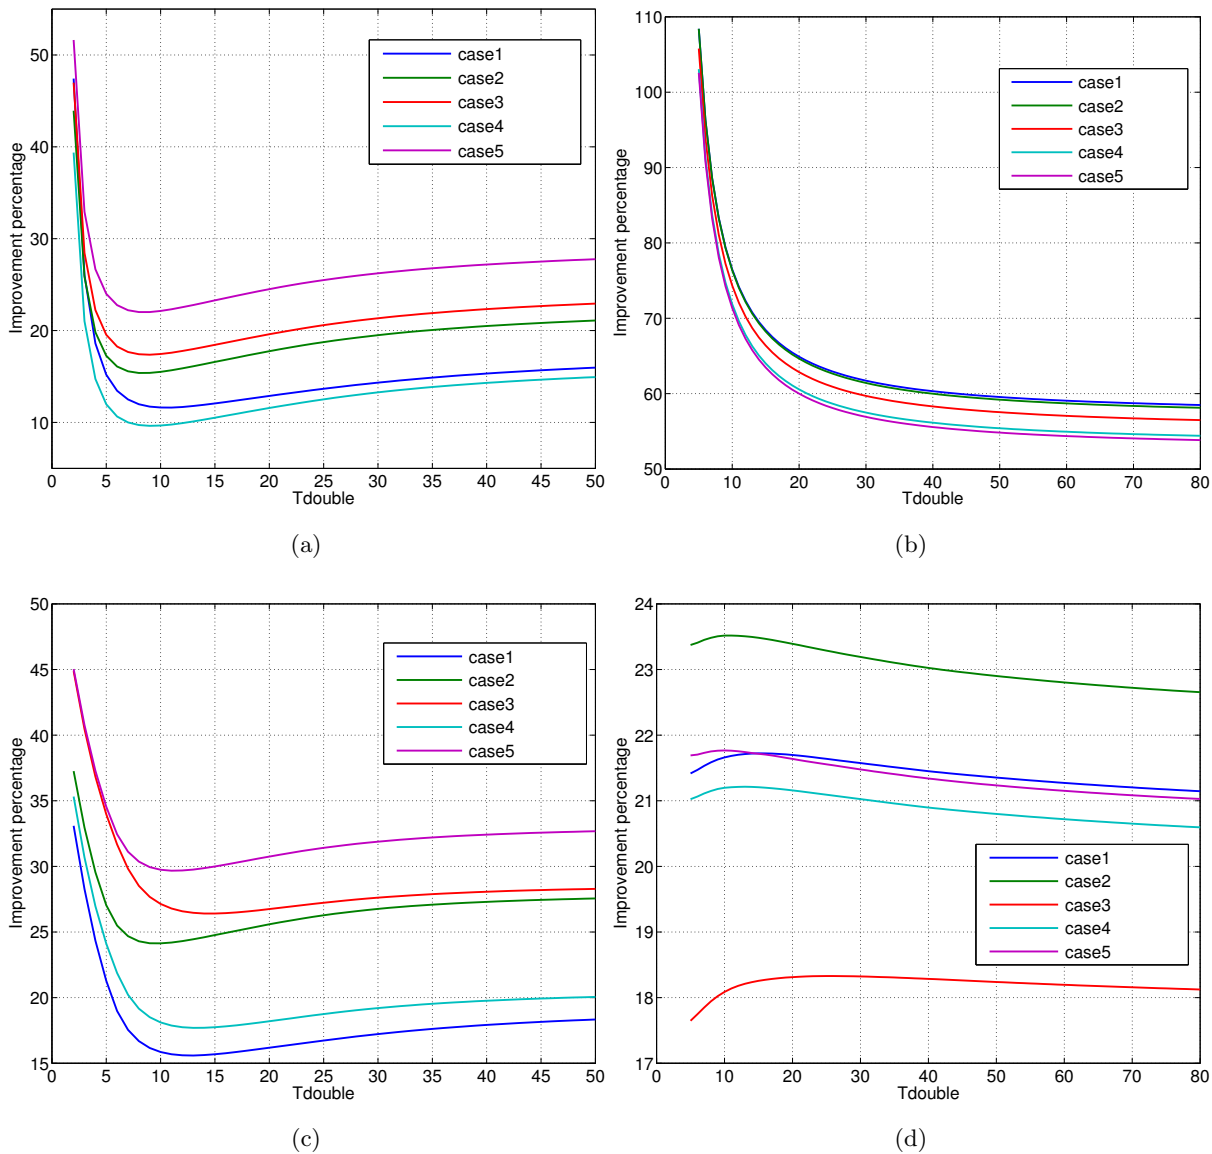


Figure 4.2: The average percentage improvement achieved by our spatiotemporally integrated approach over: (a) IMRT for head-and-neck, (b) IMRT for prostate, (c) spatiotemporally separated model in Chapter 2 for head-and-neck, and (d) spatiotemporally separated model in Chapter 2. These numbers are for $T_{\text{lag}} = 7$ days and averaged over all combinations of tumor and normal tissue α/β ratios.




Figure 4.3: The average percentage improvement achieved by our spatiotemporally integrated approach over IMRT for: (a) head-and-neck case 1, (b) prostate case 1. These numbers are for $T_{\text{lag}} = 7$ days with all normal tissue α/β ratios fixed at 3 for different tumor α/β ratios as shown in the legends.

fractionation strategies by redefining function $\tau(N)$ in Equation (4.1) as described in our previous work in Chapter 2.

Finally, one limitation of our formulation is the assumption that an identical fluence-map is used in every session; that is, it only considers equal-dosage fractionation. This is consistent with prevalent practice. However, it has been suggested in Chapter 3 and in [80] that this may not be optimal. It would be interesting to formulate and solve a more general version of our spatiotemporally integrated model that allows for the fluence-map to change across sessions. The number of variables in such a model will be an order of magnitude larger than that in our model here. We therefore defer the development of efficient algorithms for its approximate solution to future research.

Chapter 5

**A THEORETICAL STOCHASTIC CONTROL FRAMEWORK FOR
ADAPTING RADIOTHERAPY TO HYPOXIA****5.1 Introduction**

Hypoxia, that is, an inadequate supply of oxygen in living tissue, has long been known to adversely affect the outcome of radiotherapy in solid tumors and especially in head-and-neck tumors [14, 15, 25, 92–98]. Specifically, hypoxia induces radioresistance in these tumor cells and hence a higher radiation dose may be needed to achieve the same degree of cell-kill as compared to that for well-oxygenated tumor cells.

Oxygen partial pressure within tumors varies both spatially and temporally over the treatment course [14, 15, 97, 99, 100]. Moreover, there could be a considerable inter-patient variation in intra-tumor oxygen partial pressure [14, 101]. Similarly, there could be a variation in oxygen partial pressure among tumors that are otherwise histologically identical [102]. Intensity Modulated Radiation Therapy (IMRT) can in principle be used to selectively boost radiation dose to tumor subregions that are known to be hypoxic with the hope of improving tumor control in individual patients [14, 15, 21, 97, 103, 104]. However, traditionally, invasive techniques were needed to measure spatiotemporal distribution of oxygen partial pressure in solid tumors [97, 99–101, 105–107]. This curtailed the prevalence of such dose boosting in practice [105].

Non-invasive functional imaging techniques are now becoming available for directly or indirectly assessing intra-tumor oxygen partial pressure and other relevant biological information *in vivo* [13, 26, 105, 108–112]. Examples include magnetic resonance spectroscopic imaging (MRSI), single photon emission computed tomography (SPECT), and positron emission tomography (PET). For instance, ^{18}F -fluoromisonidazole (FMISO) is an extensively studied PET agent whose intracellular

retention depends on oxygen concentration. FMISO was the first radiopharmaceutical employed to quantitatively image hypoxia in human tumors with PET [102, 105, 113]. Similarly, PET imaging with the ^{18}F -2-deoxyglucose (FDG) tracer, and MRSI with choline/citrate ratios can be used to assess spatial distribution of tumor cell density (and hence to indirectly locate radioresistant tumor regions of high cell density) [21, 104, 114, 115]. Research is currently ongoing to enhance the accuracy and precision of such functional imaging techniques. There seems to be hope in the medical community that this type of imaging will become clinically viable within a decade [14, 21, 104]. This creates the future possibility of adapting IMRT treatment plans to the uncertain, spatiotemporal evolution of oxygen partial pressure and tumor cell density as inferred from functional images acquired over the treatment course for individual patients [14, 15, 21, 97, 104, 108, 116–118].

Stochastic control is the appropriate mathematical framework for rigorously formulating the problem of optimally adapting IMRT plans. Thus, the aim of this study is to estimate any potential benefits of such adaptive planning using computer simulations of computer-generated head-and-neck cancer test cases within a general, theoretical stochastic control framework that we originally envisioned in [23].

5.2 A theoretical stochastic control formalism

The conceptual idea in the abstract stochastic control formalism that we outlined in [23] is to view the tumor and normal tissues as a dynamic system whose state evolution over the treatment course is affected stochastically by the control variables chosen by the treatment planner in different treatment sessions. The planner’s goal is to optimize some quantitative measure of treatment efficacy. Any concrete implementation of this abstract idea calls for precise mathematical definitions of the tumor state dynamics, the normal tissue state dynamics, the feasible control variables, and the treatment efficacy measure. These definitions are of course problem dependent. We formalize them in detail in the specific context of hypoxia in this section.

We consider a treatment course with T sessions indexed by $t = 1, 2, \dots, T$. The

number of beamlets in the radiation field is K and we use $u^t \in \mathfrak{R}_+^K$ to denote the fluence-map, that is, the beamlet intensity vector, employed in the t th treatment session.

5.2.1 Tumor state dynamics

In our model, the tumor state includes the cell density and oxygen partial pressure in all tumor voxels as observed in functional images taken prior to a treatment session. The tumor is assumed to include n equal-volume voxels (each with volume ν) indexed by $i = 1, 2, \dots, n$. Let A denote the $n \times k$ tumor dose deposition coefficient matrix and let A_i denote its i th row that corresponds to the i th tumor voxel. That is, according to the linear dose deposition model, $A_i u^t$ is the dose deposited in the i th tumor voxel in the t th session. Let x_i^t denote the tumor cell density in tumor voxel i as assessed from a functional image acquired at the beginning of the t th session.

We model the evolution of tumor cell density in each voxel using the standard linear-quadratic (LQ) dose-response model from radiobiology [11]. Let α_i^t and β_i^t denote the radiosensitivity parameters of the LQ model in session t for tumor voxel i . The dependence of these radiosensitivity parameters on oxygen partial pressure is characterized using the standard notion of Oxygen Enhancement Ratio (OER) [14, 97, 98]. Specifically, let y_i^t denote the oxygen partial pressure in tumor voxel i as assessed from a functional image acquired at the beginning of session t . Then, we have, for $i = 1, 2, \dots, n$,

$$\alpha_i^t = \frac{\alpha}{\text{OER}_\alpha} \left(\frac{y_i^t \text{OER}_\alpha + \kappa}{y_i^t + \kappa} \right), \text{ and} \quad (5.1)$$

$$\beta_i^t = \frac{\beta}{(\text{OER}_\beta)^2} \left(\frac{y_i^t \text{OER}_\beta + \kappa}{y_i^t + \kappa} \right)^2, \quad (5.2)$$

where α and β are radiosensitivity parameters under well-oxygenated conditions, and OER_α , OER_β , and κ are other parameters of the OER model. The tumor cell dynamics are then given by the standard LQ formula

$$x_i^{t+1} = x_i^t \exp(-\alpha_i^t (A_i u^t) - \beta_i^t (A_i u^t)^2), \quad i = 1, 2, \dots, n. \quad (5.3)$$

As in [15, 94], we ignore the effect of tumor proliferation for notational simplicity; it can, however, be incorporated without any changes to our methodology.

Currently, quantitative data on voxel-by-voxel spatiotemporal evolution of oxygen partial pressure over the treatment course are not available. In addition, as stated above, there is an inter-tumor and an inter-patient uncertainty in this evolution. We therefore model the stochastic, spatiotemporal evolution of oxygen partial pressure using a random walk as an example. Specifically, we assume that the oxygen partial pressure evolves according to

$$\ln y^{t+1} = \ln y^t + \theta^t. \quad (5.4)$$

Here, $\theta^t \in \mathfrak{R}^n$ is stochastic noise that we assume to be independent across t with zero mean and with a spatial covariance matrix Σ . Thus (5.4) can also be seen as a first order vector autoregressive process. This particular choice for modeling partial pressure evolution is mainly for concreteness. It seemed more realistic than the current approach in the theoretical literature of sampling hypoxia values as independent and identically distributed random variables over space and time [119]. This random walk has the added benefit that the resulting probability distribution of oxygen partial pressure well-approximates the lognormal distribution (existing clinical studies have employed the lognormal distribution to fit oxygen partial pressure data [101, 106, 119]). Our stochastic control formalism can be easily generalized to any Markovian model of oxygen partial pressure evolution. In fact, as we shall see in Remark 5.3.1, our treatment plan optimization algorithm is model-free in that it does not actually require an explicit model of hypoxia evolution. In this sense, our random walk model here serves as a means to simulate hypoxia evolution on a computer within our theoretical framework so that we could test our optimization algorithm without real hypoxia images.

In summary, the cell density vector $x^t \triangleq (x_1^t, \dots, x_n^t) \in \mathfrak{R}_+^n$ and the oxygen partial pressure vector $y^t \triangleq (y_1^t, \dots, y_n^t) \in \mathfrak{R}_+^n$ define the tumor states in our model with dynamics given by Equations (5.3) and (5.4).

5.2.2 Control variables and normal tissue dynamics

The control variables in our stochastic control formalism correspond to the fluence-maps u^t employed in different treatment sessions. Feasible fluence-maps are subject to maximum dose, mean dose and dose-volume constraints on serial and parallel normal tissues. To explicitly model the temporal component of the problem, these constraints are expressed in terms of biologically effective dose (BED) equivalents. We also include smoothness constraints on the fluence-maps. The details of these components of our formulation are described in this section.

Let O_1, O_2, \dots, O_M denote the M different normal tissues under consideration. Let $\mathcal{M} = \{1, \dots, M\}$ be the set of indices of these normal tissues. For $m \in \mathcal{M}$, let n_m denote the number of voxels in O_m . These voxels are indexed by $j = 1, 2, \dots, n_m$. Let \mathcal{N}_m denote the set $\{1, 2, \dots, n_m\}$ of these voxels. All normal tissue voxels are assumed to have equal volume. Let A^m be the $n_m \times K$, non-negative dose deposition matrix for O_m . Let A_j^m be the j th row of this matrix; this is the row that corresponds to the j th voxel in O_m . That is, $A_j^m u$ is the dose delivered to the j th voxel in O_m in each session. Let α_m and β_m be the parameters of the LQ model for normal tissue O_m and we define $\rho_m \triangleq 1/(\alpha_m/\beta_m)$. Then, the BED corresponding to the dose delivered to the j th voxel in O_m in the t th session is given by

$$(A_j^m u^t) + \rho_m (A_j^m u^t)^2. \quad (5.5)$$

The set of feasible fluence-maps in session t is defined through BED and smoothness constraints as described next.

Maximum BED constraints for serial normal tissues

Let $\mathcal{M}_1 \subseteq \mathcal{M}$ be the set of indices of serial normal tissues for which we wish to include maximum dose constraints. These are the normal tissues whose function is hampered even when a small region is damaged by radiation. For example, maximum dose constraints are included on the spinal cord in head-and-neck cancer. Suppose for any $m \in \mathcal{M}_1$, that a total dose D_{\max}^m is known to be tolerated by each voxel in O_m if

administered in N_{conv}^m equal-dose fractions. The BED corresponding to this schedule equals

$$\text{BED}_{\text{max}}^m = D_{\text{max}}^m (1 + \rho_m (D_{\text{max}}^m / N_{\text{conv}}^m)). \quad (5.6)$$

We use the standard approach of comparing normal tissue BED. In particular, let $z_j^{t,m}$ denote the BED of doses delivered to voxel j in O_m in the first $t - 1$ sessions. Then, a dose of $(A_j^m u^t)$ can be tolerated by the j th voxel in O_m in session t if

$$z_j^{t,m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2 \leq \text{BED}_{\text{max}}^m, \quad j = 1, 2, \dots, n_m. \quad (5.7)$$

Thus, for each $m \in \mathcal{M}_1$, we define the BED vectors $z^{t,m} \triangleq (z_1^{t,m}, \dots, z_{n_m}^{t,m})$ as the state of O_m at the beginning of session t .

Mean BED constraints for parallel normal tissues

Let $\mathcal{M}_2 \subseteq \mathcal{M}$ be the set of indices of parallel normal tissues for which we wish to include mean dose constraints. These are the normal tissues where a sufficiently small portion can be damaged without affecting the organ function. For example, mean dose constraints are included on the parotid glands in IMRT for head-and-neck cancer. Suppose for any $m \in \mathcal{M}_2$, that mean dose D_{mean}^m is known to be tolerated by O_m if administered in N_{conv}^m equal-dose fractions. The BED corresponding to this mean dose is given by

$$\text{BED}_{\text{mean}}^m = D_{\text{mean}}^m (1 + \rho_m (D_{\text{mean}}^m / N_{\text{conv}}^m)). \quad (5.8)$$

Let $v^{t,m}$ denote the mean BED over all voxels in O_m of doses delivered in the first $t - 1$ sessions. These will be included among our normal tissue state variables. Then, doses $(A_j^m u^t)$ delivered to voxels $j \in \mathcal{N}_m$ in session t can be tolerated by O_m if

$$v^{t,m} + \frac{\sum_{j=1}^{n_m} (A_j^m u^t) + \rho_m \sum_{j=1}^{n_m} (A_j^m u^t)^2}{n_m} \leq \text{BED}_{\text{mean}}^m. \quad (5.9)$$

Thus, for each $m \in \mathcal{M}_2$, we define the mean BED $v^{t,m}$ as the state of O_m at the beginning of session t .

Dose-volume constraints for parallel normal tissues

Let $\mathcal{M}_3 \subseteq \mathcal{M}$ be the set of indices of normal tissues with dose-volume constraints. Examples include rectum, bladder, and lung. In particular, suppose for any $m \in \mathcal{M}_3$, that no more than a volume fraction ϕ_m of normal tissue O_m can receive a dose more than D_{dv}^m if administered in N_{conv}^m fractions. The BED of total dose D_{dv}^m administered in N_{conv}^m equal-dosage fractions is given by

$$\text{BED}_{\text{dv}}^m = D_{\text{dv}}^m (1 + \rho_m (D_{\text{dv}}^m / N_{\text{conv}}^m)). \quad (5.10)$$

Let $w_j^{t,m}$ denote the BED of doses delivered to voxel j in O_m in the first $t - 1$ sessions. Since all voxels in O_m have equal volume, the volume fraction is the same as the voxel fraction. For each $m \in \mathcal{M}_3$ and for $j = 1, 2, \dots, n_m$, we thus define binary-valued functions $f_j^m(u^t; w^{t,m})$ such that

$$f_j^m(u^t; w^{t,m}) = \begin{cases} 1 & \text{if } w^{t,m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2 > \text{BED}_{\text{dv}}^m, \\ 0 & \text{if } w^{t,m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2 \leq \text{BED}_{\text{dv}}^m. \end{cases} \quad (5.11)$$

In words, $f_j^m(u^t; w^{t,m})$ is one if the BED resulting after delivering dose by fluence-map u^t to voxel j exceeds the tolerance BED_{dv}^m ; $f_j^m(u^t; w^{t,m})$ is zero otherwise. We use the integer K_m to denote $\lfloor n_m \phi_m \rfloor$, that is, the largest integer that is at most $n_m \phi_m$. Then, the dose-volume constraints are expressed as

$$\sum_{j=1}^{n_m} f_j^m(u^t; w^{t,m}) \leq K_m, \quad m \in \mathcal{M}_3. \quad (5.12)$$

These constraints ensure that there are at most K_m voxels for which $w^{t,m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2 > \text{BED}_{\text{dv}}^m$; in other words, there are at least $n_m - K_m$ voxels for which $w^{t,m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2 \leq \text{BED}_{\text{dv}}^m$. Thus, for each $m \in \mathcal{M}_3$, we define the BED vectors $w^{t,m} \triangleq (w_1^{t,m}, \dots, w_{n_m}^{t,m})$ as the state of O_m at the beginning of session t .

Fluence-map smoothness constraints

Finally, to ensure that the intensity profile is deliverable in practice using a multi-leaf collimator, we put a smoothness constraint on each radiation field similar to Chapter

4 [83, 84]. In particular, for each radiation field, we bound the absolute relative difference between intensities of each pair of nearest neighbor beamlets by a fraction ξ . Then the smoothness constraints can be written compactly in matrix format as

$$Su \leq 0, \tag{5.13}$$

where S is a block diagonal matrix with entries $-(1 + \xi)$, $(1 - \xi)$, -1 , 0 , $+1$ at appropriate locations.

Set of feasible fluence-map policies

Recall that the states of various normal tissues at the beginning of treatment session t are denoted by $z^{t,m}$ for $m \in \mathcal{M}_1$, $v^{t,m}$ for $m \in \mathcal{M}_2$, and $w^{t,m}$ for $m \in \mathcal{M}_3$, respectively. Let $|\mathcal{M}_1|$, $|\mathcal{M}_2|$, and $|\mathcal{M}_3|$ denote the cardinalities of sets \mathcal{M}_1 , \mathcal{M}_2 , and \mathcal{M}_3 , respectively. We now define the more compact normal tissue state notation $\vec{z}^t \triangleq (z^{t,1}, \dots, z^{t,|\mathcal{M}_1|})$, $\vec{v}^t \triangleq (v^{t,1}, \dots, v^{t,|\mathcal{M}_2|})$, and $\vec{w}^t \triangleq (w^{t,1}, \dots, w^{t,|\mathcal{M}_3|})$. Thus, the combined state of all normal tissues is written as $[\vec{z}^t; \vec{v}^t; \vec{w}^t]$. The set of feasible fluence-maps in session t depends on this normal tissue state through constraints (5.7), (5.9), and (5.12). The set of fluence-maps $u^t \in \mathfrak{R}_+^K$ that satisfy these constraints as well as the smoothness constraint (5.13) is denoted by $\mathcal{U}^t([\vec{z}^t; \vec{v}^t; \vec{w}^t]) \subset \mathfrak{R}_+^K$. Then, in the standard language of the stochastic control literature [120], a feasible policy assigns a fluence-map from the set $\mathcal{U}^t([\vec{z}^t; \vec{v}^t; \vec{w}^t])$ to every possible combination of tumor and normal tissue states $[x^t; y^t; \vec{z}^t; \vec{v}^t; \vec{w}^t]$ in each treatment session $t = 1, 2, \dots, T$. We denote the set of such feasible policies by \mathcal{P} and generic policies in this set by π .

5.2.3 Bellman's equations of dynamic programming

As in [23?], we use the total number of tumor cells remaining (TNTCR) as our treatment efficacy measure. That is, our goal is to minimize the TNTCR, which is

given by

$$\text{TNTCR} \triangleq \sum_{i=1}^n \nu x_i^{T+1}. \quad (5.14)$$

Note that this perhaps less familiar objective is equivalent to the more common objective of maximizing tumor control probability (TCP); this is because the negative natural logarithm of TCP equals TNTCR. We employ the TNTCR objective here because it turns out to be algebraically less cumbersome than TCP while writing our models and our convexity proof in Lemma 5.3.3. Moreover, when the initial cell density is equal in all tumor voxels, which is an assumption we will later make, as in [94], in our simulation experiments, the improvement in TNTCR achieved by one solution method over another does not depend on the value of this initial cell density; this is not the case with TCP. This feature of TNTCR made it appealing for us as compared to the TCP. Finally, numerical calculations with TNTCR on our computer platform also seemed more reliable and less ill-conditioned than those with TCP.

Now recall that the initial state is given by $[x^1; y^1; z^1; v^1; w^1]$; in fact, note that all components of z^1 , v^1 , and w^1 are zero since no dose is delivered prior to the first session. Let $J_\pi^1([x^1; y^1; z^1; v^1; w^1])$ denote the TNTCR achieved by the end of the treatment course if policy $\pi \in \mathcal{P}$ is implemented through the treatment course starting in state $[x^1; y^1; z^1; v^1; w^1]$ at the beginning of the first session. The problem of minimizing the TNTCR can then be formulated as the stochastic control problem

$$J_*^1 \triangleq \min_{\pi \in \mathcal{P}} J_\pi^1([x^1; y^1; z^1; v^1; w^1]). \quad (5.15)$$

The standard approach for solving such problems, at least in principle, is to employ Bellman's backward recursive algorithm of dynamic programming [120]. In order to present this algorithm, we define the optimal TNTCR-to-go functions $J^t([x^t; y^t; z^t; v^t; w^t])$ for all possible states $[x^t; y^t; z^t; v^t; w^t]$ and for $t = 1, 2, \dots, T+1$. Specifically, $J^t([x^t; y^t; z^t; v^t; w^t])$ is interpreted as the minimum TNTCR reached at the end of the treatment course given that the state at the beginning of session t is $[x^t; y^t; z^t; v^t; w^t]$. These functions can in principle be obtained by solving, in the backward order $t = T, T-1, \dots, 1$, the

non-linear stochastic optimization problems

$$J^t([x^t; y^t; z^t; v^t; w^t]) = \min_{u^t \in \mathcal{U}^t([z^t; v^t; w^t])} E_{\theta^t} \left(J^{t+1}([x^{t+1}; y^{t+1}; z^{t+1}; v^{t+1}; w^{t+1}]) \right), \quad (5.16)$$

subject to

$$x_i^{t+1} = x_i^t \exp(-\alpha_i^t (A_i u^t) - \beta_i^t (A_i u^t)^2), \quad i = 1, 2, \dots, n, \quad (5.17)$$

$$\ln y^{t+1} = \ln y^t + \theta^t, \quad (5.18)$$

$$z_j^{t+1,m} = z_j^{t,m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (5.19)$$

$$v^{t+1,m} = v^{t,m} + \frac{\sum_{j=1}^{n_m} (A_j^m u^t) + \rho_m \sum_{j=1}^{n_m} (A_j^m u^t)^2}{n_m}, \quad m \in \mathcal{M}_2, \quad (5.20)$$

$$w_j^{t+1,m} = w_j^{t,m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_3, \quad (5.21)$$

for all possible states $[x^t; y^t; z^t; v^t; w^t]$ starting with the boundary condition

$$J^{T+1}([x^{T+1}; y^{T+1}; z^{T+1}; v^{T+1}; w^{T+1}]) = \sum_{i=1}^n \nu x_i^{T+1}. \quad (5.22)$$

The fluence-maps $u^t \in \mathcal{U}^t([z^t; v^t; w^t])$ that achieve the minimum in (5.16) then define an optimal policy for the stochastic control problem (5.15).

Unfortunately, for our problem, this backward algorithm is not implementable in practice because it requires that the minimization in (5.16) be performed for an *uncountable* number of states. Thus a state discretization would be necessary, which is again computationally impractical due to the curse of state-space dimensionality [121]. Moreover, even if such a state discretization were possible, each minimization in a discrete version of (5.16) would require the solution of a non-linear stochastic programming problem which itself is impractical due to the curse of action-space dimensionality [121]. Thus, as in [23], we instead consider an approximate dynamic programming technique called Certainty Equivalent Control [120] that exploits the structure of our stochastic control problem and does not require discretization.

5.3 Solution method: Certainty Equivalent Control

Certainty Equivalent Control (CEC) [120] is an easy to implement approximation technique for stochastic control problems. It begins by assuming that all future noise in the problem is replaced by its nominal values, for example, the expected values, thus leading to a deterministic control problem. This control problem is then solved to obtain an optimal sequence of controls for all future periods. However, only the control obtained for the first period is implemented while the others are discarded. The system then evolves stochastically to the next state and the process is repeated until the terminal course is completed. Thus, in one complete run of CEC, a total of T deterministic control problems are solved; the first one includes T sessions, the second includes $T - 1$ sessions, and ultimately, the last one involves only one session.

Certainty Equivalent Control

INITIALIZE: Set $t = 1$ and begin with a given initial state $[x^1; y^1; z^1; v^1; w^1]$.

DO WHILE $t \leq T$,

1. let the state at the beginning of session t be $[x^t; y^t; z^t; v^t; w^t]$;
2. fix the oxygen partial pressure vectors y^t, y^{t+1}, \dots, y^T at some “nominal” values $\hat{y}^t, \hat{y}^{t+1}, \dots, \hat{y}^T$. Substitute these in expressions (5.1) and (5.2) to obtain nominal radiosensitivity parameters $\hat{\alpha}_i^t, \hat{\alpha}_i^{t+1}, \dots, \hat{\alpha}_i^T$ and $\hat{\beta}_i^t, \hat{\beta}_i^{t+1}, \dots, \hat{\beta}_i^T$ for tumor voxels $i = 1, 2, \dots, n$;
3. for each $m \in \mathcal{M}_3$ and for $j = 1, 2, \dots, n_m$, define binary-valued functions

$$g_j^m(u^t, u^{t+1}, \dots, u^T; w^{t,m}) = \begin{cases} 1 & \text{if } w^{t,m} + \sum_{l=t}^T [(A_j^m u^l) + \rho_m (A_j^m u^l)^2] > \text{BED}_{\text{dv}}^m, \\ 0 & \text{if } w^{t,m} + \sum_{l=t}^T [(A_j^m u^l) + \rho_m (A_j^m u^l)^2] \leq \text{BED}_{\text{dv}}^m; \end{cases} \quad (5.23)$$

4. solve the deterministic optimization problem

$$(P_t) \min \sum_{i=1}^n \nu x_i^t \prod_{l=t}^T \exp(-\hat{\alpha}_i^l(A_i u^l) - \hat{\beta}_i^l(A_i u^l)^2) \quad (5.24)$$

subject to

$$\sum_{l=t}^T [(A_j^m u^l) + \rho_m (A_j^m u^l)^2] \leq \text{BED}_{\max}^m - z_j^{t,m}, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (5.25)$$

$$\sum_{l=t}^T \left[\sum_{j=1}^{n_m} (A_j^m u^l) + \rho_m \sum_{j=1}^{n_m} (A_j^m u^l)^2 \right] \leq n_m \text{BED}_{\text{mean}}^m - n_m v^{t,m}, \quad m \in \mathcal{M}_2, \quad (5.26)$$

$$\sum_{j=1}^{n_m} g_j^m(u^t, u^{t+1}, \dots, u^T; w^{t,m}) \leq K_m, \quad m \in \mathcal{M}_3, \quad (5.27)$$

$$S u^l \leq 0, \quad l = t, t+1, \dots, T, \quad (5.28)$$

$$u^l \geq 0, \quad l = t, t+1, \dots, T; \quad (5.29)$$

to obtain an optimal sequence of fluence-maps $u_*^t, u_*^{t+1}, \dots, u_*^T$;

5. discard $u_*^{t+1}, u_*^{t+2}, \dots, u_*^T$ and use fluence-map u_*^t in session t ;

6. calculate α_i^t and β_i^t via expressions (5.1)-(5.2) using the imaged hypoxia values y_i^t , for $i = 1, 2, \dots, n$; update tumor cell density and normal tissue states using u_*^t as

$$x_i^{t+1} = x_i^t \exp(-\alpha_i^t(A_i u_*^t) - \beta_i^t(A_i u_*^t)^2), \quad i = 1, 2, \dots, n, \quad (5.30)$$

$$z_j^{t+1,m} = z_j^{t,m} + (A_j^m u_*^t) + \rho_m (A_j^m u_*^t)^2, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (5.31)$$

$$v^{t+1,m} = v^{t,m} + \frac{\sum_{j=1}^{n_m} (A_j^m u_*^t) + \rho_m \sum_{j=1}^{n_m} (A_j^m u_*^t)^2}{n_m}, \quad m \in \mathcal{M}_2, \quad (5.32)$$

$$w_j^{t+1,m} = w_j^{t,m} + (A_j^m u_*^t) + \rho_m (A_j^m u_*^t)^2, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_3; \quad (5.33)$$

7. sample θ^t from its probability distribution and set $\ln y^{t+1} = \ln y^t + \theta^t$;

8. update $t \leftarrow t + 1$.

END DO

Remark 5.3.1. *If this CEC approach were implemented in practice, the fluence-map u_*^t would be employed in the t th session and then the imaged oxygen partial pressure vector y^t evolves stochastically to a new state y^{t+1} that is again imaged at the beginning of the $t + 1$ st session. Consequently, an explicit model for oxygen partial pressure evolution, such as our random walk model, is **not** needed. Specifically, Step 7 in the above algorithm is then not needed; this step is needed in our computer simulation of this physical process based on our stochastic model of oxygen partial pressure evolution.*

Remark 5.3.2. *We emphasize that **only** one cell density image is needed in order to implement the CEC approach above in practice. This cell density image is acquired at the beginning of the treatment course wherein the planner can observe the initial cell density x^1 . All subsequent cell densities are then calculated, when needed, via formula (5.30) using radiosensitivity parameters α_i^t and β_i^t obtained by the OER expressions (5.1)-(5.2) using the imaged hypoxia values y_i^t , for $i = 1, 2, \dots, n$. In fact, if we assume that the initial tumor cell density is homogeneous across all tumor voxels, that is, if $x_i^1 = c$, for $i = 1, 2, \dots, n$, where c is some fixed cell density value, then even an initial cell density image is not needed. This is because c factors out from all our objective functions and thus has no effect on the fluence-maps obtained by our solution method. As such, the value of c is not needed anywhere by the above algorithm. It is for this reason that we call the above algorithm “**CEC with hypoxia images**” in Section 5.4 below.*

It remains to choose nominal values of oxygen partial pressure vectors in Step 2 of the loop over sessions t and then to devise an efficient procedure for solving problem (P_t) . We do this in the next section.

5.3.1 Efficient solution of problem (P_t)

Exact solution of (P_t) , as formulated in the above CEC algorithm, is computationally intractable. It is well-known in the IMRT literature that dose-volume constraints are some of the most difficult to handle [6]. Even without the dose-volume constraints (5.27), (P_t) is a non-convex problem because although the constraints are convex, the objective function is in general not. In realistic instances of (P_t) , the number of beamlets K is likely to equal a few thousand [6]. With $T = 35$ (this is the number of fractions typical in current treatment guidelines [1]), this implies that (P_t) can include about a hundred thousand variables. Also, the number of constraints can be in the tens of thousands depending on the total number of normal tissue voxels [6]. We therefore propose an efficient procedure for approximate solution of (P_t) .

The first step is to choose nominal values for oxygen partial pressures. Since the partial pressure map y^t is already available in the state as observed in the functional image acquired prior to treatment session t , we set $\hat{y}^t = y^t$. Moreover, if the treatment planner were not to re-plan in subsequent sessions, it is natural to plan the remaining sessions assuming that oxygen partial pressures in the subsequent sessions will be as seen in the image at the beginning of session t . As a result, we also set $\hat{y}^{t+1} = \hat{y}^{t+2} = \dots = \hat{y}^T = y^t$. Now, given that the nominal values of oxygen partial pressures in all remaining sessions are chosen to be identical, and in fact equal to y^t , we denote them by \hat{y} for simplicity. This also implies that the radiosensitivity parameters $\hat{\alpha}_i^t, \hat{\alpha}_i^{t+1}, \dots, \hat{\alpha}_i^T$ and $\hat{\beta}_i^t, \hat{\beta}_i^{t+1}, \dots, \hat{\beta}_i^T$ can be simply denoted by $\hat{\alpha}_i$ and $\hat{\beta}_i$, where, from formulas (5.1) and (5.2), we have,

$$\hat{\alpha}_i = \frac{\alpha}{\text{OER}_\alpha} \left(\frac{\hat{y}_i \text{OER}_\alpha + \kappa}{\hat{y}_i + \kappa} \right), \text{ and} \quad (5.34)$$

$$\hat{\beta}_i = \frac{\beta}{(\text{OER}_\beta)^2} \left(\frac{\hat{y}_i \text{OER}_\beta + \kappa}{\hat{y}_i + \kappa} \right)^2, \quad (5.35)$$

for tumor voxels $i = 1, 2, \dots, n$. Now, since the tumor's dose-response parameters are time-invariant, it is natural, although not necessarily optimal (see Chapter 3 and

[80]), to search for equal-dosage treatment plans. That is, we set $u^t = u^{t+1} = \dots = u^T$ and simply call these fluence-maps u . Note that this thought is also consistent with the current practice of using identical fluence-maps in all sessions. This simplifies problem (P_t) to

$$(Q_t) \min \sum_{i=1}^n \nu x_i^t \exp(-(T-t+1)(\hat{\alpha}_i(A_i u) + \hat{\beta}_i(A_i u)^2)) \quad (5.36)$$

subject to

$$(T-t+1)[(A_j^m u) + \rho_m (A_j^m u)^2] \leq \text{BED}_{\max}^m - z_j^{t,m}, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (5.37)$$

$$(T-t+1)\left[\sum_{j=1}^{n_m} (A_j^m u) + \rho_m \sum_{j=1}^{n_m} (A_j^m u)^2\right] \leq n_m \text{BED}_{\text{mean}}^m - n_m v^{t,m}, \quad m \in \mathcal{M}_2, \quad (5.38)$$

$$\sum_{j=1}^{n_m} g_j^m(u, u, \dots, u; w^{t,m}) \leq K_m, \quad m \in \mathcal{M}_3, \quad (5.39)$$

$$Su \leq 0, \quad (5.40)$$

$$u \geq 0. \quad (5.41)$$

Moreover, since $A_j^m u \geq 0$, constraints (5.37) can be equivalently rewritten as linear constraints wherein the right hand side is obtained by solving a quadratic equation.

This yields the equivalent problem

$$(Q_t) \min \sum_{i=1}^n \nu x_i^t \exp(-(T-t+1)(\hat{\alpha}_i(A_i u) + \hat{\beta}_i(A_i u)^2)) \quad (5.42)$$

subject to

$$A_j^m u \leq \frac{-1 + \sqrt{1 + 4\rho_m(\text{BED}_{\max}^m - z_j^{t,m})/(T-t+1)}}{2\rho_m}, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (5.43)$$

$$(T-t+1) \left[\sum_{j=1}^{n_m} (A_j^m u) + \rho_m \sum_{j=1}^{n_m} (A_j^m u)^2 \right] \leq n_m \text{BED}_{\text{mean}}^m - n_m v^{t,m}, \quad m \in \mathcal{M}_2, \quad (5.44)$$

$$\sum_{j=1}^{n_m} g_j^m(u, u, \dots, u; w^{t,m}) \leq K_m, \quad m \in \mathcal{M}_3, \quad (5.45)$$

$$Su \leq 0, \quad (5.46)$$

$$u \geq 0. \quad (5.47)$$

The dose-volume constraints (5.45) in problem (Q_t) are still difficult to tackle. We therefore propose a constraint generation approach to surmount this difficulty.

Constraint generation method

In particular, we first solve (Q_t) but *without* constraints (5.45). This problem is given by

$$(\hat{Q}_t) \min \sum_{i=1}^n \nu x_i^t \exp(-(T-t+1)(\hat{\alpha}_i(A_i u) + \hat{\beta}_i(A_i u)^2)) \quad (5.48)$$

subject to

$$A_j^m u \leq \frac{-1 + \sqrt{1 + 4\rho_m(\text{BED}_{\max}^m - z_j^{t,m})/(T-t+1)}}{2\rho_m}, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (5.49)$$

$$(T-t+1) \left[\sum_{j=1}^{n_m} (A_j^m u) + \rho_m \sum_{j=1}^{n_m} (A_j^m u)^2 \right] \leq n_m \text{BED}_{\text{mean}}^m - n_m v^{t,m}, \quad m \in \mathcal{M}_2, \quad (5.50)$$

$$Su \leq 0, \quad (5.51)$$

$$u \geq 0. \quad (5.52)$$

Suppose \hat{u} is an optimal solution to problem (\hat{Q}_t) . Then, for each $m \in \mathcal{M}_3$, we find $n_m - K_m$ voxels that receive the smallest doses among the n_m voxels in O_m under fluence-map \hat{u} . Let subset $\mathcal{N}_m^t(\hat{u}) \subseteq \mathcal{N}_m$ denote this group of voxels. We then re-solve (Q_t) but this time by replacing the dose-volume constraints with tolerance limits on all voxels in the set $\mathcal{N}_m^t(\hat{u})$. Finally, we note that these tolerance limits can be equivalently re-written as linear constraints whose right hand side is obtained by

solving a quadratic equation. That is, this problem is given by

$$(\bar{Q}_t) \min \sum_{i=1}^n \nu x_i^t \exp(-(T-t+1)(\hat{\alpha}_i(A_i u) + \hat{\beta}_i(A_i u)^2)) \quad (5.53)$$

subject to

$$A_j^m u \leq \frac{-1 + \sqrt{1 + 4\rho_m(\text{BED}_{\max}^m - z_j^{t,m})/(T-t+1)}}{2\rho_m}, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (5.54)$$

$$(T-t+1) \left[\sum_{j=1}^{n_m} (A_j^m u) + \rho_m \sum_{j=1}^{n_m} (A_j^m u)^2 \right] \leq n_m \text{BED}_{\text{mean}}^m - n_m v^{t,m}, \quad m \in \mathcal{M}_2, \quad (5.55)$$

$$A_j^m u \leq \frac{-1 + \sqrt{1 + 4\rho_m(\text{BED}_{\text{dv}}^m - w_j^{t,m})/(T-t+1)}}{2\rho_m}, \quad j \in \mathcal{N}_m^t(\hat{u}), \quad m \in \mathcal{M}_3, \quad (5.56)$$

$$Su \leq 0, \quad (5.57)$$

$$u \geq 0. \quad (5.58)$$

In the next section, we describe a log-barrier interior point algorithm for solving (\hat{Q}_t) and (\bar{Q}_t) .

Log-barrier interior point algorithm for solving (\hat{Q}_t) and (\bar{Q}_t)

Since problems (\hat{Q}_t) and (\bar{Q}_t) are identical in structure, we present the interior point algorithm only for (\hat{Q}_t) . Before proceeding, we equivalently re-write (\hat{Q}_t) in a more compact and convenient form. In particular, we use numbers b_j^m , for $j = 1, 2, \dots, n_m$ and $m \in \mathcal{M}_1$ to denote the right hand sides in constraint (5.49). Similarly, we use positive definite matrices Q^m , matrices P^m , and numbers r^m to compactly write the

quadratic constrains (5.50). That is, we have,

$$(\hat{Q}_t) \min \sum_{i=1}^n \nu x_i^t \exp(-(T-t+1)(\hat{\alpha}_i(A_i u) + \hat{\beta}_i(A_i u)^2)) \quad (5.59)$$

subject to

$$A_j^m u - b_j^m \leq 0, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (5.60)$$

$$u' Q^m u + P^m u - r^m \leq 0, \quad m \in \mathcal{M}_2, \quad (5.61)$$

$$S u \leq 0, \quad (5.62)$$

$$u \geq 0. \quad (5.63)$$

Lemma 5.3.3. *Problem (\hat{Q}_t) as formulated in (5.59)-(5.63) is convex if $\alpha^2 \geq 2\beta$ and $OER_\alpha \leq OER_\beta$. This convexity property also holds for (\bar{Q}_t) .*

Proof. The constraints in (\hat{Q}_t) are linear and convex quadratic. So it only remains to show that the objective function is convex if $\alpha^2 \geq 2\beta$ and $OER_\alpha \leq OER_\beta$. That is, we need to show that $\sum_{i=1}^n \nu x_i^t \exp(-(T-t+1)(\hat{\alpha}_i(A_i u) + \hat{\beta}_i(A_i u)^2))$ is convex. Since a non-negative linear combination of convex functions is convex, it suffices to show that $\exp(-(T-t+1)(\hat{\alpha}_i(A_i u) + \hat{\beta}_i(A_i u)^2))$ is convex for each i . To prove this, we first establish an intermediate claim below.

Claim: For all real numbers $a \geq 0$ and $b \geq 0$, the single variable function $h(z) = \exp(-az - bz^2)$ is convex over $z \geq 0$ if $a^2 \geq 2b$.

Proof of Claim: Observe that

$$\frac{dh}{dz} = -(a + 2bz) \exp(-az - bz^2), \quad \text{and} \quad (5.64)$$

$$\frac{d^2h}{dz^2} = (a + 2bz)^2 \exp(-az - bz^2) - 2b \exp(-az - bz^2) = \exp(-az - bz^2)(a^2 + 4abz + 4b^2z^2 - 2b). \quad (5.65)$$

Therefore, because $a \geq 0$ and $b \geq 0$, $\frac{d^2h}{dz^2}$ is non-negative (and hence $h(\cdot)$ is convex) over $z \geq 0$ when $a^2 \geq 2b$. Since composition of a convex function with an affine mapping is convex (see Section in [64]), and since $A_i u \geq 0$ for all $u \geq 0$, we have that $\exp(-(T-t+1)(\hat{\alpha}_i(A_i u) + \hat{\beta}_i(A_i u)^2))$ is convex over $u \geq 0$ for each i if $\hat{\alpha}_i^2 \geq 2\hat{\beta}_i$.

From formulas (5.34) and (5.35), we see that $\hat{\alpha}_i^2 \geq 2\hat{\beta}_i$ if

$$\frac{\alpha^2(\hat{y}_i \text{OER}_\alpha + \kappa)^2}{\text{OER}_\alpha^2} \geq \frac{2\beta(\hat{y}_i \text{OER}_\beta + \kappa)^2}{\text{OER}_\beta^2}. \quad (5.66)$$

After algebraic simplification, this yields that $\hat{\alpha}_i^2 \geq 2\hat{\beta}_i$ if $\alpha^2 \geq 2\beta$ and $\text{OER}_\alpha \leq \text{OER}_\beta$. \square

Fortunately, this sufficient condition for convexity is met for the values of $\alpha, \beta, \text{OER}_\alpha, \text{OER}_\beta$ currently available in the clinical literature for head-and-neck cancer. Consequently, as discussed in [64], the log-barrier interior point algorithm with Newton's method with backtracking line search is particularly suitable for solving (\hat{Q}_t) and (\bar{Q}_t) . Different components of this algorithm are described in various sections of [64] and we refer the interested reader to that book for a rigorous convergence analysis of such algorithms. We present this algorithm here using our notation for completeness.

For any log-barrier parameter $\mu > 0$, the log-barrier function for problem (5.59)-(5.63) is given by

$$\begin{aligned} \phi(u; \mu) \triangleq & \sum_{i=1}^n \nu x_i^t \exp(-(T-t+1)(\hat{\alpha}_i(A_i u) + \hat{\beta}_i(A_i u)^2)) - (1/\mu) \left(\sum_{m \in \mathcal{M}_1} \sum_{j=1}^{n_m} \ln(b_j^m - A_j^m u) + \right. \\ & \left. \sum_{m \in \mathcal{M}_2} \ln(r^m - P^m u - u' Q^m u) + \sum_{k=1}^K \ln u_k + \sum_{l=1}^L \ln(-S_l u) \right). \end{aligned} \quad (5.67)$$

The idea (see Algorithm 11.1 on page 569 of [64]) then is to use Newton's procedure (Algorithm 9.5 on page 487 of [64]) to solve a sequence of unconstrained problems $\min \phi(u; \mu_\eta)$ for a strictly increasing and divergent sequence of parameters $\{\mu_\eta\}_{\eta=1}^\infty$. We use $\nabla \phi(u; \mu)$ and $\nabla^2 \phi(u; \mu)$ to denote the gradient and Hessian of $\phi(u; \mu)$ with respect to the decision variables $u \in \Re^K$. The step-size δ in Newton's method is chosen by performing a backtracking line search (Algorithm 9.2 on page 464 of [64]) with parameters $0 < p < 0.5$ and $0 < q < 1$. Finally, note that the total number of constraints in (\hat{Q}_t) is $K + L + |\mathcal{M}_2| + \sum_{m \in \mathcal{M}_1} n_m$ and we use C to denote this integer in our stopping condition below.

Log-barrier interior point algorithm

INITIALIZE: Set $\eta = 1$ and begin with an initial guess fluence-map u in the interior of the feasible region of (\hat{Q}_t) . Set tolerance to ϵ and pre-select a strictly increasing sequence of log-barrier parameters $\mu_\eta \rightarrow \infty$.

DO WHILE $C/\mu_\eta \geq \epsilon$

$\lambda^2 \triangleq \infty$.

DO WHILE $\lambda^2/2 > \epsilon$

$\Delta u \triangleq -\nabla^2 \phi(u; \mu_\eta) \nabla \phi(u; \mu_\eta)$, and $\lambda^2 \triangleq -(\nabla \phi(u; \mu_\eta))' \Delta u$.

$\delta = 1$.

DO WHILE $\phi(u + \delta \Delta u; \mu_\eta) > \phi(u) + p\delta \nabla \phi(u; \mu_\eta) \Delta u$

$\delta = q\delta$.

END DO

$u \triangleq u + \delta \Delta u$.

END DO

$\eta \leftarrow \eta + 1$.

END DO

TERMINATE: Report fluence-map u as an optimal solution to (\hat{Q}_t) .

5.3.2 A streamlined version of CEC

Based on the above discussion, we propose the following simplified version of CEC for approximate solution of our stochastic control problem.

Streamlined CEC

INITIALIZE: Set $t = 1$ and begin with a given initial state $[x^1; y^1; z^1; v^1; w^1]$.

DO WHILE $t \leq T$,

1. let the state at the beginning of session t be $[x^t; y^t; z^t; v^t; w^t]$;
2. fix the oxygen partial pressure vectors y^t, y^{t+1}, \dots, y^T all at y^t and call this \hat{y} . Substitute \hat{y} in Equations (5.34) and (5.35) to obtain nominal radiosensitivity parameters $\hat{\alpha}_i$ and $\hat{\beta}_i$ for tumor voxels $i = 1, 2, \dots, n$;

3. solve problem (\hat{Q}_t) using the log-barrier interior point algorithm that uses Newton's method along with a line search procedure (see Section 5.3.1 above) to obtain an optimal fluence-map \hat{u} and, for each $m \in \mathcal{M}_3$, let $\mathcal{N}_m^t(\hat{u}) \subseteq \mathcal{N}_m$ denote the subset of $n_m - K_m$ voxels that receive the smallest doses among the n_m voxels in O_m under fluence-map \hat{u} ;
4. solve problem (\bar{Q}_t) using the log-barrier interior point algorithm that uses Newton's method along with a line search procedure (see Section 5.3.1 above) to obtain an optimal fluence-map u_* and use it in session t ;
5. calculate α_i^t and β_i^t via expressions (5.1)-(5.2) using the imaged hypoxia values y_i^t , for $i = 1, 2, \dots, n$; update tumor cell density and normal tissue states using u_* as

$$x_i^{t+1} = x_i^t \exp(-\alpha_i^t(A_i u_*) - \beta_i^t(A_i u_*)^2), \quad i = 1, 2, \dots, n, \quad (5.68)$$

$$z_j^{t+1,m} = z_j^{t,m} + (A_j^m u_*) + \rho_m (A_j^m u_*)^2, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (5.69)$$

$$v^{t+1,m} = v^{t,m} + \frac{\sum_{j=1}^{n_m} (A_j^m u_*) + \rho_m \sum_{j=1}^{n_m} (A_j^m u_*)^2}{n_m}, \quad m \in \mathcal{M}_2, \quad (5.70)$$

$$w_j^{t+1,m} = w_j^{t,m} + (A_j^m u_*) + \rho_m (A_j^m u_*)^2, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_3; \quad (5.71)$$

6. sample θ^t from its probability distribution and set $\ln y^{t+1} = \ln y^t + \theta^t$;
7. update $t \leftarrow t + 1$.

END DO

This CEC method assumes that a hypoxia image is acquired at the beginning of every treatment session. Another alternative is to not use hypoxia images but instead acquire cell density images at the beginning of every session. Note that, after a few treatment sessions, hypoxia information is at least indirectly included in the cell density image because hypoxic regions are radioresistant, and hence, *ceteris paribus*, are likely to have a higher cell density. One change that is needed to the

above CEC algorithm in this case is in Step 2 of the DO WHILE loop over $t \leq T$ — since a hypoxia image is not acquired, we do not know the oxygen partial pressure values y^t and hence we cannot employ Equations (5.1) and (5.2) to calculate the values of $\hat{\alpha}_i$ and $\hat{\beta}_i$; we thus instead simply use the constant values α and β under well-oxygenated conditions. Another change is in Step 5 — since the values of α_i^t and β_i^t cannot be calculated, we cannot use the update (5.68) to *calculate* the next cell density vector x^{t+1} . Instead, this cell density is directly *observed* in an image acquired at the beginning of the $t + 1$ st session. Consequently, we call this alternative approach “**CEC with cell density images**” below.

In the next section, we apply the CEC methodology developed here to five head-and-neck cancer test cases.

5.4 Results

We first describe the five test cases that were used in our simulation experiments. All test cases were generated using our in-house software Phantom Creator (PhanC) written in MATLAB (documentation provided as supplementary material not for publication). These test cases were three-dimensional and were carefully developed to be representative of clinical scenarios in terms of geometry. Specifically, our test cases were similar to those we used in Chapter 4, which were in turn similar in structure to those in [6, 85–88]. All cases used seven equally spaced coplanar beams and the beamlet resolution was $5 \times 5 \text{ mm}^2$. All voxels were $5 \times 5 \times 5 \text{ mm}^3$ leading to $\nu = 125 \text{ mm}^3$. We acknowledge that this voxel size is somewhat larger than what is used in the clinic today. This size helped us efficiently run hundreds of computer simulations to derive qualitative insights that we believe should hold independently of the voxel size. Another practical issue in any future clinical implementation is that the voxel resolution used in IMRT planning is likely to be finer than what is detectable with reasonable accuracy in a functional image. Since our tumor states correspond to voxel-by-voxel information acquired from functional images, we preferred to employ a coarser voxel resolution so as to avoid using a tedious notation that would account

for two different resolutions in our problem formulation.

All cases included spinal cord, brainstem, left and right parotids and unspecified normal tissue between these critical organs. The total number of voxels in the head-and-neck target and in the normal tissues, and the total number of beamlets is shown in Table 5.1 below.

case #	# of beamlets (k)	# of tumor voxels (n)	# of normal tissues voxels
1	1680	9096	17040
2	1485	6384	15360
3	1500	7248	16536
4	1584	7350	13850
5	1746	8640	21720

Table 5.1: Description of the geometry used in head-and-neck cancer cases.

The conventional fractionation schedule was assumed to include $N_{\text{conv}} = 35$ fractions. We included maximum dose constraints for spinal cord, brainstem and unspecified normal tissue. A dose-volume constraint for unspecified normal tissue was also added. Mean dose constraints were used for left and right parotids. The tolerance dose values for various normal tissues were similar to [1, 70–72] and are listed in Table 5.2 below.

Dose-response parameter estimates were chosen from the existing clinical literature [38, 44, 47–49, 51, 65, 75]. Head-and-neck tumor was assumed to exhibit $\alpha = 0.35 \text{ Gy}^{-1}$ and $\beta = 0.035 \text{ Gy}^{-2}$ under well-oxygenated conditions. The α/β ratios for all normal tissues were fixed at 3 Gy. Parameters OER_α , OER_β , and κ were fixed at 2.5, 3, and 3.28 mm Hg, respectively, as in [14, 97, 98]. Note that the two sufficient conditions for convexity in Lemma 5.3.3 hold for this choice of parameters. All computer simulations were performed on a 3.1 GHz iMac desktop with 16 GB RAM.

The stochastic vectors θ_t were assumed to be independent across treatment sessions

Normal tissue	D_{\max} (Gy)	D_{mean} (Gy)	D_{dv} (Gy), ϕ
Spinal cord	45	N/A	N/A
Brainstem	50	N/A	N/A
Left and right parotids	N/A	28	N/A
Unspecified normal tissue	77	N/A	70, 0.05

Table 5.2: Tolerance doses for various normal tissues in our head-and-neck test cases where the dose is administered in $N_{\text{conv}} = 35$ equal-dose fractions. Recall that for dose-volume type constraints no more than a volume fraction ϕ of the normal tissue can receive dose more than D_{dv} .

and to have a multivariate normal distribution with zero mean. Intuitively, tumor voxels that are close to each other should have similar oxygen partial pressures. This was captured by using a distance-based covariance matrix Σ for each θ_t . Suppose D_{ij} denotes the distance between tumor voxels i and j . We used two standard covariance functions from spatiotemporal statistics [122] to define matrix Σ : the exponential function and the rational-quadratic function. That is, the entries σ_{ij} of matrix Σ were obtained by

$$\sigma_{ij} = \exp(-D_{ij}/\sigma) \quad (5.72)$$

for the exponential case, and by

$$\sigma_{ij} = \left(1 + \frac{D_{ij}^2}{2\sigma}\right)^{-\sigma} \quad (5.73)$$

for the rational-quadratic case, where $\sigma > 0$ is a parameter of these functions. Note that other researchers have also used different distance-based (deterministic) models of hypoxia evolution [14, 118], and our thought process here is similar to theirs. Finally, the initial cell density was fixed at the normalized value of 1 in all tumor voxels since our results below do not depend on this value as long as it is invariant across voxels. This assumption of homogeneous initial cell density is also made in [94], and in fact can be easily relaxed without any change to our solution procedure

here.

We first compare the TNTCR values attained by three different methods for all five test cases. The first method, which we see as the base case or the conventional approach, is static; that is, it does not acquire any images, and hence does not adapt to hypoxia either directly or indirectly. In particular, we simply solved problem Q_1 once at the beginning of the treatment session with tumor radiosensitivity values fixed at $\alpha = 0.35 \text{ Gy}^{-1}$ and $\beta = 0.035 \text{ Gy}^{-2}$, and used the resulting fluence-map vector in all sessions. The second method employed CEC with hypoxia images and the third method employed CEC with cell density images as described above. For each test case, we performed two sets of 30 simulations each. For one set, we used the exponential covariance function and for the second set, we used the rational-quadratic function. The σ value was fixed at 0.1 for all simulations (we also tried a few other values of σ , but our qualitative conclusions were identical in all cases, so we report results for only one value here). We compared the average TNTCR over 30 simulations for these three methods. The percentage improvements of the two CEC methods over the static method were computed as follows:

$$\text{Percentage improvement} = \frac{\sum_{\text{sim}=1}^{30} \text{TNTCR}_{\text{sim}}^{\text{static}} - \sum_{\text{sim}=1}^{30} \text{TNTCR}_{\text{sim}}^{\text{CEC}}}{\sum_{i=1}^{30} \text{TNTCR}_{\text{sim}}^{\text{static}}} \times 100. \quad (5.74)$$

The improvement numbers are reported in Table 5.3. Figures 5.1(a) and (b) show the improvement in TNTCR for each simulation for all cases, for the exponential and rational-quadratic covariance functions, respectively, when only hypoxia images are used. Similarly, Figures 5.2(a) and (b) show the improvement in TNTCR for all cases, for the exponential and rational-quadratic covariance functions, respectively, when only cell density images are used.

Table 5.3, and Figures 5.1 and 5.2 show that adapting the treatment plan to the spatiotemporal evolution of hypoxia, either directly via hypoxia images or indirectly

Percentage improvement in TNTCR				
		Exponential covariance, $\sigma = 0.1$		rational-quadratic covariance, $\sigma = 0.1$
case #	hypoxia images	cell density images	hypoxia images	cell density images
1	14%	13%	24%	22%
2	63%	62%	66%	63%
3	60 %	59%	59%	58%
4	13%	11%	26%	24%
5	45%	42%	46%	42%

Table 5.3: Improvement in average TNTCR obtained by CEC over the static approach for exponential and rational-quadratic covariance functions.

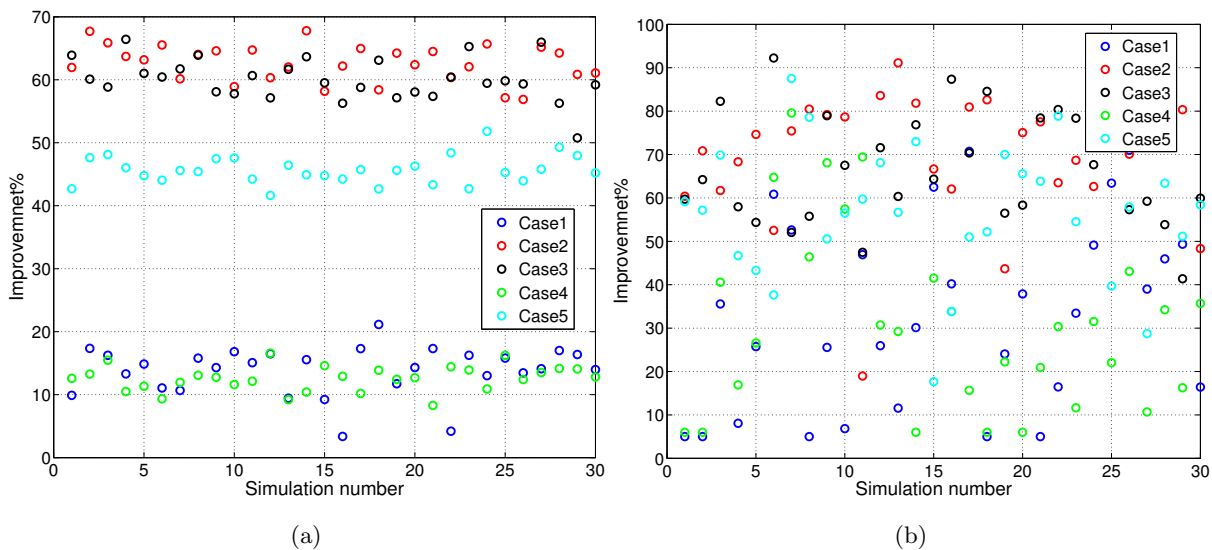


Figure 5.1: Improvement in TNTCR obtained by CEC using hypoxia images over the static approach for: (a) exponential covariance function, (b) rational-quadratic covariance function.

via cell density images, led to an improvement in average TNTCR for both covariance functions; in fact, adaptive treatment led to a better TNTCR in *every* simulation.

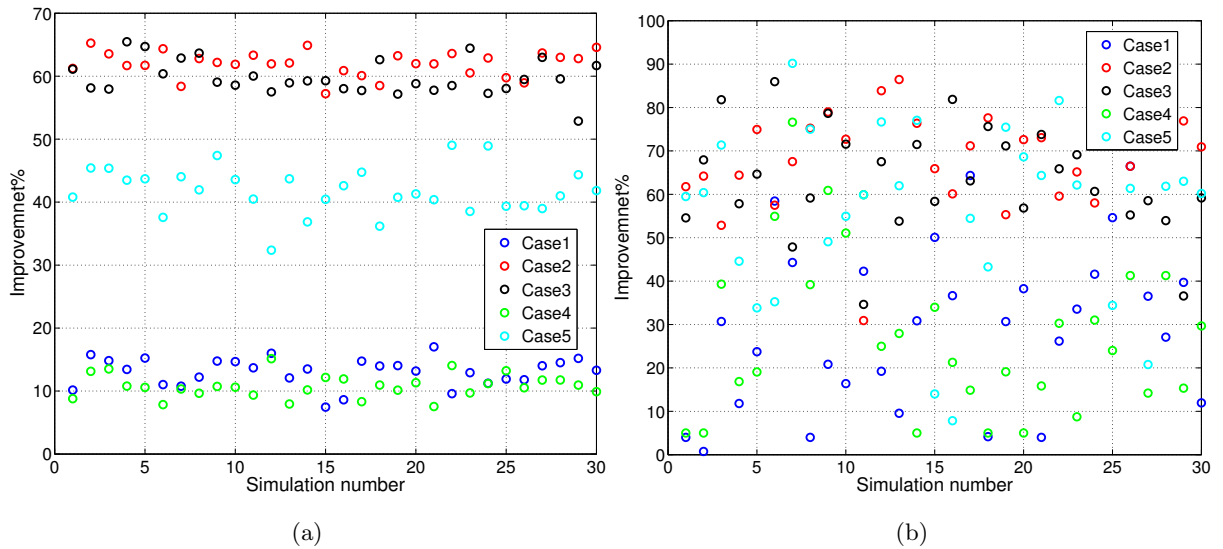


Figure 5.2: Improvement in TNTCR obtained by CEC with cell density images over the static approach for: (a) exponential covariance function, (b) rational-quadratic covariance function.

The improvements were essentially identical no matter whether hypoxia images were used or cell density images were used. This seems to corroborate our original belief that sufficient information about the spatiotemporal evolution of hypoxia might be included in cell density images. A potential clinical implication of this could be that the treatment planner might be able rely on other logistical and technological considerations in choosing which images to use in future practice.

There was a marked difference in the improvement achieved by CEC, with both covariance functions, for cases 1 and 4 as compared to cases 2, 3 and 5. Specifically, the improvement for cases 1 and 4 was smaller than that for cases 2, 3 and 5. We believe that this can be explained based on the difference in the geometry of these cases. The spinal cord and brainstem were much closer to the tumor in cases 1 and 4 than in cases 2, 3 and 5. This implies that cases 2, 3 and 5 inherently had more leeway for improvement via better treatment planning; as such, cases 2, 3 and 5 were

“easier”.

A careful comparison of Figure 5.1(a) with Figure 5.1(b) shows that the variance in improvements across different simulations was higher with the rational-quadratic covariance function than the exponential one. Similarly for Figures 5.2(a) and (b). We believe that this is rooted in the fact that the rational-quadratic covariance decreases with increasing distance at a slower rate than the exponential function. This means that the oxygen partial pressure in distant voxels tends to be more correlated in the rational-quadratic case thus leading to a higher variance in improvements.

We wanted to study the root-cause of the improvements achieved by CEC over the static method. Intuitively, it should be that CEC is better able to spatiotemporally redistribute dose based on observed images. This is verified by Figures 5.3, 5.4, 5.5, and 5.6. Specifically, Figures 5.3 and 5.4 show that the average dose (over all tumor voxels) delivered by CEC in various treatment sessions was different and in fact was higher in almost every simulation than that delivered by the static method. The CEC dose showed an upward trend over treatment sessions. This temporal dose boost can be explained, for example in the case of cell density images, as follows. Both the CEC and static methods start out with an identical dose distribution in the first treatment session. As treatment progresses, the CEC method, which gets the opportunity to re-plan, discovers that there is room for higher doses that will reduce the TNTCR while still protecting the normal tissues; it thus administers these higher doses in later sessions. Figures 5.5 and 5.6 further elaborate this point by contrasting the spatial distribution of the doses delivered by CEC in each simulation against that delivered by the static method, for cases 3 and 2 with the rational-quadratic and exponential covariance functions, respectively.

Finally, we employed a random walk model to simulate the spatiotemporal evolution of oxygen partial pressure in this paper. In lieu of actual hypoxia images, it is therefore important to validate that this simulation method generates a reasonable

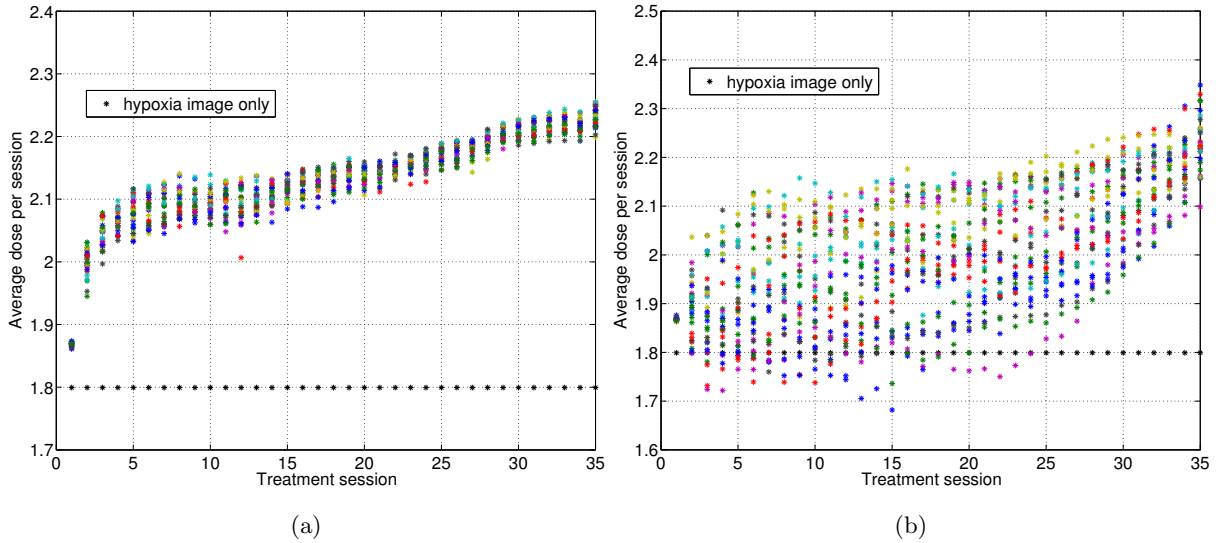


Figure 5.3: The average (over all tumor voxels) dose (Gy) per session delivered by CEC with hypoxia images in each one of the 30 simulations over 35 treatment sessions in case 1: (a) exponential covariance function (b) rational-quadratic covariance function. The solid flat line shows the constant dose (Gy) delivered by the static method in every session.

distribution of hypoxia. Toward this end, as an example, we show histograms of the hypoxia distribution generated over different treatment sessions in one simulation for case 2 with the exponential covariance function in Figure 5.7. The figure shows that the shape of the hypoxia distribution does indeed look lognormal, and in fact, as the treatment progresses, matches with the shapes shown in [101, 106, 118, 119] based on clinical literature as well as simulation studies. Histograms of the corresponding α_i^t and β_i^t values are also shown in Figures 5.8 and 5.9. The spikes in the histograms near 100mm Hg are somewhat artificial because we truncated the partial pressure values at 100mm Hg. We believe that this truncation has essentially no effect on our results since the radiosensitivity parameters α_i^t and β_i^t asymptotically and monotonically approach their well-oxygenated values of 0.35 Gy^{-1} and 0.035 Gy^{-2} , respectively, as the oxygen partial pressure increases. In fact, these parameters are already close to their

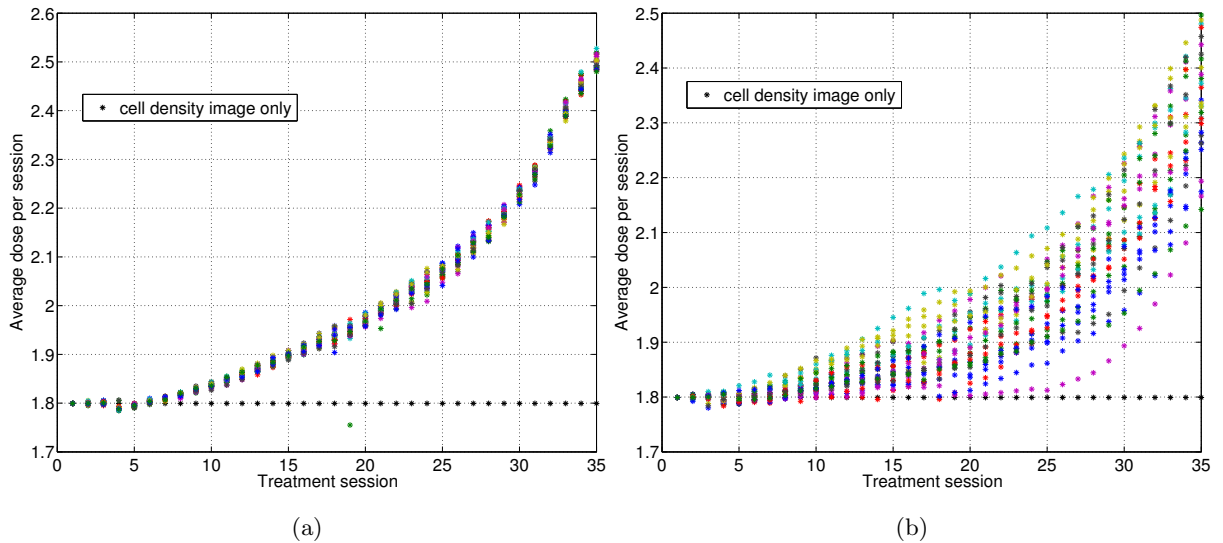


Figure 5.4: The average (over all tumor voxels) dose (Gy) per session delivered by CEC with cell density images in each one of the 30 simulations over 35 treatment sessions in case 1: (a) exponential covariance function (b) rational-quadratic covariance function. The solid flat line shows the constant dose (Gy) delivered by the static method in every session.

well-oxygenated values at 100mm Hg. The specific value of 100 was chosen from the histograms shown in [94, 119].

5.5 Conclusions and discussion

We applied the general stochastic control formalism for dynamic radiotherapy planning from [23] to hypoxia. We developed in detail a concrete implementation of an approximate control method called CEC by specifically tailoring it to our problem. This method involved the solution of a sequence of convex optimization problems; dose-volume constraints were tackled via a simple constraint generation procedure. We implemented two versions of this CEC method. One used (simulated) hypoxia images and hence adapted to the spatiotemporal evolution of hypoxia directly; the

other used (simulated) cell density images and thus adapted indirectly. The evolution of oxygen partial pressure was simulated using a first order vector autoregressive process with a distance-based covariance matrix from the literature in statistics. This simulation approach led to partial pressures that followed a lognormal distribution as observed in the clinical literature. The effect of hypoxia on tumor-radiosensitivity was modeled using the well-known OER approach from radiobiology.

We performed computer simulations to quantify any potential benefits of dynamically redesigning the fluence-maps based on hypoxia or cell density information acquired over the treatment course. Our simulations suggest that such dynamic planning could offer an improvement over static planning. This conclusion is similar in essence to t [15, 94] using dynamic planning to treat a canine patient and also to that in [14, 118] using more stylized (spherical) models of hypothetical tumors.

Significant challenges would need to be overcome, however, before any future clinical implementation of our theoretical approach. Potential benefits of adaptive planning are likely to depend crucially on the accuracy of the functional images used. There is hope that functional imaging technology will continue to improve over the next few decades; this could impact the value of adaptive planning. Ultimately, clinical trials will need to be run to assess the benefits of adapting to hypoxia. We hope that computer simulation studies, such as the one we presented here, would help guide the design of such trials.

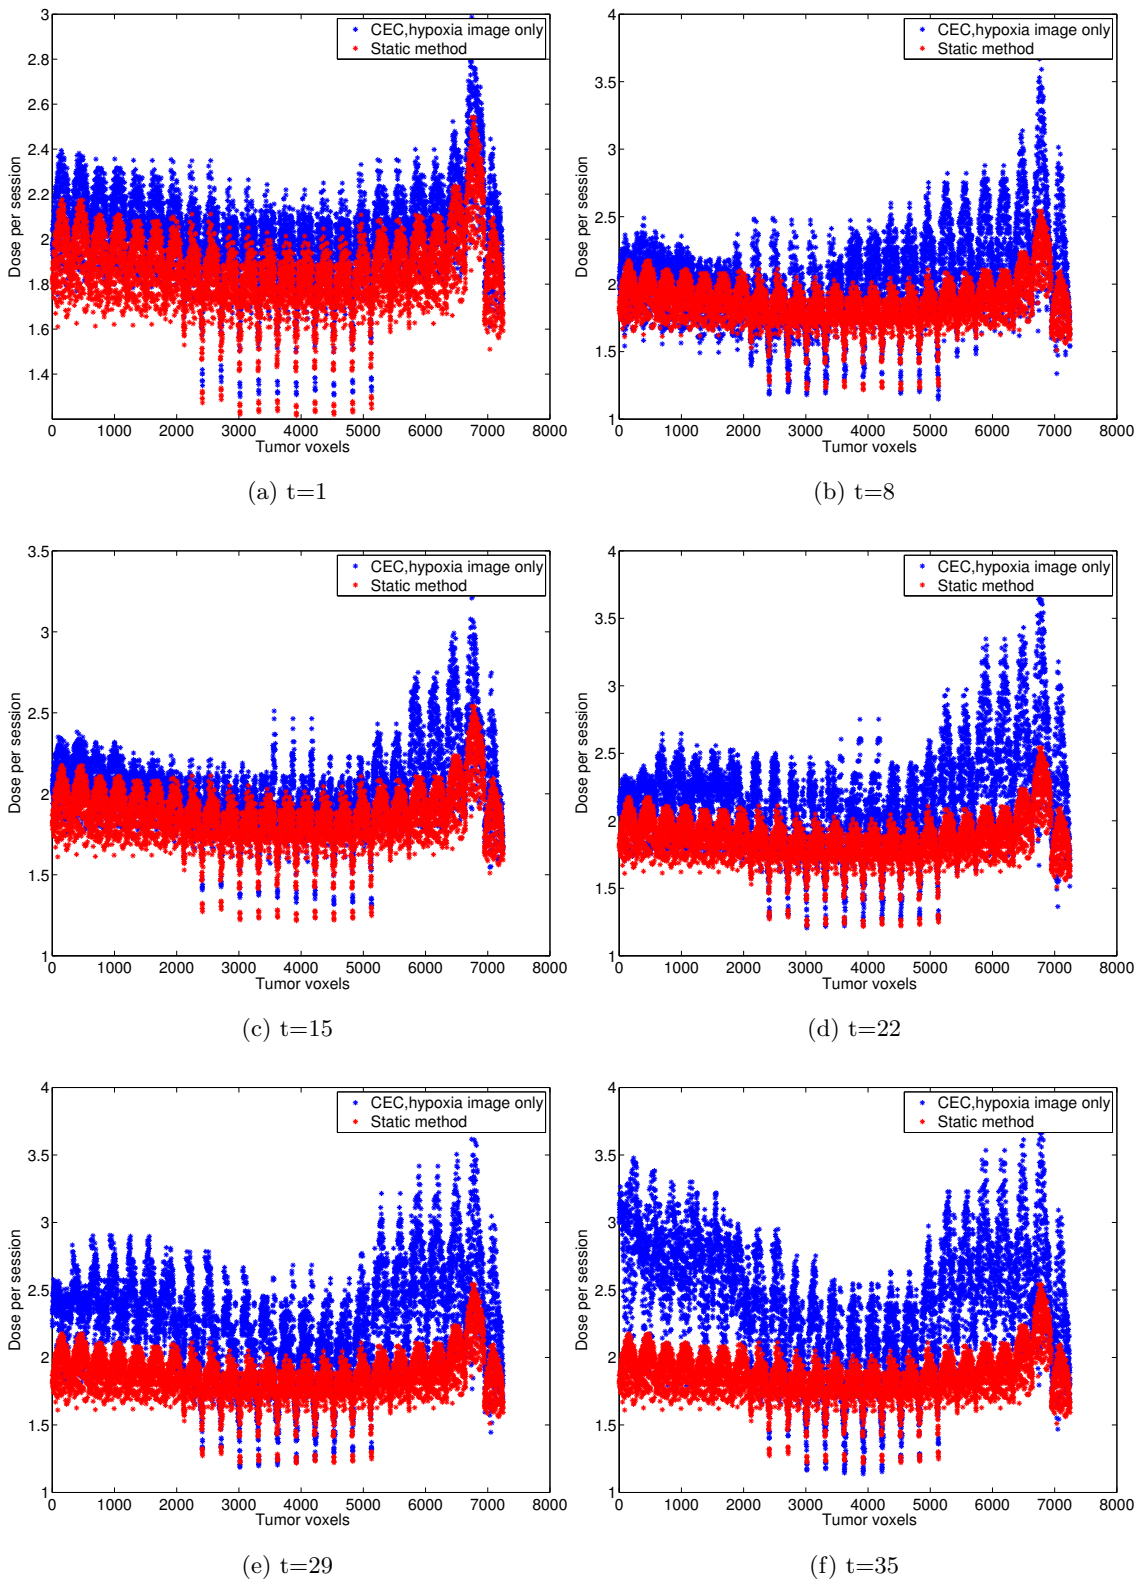


Figure 5.5: The dose (Gy) delivered by CEC with hypoxia images and by the static method in various treatment sessions for case 3 with the rational-quadratic covariance function.

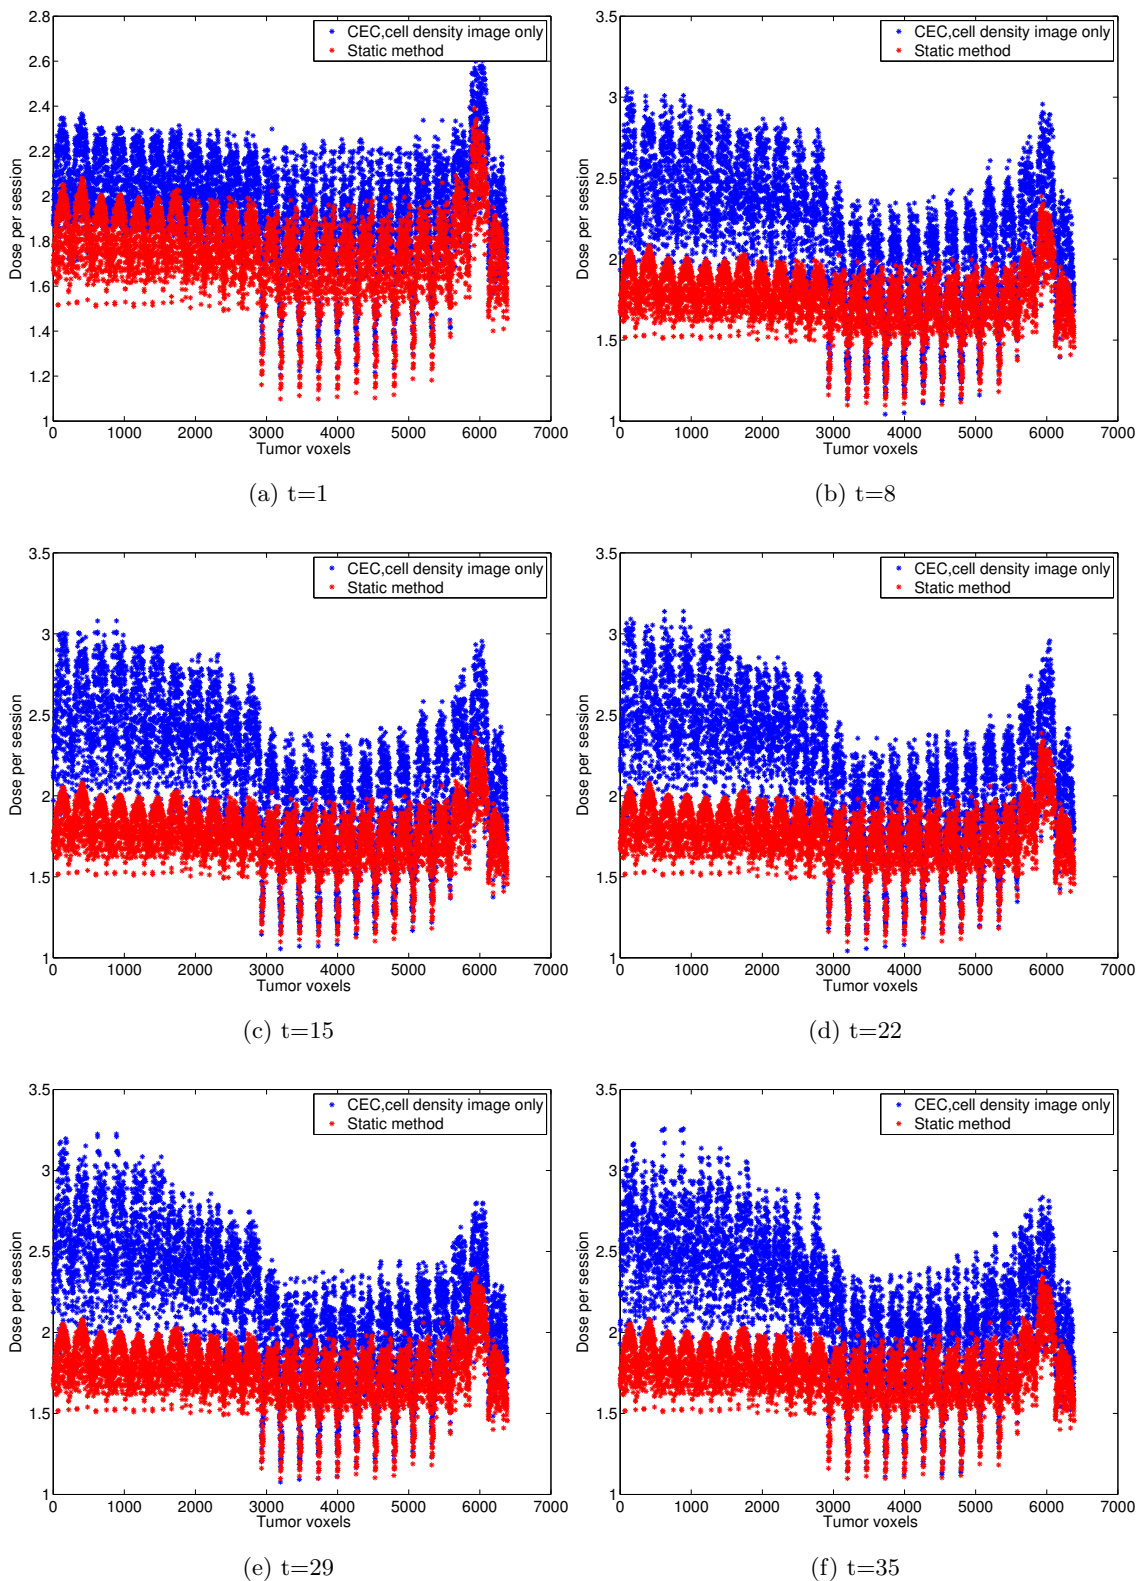


Figure 5.6: The dose (Gy) delivered by CEC with cell density images and by the static method in various treatment sessions for case 2 with the exponential covariance function.




Figure 5.7: A histogram of oxygen partial pressure in different tumor voxels for various treatment sessions for case 2 with the exponential covariance function.




Figure 5.8: A histogram of tumor alpha values (Gy^{-1}) over different tumor voxels for various treatment sessions for case 2 with the exponential covariance function.




Figure 5.9: A histogram of tumor beta values (Gy^{-2}) over different tumor voxels for various treatment sessions for case 2 with the exponential covariance function.

Chapter 6

OPTIMAL LEARNING OF TUMOR RESPONSE

6.1 Introduction

In all previous chapters of this dissertation, as well as in other existing work in this area [22–24], the overarching assumption is that the distribution of uncertainty in tumor-response parameters is known to the treatment planner, perhaps based on population estimates. In the radiation oncology community, there has been a recent surge of excitement generated by the future possibility of using tumor-images acquired over the treatment course to learn patient-specific parameter distributions [123]. The ultimate goal in this context is to learn a tumor’s response *while* optimally dosing a patient over the treatment course. In this chapter, we provide the first mathematical framework as an initial step toward this goal. Our work here falls within the fast-growing field of optimal learning, where a decision-maker chooses good-quality actions while learning his/her environment; we refer the reader to [124] for a survey/tutorial.

We continue to use the same notation as in Chapter 5. Recall that the tumor-response parameters for voxel i in session t are given by α_i^t and β_i^t , and we use $\rho_i^t \triangleq \beta_i^t/\alpha_i^t$. There is a large body of clinical literature that reports good population-based estimates of this so-called α over β ratio. So we assume that this ratio is a fixed constant r but the patient-specific uncertainty in response is captured by the radiosensitivity parameter α_i^t . Then, the tumor-cell dynamics in session t are written as

$$x_i^{t+1} = x_i^t \exp(-\alpha_i^t[(A_i u^t) + r(A_i u_t)^2]), \quad i = 1, 2, \dots, n. \quad (6.1)$$

As is prevalent in the existing literature [22–24], we assume that α_i^t are random variables that are independent and identically distributed across t and i . For concreteness, we assume that these random variables follow a Categorical distribution with k cat-

egories denoted by the vector $\vec{b} \triangleq (b_1, b_2, \dots, b_k)$ and the corresponding probability vector $\vec{p} \triangleq (p_1, p_2, \dots, p_k)$. That is, the probability that α_i^t takes a value b_j equals p_j , for $j = 1, 2, \dots, k$. Now, if this probability vector were known to the treatment planner, the Bellman's equations for the stochastic control problem of minimizing the expected TNTCR would be given by

$$J^t([x^t; z^t; v^t; w^t]) = \min_{u^t \in \mathcal{U}^t([z^t; v^t; w^t])} E_{\vec{p}} \left(J^{t+1}([x^{t+1}; z^{t+1}; v^{t+1}; w^{t+1}]) \right), \quad (6.2)$$

subject to

$$x_i^{t+1} = x_i^t \exp(-\alpha_i^t [(A_i u^t) + r(A_i u^t)^2]), \quad i = 1, 2, \dots, n, \quad (6.3)$$

$$z_j^{t+1, m} = z_j^{t, m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (6.4)$$

$$v^{t+1, m} = v^{t, m} + \frac{\sum_{j=1}^{n_m} (A_j^m u^t) + \rho_m \sum_{j=1}^{n_m} (A_j^m u^t)^2}{n_m}, \quad m \in \mathcal{M}_2, \quad (6.5)$$

$$w_j^{t+1, m} = w_j^{t, m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_3, \quad (6.6)$$

for all possible states $[x^t; z^t; v^t; w^t]$ starting with the boundary condition

$$J^{T+1}([x^{T+1}; z^{T+1}; v^{T+1}; w^{T+1}]) = \sum_{i=1}^n \nu x_i^{T+1}. \quad (6.7)$$

Again, for the reasons discussed in Chapter 5, this problem is not computationally tractable. Consequently, a much simpler, yet still intractable, version of this problem was studied using CEC in [22, 23]. Here, our goal is to study a more difficult version of this problem, where the distribution \vec{p} is not known to the decision maker.

Specifically, we pursue a Bayesian approach to tackle this optimal learning problem. We assume that the treatment planner has a Dirichlet prior with hyperparameter vector generically written as $\vec{a} = (a_1, a_2, \dots, a_k)$ on the unknown Categorical distribution \vec{p} . A key benefit of this approach is that the Dirichlet distribution is conjugate to the Categorical distribution. This means that the posterior distribution, after observing each tumor-voxel's response, is also Dirichlet and furthermore, the posterior hyperparameters are obtained via a simple update formula. In particular, if the hyperparameters at the beginning of session t are given by $\vec{a}_t = (a_{t,1}, a_{t,2}, \dots, a_{t,k})$, and

if the realized values of the random variables α_i^t in that session are such that c_j of the values belong to category j for $j = 1, 2, \dots$, then the posterior hyperparameters are simply given by $\vec{a}_{t+1} = (a_{t,1} + c_1, a_{t,2} + c_2, \dots, a_{t,k} + c_k)$. Note that the realized values of the random variables α_i^t are not observed but rather are *calculated* by inverting Equation (6.1) given $x_{t,i}$, $x_{t+1,i}$, and u_t . In this calculation process, which we call “information update”, x_t and x_{t+1} are observed in images and u_t is chosen by the decision maker. As such, this conjugate prior property allows us to simply use the hyperparameters as the information state in a Bayesian stochastic control problem whose Bellman’s equations are written as

$$J^t([\vec{a}_t; x^t; z^t; v^t; w^t]) = \min_{u^t \in \mathcal{U}^t([z^t; v^t; w^t])} E_{\vec{a}_t} \left(J^{t+1}([\vec{a}_{t+1}; x^{t+1}; z^{t+1}; v^{t+1}; w^{t+1}]) \right), \quad (6.8)$$

subject to

$$x_i^{t+1} = x_i^t \exp(-\alpha_i^t [(A_i u^t) + r(A_i u^t)^2]), \quad i = 1, 2, \dots, n, \quad (6.9)$$

$$\vec{a}_{t+1} = \vec{a}_t + \vec{c}, \quad (6.10)$$

$$z_j^{t+1,m} = z_j^{t,m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_1, \quad (6.11)$$

$$v^{t+1,m} = v^{t,m} + \frac{\sum_{j=1}^{n_m} (A_j^m u^t) + \rho_m \sum_{j=1}^{n_m} (A_j^m u^t)^2}{n_m}, \quad m \in \mathcal{M}_2, \quad (6.12)$$

$$w_j^{t+1,m} = w_j^{t,m} + (A_j^m u^t) + \rho_m (A_j^m u^t)^2, \quad j = 1, 2, \dots, n_m, \quad m \in \mathcal{M}_3, \quad (6.13)$$

for all possible states $[\vec{a}_t; x^t; z^t; v^t; w^t]$ starting with the boundary condition

$$J^{T+1}([\vec{a}_{T+1}; x^{T+1}; z^{T+1}; v^{T+1}; w^{T+1}]) = \sum_{i=1}^n \nu x_i^{T+1}. \quad (6.14)$$

Since the original clairvoyant problem (6.2) is computationally intractable, its Bayesian learning modification in (6.8) is also intractable. We therefore again resort to CEC for its approximate solution. Unlike the problem studied in [22, 23] and the problem studied in Chapter 5, the present optimal learning problem has two sources of uncertainty — one in the tumor-response parameter and the other in the treatment planner’s belief about the distribution of this parameter. In each session, our CEC

approach derives a deterministic nonlinear problem by replacing both these uncertainties with their expected values. Specifically, in session t , the treatment planner chooses the fluence-map u_t by assuming that the components p_1, p_2, \dots, p_k of the Categorical probability vector \vec{p} equal the corresponding component-expectations of the current Dirichlet belief, and that the tumor-response parameter equals the expected value of this resulting Categorical distribution. That is, the treatment planner solves a $T - t + 1$ -session deterministic nonlinear problem by assuming that for all subsequent treatment sessions $l = t, t + 1, \dots, T$, $\alpha_{l,i} = \sum_{j=1}^k p_j b_j$, where $p_j = \frac{a_{t,j}}{\sum_{j=1}^k a_{t,j}}$ for $j = 1, 2, \dots, k$. This deterministic nonlinear problem is identical in structure to the one in Chapter 5 and we therefore solve it using the interior point methodology described in detail there. The fluence-map so obtained is implemented and then the system stochastically evolves to the next state. The process continues until the end of the treatment session.

6.2 Numerical results

In this section, we use test case 1 from Chapter 5 to numerically test the performance of the aforementioned CEC approach. This case includes spinal cord, brainstem, left and right parotids and unspecified normal tissue between these critical organs. The tolerance doses for various normal tissues for this case were similar to the ones in Chapter 5. The α/β ratios for all normal tissues were fixed at 3 Gy and the α/β for tumor is fixed at $r = 10$. We assume that α_i^t variables follow a Categorical distribution with $k = 15$ categories with $\vec{b} = (0.1, 0.11, 0.12, \dots, 0.6)$. We used three probability vectors for the categorical distribution. The first one is based on a normal distribution between 0.1 and 0.5 with 0.35 and 0.25/3 for mean and standard deviation, respectively. The second and third probability factors are based on beta distributions, Beta(6,8) and Beta(2,6). Figure 6.1 shows this probability vector values.

We compare the TNTCR values attained by two different methods for this head-and-neck test case. In the first method, we assume that the probability vector for

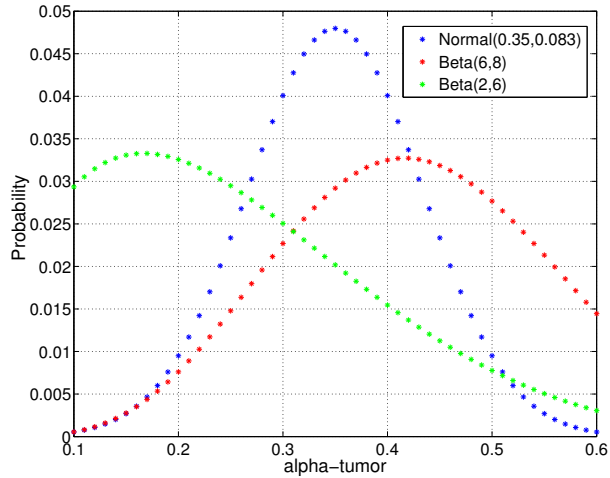


Figure 6.1: Three different probability vectors for tumor dose response parameter

the categorical distribution is known. In the second method, we learn the probability vector while optimizing dose over the treatment course. We compared the average TNTCR over 10 simulations for these two methods. Figure 6.3 shows the process of learning the probability vectors over 35 treatment sessions for one simulation.

The percentage difference between the TNTCR of both methods is shown in Figure 6.3. This percentage difference was computed as follows:

$$\text{Percentage difference} = \frac{\left| \sum_{\text{sim}=1}^{10} \text{TNTCR}_{\text{sim}}^{\text{CEC}} - \sum_{\text{sim}=1}^{10} \text{TNTCR}_{\text{sim}}^{\text{known}} \right|}{\sum_{i=1}^{10} \text{TNTCR}_{\text{sim}}^{\text{known}}} \times 100. \quad (6.15)$$

The TNTCR achieved while learning the distribution of the tumor dose-response parameter is only about 7.7% different on average over 10 simulations than the TNTCR when we know the distribution of the tumor dose-response parameter.

6.3 Conclusions

We developed a framework based on Bayesian dynamic programming to learn a patient-specific distribution of tumor-response uncertainty over the treatment course

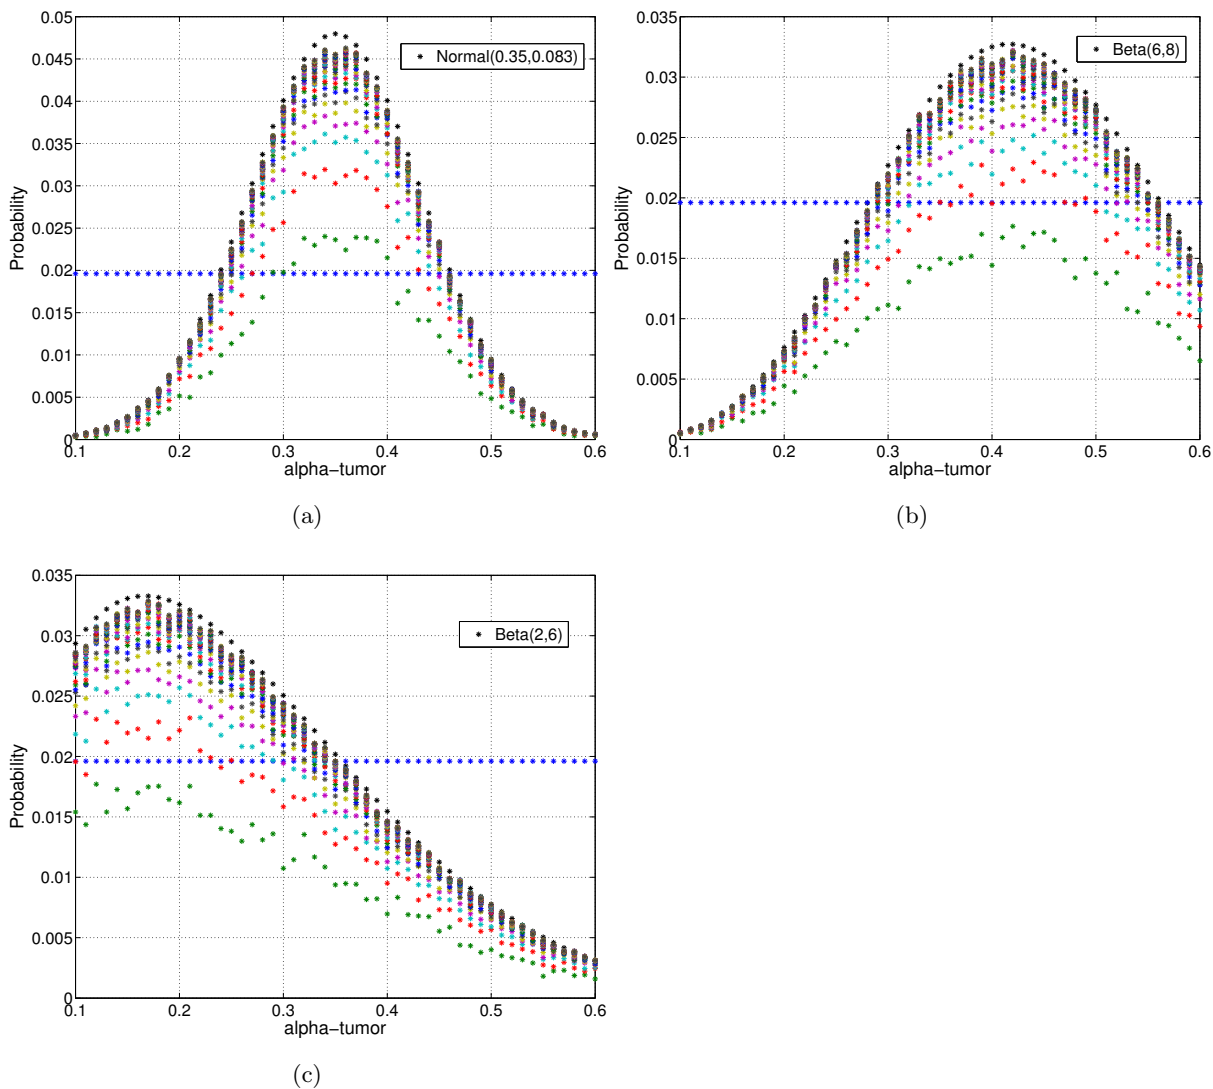



Figure 6.2: Learning the probability vector over 35 treatment sessions for normal and beta distributions

while simultaneously optimizing fluence-maps. Specifically, we assumed that a particular tumor-response parameter has a Categorical distribution and that the decision maker uses a Dirichlet prior on this Categorical distribution. We then exploited the conjugate property of the Categorical-Dirichlet pair to formulate a Bayesian dynamic



(a)

Figure 6.3: Percentage difference between the TNTCRs for 10 simulations for normal and beta distributions

programming formulation of our optimal learning problem. We then extended the approximate control scheme from Chapter 5 to efficiently tackle this problem. We showed via computer simulations on one head-and-neck test case that we were able to efficiently learn the tumor-response parameter distribution over the treatment course. Our simulations suggest that the TNTCR attained by our learning approach is close to that obtained by a hypothetical, idealized, clairvoyant algorithm.

Chapter 7

FUTURE WORK

There are several potentially fruitful directions for extending the research reported in this dissertation. We briefly outline them in this chapter.

It was shown in Chapter 3, in the context of our spatiotemporally separated model, that it may be suboptimal to use an equal-dosage or a single-dosage solution. It was recently shown by example in [80], in the context of a spatiotemporally integrated model, that equal-dosage solutions could be suboptimal. Unfortunately, a spatiotemporally integrated formulation of a problem that allows the fluence-maps to change across fractions is computationally difficult to solve. This is because the number of variables, which number already in the thousands in the equal-dosage formulation, is dozens of times larger in the non-stationary formulation. It would thus be interesting to devise approximate solution techniques to solve such problems by extending our convex programming method from Chapter 4.

The α and β values for both the tumor and the normal tissues were assumed to be known in Chapter 2, 3, and 4. It would be interesting to relax this assumption by formulating robust counterparts of these models. Specifically, it would be interesting to study whether the two-variable linear programming formulation from Chapter 3 can be utilized in exact or approximate solution of its robust counterpart. Moreover, it might be worthwhile to investigate a robust formulation and corresponding solution methods for our spatiotemporally integrated model from Chapter 4.

Lastly, our formulations of the optimal fractionation problem in Chapter 2, 3, 4 focussed on the deterministic case. When tumor-response is uncertain, it will be interesting to consider optimal stopping variants of these formulations. The basic groundwork for this idea was laid in [22] and its benefits were demonstrated using sim-

ulations on small-scale examples. In the future, it may be fruitful to implement these ideas on more full-scale problems as in Chapter 4, and in fact, compare the tumor-biological effect attained by such optimal stopping extensions with those achieved by the deterministic approach.

Appendix A

PHANC: A MATLAB SOFTWARE FOR CREATING PHANTOM TEST CASES ON A COMPUTER

This document describes how to generate the dose coefficient matrix A in PhanC software. Accurate dose calculation in the region of interest within a body resulting from a treatment field requires significant computation time, typically in the order of minutes. Beamlet intensity (fluence map) optimization involves a number of dose calculations at various optimization stages, and therefore it is widely used to approximate the dose calculation using a dose deposition coefficient matrix A in order to speed the optimization process. Suppose that there are N voxels in the region of interest and K beamlets in the treatment fields. A voxel is a small 3D segment within a body and a beamlet is a small 2D segment of treatment fields. Then the dose D is approximated by

$$D = Au, \tag{A.1}$$

$$D \in \mathbb{R}_+^N, A \in \mathbb{R}_+^{N \times K}, u \in \mathbb{R}_+^K, \tag{A.2}$$

where u_j is the j th beamlet intensity, and A_{ij} is the dose deposited in voxel i from the j th beamlet with unit intensity. Therefore, D_i is the dose deposited in voxel i from all K beamlets. It should be noted that A_{ij} does not represent the actual dose from such a small beamlet, where the lateral electronic equilibrium lacks. The accurate dose in voxel i cannot be obtained by summing the dose from each small beamlet in the treatment field since the beam profile of an individual beamlet is different from that of a broad beam. Our goal is to achieve a reasonable D using the above linear approximation with matrix A . Accurate dose calculation must follow for clinical treatment plans, which is beyond the scope of this document.

A.1 Input data

The starting point in to create the ultimate matrix A is to import and load an image of the internal anatomies including the cancerous targets, and nearby unspecified normal tissue and organ(s)-at-risk (OAR). OARs are healthy anatomies, such as the spinal cord or heart, that are close to the cancerous regions and need to be spared as much as possible. The required geometrical information is obtained using an anatomical image such as Computed Tomography (CT) or a Magnetic Resonance Imaging (MRI) scan. Once a CT or MRI scan is taken, a physician can contour the relevant tumor volume and delineate the critical organs close to the cancerous regions. This geometry information of the region of interest in patient's body can be stored as a 3D Geometry matrix, G , which is the first input for PhanC. Each entry in G represents a voxel within patient's body. Value of each entry in matrix G shows the voxel type, which can be either tumor, organ at risk or unspecified normal tissue between organs. These values are defined as following:

- Tumor voxels: -1 should be used in G to represent the tumor voxels.
- OAR voxels: Integer numbers from 1 to the number of OARs should be used to show different OARs. For example, in head and neck cancer, if we consider four ORAs: spinal cord, brain stem, right and left parotid, then integer numbers 1 to 4 should be used to show the voxels belong respectively to these organs.
- Unspecified normal tissue voxels: To show the voxels between different organs (OARs and tumor), -4 should be used.

Construction of matrix G is the initial step in providing the inputs for PhanC and constructing the ultimate matrix A . Once the user constructs the matrix G with above mentioned specification, it should be saved under the name and file extension " $G.mat$ " in the main directory(This file will be used later while running the code).

The rest of the inputs are provided step-by-step via the Graphic User Interface (GUI). The instruction for filling each row of this GUI is as follows:

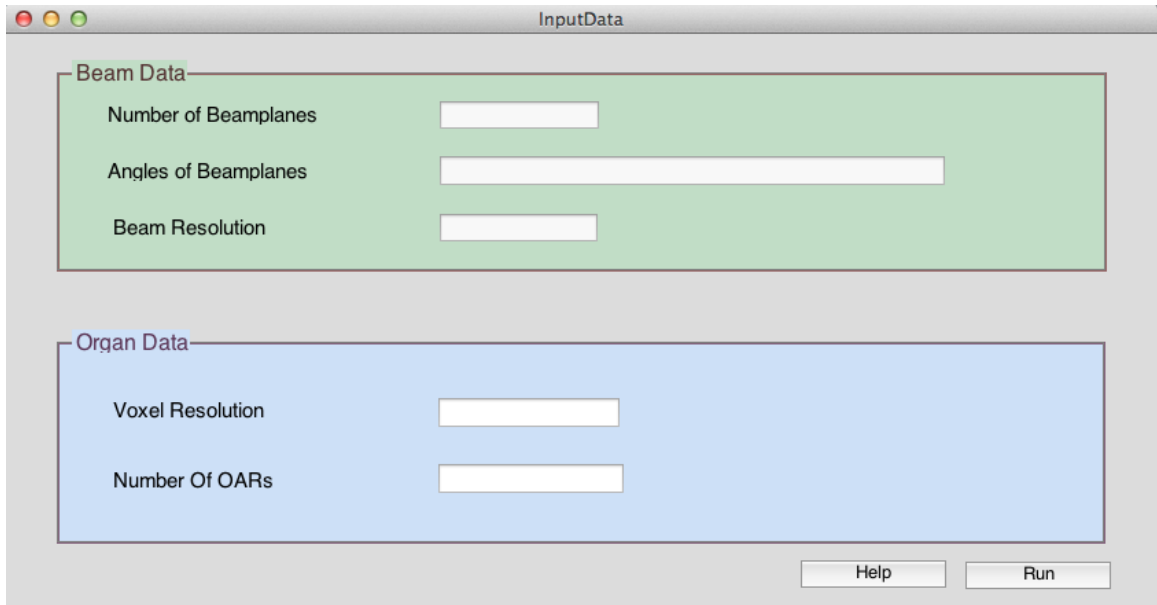


Figure A.1: PhanC Graphic User Interface (GUI)

- **Number of beamplanes:** In this box, user should insert number of radiation fields that is determined to deliver the radiation.
- **Angle of beamplanes:** This box is for the angles from which to deliver the radiation. The input could be just one number or a vector of numbers depending on the number of beam planes. The user should enter the angles as a coefficient of π . As an example, if the chosen angles are $[\pi/4, \pi/2, \pi, 3\pi/2]$, the input should be $1/4, 1/2, 1, 3/2$.
- **Beam resolution:** This input represents the length of each beamlet in *cm*. For example, if the size of each beamlet is $5 \times 5 \text{ mm}^2$, the input should be 0.5.

- **Voxel resolution:** Similar to the previous box, this box is for each voxel size. For instance, if the size of each voxel is $3 \times 3 \times 3 \text{ mm}^3$, the input should be 0.3.
- **Number of OARs:** Finally, the last box is to enter the number of organ(s)-at-risk which user has included for this case.

A.2 Method

Consider a geometry shown in Figure A.2, where d_1 represents the distance from the source to the body surface and d_2 is the distance from the source to the voxel of interest, i . Then A_{ij} is given by

$$A_{ij} = TMR(d_2 - d_1) \times DF(d_2) \times OF(S_{eq}) \times SF, \quad (\text{A.3})$$

where each factor in the above equation is as explained below. Note that it is normalized so that it yields 1cGy at d_{max} (depth at which the dose achieves its maximum along the beam central axis), where the distance from the 10cm x 10cm radiation field to d_{max} is set to 100cm. This is identical to the calibration condition for the linear accelerators in practice.




Figure A.2: Geometry to calculate the dose deposition coefficients, A_{ij} , in matrix A

- TMR: Tissue maximum ratio (TMR) is the ratio of the dose at an arbitrary depth to the dose at d_{max} from 10cm x 10cm radiation field size. The distance from the radiation field to the point of measurement is constant at 100cm. Therefore, TMR is a function of depth and photon energy. Note that TMR is normalized so that it is unity at d_{max} .
- DF: Distance factor (DF) is to account for the fact that the dose is inversely related to the square of distance since TMR is measured at 100cm whereas the dose calculation is done at d_2 from the source. It is defined as $DF = \left(\frac{100}{d_2}\right)^2$.
- OF: Output factor (OF) is the ratio of the dose from an arbitrary field size to the dose from 10cm x 10cm field size at d_{max} . An arbitrary field size is converted to its equivalent square (S_{eq}) with which the OF was measured. S_{eq} is the square that has the same $\frac{4 \times Area}{Perimeter}$ as the original arbitrary field size, see Table A.1.
- SF: Scaling factor (SF) is applied to account for the penumbra of the treatment field, and is dependent on the depth of a voxel, i.e., $d_2 - d_1$. Let E_1 and E_2 denote the set of beamlets that are located outermost and second outermost respectively. All other beamlets are considered as interior beamlets denoted by I as shown in Figure A.3. For each voxel i , primary beamlets, P_i , are defined as the set of beamlets that are within the specified distance from the orthogonal projection of a voxel i to the beam plane.

For the beamlet $j \in I$, SF is unity if j is the primary beamlet of the voxel i , and zero otherwise. For the beamlets in E_1 and E_2 , SF is provided by the measured penumbra profiles from 10cm x 10cm radiation field at various depths ($d_2 - d_1$). In profiles, x axis represents the distance from the beamlet j (outer side for $j \in E_1$ and inner side for $j \in E_2$ - see Figure A.3(c)) to the orthogonal projection of the center of the voxel i to the beam plane, and y axis is the relative dose normalized to the dose at the center of the field. Using E_1 and E_2

ensures that the edge of the field corresponds to 50% of the dose at the center of the field, and the dose outside the radiation field is approximated using the measured profile.

All measured profiles are obtained from the data using University of Washington Medical Center, Elekta Synergy C machine with 6 MV.

In summary, SF is given by

$$SF(i, j, d_2 - d_1) = \left\{ \begin{array}{ll} 1 & \forall j \in I, \text{ and } j \in P_i \\ 0 & \forall j \in I, \text{ and } j \notin P_i \\ \text{profile (See Figure A.4, blue curve)} & \forall j \in E_1, \text{ and } j \in P_i \\ 0 & \forall j \in E_1, \text{ and } j \notin P_i \\ \text{profile (See Figure A.4, red curve)} & \forall j \in E_2 \end{array} \right\}.$$




Figure A.3: Definition of (a) interior and exterior beamlets (b) primary beamlets P_i for voxel i (c) x-axis in profiles




Figure A.4: Measured penumbra from 10cm x 10cm field at various depths

A.3 Validation

The cube phantom with a size of $20 \times 20 \times 20$ (cm³) was created. The isodose distributions at 1.5 (6MV d_{max}), 5, 10, and 15 (cm) depths from the field size of 10×10 , 20×20 , 30×30 (cm²) with unit radiation intensities were calculated using PhanC and compared with the measured profiles.

Plots should show “calculated (dashed line)” vs “measured (solid line)” in the same plot - errors are estimated by the maximum difference in percentage (flat area)

or the distance to match (high gradient area)

Field size, cm	Output factor (OF) at d_{max}
2	0.913
3	0.933
4	0.948
5	0.957
6	0.970
7	0.976
8	0.987
9	0.992
10	1.000
11	1.003
12	1.011
13	1.015
14	1.021
15	1.024
20	1.044
25	1.053
30	1.062
35	1.065
40	1.066

Table A.1: OF is the ratio of the dose from an arbitrary field size to the dose from 10cm x 10cm field size at d_{max}

BIBLIOGRAPHY

- [1] Lawrence B. Marks, Ellen D. Yorke, Andrew Jackson, Randall K. Ten Haken, Louis S. Constine, Avraham Eisbruch, Søren M. Bentzen, Jiho Nam, and Joseph O. Deasy. Use of normal tissue complication probability models in the clinic. *International Journal of Radiation Oncology Biology Physics*, 76(3):S10–S19, 2010.
- [2] B Stewart and C P Wild. World cancer report. Technical report, International Agency for Research on Cancer, 2014.
- [3] National Cancer Institute. Radiation therapy for cancer. <http://www.cancer.gov/cancertopics/factsheet/Therapy/radiation>, June 2010.
- [4] S. Soderstrom and A. Brahme. Optimization of the dose delivery in a few field techniques using radiobiological objective functions. *Med Phys*, 1993.
- [5] T. Bortfeld and W. Schlegel. Optimization of beam orientations in radiation therapy: some theoretical considerations. *Physics in Medicine and Biology*, 38:291–304, 1993.
- [6] H E Romeijn, R K Ahuja, J F Dempsey, and A Kumar. A new linear programming approach to radiation therapy treatment planning problems. *Operations Research*, 54(2):201–216, 2006.
- [7] Michael C. Ferris, Jinho Lim, and David M. Shepard. An optimization approach for radiosurgery treatment planning. *SIAM Journal on Control and Optimization*, 13(3):921–937, 2003.
- [8] Eva K. Lee, Tim Fox, and Ian Crocker. Integer programming applied to

- intensity- modulated radiation therapy treatment planning. *Annals of Operations Research*, 119:165–181, 2003.
- [9] D M Shepard, M C Ferris, G H Olivera, and T R Mackie. Optimizing the delivery of radiation therapy to cancer patients. *SIAM Review*, 41(4):721–744, 1999.
- [10] R. Reemtsen and M. Alber. Continuous optimization of beamlet intensities for photon and proton radiotherapy. *Technical report, Brandenburgisch Technische Universit at Cottbus*, 2004.
- [11] E J Hall and A J Giaccia. *Radiobiology for the Radiologist*. Lippincott Williams & Wilkins, Philadelphia, Pennsylvania, USA, 2005.
- [12] J R Bading and A F Shields. Imaging of cell proliferation: status and prospects. *Journal of Nuclear Medicine*, 49(6):64S–80S, 2008.
- [13] S M Eschmann, F Paulsen, C Bedeshem, H-J Machulla, T Hehr, M Bamberg, and R Bares. Hypoxia-imaging with 18-F-Misonidazole and PET: Changes of kinetics during radiotherapy of head-and-neck cancer. *Radiotherapy and Oncology*, 83(3):406–410, 2007.
- [14] C P South, M Partridge, and P M Evans. A theoretical framework for prescribing radiotherapy dose distributions using patient-specific biological information. *Medical Physics*, 35(10):4599–4611, 2008.
- [15] A Sovik, E Malinen, H K Skogmo, Soren M Bentzen, O S Bruland, and D R Olsen. Radiotherapy adapted to spatial and temporal variability in tumor hypoxia. *International Journal Radiation Oncology Biology Physics*, 68(5):1496–1504, 2007.
- [16] L Barazzuol, N G Burnet, R Jena, B Jones, S J Jefferies, and N F Kirkby. A mathematical model of brain tumour response to radiotherapy and

- chemotherapy considering radiobiological aspects. *Journal of Theoretical Biology*, 262(3):553–565, 2010.
- [17] L G Hanin. A stochastic model of tumor response to fractionated radiation: limit theorems and rate of convergence. *Mathematical Biosciences*, 191:1–17, 2004.
- [18] G Powathil, M Kohandel, S Sivaloganathan, A Oza, and M Milosevic. Mathematical modeling of brain tumors: effects of radiotherapy and chemotherapy. *Physics in Medicine and Biology*, 52:3291–3306, 2007.
- [19] U Schneider. Mechanistic model of radiation-induced cancer after fractionated radiotherapy using the linear-quadratic formula. *Medical Physics*, 36(4):1138–1143, 2009.
- [20] N Stavreva, P Stavrev, and B G Fallone. Probability dynamics of a repopulating tumor in case of fractionated external radiotherapy. *Physica Medica*, 25(4):181–191, 2009.
- [21] R D Stewart and X A Li. BGRT: Biologically guided radiation therapy — the future is fast approaching. *Medical Physics*, 34(10):3739–3751, 2007.
- [22] M Kim. *A Mathematical Framework for Spatiotemporal Optimality in Radiation Therapy*. PhD thesis, University of Washington, Seattle, Washington, USA, 2010.
- [23] Minsun Kim, Archis Ghate, and Mark H. Phillips. A stochastic control formalism for dynamic biologically conformal radiation therapy. *European Journal of Operational Research*, 219(3):541 – 556, 2012.
- [24] M Kim, A Ghate, and M H Phillips. A markov decision process approach to temporal modulation of dose fractions in radiation therapy planning. *Physics in Medicine and Biology*, 54:4455–4476, 2009.

- [25] C C Ling, J Humm, S Larson, H Amols, Z Fuks, S Leibel, and J A Koutcher. Towards multidimensional radiotherapy (MD-CRT): biological imaging and biological conformity. *International Journal of Radiation Oncology Biology Physics*, 47(3):551–560, 2000.
- [26] J L Tatum, G J Kelloff, R J Gillies, J M Arbeit, J M Brown, and K S Chao KS *et al.* Hypoxia: importance in tumor biology, noninvasive measurement by imaging, and value of its measurement in the management of cancer therapy. *International Journal of Radiation Oncology Biology Physics*, 82, 2006.
- [27] S Webb. *Contemporary IMRT: Developing Physics and Clinical Implementation*. IOP Publishing, Bristol, UK, 2010.
- [28] Chandra Burman, Chen-Shou Chui, Gerald Kutcher, Steven Leibel, Michael Zelefsky, Thomas LoSasso, Spiridon Spirou, Qiuwen Wu, Jie Yang, J Stein, Radhe Mohan, Zvi Fuks, and C C Ling. Planning, delivery, and quality assurance of intensity-modulated radiotherapy using dynamic multileaf collimator: A strategy for large-scale implementation for the treatment of carcinoma of the prostate. *International Journal of Radiation Oncology Biology Physics*, 39(4):863 – 873, 1997.
- [29] M Ehrgott, C Guler, H W Hamacher, and L Shao. Mathematical optimization in intensity modulated radiation therapy. *4OR*, 6(3):199–262, 2008.
- [30] Avraham Eisbruch. Intensity-modulated radiotherapy of head-and-neck cancer: encouraging early results. *International Journal of Radiation Oncology Biology Physics*, 53(1):1–3, 2002.
- [31] M Langer, E K Lee, J O Deasy, R L Rardin, and J A Deye. Operations research applied to radiotherapy, an NCI-NSF-sponsored workshop February 7-9, 2002. *International Journal of Radiation Oncology Biology Physics*, 57(3):762–768, 2003.

- [32] H Rodney Withers. Biologic basis for altered fractionation schemes. *Cancer*, 55(S9):2086–2095, 1985.
- [33] Anesa Ahamad, David I. Rosenthal, and K. Kian Ang. *Squamous cell head and neck cancer : recent clinical progress and prospects for the future*. Current Clinical Oncology. Humana Press, Totowa, New Jersey, USA, 2005.
- [34] Giorgio Arcangeli, Biancamaria Saracino, Sara Gomellini, Maria Grazia Petronari, Stefano Arcangeli, Steno Sentinelli, Simona Marzi, Valeria Landoni, Jack Fowler, and Lidia Strigari. A prospective phase iii randomized trial of hypofractionation versus conventional fractionation in patients with high-risk prostate cancer. *International Journal of Radiation Oncology Biology Physics*, 78(1):11–18, 2010.
- [35] Karen K. Fu, Thomas F. Pajak, Andy Trotti, Christopher U. Jones, Sharon A. Spencer, Theodore L. Phillips, Adam S. Garden, John A. Ridge, Jay S. Cooper, and K.Kian Ang. A radiation therapy oncology group (RTOG) phase iii randomized study to compare hyperfractionation and two variants of accelerated fractionation to standard fractionation radiotherapy for head and neck squamous cell carcinomas: first report of RTOG 9003. *International Journal of Radiation Oncology Biology Physics*, 48(1):7 – 16, 2000.
- [36] A S Garden. Altered fractionation for head and neck cancer. *Oncology*, 15(10):1326–1332, 2001.
- [37] J C Horiot, R Le Fur, T N’Guyen, C Chenal, S Schraub, S Alfonsi, G Gardani, W Van Den Bogaert, S Danczak, and M Bolla. Hyperfractionation versus conventional fractionation in oropharyngeal carcinoma: final analysis of a randomized trial of the eortc cooperative group of radiotherapy. *Radiotherapy Oncology*, 25(4):231–241, December 1992.
- [38] K F Ho, J F Fowler, A J Sykes, B K Yap, L W Lee, and N J Slevin. Imrt

- dose fractionation for head and neck cancer: Variation in current approaches will make standardisation difficult. *Acta Oncologica*, 3(431-439):2009, 48.
- [39] JC Horiot, P. Bontemps, Den Bogaert W Van, Fur R. Le, Den Weijngaert D Van, M. Bolla, J. Bernier, A. Lusinchi, M. Stuschke, J. Lopez-Torrecilla, AC Begg, M. Pierart, and L. Collette. Accelerated fractionation (AF) compared to conventional fractionation (CF) improves loco-regional control in the radiotherapy of advanced head and neck cancers: Results of the eortc 22851 randomized trial. *Radiotherapy and Oncology : Journal of the European Society for Therapeutic Radiology and Oncology*, 44(2):111–121, 1997.
- [40] Hosam A. Kader, Aminudin R. Mydin, Matthew Wilson, Cheryl Alexander, Jeevin Shahi, Irvin Pathak, Jonn S. Wu, and Pauline T. Truong. Treatment outcomes of locally advanced oropharyngeal cancer: A comparison between combined modality radio-chemotherapy and two variants of single modality altered fractionation radiotherapy. *International Journal of Radiation Oncology Biology Physics*, 80(4):1030 – 1036, 2011.
- [41] Simona Marzi, Biancamaria Saracino, Maria Petrongari, Stefano Arcangeli, Sara Gomellini, Giorgio Arcangeli, Marcello Benassi, and Valeria Landoni. Modeling of alpha/beta for late rectal toxicity from a randomized phase ii study: conventional versus hypofractionated scheme for localized prostate cancer. *Journal of Experimental & Clinical Cancer Research*, 28(1):117–124, 2009.
- [42] S Rockwell. Experimental radiotherapy: a brief history. *Radiation Research*, 150(Supplement):S157–S169, November 1998.
- [43] A. Trotti, K.K. Fu, T.F. Pajak, C.U. Jones, S.A. Spencer, T.L. Phillips, A.S. Garden, J.A. Ridge, J.S. Cooper, and K.K. Ang. Long term outcomes of RTOG 90–03: A comparison of hyperfractionation and two variants of accelerated fractionation to standard fractionation radiotherapy for head and neck squamous

- cell carcinoma. *International Journal of Radiation Oncology Biology Physics*, 63(Supplement 1):S70–S71, 2005.
- [44] J F Fowler. How worthwhile are short schedules in radiotherapy?: A series of exploratory calculations. *Radiotherapy and Oncology*, 18(2):165–181, 1990.
- [45] J F Fowler and M A Ritter. A rationale for fractionation for slowly proliferating tumors such as prostatic adenocarcinoma. *International Journal of Radiation Oncology Biology Physics*, 32(2):521–529, 1995.
- [46] B Jones, L T Tan, and R G Dale. Derivation of the optimum dose per fraction from the linear quadratic model. *The British Journal of Radiology*, 68(812):894–902, 1995.
- [47] J F Fowler. Biological factors influencing optimum fractionation in radiation therapy. *Acta Oncologica*, 40(6):712–717, 2001.
- [48] J F Fowler. Is there an optimal overall time for head and neck radiotherapy? a review with new modeling. *Clinical Oncology*, 19(1):8–27, 2007.
- [49] J F Fowler. Optimum overall times ii: Extended modelling for head and neck radiotherapy. *Clinical Oncology*, 20(2):113–126, 2008.
- [50] C I Armpilia, R G Dale, and B Jones. 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, 2004.
- [51] Yong Yang and Lei Xing. Optimization of radiotherapy dose-time fractionation with consideration of tumor specific biology. *Medical Physics*, 32(12):3666–3677, 2005.
- [52] David J. Brenner, Lynn R. Hlatky, Philip J. Hahnfeldt, Eric J. Hall, and Rainer K. Sachs. A convenient extension of the linear-quadratic model to include

- redistribution and reoxygenation. *International Journal of Radiation Oncology Biology Physics*, 32(2):379–390, 1995.
- [53] A Bertuzzi, C Bruni F Papa, and C Sinisgalli. Optimal solution for a cancer radiotherapy problem. *Journal of Mathematical Biology*, 66(1-2):311–349, 2013.
- [54] Masahiro Mizuta, Seishin Takao, Hiroyuki Date, Naoki Kishimoto, Kenneth L. Sutherland, Rikiya Onimaru, and Hiroki Shirato. 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, 2012.
- [55] Jan Unkelbach, David Craft, Ehsan Saleri, Jagdish Ramakrishnan, and Thomas Bortfeld. The dependence of optimal fractionation schemes on the spatial dose distribution. *Physics in Medicine and Biology*, 58(1):159–167, 2013.
- [56] R Jeraj and P Keall. Monte carlo-based inverse treatment planning. *Physics in Medicine and Biology*, 44(8):1885–1896, 1999.
- [57] J V Siebers, S D Tong, M Lauterbach, Q W Wu, and R Mohan. Acceleration of dose calculations for intensity-modulated radiotherapy. *Medical Physics*, 28(6):903–910, 2001.
- [58] S V Spirou and C S Chui. A gradient inverse planning algorithm with dose-volume constraints. *Medical Physics*, 25(3):321–333, 1998.
- [59] Z Tian, M Zarepisheh, X Jia, and S B Jiang. The fixed-point iteration method for imrt optimization with truncated dose deposition coefficient matrix. Technical report, arxiv, 2013.
- [60] S Webb and M Oldham. 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, 1996.

- [61] H Keller, G Meier, A Hope, and M Davison. Fractionation schedule optimization for lung cancer treatments using radiobiological and dose distribution characteristics. *Medical Physics*, 39(6):3811–3811, 2012.
- [62] Thomas Bortfeld, Jagdish Ramakrishnan, John N Tsitsiklis, and Jan Unkelbach. Optimization of radiotherapy fractionation schedules in the presence of tumor repopulation. http://pages.discovery.wisc.edu/~jramakrishnan/BRT2013_repop.pdf, December 2013.
- [63] Zhi-Quan Luo, Wing-Kin Ma, A. M. C So, and Yinyu Ye. Semidefinite relaxation of quadratic optimization problems. *IEEE Signal Processing Magazine*, 27(3):20–34, 2010.
- [64] S Boyd and L Vandenberghe. *Convex Optimization*. Cambridge University Press, Cambridge, UK, 2004.
- [65] M V Williams, J Denekamp, and J F Fowler. A review of alpha/beta ratios for experimental tumors: implications for clinical studies of altered fractionation. *International Journal of Radiation Oncology Biology Physics*, 11(1):87–96, January 1985.
- [66] W C Waterhouse. Do symmetric problems have symmetric solutions? *The American Mathematical Monthly*, 90(6):378–387, 1983.
- [67] A Dasu. Is the alpha/beta value for prostate tumours low enough to be safely used in clinical trials? *Clinical Oncology*, 19(5):289–301, 2007.
- [68] J F Fowler, R Chappell, and M Ritter. Is alpha/beta for prostate tumors really low? *International Journal of Radiation Oncology Biology Physics*, 50(4):1021–1031, 2001.
- [69] J Z Wang, M Guerrero, and X A Li. How low is the alpha/beta ratio for prostate cancer? *International Journal of Radiation Oncology Biology Physics*, 55(1):194–203, 2003.

- [70] Emami B, Lyman J, Brown A, Coia L, Goiten M, and et al. Munzenrider JE. Tolerance of normal tissue to therapeutic radiation. *International Journal of Radiation Oncology Biology Physics*, 21(1):109–122, 1991.
- [71] T S Kehwar. Analytical approach to estimate normal tissue complication probability using best fit of normal tissue tolerance doses into the ntcp equation of the linear quadratic model. *Journal of Cancer Research and Therapy*, 1(3):168–179, 2005.
- [72] Panayiotis Mavroidis, Brigida Costa Ferreira, and Maria do Carmo Lopes. Response-probability volume histograms and iso-probability of response charts in treatment plan evaluation. *Medical Physics*, 38(5):2382–2397, 2011.
- [73] M Grant and S Boyd. CVX: MATLAB software for disciplined convex programming (web page and software), 2009.
- [74] J Ananworanich, U Siangphoe, P Cardiello, W Apateerapong, B Hirschel, A Mahanontharit, S Ubolyam, D Cooper, P Phanuphak, and K Ruxrungtham. Highly active antiretroviral therapy (HAART) retreatment in patients on CD4-guided therapy achieved similar virologic suppression compared with patients on continuous HAART - The HIV Netherlands Australia Thailand Research Collaboration 001.4 Study. *Journal of Acquired Immune Deficiency Syndromes*, 39(5):523–529, 2005.
- [75] X S Qi, Q Yang, S P Lee, X A Li, and D Wang. An estimation of radiobiological parameters for head-and-neck cancer cells and the clinical implications. *Cancers*, 4:566–580, 2012.
- [76] D J Brenner and Eric J Hall. Fractionation and protraction for radiotherapy of prostate carcinoma. *International Journal of Radiation Oncology Biology Physics*, 43(5):1095–1101, 1999.

- [77] David J. Brenner. Fractionation and late rectal toxicity. *International Journal of Radiation Oncology Biology Physics*, 60(4):1013–1015, 2004.
- [78] M Gao, N A Mayr, Z Huang, H Zhang, and J Z Wang. When tumor repopulation starts? the onset time of prostate cancer during radiation therapy. *Acta Oncologica*, 49(8):1269–1275, 2010.
- [79] Karin M.G. Haustermans, Ingrid Hofland, Hein Van Poppel, Raymond Oyen, Wim Van de Voorde, Adrian C. Begg, and Jack F. Fowler. Cell kinetic measurements in prostate cancer. *International Journal of Radiation Oncology Biology Physics*, 37(5):1067–1070, 1997.
- [80] J Unkelbach, C Zeng, and M Engelsman. Simultaneous optimization of dose distributions and fractionation schemes in particle radiotherapy. *Medical Physics*, 40(9):091702, 2013.
- [81] B Choi and J O Deasy. The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning. *Physics in Medicine and Biology*, 47(20):3579–3589, 2002.
- [82] A Niemierko. Reporting and analyzing dose distributions: a concept of equivalent uniform dose. *Medical Physics*, 24(1):103–110, 1997.
- [83] Sebastiaan Breedveld, Pascal R M Storchi, Marleen Keijzer, and Ben J M Heijmen. Fast, multiple optimizations of quadratic dose objective functions in imrt. *Physics in Medicine and Biology*, 51(14):3569–3579, 2006.
- [84] S Webb, D J Convery, and P M Evans. Inverse planning with constraints to generate smoothed intensity-modulated beams. *Physics in Medicine and Biology*, 43(10):2785–2794, 1998.
- [85] X Gu, D Choi, C Men, H Pan, A Majumdar, and S B Jiang. GPU-based ultrafast dose calculation using a finite size pencil beam model. *Physics in Medicine and Biology*, 54:6287–6297, 2009.

- [86] C Men, X Gu, D Choi, A Majumdar, Z Zheng, K Mueller, and S B Jiang. GPU-based ultrafast IMRT plan optimization. *Physics in Medicine and Biology*, 54:6565–6573, 2009.
- [87] H E Romeijn and J F Dempsey. Intensity modulated radiation therapy treatment plan optimization. *TOP*, 16(2):215–243, 2008.
- [88] E Salari and H E Romeijn. Quantifying the trade-off between IMRT treatment plan quality and delivery efficiency using direct aperture optimization. *INFORMS Journal on Computing*, 24(4):518–533, 2012.
- [89] J F Fowler. Non-standard fractionation in radiotherapy. *International Journal of Radiation Oncology Biology Physics*, 10(5):755–759, 1984.
- [90] M Joiner and A van der Kogel. *Basic Clinical Radiobiology*. Hodder Arnold, London, UK, fourth edition, 2009.
- [91] H D Thames, S M Bentzen, I Turesson, M Overgaard, and W van den Bogaert. Fractionation parameters for human tissues and tumors. *International Journal Radiation Biology*, 56(5):701–710, 1989.
- [92] W M Harriss. *Monte Carlo modelling of tumor growth, hypoxia and radiotherapy in head and neck squamous cell carcinoma*. PhD thesis, University of Adelaide, Australia, 2011.
- [93] W M Harriss-Phillips, E Bezak, E K Yeoh, and F Ranzcr. Monte carlo radiotherapy simulations of accelerated repopulation and reoxygenation for hypoxic head and neck cancer. *The British Journal of Radiology*, 84(1006):903–918, 2011.
- [94] Eirik Malinen, Åste Søvik, Dimitre Hristov, Øyvind S Bruland, and Dag Rune Olsen. Adapting radiotherapy to hypoxic tumours. *Physics in Medicine and Biology*, 51(19):4903, 2006.

- [95] M Nordmark, S M Bentzen, V Rudat, D Brizel, E Lartigau, P Stadler, Axel Becker, M Adam, M Molls, J Dunst, D J Terris, and J Overgaard. Prognostic value of tumor oxygenation in 397 head and neck tumors after primary radiation therapy. an international multi-center study. *Radiotherapy and Oncology*, 77(1):18–24, 2005.
- [96] M Nordmark, M Overgaard, and J Overgaard. Pretreatment oxygenation predicts radiation response in advanced squamous cell carcinoma of the head and neck. *Radiotherapy and Oncology*, 41(1):31–39, 1996.
- [97] A Sovik, E Malinen, O S Bruland, S M Bentzen, and D R Olsen. Optimization of tumor control probability in hypoxic tumors by radiation dose redistribution: a modeling study. *Physics in Medicine and Biology*, 52(2):499–513, 2007.
- [98] B G Wouters and J M Brown. Cells at intermediate oxygen levels can be more important than the “hypoxic fraction” in determining tumor response to fractionated radiotherapy. *Radiation Research*, 147:541–550, 1997.
- [99] H Lyng, G Tanum, J F Evensen, and E K Rofstad. Changes in oxygen tension during radiotherapy of head and neck tumours. *Acta Oncologica*, 38(8):1037–1042, 1999.
- [100] M Nordmark, S Bentzen, and J Overgaard. Measurement of human tumour oxygenation status by a polarographic needle electrode. an analysis of inter- and intratumour heterogeneity. *Acta Oncologica*, 33(4):383–389, 1994.
- [101] M F Adam, E C Gabalski, D A Bloch, J W Oehlert, J M Brown, A A Elsaid, H A Pinto, and D J Terris. Tissue oxygen distribution in head and neck cancer patients. *Head and Neck*, 21(2):146–153, 1999.
- [102] J S Rasey, W J Koh, M L Evans, L M Peterson, T K Lewellen, M M Graham, and K A Krohn. Quantifying regional hypoxia in human tumors with positron emission tomography of [18f]fluoromisonidazole: A pretherapy study

- of 37 patients. *International Journal of Radiation Oncology Biology Physics*, 36(2):417–428, 1996.
- [103] M Alber, F Paulsen, S M Eschmann, and H J Machulla. On biologically conformal boost dose optimization. *Physics in Medicine and Biology*, 48(2):N31–N35, 2003.
- [104] Y Yang and L Xing. Towards biologically conformal radiation therapy (BCRT): selective IMRT dose escalation under the guidance of spatial biology distribution. *Medical Physics*, 32(6):1473–1484, 2005.
- [105] S S Foo, D F Abbott, N Lawrentschuk, and A M Scott. Functional imaging of intratumoral hypoxia. *Molecular Imaging and Biology*, 6(5):291–305, 2004.
- [106] E Lartigau, A Lusinchi, P Weeger, P Wibault, B Luboinski, F Eschwege, and M Guichard. Variations in tumor oxygen tension during accelerated radiotherapy of head and neck carcinoma. *European Journal of Cancer*, 34(6):856–861, 1998.
- [107] P Vaupel, K Schlenger, and C Knoop M Hockel. Oxygenation of human tumors: evaluation of tissue oxygenation distribution in breast cancers by computerized O_2 tension measurements. *Cancer Research*, 51(12):3316–3322, 1991.
- [108] M. R. Horsman, L. S. Mortensen, J. B. Petersen, M. Busk, and J. Overgaard. Imaging hypoxia to improve radiotherapy outcome. *Nature Reviews Clinical Oncology*, 9(12):674–687, 2012.
- [109] G Komar, M Seppanen, O Eskola, P Lindholm, T J Gronroos, S Forsback, H Sipila, S M Evans, O Solin, and H Minn. 18F-EF5: A new PET tracer for imaging hypoxia in head and neck cancer. *Journal of Nuclear Medicine*, 49 (12), 2008.
- [110] H Minn, T J Gronroos, G Komar, O Eskola, K Lehtia, J Tuomela, M Seppanen,

- and O Solin. Imaging of tumor hypoxia to predict treatment sensitivity. *Current Pharmaceutical Design*, 14 (28), 2008.
- [111] J G Rajendran and D A Mankoff. “Positron Emission Tomography imaging of blood flow and hypoxia in tumors”, in *In vivo imaging of cancer therapy*, chapter 4, pages 47–72. Humana Press, first edition, 2007.
- [112] X Sun, G Niu, N Chan, B Shen, and X Chen. Tumor hypoxia imaging. *Molecular Imaging and Biology*, 13(3):399–410, 2011.
- [113] W J Koh, J S Rasey, M L Evans, J R Grierson, T K Lewellen, M M Graham, K A Krohn, and T W Griffin. Imaging of hypoxia in human tumors with [f-18]fluoromisonidazole. *International Journal of Radiation Oncology Biology Physics*, 22(1):199–212, 1992.
- [114] A Dasu. Treatment planning optimisation based on imaging tumour proliferation and cell density. *Acta Oncologica*, 47(7):1221–1228, 2008.
- [115] S M Zhou, T Z Wong, and L B Marks. Using FDG-PET activity as a surrogate for tumor cell density and its effect on equivalent uniform dose calculation. *Medical Physics*, 31(9):2577–2583, 2004.
- [116] D Thorwarth and M Alber. Implementation of hypoxia imaging into treatment planning and delivery. *Radiotherapy and Oncology*, 97(2):172–175, 2010.
- [117] D Thorwarth, S-M Eschmann, F Paulsen, and M Alber. Hypoxia dose painting by numbers: A planning study. *International Journal of Radiation Oncology Biology Physics*, 68(1):291–300, 2007.
- [118] I Toma-Dasu, A Dasu, and A Brahme. Dose prescription and optimisation based on tumour hypoxia. *Acta Oncologica*, 48, 2009.
- [119] W Tuckwell, E Bezak, E Yeoh, and L Marcu. Efficient Monte Carlo modelling of individual tumour cell propagation for hypoxic head and neck cancer. *Physics in Medicine and Biology*, 53(17):4489–4507, 2008.

- [120] D P Bertsekas. *Dynamic Programming and Optimal Control*, volume 1 and 2. Athena Scientific, Nashua, NH, third edition, 2007.
- [121] W Powell. *Approximate Dynamic Programming*. John Wiley and Sons, New Jersey, 2007.
- [122] N Cressie and C K Wikle. *Statistics for spatiotemporal data*. Wiley, Hoboken, New Jersey, USA, 2011.
- [123] Y R Lawrence, M Werner-Wasik, and A P Dicker. Biologically conformal treatment: biomarkers and functional imaging in radiation oncology. *Future Oncology*, 4(5):689–704, 2008.
- [124] W B Powell and P Frazier. Optimal learning. In *Tutorials in Operations Research*, pages 213–245. Institute for Operations Research and the Management Sciences, 2008.