

Optimization Problems in Radiation Therapy Treatment Planning

by

Troy C. Long

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Industrial and Operations Engineering)
in The University of Michigan
2015

Doctoral Committee:

Professor H. Edwin Romeijn, Co-Chair
Associate Professor Marina A. Epelman, Co-Chair
Professor Jon Lee
Assistant Professor Martha M. Matuszak
Professor Randall K. Ten Haken

© Troy C. Long 2015
All Rights Reserved

For Brady, Nancy, and Larry

ACKNOWLEDGEMENTS

While working with Prof. Scott Mason and Prof. Chase Rainwater at Arkansas, they mentioned to me, “you should talk with Prof. Edwin Romeijn up at Michigan... you two might get along.” After meeting with him at INFORMS over orange juice and coffee, talking about building an optimization toolbox and radiation therapy, working with him at Michigan became my new goal. I’m thankful to my U of A advisors for guiding me towards such a fulfilling experience here at Michigan.

Five years later, my dissertation has come together. This would not have happened without the guidance of Prof. Edwin Romeijn, Prof. Marina Epelman, and Prof. Martha Matuszak. They were the Tripitaka to my Monkey-like levels of chaotic research fervor (see *Waley*, 1961). I’m also grateful for the support of Prof. Randy Ten Haken, Prof. Steve Jiang, and Prof. Thomas Bortfeld for supporting my collaborative research efforts during my studies.

In the IOE building, I’d like to give a heartfelt shout-out to Chris and Tina for always patiently helping me and attending to my onslaught of questions and requests. I had warned them that I was a bit of an administrative nightmare when I arrived, and they handled that realization with poise and skill.

My ~~comrades~~ fellow students will always hold a special place in my experience here, and helped me through the times when light at the end of the PhD tunnel seemed most dim. Special thanks to Victor, Ilbin, Jivan, Zohar, and Gregg for putting up with my distractions and always offering their support.

Lastly, I’d like to thank my family, Brady, Larry, and Nancy Long for, well, everything.

TABLE OF CONTENTS

DEDICATION	ii
ACKNOWLEDGEMENTS	iii
LIST OF FIGURES	vii
LIST OF TABLES	ix
LIST OF APPENDICES	x
LIST OF ABBREVIATIONS	xi
ABSTRACT	xiii
 CHAPTER	
I. Introduction	1
1.1 Radiation Therapy Treatment Planning	1
1.2 Traditional Treatment Planning Models	3
1.2.1 Fluence Map Optimization	3
1.2.2 Direct Aperture Optimization	4
1.2.3 Modeling Criteria	5
1.3 Chapter Summaries	5
1.3.1 Sensitivity Analysis in Lexicographic Ordering, Chapter II	5
1.3.2 Beam Orientation Optimization, Chapter III	6
1.3.3 Non-coplanar Volumetric Modulated Arc Therapy, Chapter	
IV	6
1.3.4 Adaptive Treatment Planning for Lung Cancer, Chapter V	7
 II. Sensitivity Analysis in Lexicographic Ordering for Radiation Ther-	
apy Treatment Planning	8
2.1 Introduction	8
2.2 Model and Methodology	11
2.2.1 Bi-criteria Treatment Planning Model	11

2.2.2	Sensitivity Analysis in Lexicographic Optimization Methodology	12
2.2.3	Treatment Plan Evaluation Criteria	17
2.3	Illustration of SALO on Clinical Cases	18
2.3.1	Data and computations	18
2.3.2	Results and Discussion	19
2.3.3	SALO Methodology Discussion	23
2.4	SALO Extensions and Improvements	26
2.4.1	Pros and Cons of SALO	26
2.4.2	SALO with Automation	28
2.4.3	SALOA Model	28
2.4.4	Solution Methodology	30
2.4.5	SALOA Results	34
2.5	Conclusions	34
 III. Beam Orientation Optimization		36
3.1	Introduction	36
3.2	Treatment Planning Models	37
3.2.1	Traditional FMO Model	38
3.2.2	Integrated BOO and FMO Model	38
3.3	Solution Methodology for $P_{\text{FMO}}^{\text{BOO}}$	39
3.3.1	KKT Motivation	40
3.3.2	Steepest Descent Motivation	43
3.3.3	Second-order Motivation	44
3.4	Application	46
3.4.1	Assessing the BOO Methodology	46
3.4.2	BOO Results	48
3.4.3	Methodology Discussion	51
3.5	Bounding on the Optimal Objective Value	55
3.5.1	Bounding $\mathcal{F}(\mathcal{B})$ Using an FMO Model	55
3.5.2	Bounding $\mathcal{F}(\mathcal{B})$ Using a DAO Model	59
3.5.3	Equivalence of Pricing Problem Constraints	66
3.5.4	Considerations and Potential Bounding Uses	67
3.6	Clinical Applicability and Future Research	68
3.6.1	Potential Clinical Implementation	68
3.6.2	Future Work	69
3.7	Acknowledgements	69
 IV. Non-coplanar Volumetric Modulated Arc Therapy		70
4.1	Introduction	70
4.2	Treatment Planning Models	71
4.2.1	Traditional VMAT Model	72
4.2.2	Non-coplanar VMAT Full Problem	72

4.2.3	Master Problem	74
4.3	Heuristic Methodology for NCP VMAT	74
4.3.1	Restricted Master Problem	75
4.3.2	Pricing Problem	75
4.3.3	Solving the PP	79
4.4	Application	81
4.4.1	Brain Case Details	81
4.4.2	Model Assumptions	81
4.4.3	Results	82
4.5	Conclusions and Future Work	85
4.6	Acknowledgements	87
V.	Adaptive Treatment Planning for Lung Cancer	88
5.1	Introduction	88
5.1.1	Adaptation Motivation	88
5.2	Lung Cancer Setting	89
5.2.1	Local Tumor Control	90
5.2.2	Radiation Induced Lung Toxicity	90
5.2.3	Bounding P_{RILT} in Pre-Treatment Planning Models	92
5.3	Planning Strategies and Models	94
5.3.1	Non-adaptive Treatment Planning Models	94
5.3.2	Framework for Adaptive Treatment Planning Models	96
5.3.3	Stochastic Adaptive Treatment Planning Models	97
5.3.4	Approximate Adaptive Treatment Planning Model	100
5.4	Results and Discussion	102
5.4.1	Discretizing the Biomarker Ratio Distribution	103
5.4.2	Sanity Check: Flexibility in stage-one dose distribution	104
5.4.3	Analysis for a Two-scenario Patient Population	104
5.4.4	The Impact of Treatment Adaptation to Exact Biomarker Value, and of Adaptation Timing	106
5.4.5	When to Use These Techniques	110
5.5	Conclusions and Future Research	112
5.6	Acknowledgements	113
	APPENDICES	114
	BIBLIOGRAPHY	124

LIST OF FIGURES

Figure

1.1	Illustration of the stages a patient goes through from diagnosis to treatment	2
1.2	(left) Sliding leaves in a multi-leaf collimator, (right) VMAT treatment delivery machine, gantry ghosting through delivery arc (photo credits to Varian Medical Systems)	2
2.1	Example upper bounds generated by the sandwich algorithm	16
2.2	SALO progression for the brain case, scenario A: (a) stage 1, (b) stage 2, (c) stage 3, and (d) stage 4	22
2.3	SALO progression for the brain case, scenario B: (a) stage 1, (b) stage 2, (c) stage 3, and (d) stage 4	23
2.4	Brain case: DVHs for treatment plan for (a) scenario A versus LO and (b) scenario A versus scenario B	24
2.5	SALO progression for the prostate case, scenario A: (a) stage 1, (b) stage 2, (c) stage 3, and (d) stage 4	25
2.6	SALO progression for the prostate case, scenario B: (a) stage 1, (b) stage 2, (c) stage 3, and (d) stage 4	26
2.7	Prostate case: DVHs for treatment plan for (a) scenario A versus LO and (b) scenario A versus scenario B	27
2.8	Critical region of clinical relevance for G_1 and G_2	30
2.9	Example endpoint locations and possible tradeoff curves for major scenario (I) with I_i : Red Square I_{ii} : Blue Diamond, I_{iii} : Green Octagon	32
2.10	Example endpoint locations and possible tradeoff curves for major scenario (II) with II_i : Red Square, II_{ii} : Blue Diamond	33
2.11	Feasible region for the convex tradeoff curve for minor scenario II_i given endpoint locations	33
3.1	Predictive results for steepest descent and second-order methodologies scaled to explicit $\Delta\mathcal{F}(\mathcal{B}', b)$ calculation for $ \mathcal{B}' = 2$ on a coplanar head-and-neck case case	49
3.2	Predictive results for steepest descent and second-order methodologies scaled to explicit $\Delta\mathcal{F}(\mathcal{B}', b)$ calculation for $ \mathcal{B}' = 4$ on a coplanar prostate case case	50
3.3	Objective function value vs. added number of beams for a coplanar prostate case with 60 beams	50
3.4	Objective function value vs. added number of beams for a brain case with 60 coplanar beams and 394 non-coplanar beams	51

3.5	Objective function value vs. added number of beams for a liver case with 56 non-coplanar beams	51
3.6	Dosewash of 7 equispaced beams	52
3.7	Dosewash of 7 beams selected using the second order method	52
3.8	DVH for the patient from Figures 3.6 and 3.7	53
3.9	Computation time for calculating beam score, $ \mathcal{B} = 60$	54
4.1	Illustration of $g = 2$ and $c = 1$ values between two red control points	78
4.2	36 control point NCP VMAT (solid) vs. coplanar VMAT (dashed)	83
4.3	36 control point NCP VMAT (dashed) vs. IMRT (solid)	83
4.4	Illustration of potential control points (blue), 36-point NCP VMAT trajectory (red), and interpolated 180-point NCP VMAT trajectory (green)	84
4.5	180 control point VMAT along 36 control point NCP VMAT path (solid) vs. 36 control point NCP VMAT (dashed)	85
4.6	180 control point VMAT along 36 control point NCP VMAT path (dashed) vs. 180 control point IMRT (solid)	85
4.7	180 control point VMAT along 36 control point NCP VMAT path (solid) vs. 180 control point coplanar VMAT (dashed)	86
4.8	(top) 180 control point coplanar VMAT dosewash, (bottom) 180 control point VMAT along 36 control point NCP VMAT path dosewash	87
5.1	Population based dependence of P_{RILT} on mean lung dose (or expected P_{RILT}).	91
5.2	P_{RILT} as a function of mean lung dose for the population (expected P_{RILT}) and for sensitive and resistant patients.	92
5.3	Sorted differences in first-stage dose distribution vectors (a) $z_{\text{rob}}^{\text{NA}} - z_{\text{rob}}^{\text{AA}}$ lung, (b) $z_{\text{rob}}^{\text{NA}} - z_{\text{rob}}^{\text{AA}}$ PTV, (c) $z_{\text{rob}}^{\text{NA}} - z_{\text{rob}}^{\text{SA}}$ lung, and (d) $z_{\text{rob}}^{\text{NA}} - z_{\text{rob}}^{\text{SA}}$ PTV.	105
5.4	Feasible region with respect to mean lung dose for the second stage dose distribution for the different methods (WC Bound = “rob” strategies)	106
5.5	P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 2, NA - pop and rob	107
5.6	P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 2, SA - pop	107
5.7	P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 2, SA - rob	108
5.8	P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 4, NA - pop and rob	109
5.9	P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 4, SA - pop	109
5.10	P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 4, SA - rob	110
5.11	(top) Treatment outcomes for non-adaptive plans, and plans with limited (scenario-based only) adaptation, vs. patient’s biomerker value; Realized P_{RILT} (top) and P_{LTC} (bottom) realized.	111
5.12	For $w = 0.33$, (left) realized P_{RILT} vs. biomarker ratio realization for re-optimized plans, (right) realized P_{LTC} vs. biomarker ratio realization for reoptimized plans; all models on nearly the same lines	112
5.13	P_{RILT} vs. P_{LTC} for $w =$ (a) 0.33, (b) 0.67, (c) 0.90, and (d) 0.95	112

LIST OF TABLES

Table

2.1	Structures with corresponding number of voxels and gEUD-parameters for the two clinical test cases	20
2.2	Number of solver calls and physician input queries for SALO and SALOA	34
4.1	Parameters tested over for the initial set of runs	82
5.1	Parameter estimates for calculating the probability of local tumor control .	90
5.2	Parameter estimates for P_{RILT}	91
5.3	The sizes of clinical cases: number of voxels and number of beamlets. The clinical number is the identifier for the case in UMPlan, the University of Michigan’s in-house treatment planning system, included here for figure references and in case data needs to be retrieved.	102
5.4	The number of variables, constraints, and non-zeros as well as run time ranges in minutes for case 3.	102
5.5	Scenario-specific biomarker ratio realizations for different numbers of quantiles	104
5.6	Realized P_{RILT} and P_{LTC} for $\alpha = 0.85$	108

LIST OF APPENDICES

Appendix

A.	Extending the Feasible Aperature Region by Restricting Gantry Travel Time	115
B.	Derivations for P_{RILT} and P_{LTC}	117
C.	Full Linear Models for Stochastic Adaptive Lung Formulations	121

LIST OF ABBREVIATIONS

AA	approximate adaptive
BEV	beams-eye-view
BLAS	basic linear algebra subprograms
BOO	beam orientation optimization
CERR	Computational Environment for Radiotherapy Research
CORT	common optimization for radiation therapy
CVaR	conditional value-at-risk
DAO	direct aperture optimization
DVH	dose-volume histogram
EUD	equivalent uniform dose
FMO	fluence map optimization
gEUD	generalized equivalent uniform dose
GPU	graphical processing unit
IMRT	intensity modulated radiation therapy
KKT	Karush–Kuhn–Tucker
LO	lexicographic optimization
LP	linear program
LS	leaf sequencing
MCO	multi-criteria optimization
MLC	multileaf collimator

NA non-adaptive
NCP non-coplanar
NTCP normal tissue complication probability
OAR organ-at-risk
pBEV pseudo-beams-eye-view
PP pricing problem
PTV planning target volume
RMP restricted master problem
RILT radiation-induced lung toxicity
SA stochastic adaptive
SALO sensitivity analysis in lexicographic ordering
SALOA sensitivity analysis in lexicographic ordering with automation
TCP tumor control probability
VMAT volumetric modulated arc therapy

ABSTRACT

Optimization Problems in Radiation Therapy Treatment Planning

by

Troy Long

Chairs: Edwin Romeijn, Marina Epelman

Radiation therapy is one of the most common methods used to treat many types of cancer. External beam radiation therapy and the models associated with developing a treatment plan for a patient are studied. External beams of radiation are used to deliver a highly complex so-called dose distribution to a patient that is designed to kill the cancer cells while sparing healthy organs and normal tissue. Treatment planning models and optimization are used to determine the delivery machine instructions necessary to produce a desirable dose distribution. These instructions make up a treatment plan. This thesis studies four problems in radiation therapy treatment plan optimization.

First, treatment planners generate a plan with a number of competing treatment plan criteria. The relationship between criteria is not known a priori. A methodology is developed for physicians and treatment planners to efficiently navigate a clinically relevant region of the Pareto frontier generated by trading off these different criteria in an informed way. Second, the machine instructions for intensity modulated radiation therapy, a common treatment modality, consist of the locations of the external beams and the non-uniform intensity profiles delivered from each of these locations. These decisions are traditionally made with separate, sequential models. These decisions are integrated into a single model and propose a heuristic solution methodology. Third, volumetric modulated arc therapy (VMAT), a treatment modality where the beam travels in a coplanar arc around the patient while continuously delivering radiation, is a popular topic among optimizers studying treatment planning due to the difficult nature of the problem and the lack of a universally accepted treatment planning method. While current solution methodologies assume a pre-determined coplanar path around the patient, that assumption is relaxed and the generation

of a non-coplanar path is integrated into a VMAT planning algorithm. Fourth, not all patient information is available when developing a treatment plan pre-treatment. Some information, like a patient's sensitivity to radiation, can be realized during treatment through physiological tests. Methodologies of pre-treatment planning considering adaptation to new information are studied.

CHAPTER I

Introduction

1.1 Radiation Therapy Treatment Planning

Radiation therapy is one of the most common methods used to treat many different types of cancer. There are different techniques for exposing cancerous tissue to radiation, and in these studies we will be investigating external beam radiation therapy. With this technique, external beams of radiation are used to deliver a highly complex so-called dose distribution to a patient that is designed to kill the cancer cells while sparing healthy organs and normal tissue. Treatment planning models and optimization are used to determine the delivery machine instructions necessary to produce a desirable dose distribution. This thesis will focus on models that address some of the shortcomings with current treatment planning and offer clinically applicable techniques for solving these problems.

After a patient is diagnosed with cancer, a physician will determine the patient's course of treatment (see Figure 1.1). If radiation therapy is included in this treatment, a patient will have detailed imaging performed to allow treatment planners to construct a 3D model of the patient within a treatment planning system, a software package that assists treatment planners in determining machine instructions for delivery. After imaging, a physician will manually contour each critical structure and target within the patient. The system then uses this geometric information to simulate radiation delivered to the patient to estimate the effects of delivering radiation from a certain angle at a certain intensity. This information will be used to relate machine instructions to dose received by the patient. After the machine instructions and corresponding dose distribution are decided upon by the physician and the treatment planner, the patient is treated. Patients receive radiation over a number of days, each session delivering a portion of the total treatment. These sessions are referred to as fractions. During these fractions, new patient information may be observed and prompt adapting the plan to the new information. After the patient completes a full course of treatment fractions, the patient is imaged and the diagnosis is reassessed.




Figure 1.1: Illustration of the stages a patient goes through from diagnosis to treatment

The main treatment modalities we consider are intensity modulated radiation therapy (IMRT) and volumetric modulated arc therapy (VMAT). In both modalities, radiation is delivered through the rectangular opening on the gantry and through an multileaf collimator (MLC) (see Figure 1.2) that uses sliding tungsten leaves to block radiation and form shapes called apertures. IMRT utilizes a fixed number of stationary beam locations to deliver radiation through multiple apertures of varying intensity at each location of varying intensities. The primary decisions that are needed to define an IMRT treatment plan are the beam directions, aperture shapes, and aperture intensities. VMAT, a more advanced and modern modality, employs a beam that rotates in an arc continuously with dynamically changing apertures and intensity (see Figure 1.2). The primary decisions that are needed to define a VMAT treatment plan are the MLC leaf positions along the arc, the dose rate, and the speed of the gantry.




Figure 1.2: (left) Sliding leaves in a multi-leaf collimator, (right) VMAT treatment delivery machine, gantry ghosting through delivery arc (photo credits to Varian Medical Systems)

1.2 Traditional Treatment Planning Models

We present the general structure of treatment plan optimization models for IMRT without the project-specific minutiae. The VMAT treatment planning model only applies to chapter IV and thus will not be included here. Clinically motivated changes to these models as well as corresponding solution methodologies will be explored in the following chapters.

As stated earlier, the decisions to be made in IMRT are the beam directions, aperture shapes, and aperture intensities. Of these decisions, let us assume that the beam directions are selected and fixed through some process (see chapter III). The remaining decisions are the aperture shapes and aperture intensities. There are two main ways to model and solve for these values: fluence map optimization (FMO) followed by leaf sequencing (LS) and direct aperture optimization (DAO). FMO generates a non-uniform intensity profile for each beam and the post-processing LS step generates the apertures and corresponding intensities necessary to deliver a plan that produces a dose distribution close to that generated by the original FMO intensity profile. DAO is an iterative procedure that actively generates deliverable apertures to be used in the treatment planning model.

There is a core set of notation that will be used throughout this thesis. We present that notation here, and any deviations or additions will be updated within the associated chapters. Common to all presented treatment planning models, the patient geometry is discretized into small volumes called voxels represented by set \mathcal{V} . Let z_j be the dose received by voxel $j \in \mathcal{V}$. Let \mathcal{B} be the set of potential discrete beam locations around the body.

1.2.1 Fluence Map Optimization

In FMO, we need to define the sets of beam locations and beamlets. Let \mathcal{B}' be the set of active (i.e., used in the treatment plan) beam locations. These locations can be in a coplanar arc around the patient or in a non-coplanar arrangement. Non-coplanar arrangements are often referred to as the 4π space around a patient. In order to represent the FMO intensity profile, each beam $b \in \mathcal{B}'$ is discretized into a grid of beamlets \mathcal{N}_b . Let decision variable x_{bi} be the intensity for beamlet $i \in \mathcal{N}_b$ for beam $b \in \mathcal{B}'$. Sometimes beamlet intensity is presented as x_i for $i \in \mathcal{N}$ when beam-specificity is not relevant to the problem (i.e., $i \in \mathcal{N} = \cup_{b \in \mathcal{B}'} \mathcal{N}_b$). To relate the beamlet intensities to the voxel doses, let D_{bij} be the dose delivered to voxel $j \in \mathcal{V}$ from beamlet $i \in \mathcal{N}_b$ in beam $b \in \mathcal{B}'$ at unit intensity. With objective function $F(z)$, representing treatment plan evaluation criteria as a function of dose, and set of deliverable treatment plans Z , we can construct the basic FMO model as follows:

$$\underset{x,z}{\text{minimize}} \quad F(z) \tag{1.1}$$

$$\text{subject to} \quad z_j = \sum_{b \in \mathcal{B}'} \sum_{i \in \mathcal{N}_b} D_{bij} x_{bi} \quad \forall j \in \mathcal{V} \tag{1.2}$$

$$z \in Z \tag{1.3}$$

$$x_{bi} \geq 0 \quad \forall b \in \mathcal{B}', \forall i \in \mathcal{N}_b. \tag{1.4}$$

LS occurs after FMO to generate apertures and aperture intensities. For the purposes of this thesis, LS is assumed to occur without incident and thus is not considered when modeling using FMO (see *Xia and Verhey, 1998* for more information on LS).

1.2.2 Direct Aperture Optimization

In DAO, the variables of aperture intensities are directly represented in the model. Let \mathcal{A}_b be the set of deliverable apertures on beam $b \in \mathcal{B}'$ and let y_{bk} be the intensity of aperture $k \in \mathcal{A}_b$ on beam $b \in \mathcal{B}'$. To relate aperture intensity to dose received, let \mathcal{D}_{bkj} be the dose delivered to voxel $j \in \mathcal{V}$ from aperture $k \in \mathcal{A}_b$ at beam $b \in \mathcal{B}'$. The full treatment planning model looks as follows:

$$\underset{y,z}{\text{minimize}} \quad F(z) \tag{1.5}$$

$$\text{subject to} \quad z_j = \sum_{b \in \mathcal{B}'} \sum_{k \in \mathcal{A}_b} \mathcal{D}_{bkj} y_{bk} \quad \forall j \in \mathcal{V} \tag{1.6}$$

$$z \in Z \tag{1.7}$$

$$y_{bk} \geq 0 \quad \forall b \in \mathcal{B}', \forall k \in \mathcal{K}_b. \tag{1.8}$$

While the model structure seems similar to the FMO model, there are many more deliverable apertures than beamlets for each beam (i.e., $|\mathcal{N}_b| \ll |\mathcal{A}_b|$). Therefore, sophisticated solution methodologies are used to solve this problem (see, e.g., *Romeijn et al., 2005* for a column-generation approach). Sometimes the $b \in \mathcal{B}'$ subscript is dropped when beam-specificity is not needed. Due to the large number of \mathcal{D}_{kj} values, \mathcal{D}_{kj} values are constructed as needed using a beamlet approximation. Let $i \in \mathcal{N}_k$ be the set of beamlets that make up aperture $k \in \mathcal{A}$. We can then calculate \mathcal{D}_{kj} as the following for all $k \in \mathcal{A}$ and $j \in \mathcal{V}$:

$$\mathcal{D}_{kj} = \sum_{i \in \mathcal{N}_k} D_{ij}. \tag{1.9}$$

1.2.3 Modeling Criteria

As mentioned before, the main goals of radiation therapy treatment planning are to deliver a dose distribution that eradicates the cancerous tissue while sparing healthy tissue and functioning organs. In the treatment planning models in Section 1.2, the objective of the models is some function of dose, $F(z)$. How $F(z)$ is constructed, however, is not a well-established, clear procedure in the clinic. There are a multitude of treatment plan evaluation functions used to describe the effects of radiation in the body. There are functions that give probabilities of control or complications for tumors and non-cancerous tissue, respectively. Some examples include tumor control probability (TCP) and normal tissue complication probability (NTCP) (see, e.g., *Webb and Nahum, 1993*; *Warkentin et al., 2004*; *Marks et al., 2010*). Other criteria include voxel-based evaluation functions and dose-volume functions (see *Shepard et al., 1999* for an overview), metrics that convert non-uniform dose to a structure to the radiobiologically equivalent uniform dose (see equivalent uniform dose (EUD) in *Niemierko, 1999*), and proxy metrics for some of those listed earlier (e.g., conditional value-at-risk (CVaR) in *Romeijn et al., 2003* and linearized EUD in *Thieke et al., 2002*). Some are “easy” objectives (e.g., piecewise quadratic voxel-based penalties are convex) while others are extremely difficult to incorporate explicitly into the treatment planning model (e.g., dose-volume objectives are non-convex).

On top of it being unclear which functions to use in the objective, the relative importance of different objectives is also unclear a priori (see chapter II for more information). Multiple objectives will be given for planners to improve upon or satisfy, but how these competing objectives behave is not immediately clear. Traditionally, physicians and treatment planners work together in an iterative, trial-and-error fashion to select a combination of evaluation functions and relative importances to arrive at a desirable treatment plan.

1.3 Chapter Summaries

1.3.1 Sensitivity Analysis in Lexicographic Ordering, Chapter II

When presented with a radiation therapy case, physicians approach the treatment planning process with a number of goals for the treatment outcome of varying importance. For example, sparing the brainstem of lethal radiation will be much more critical than reducing hotspots in normal tissue surrounding the brainstem. In this process, the physicians develop a hierarchy of criteria they believe to be significant in determining a treatment plan. Two questions arise when considering different criteria of unclear relative importance. First, what tradeoff information between competing objectives in the model might prove beneficial to

treatment planners and physicians? Sometimes yielding a small amount on the most important treatment goals may provide significant benefits in overall treatment plan quality, and identifying this tradeoff can be non-trivial. Second, how can these tradeoffs be calculated efficiently and presented in a manner valuable to physicians? Because the calculations necessary to produce accurate representations of these relationships could be lengthy, techniques to quicken the process without losing accuracy can help maintain clinical feasibility. We address both of these questions in chapter II, as well as a several improvements that increase the efficiency of the initially developed methodology. To improve the methodology, we incorporate additional bounds on the relevant region of the generated Pareto frontier and automate some of the proposed interactive steps.

1.3.2 Beam Orientation Optimization, Chapter III

In IMRT beam orientation optimization, a small number of beam positions must be identified that allow for both high treatment plan quality and efficient deliverability, creating a large-scale combinatorial optimization problem. Traditionally, beam selection and FMO are approached as separate, sequential models. We integrate these models and develop efficient and effective methods for selecting high-quality coplanar or non-coplanar beam orientations for IMRT treatments. The proposed methods explicitly incorporate the effects of beam selection on the quality of the resulting optimal dose distribution. To this end, we propose a greedy heuristic framework for solving the integrated model. The algorithm iteratively adds beams to the model according to a dynamically updated attractiveness measure for each remaining candidate beam. We consider measures that are based explicitly on the optimal dose distribution corresponding to the currently selected set of beams. Several specific attractiveness measures are proposed that use either first-order or both first and second-order information. Performance of the algorithm is assessed on clinical data demonstrating a high quality dose distributions found with fewer beams than traditional methods. Theoretical work is also presented aimed at finding methods for bounding optimal solutions to the integrated beam orientation optimization and fluence map optimization treatment planning model.

1.3.3 Non-coplanar Volumetric Modulated Arc Therapy, Chapter IV

VMAT is rapidly emerging as a method for delivering radiation therapy treatments to cancer patients that are of comparable quality to IMRT but much more efficient. In VMAT, the beam moves along a coplanar arc (i.e., fixed couch position with a moving gantry) while delivering radiation through dynamically changing apertures. The decisions here are the

(usually fixed) coplanar arc path, the movements of the MLC along the arc, the speed of the gantry, and the dynamic dose rate of the beam. When making these decisions, we must consider machine constraints such as MLC leaf movement restrictions, dose rate limits, and gantry acceleration bounds. The next step in taking advantage of the flexibility of the delivery machines is to consider non-coplanar paths in the 4π space around the patient as well. We propose a constructive approach that employs both column generation and routing heuristics to simultaneously determine the non-coplanar path of the beam, MLC movements, and dose rate. A proof-of-concept implementation is presented on a brain cancer case.

1.3.4 Adaptive Treatment Planning for Lung Cancer, Chapter V

When developing a personalized treatment, planners use a physician's treatment planning goals based on patient information obtained pre-treatment. However, over the course of treatment these patient characteristics may change, evolving the planner's initial understanding of patient characteristics and causing inaccuracies in the treatment plan. The pre-treatment plan is often reoptimized mid-treatment to these new data. For example, certain biomarker data obtained during treatment after several fractions of radiation have been delivered has been shown to be predictive of a patient's predisposition to radiation-induced lung toxicity. We develop several methodologies for pre-treatment planning when adaptation is likely to occur. A two-stage stochastic treatment plan optimization model to explicitly consider future patient-specific biomarker information is developed along with a less-computationally-intensive heuristic methodology. Recommendations for building flexibility for future adaptation into initially-delivered dose distributions are presented.

Research Collaborators Chapter-specific acknowledgement sections will identify research collaborators for each project.

Research Funding All projects partially supported by NSF GRFP between 2012 and 2015. Funding for projects in Chapters II and V partially from grant number NIH-P01-CA59827. Two Rackham Student Research Grants were awarded for partial support on Chapters III (2012) and IV (2014).

CHAPTER II

Sensitivity Analysis in Lexicographic Ordering for Radiation Therapy Treatment Planning

This chapter discusses a project we worked on in 2010-2012. This project allowed us to get into both the modeling side of radiation therapy treatment planning as well as the computational aspects of treatment plan optimization. Sections 2.1 to 2.3 of the following text is largely pulled from the paper published in *Medical Physics* (see *Long et al.*, 2012) along with some additional work done with a few undergraduate student researchers in Section 2.4.

2.1 Introduction

When addressing a radiation therapy case, a physician generally presents the treatment planner with a number of dosimetric goals of varying importance. Although the general objective is to deliver a prescribed radiation dose to the target(s) while simultaneously sparing critical structures, a major challenge remains how to make the unavoidable tradeoffs between these conflicting goals. The literature on radiation therapy treatment planning contains a multitude of evaluation criteria that can be used to quantify various properties of a treatment plan. Because treatment planning is generally a time-consuming endeavor which has to be performed for individual patients, providing a treatment planner with tools that allow for an efficient assessment of the interplay and tradeoffs between conflicting treatment plan evaluation criteria is essential to an efficient and effective treatment planning process.

Traditionally, radiation therapy treatment planning is based on optimization models containing a single objective function to be optimized subject to a set of hard constraints on the treatment plan. The objective function is typically a simple weighted sum of individual treatment plan evaluation criteria (see *Breedveld et al.*, 2009). Since there is no formal basis for choosing a priori values for these weights, their values are usually updated manually by the treatment planner in an iterative fashion in order to arrive at a clinically desirable

treatment plan. Occasionally this method yields acceptable results quickly, but in general this approach is inefficient and may lead to inferior treatment plans.

A modern technique for exploring the tradeoffs between treatment plan evaluation criteria is based on multi-criteria optimization (MCO) (see, e.g., *Küfer et al.*, 2009 for a recent overview of this area). In this approach, the goal is to approximate the Pareto frontier containing all efficient treatment plans, i.e., treatment plans with the property that it is not possible to improve the plan with respect to one of the criteria without deteriorating the plan with respect to at least one other. While there are many methods for generating this frontier (see, e.g., *Ruzika and Wiecek*, 2005), a common technique is to solve a sequence of single-objective optimization problems, each using an appropriately chosen set of weights for the individual criteria. When all criteria are convex functions of the dose distribution delivered to the patient, each of the corresponding solutions will represent a point on the Pareto frontier. If the number of such solutions is large enough to allow the Pareto frontier to be accurately approximated (typically using interpolation), the treatment planner can assess the tradeoffs between competing objectives by navigating the frontier and use this information to select a treatment plan. Using MCO as a means of quantifying tradeoffs is conceptually attractive, in the sense that it provides the treatment planner with complete and comprehensive tradeoff information on all criteria. However, the number of competing criteria can be large (say, on the order of 10–25 in a typical clinical setting), which means that the Pareto frontier is embedded in a correspondingly high-dimensional space. Many solutions may then be required to accurately approximate the Pareto frontier, which reduces the efficiency of the methodology (see *Hong et al.*, 2008). Moreover, visualizing and interpreting the plethora of tradeoffs can prove difficult (see *Craft and Monz*, 2010). Of course a reduction in the number of criteria or data reduction in the form of a coarser representation of patient geometry and/or capabilities of the delivery equipment may mitigate these drawbacks, but this may affect the accuracy of the frontier or the quality of the tradeoff information (see, e.g., *Craft and Bortfeld*, 2008).

A key observation is that the full Pareto frontier identified by MCO will likely contain many tradeoff regions that are clinically unacceptable or irrelevant. This not only complicates the navigation process as outlined above, it also means that a large amount of time may be spent identifying such uninteresting tradeoffs. It therefore seems appropriate to explicitly incorporate better a priori clinical information on priorities associated with the different criteria into the treatment plan optimization process. One such approach is lexicographic optimization (LO), which is sometimes also referred to as prioritized optimization (see *Wilkins et al.*, 2007; *Jee et al.*, 2007; *Clark et al.*, 2008). This is a multi-stage approach that is based on a complete ranking or prioritization of treatment planning goals. In its

purest form it starts by optimizing the highest ranked criterion. The optimal value to this problem is then used to constrain the value of the corresponding criterion in subsequent optimization models. In particular, in the following stage the second criterion on the prioritized list is optimized subject to the value of the first criterion being optimal. Following that, the third is optimized subject to the value of the first and second criteria’s solutions from the previous model. This approach is then continued for each criterion on the list, and the solution to the final optimization problem in the sequence is the optimal treatment plan with respect to the prioritized list of criteria. LO is computationally efficient and provides a clear, systematic approach. In contrast with MCO, LO does not rely on interaction with the treatment planner (once the prioritization is fixed). However, much flexibility is sacrificed in the wake of the computational and structural benefits. In particular, a notable drawback of using an LO approach is that the treatment planner may be unaware of opportunities that may exist to improve a treatment plan. In terms of MCO, the LO approach can be interpreted as confining the treatment planner’s view to a specific extreme solution on the full Pareto frontier of inter-criterion tradeoffs. If a minor sacrifice in high-priority criteria could yield meaningful benefits with respect to lower-priority criteria, the pure LO approach would not recognize or identify this opportunity. In order to introduce some flexibility into the process one might relax the optimality constraint on high-priority criteria and instead require previously optimized criteria to remain “near-optimal.” Since tradeoffs are not characterized and assessed explicitly, it is not clear how to quantify the concept of near-optimality, nor how to predict the consequences of allowing a deviation from optimality. In contrast, our method will provide an interactive way for the user to select the relaxation based on a formal sensitivity analysis.

In this chapter, we propose a systematic approach, *sensitivity analysis in lexicographic ordering (SALO)*, which combines the benefits of MCO (flexibility and comprehensiveness) and LO (efficiency and clinical focus) while avoiding their pitfalls. Similar to LO, it incorporates clinical information through a prioritized list of treatment plan evaluation criteria. However, in contrast with LO, it uses this information to efficiently navigate a clinically interesting and relevant segment of the Pareto efficient frontier in an interactive and iterative fashion. In Section 2.2 of this chapter we will provide a formal and detailed description of the SALO approach. In Section 2.3 we will then illustrate the approach on two clinical cases and discuss the performance of the algorithm. In Section 2.4 we will discuss some improvements made to the model and conclude in Section 2.5.

2.2 Model and Methodology

The goal of the SALO approach is to provide local information on the shape of the Pareto frontier to treatment planners for use as a decision making aid, based on clinical preferences represented via a prioritized list of treatment plan evaluation criteria. This local information takes the form of a two-dimensional Pareto frontier that, in each stage, characterizes the tradeoff between two consecutive criteria while (i) constraining higher priority criteria to values that have been established earlier in the process; and (ii) temporarily ignoring lower priority criteria. The treatment planner can then examine this tradeoff curve and select a point that appropriately captures the tradeoff between the two criteria currently under consideration. This point then defines a bound for the criterion that has the higher priority.

2.2.1 Bi-criteria Treatment Planning Model

Optimization models for radiation therapy treatment planning are usually classified as “beamlet-based” (yielding an optimal fluence map, which subsequently needs to be converted into a deliverable plan in a leaf-sequencing stage) and “aperture-based.” We have chosen to use the latter, DAO, approach (see, e.g., *Shepard et al.*, 2002; *Preciado-Walters et al.*, 2004; *Romeijn et al.*, 2005; *Men et al.*, 2007) since it not only eliminates the need for a leaf-sequencing stage but can also allow for a more efficient implementation and solution since an instance of the DAO model is typically much smaller and hence can be solved more rapidly than an instance of a beamlet-based FMO problem. This is particularly important since many of these problem instances will need to be solved during the course of the SALO procedure. However, if desired, the general SALO approach can be applied to a more traditional FMO model with only minor modifications.

The notation in this chapter is consistent with notation presented in 1.2.1. For simplicity, we will let $y = (y_k : k \in \mathcal{A})$ and $z = (z_j : j \in \mathcal{V})$ denote the corresponding vectors. Convexity of the set of feasible treatment plans, Z , is important for tractability of our approach, and we usually expect this set to contain only simple lower and upper bound constraints on the individual voxel doses. In principle other hard constraints on treatment plan evaluation criteria could be included as well, although we envision those tradeoffs to be made in the actual SALO procedure rather than by a priori excluding certain dose distributions.

The treatment plan evaluation criteria are given as functions of the dose distribution: $G_\ell : \mathbb{R}^{|\mathcal{V}|} \rightarrow \mathbb{R}$ ($\ell = 1, \dots, L + 1$), where we assume that the criteria are indexed in order of decreasing priority. For mathematical convenience we will assume that these criteria are such that smaller values are preferred to larger values. In addition, we will generally assume that they are all convex functions. Note that, in a multicriteria context, many

common treatment plan evaluation criteria, such as voxel-based penalty functions, EUD, generalized equivalent uniform dose (gEUD), TCP, NTCP, or CVaR, are either convex or can equivalently be replaced by convex ones (see *Romeijn et al.*, 2004). Our proposed approach could in principle be generalized to accommodate a nonconvex set Z and/or truly nonconvex criteria, such as traditional Dose-Volume Histogram (DVH) constraints, albeit at the expense of computational efficiency. The last criterion, G_{L+1} , is typically chosen in order to minimize total dose delivered to the patient while maintaining treatment plan quality with respect to all previously considered criteria. The SALO approach then interactively searches for a treatment plan by solving a sequence of bi-criteria optimization models of the following form (for $\ell = 1, \dots, L$), referred to as *stages* of the procedure:

$$\underset{(y,z)}{\text{minimize}} \{G_\ell(z), G_{\ell+1}(z)\}$$

subject to (P^(ℓ))

$$z_j = \sum_{k \in \mathcal{A}} \mathcal{D}_{kj} y_k \quad \text{for } j \in \mathcal{V} \quad (2.1)$$

$$G_{\ell'}(z) \leq \bar{G}_{\ell'} \quad \text{for } \ell' = 1, \dots, \ell - 1 \quad (2.2)$$

$$y_k \geq 0 \quad \text{for } k \in \mathcal{A} \quad (2.3)$$

$$z \in Z, \quad (2.4)$$

where $\bar{G}_{\ell'}$ is an upper bound on treatment plan evaluation criterion $G_{\ell'}$ that is set in the earlier stage ℓ' of the procedure by solving the prior bi-criteria optimization problem (P^(ℓ')) (for $\ell' = 1, \dots, \ell - 1$) and making the corresponding tradeoff.

Due to the convexity of the criterion functions, the solution to the bi-criteria optimization problem (P^(ℓ)) can be found by solving single-criterion optimization problems with an objective function of the form:

$$\alpha G_\ell(z) + (1 - \alpha) G_{\ell+1}(z), \quad (2.5)$$

for all $\alpha \in [0, 1]$.

2.2.2 Sensitivity Analysis in Lexicographic Optimization Methodology

If the set of all deliverable apertures were manageable we could directly apply the approach outlined above. Unfortunately, in general the cardinality of the set \mathcal{A} is very large and the optimization problems (P^(ℓ)) cannot be solved explicitly. One potential approach would be to generate high-quality apertures “on the fly” according to a column genera-

tion approach that has been proposed for solving single-criterion DAO problems (see, e.g., *Preciado-Walters et al., 2004; Romeijn et al., 2005; Men et al., 2007*). However, this would mean that the set of apertures considered in later stages of the algorithm is different from (in fact, larger than) the set of apertures allowed in earlier stages. This means that the tradeoffs between the higher priority, and hence clinically more important, criteria are based on a more limited set of apertures. Intuitively it would seem more attractive to base the more important (or, in fact, all) tradeoff decisions on the most accurate representation of the optimization model rather than the least accurate, which makes a straightforward application of this idea undesirable.

In order to address this issue we propose to start the SALO procedure with an initial phase in which a high-quality pool of apertures is generated, which is then kept fixed throughout the L stages of the SALO procedure. This does not only improve the computational efficiency of the approach, but also ensures that all decisions are made based on consistent input and information. However, it is clear that, in this process, the tradeoff decisions are not made with respect to the full information regarding all deliverable apertures. We therefore also propose a final phase in which a full DAO model is solved to identify a new set of apertures that minimizes the last criterion, G_{L+1} , subject to all bounds imposed on criteria G_1, \dots, G_L . Clearly, this final phase could also take other considerations, such as treatment delivery efficiency, into account. In summary, we propose a SALO procedure that proceeds in three phases:

Sensitivity Analysis in Lexicographic Optimization

Phase 1. Generation of a clinically relevant aperture pool of computationally manageable cardinality.

Phase 2. Generation of patient-specific treatment planning goals \bar{G}_ℓ ($\ell = 1, \dots, L$) by solving a sequence of bi-criteria optimization problems $(P^{(1)}), \dots, (P^{(L-1)})$.

Phase 3. Generation of final treatment plan that satisfies the patient-specific treatment planning goals while minimizing an overall single objective function.

In the remainder of this section we will discuss these three stages in more detail.

2.2.2.1 Aperture Pool Generation

We generate an aperture pool by solving a traditional single-criterion treatment plan optimization model based on the treatment plan evaluation criteria:

$$\underset{(y,z)}{\text{minimize}} \quad \sum_{\ell=1}^L \alpha_{\ell} G_{\ell}(z)$$

subject to (P)

$$z_j = \sum_{k \in \mathcal{A}} \mathcal{D}_{kj} y_k \quad \text{for } j \in \mathcal{V} \quad (2.6)$$

$$y_k \geq 0 \quad \text{for } k \in \mathcal{A} \quad (2.7)$$

$$z \in Z, \quad (2.8)$$

where $\alpha_{\ell} \geq 0$ ($\ell = 1, \dots, L$) are nonnegative criterion weights. The set of weights used in the aperture generation phase could be based on experience with other, similar, patient cases. Alternatively, we could use a sequence of criterion weights, allowing for the generation of apertures that are attractive with respect to a variety of tradeoffs. We will denote the set of apertures in the pool by $\overline{\mathcal{A}}$.

2.2.2.2 Solving the Bi-criterion Optimization Problem

We use the so-called Sandwich Algorithm (see, e.g., *Hoffmann et al.*, 2006; *Küfer et al.*, 2009; *Ehrgott et al.*, 2010) to approximate the Pareto frontier at a given stage of the SALO procedure. This algorithm, which tries to balance clinical accuracy and computational efficiency, applies when all treatment plan evaluation criteria are convex, and is particularly efficient in the bi-criteria case. The idea behind this algorithm is to approximate the entire Pareto frontier by constructing both an upper (“conservative”) and a lower (“optimistic”) bound on the frontier based on a finite set of points on the frontier. This is done by solving a sequence of optimization problems of the form (P^(ℓ)) with objective function of the form (2.5) for different values of α . The optimal solutions to these problems yield points on the Pareto frontier. For convenience, let $z^*(\alpha; \ell)$ denote an optimal solution to (P^(ℓ)) when parameter α is used. Then let

$$G_{\ell}^*(\alpha; \ell) = G_{\ell}(z^*(\alpha; \ell)) \quad (2.9)$$

$$G_{\ell+1}^*(\alpha; \ell) = G_{\ell+1}(z^*(\alpha; \ell)) \quad (2.10)$$

$$G_{\ell, \ell+1}^*(\alpha; \ell) = \alpha G_{\ell}^*(\alpha; \ell) + (1 - \alpha) G_{\ell+1}^*(\alpha; \ell). \quad (2.11)$$

The Sandwich Algorithm then determines upper and lower bounds on the Pareto frontier as follows:

Upper bound Using simple linear interpolation of a set of Pareto efficient solutions we obtain a piecewise-linear and convex function which is well-known to form an upper bound on the Pareto frontier. This follows immediately from the fact that the line segment connecting any two points of the form $(G_\ell^*(\alpha; \ell), G_{\ell+1}^*(\alpha; \ell))$ for different values of α is guaranteed to be entirely on or above the Pareto frontier. More formally, such a line segment can be characterized as

$$\{(\lambda G_\ell^*(\alpha'; \ell) + (1 - \lambda)G_\ell^*(\alpha''; \ell), \lambda G_{\ell+1}^*(\alpha'; \ell) + (1 - \lambda)G_{\ell+1}^*(\alpha''; \ell)) : \lambda \in [0, 1]\} \quad (2.12)$$

where $0 < \alpha' \neq \alpha'' < 1$. The sandwich algorithm process can be seen in Figure 2.1 as an example between the EUD to the bladder and the EUD to the rectum. Each iteration generates points in between existing points, effectively doubling the precision), and adds upper bounds on the curve using equation (2.12), shown as piecewise linear functions. The lower bounds are not shown for simplicity, but the true curve must pass through all generated points along the curve.

It is interesting to note that due to the convexity of G_ℓ^* and $G_{\ell+1}^*$ we can find an even better bound on the Pareto frontier by, instead of interpolating the pairs of optimal objective function values for different values of α , interpolating the optimal treatment plans (or, equivalently, optimal dose distributions) for different values of α . In other words, the curve of the form:

$$\{(G_\ell(\lambda z^*(\alpha; \ell) + (1 - \lambda)z^*(\alpha'; \ell)), G_{\ell+1}(\lambda z^*(\alpha; \ell) + (1 - \lambda)z^*(\alpha'; \ell))) : \lambda \in [0, 1]\}, \quad (2.13)$$

where $0 < \alpha \neq \alpha' < 1$ is guaranteed to not only be entirely on or above the Pareto frontier, but also entirely on or below the curve in equation (2.12). While the precise values of the upper bound do not directly factor into our algorithm, we recommend using equation (2.13) for upper bound estimation.

Lower bound A lower bound can be determined by observing that for some fixed α and associated $G_{\ell, \ell+1}^*(\alpha; \ell)$ a line given by the following:

$$\{(g_\ell, g_{\ell+1}) : \alpha g_\ell + (1 - \alpha)g_{\ell+1} = G_{\ell, \ell+1}^*(\alpha; \ell)\}, \quad (2.14)$$

is entirely on or below the Pareto frontier (where $(g_\ell, g_{\ell+1})$ denotes a point in \mathbb{R}^2). As α changes, lines satisfying this equation lie tangent to the true curve. This follows since (i) the Pareto frontier is convex by convexity of the criteria functions G_ℓ and $G_{\ell+1}$; and (ii) the

point $(G_{\ell}^*(\alpha; \ell), G_{\ell+1}^*(\alpha; \ell))$ lies both on the efficient frontier and the line (2.14). This means that the upper envelope of these lines over a collection of different values of $\alpha \in (0, 1)$ is a piecewise-linear convex function that is entirely on or below the Pareto frontier as well.

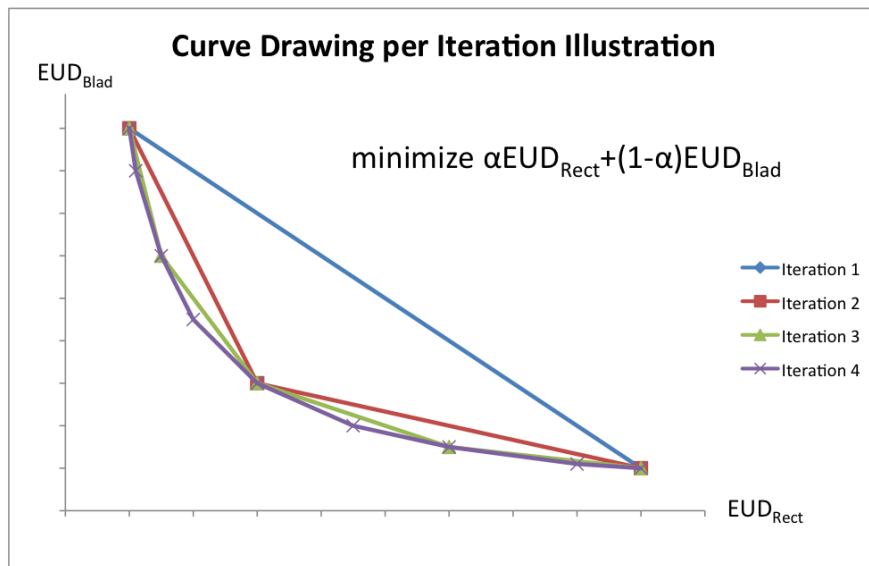


Figure 2.1: Example upper bounds generated by the sandwich algorithm

Choosing objective function weights α There are different ways in which the set of values for α to be used at a particular SALO stage can be determined. In an interactive implementation, the treatment planner could indicate which value to use with the goal of refining the approximation of the Pareto frontier in the clinically most relevant or interesting areas. In an automated setting this can be done by measuring the discrepancy between the upper and lower bounds, and choosing that value of α where the discrepancy is largest. Since the bounds are themselves curves, different discrepancy measures can be used, and each of them will yield a different sequence of values for α and a different bounding of the frontier. However, with careful design, we can ensure that the lower and upper bounds both converge to the Pareto frontier as the number of values for α increases (see *Craft et al.*, 2006).

2.2.2.3 Final Treatment Plan Optimization

This full model clearly has a feasible solution (by construction), i.e., by generating a new set of apertures from scratch we know it will be possible to achieve all previously identified treatment planning goals. The DAO column generation procedure can therefore be initialized with the final solution obtained by the SALO procedure. However, if there is an additional goal related to treatment plan delivery efficiency, such as limiting the beam-on-time or number of apertures, it may be preferable to discard the original aperture pool

and start the DAO algorithm from scratch. Since a feasible solution is needed to start the procedure, the algorithm is then started by first solving an auxiliary problem of the following form:

$$\begin{aligned}
& \underset{(y,z)}{\text{minimize}} && \sum_{\ell=1}^L \max \{G_{\ell}(z) - \bar{G}_{\ell}, 0\} \\
& \text{subject to} && \\
& z_j = \sum_{k \in \mathcal{A}} \mathcal{D}_{kj} y_k && \text{for } j \in \mathcal{V} \\
& y_k \geq 0 && \text{for } k \in \mathcal{A} \\
& z \in Z. &&
\end{aligned} \tag{I}$$

Any feasible solution to this problem with objective function value 0 is a feasible solution to the actual problem from which the DAO algorithm can be started.

2.2.3 Treatment Plan Evaluation Criteria

For our experiments we have chosen to use common measures of gEUD for each target and major critical structure as our main treatment plan evaluation criteria. Letting \mathcal{S} denote the set of structures and \mathcal{V}_s the set of voxels in structure $s \in \mathcal{S}$, the gEUD corresponding to the dose distribution in structure $s \in \mathcal{S}$ is given by the following:

$$\text{gEUD}_s(z; a_s) = \left(\frac{1}{|\mathcal{V}_s|} \sum_{j \in \mathcal{V}_s} z_j^{a_s} \right)^{\frac{1}{a_s}} \tag{2.15}$$

where $1 \leq a_s \leq \infty$ if s is a critical structure while $-\infty \leq a_s \leq 1$ if s is a target, and where

$$\text{gEUD}_s(z; a_s) = \begin{cases} \max_{j \in \mathcal{V}_s} z_j & \text{if } a_s = \infty \\ \frac{1}{|\mathcal{V}_s|} \sum_{j \in \mathcal{V}_s} z_j & \text{if } a_s = 1 \\ \min_{j \in \mathcal{V}_s} z_j & \text{if } a_s = -\infty, \end{cases}$$

(see, e.g., *Niemierko, 1999; Choi and Deasy, 2002*). For the sake of computational efficiency we have chosen to use an approximation of gEUD given by a convex combination of mean and maximum dose for critical structures and of mean and minimum dose for targets (see *Craft et al., 2005; Thieke et al., 2002*). In particular, we choose treatment plan evaluation

criteria functions G_ℓ of the form:

$$\gamma_s \text{gEUD}_s(z; 1) + (1 - \gamma_s) \text{gEUD}_s(z; \infty) = \gamma_s \frac{1}{|\mathcal{V}_s|} \sum_{j \in \mathcal{V}_s} z_j + (1 - \gamma_s) \max_{j \in \mathcal{V}_s} z_j \quad (2.16)$$

if s is a critical structure, and,

$$\gamma_s \text{gEUD}_s(z; 1) + (1 - \gamma_s) \text{gEUD}_s(z; -\infty) = \gamma_s \frac{1}{|\mathcal{V}_s|} \sum_{j \in \mathcal{V}_s} z_j + (1 - \gamma_s) \min_{j \in \mathcal{V}_s} z_j, \quad (2.17)$$

if s is a target, where in both cases $\gamma_s \in [0, 1]$. (Note that, unlike the generic model, target-related criteria will be maximized rather than minimized). These gEUD approximations pertaining to targets are concave, thus taking the negative and minimizing fits into our generic convex problem formulation. The advantage of using these approximations is that our optimization problems can be formulated and solved as linear programs.

As our final criterion we have chosen to minimize the sum of all voxel doses:

$$G_{L+1}(z) = \sum_{j \in \mathcal{V}} z_j. \quad (2.18)$$

Moreover, we assume that no tradeoff takes place between gEUD-criterion G_L and this final criterion G_{L+1} , so that, in problem $(P^{(L)})$, we limit ourselves to $\alpha = 1$ in the corresponding objective function (2.5).

2.3 Illustration of SALO on Clinical Cases

2.3.1 Data and computations

We illustrate our SALO procedure on two clinical cases from different sites: brain and prostate. The brain cancer case has 8 beams and 575 beamlets while the prostate cancer case has 7 beams and 796 beamlets. In both cases the beamlets were of dimension 5×5 mm², and only beamlets whose primary trajectory intersected with the target(s) were included in the model. We chose the set Z to be of the form

$$Z = \{z \in \mathbb{R}^{|\mathcal{V}|} : \underline{z}_s \leq z_j \leq \bar{z}_s, j \in \mathcal{V}_s, s \in \mathcal{S}\}, \quad (2.19)$$

for \underline{z}_s and \bar{z}_s chosen by physicians. The weighting parameters γ_s in the gEUD-approximation described in Section 2.2.3 were found by evaluating both the gEUD and its approximation on a clinically acceptable dose distribution, where suitable values of the gEUD parameters a_s

were taken from *Burman et al. (1991)*; *Lawrence et al. (2010)*; *Mayo et al. (2010)*; *Michalski et al. (2010)*; *Roach, III et al. (2010)*; *Viswanathan et al. (2010)*, and clinical practice at the University of Michigan Department of Radiation Oncology. Values of γ_s were tweaked until the approximate gEUD values from equations (2.16) and (2.17) were the same for the gEUD values from (2.15) for clinically acceptable dose distributions. Table 2.1 provides, for each of the two cases, the structures and two prioritization scenarios, as well as the number of voxels $|\mathcal{V}_s|$ in each structure $s \in \mathcal{S}$, the gEUD parameters a_s and γ_s , and the dose upper and lower bounds \bar{z}_s and \underline{z}_s (where the latter is 0 if omitted). The structures without prioritization values will be addressed in the final treatment plan optimization, but not in the interactive portion of the algorithm.

The optimization problems were all solved on a Mac Pro 4,1 with a single 2.93 GHz Quad-Core Intel Xeon processor and 12 GB DDR3 memory at 1066 MHz. All model generation code was written in C++ and executed in Xcode, and the primal simplex method of CPLEX 12.2 was used as the solver. Since the coefficients \mathcal{D}_{kj} are too numerous to be precomputed and stored, and since the column generation relies on an efficient representation of these coefficients, we make the common assumption that these coefficients can be expressed in terms of so-called beamlet-based dose deposition coefficients,

$$\mathcal{D}_{kj} = \sum_{i \in \mathcal{N}_k} D_{ij}, \quad (2.20)$$

where $\mathcal{N}_k \subseteq \mathcal{N}$ is the subset of beamlets that is exposed in deliverable aperture $k \in \mathcal{A}$, \mathcal{N} is a set of beamlets that discretizes the beams used for treatment, and D_{ij} is the dose delivered to voxel $j \in \mathcal{V}$ from beamlet $i \in \mathcal{N}$ at unit intensity. Storing the (nonzero) coefficients D_{ij} for all $j \in \mathcal{V}$ and $i \in \mathcal{N}$ is manageable in a sparse format. The prostate case used 7 beam directions and the brain case used 9 beam directions. Finally, in Phase 1 of the SALO procedure we generated a pool of $|\bar{\mathcal{A}}| = 100$ apertures.

2.3.2 Results and Discussion

In this section we will describe in detail how the treatment planning process based on the SALO procedure would proceed in the two clinical cases that were solved with a radiation oncologist present. We will illustrate the SALO procedure by going through all steps that a treatment planner would take when developing a treatment plan. For both cases we will show two examples of potential a priori clinical priority lists, along with a potential sequence of decisions made by the treatment planner, for a total of four SALO applications. Furthermore, for ease of exposition we limit ourselves in both cases to a relatively small set of criteria. The tradeoff decisions made during the SALO process were made by a radiation oncologist.

Table 2.1: Structures with corresponding number of voxels and gEUD-parameters for the two clinical test cases

Site	Priority		Structure (s)	$ \mathcal{V}_s $	a_s	γ_s	\bar{z}_s	\underline{z}_s	
	A	B							
brain	1	2	PTV	6,318	-15	0.9	63	56	
		2	1	chiasm	216	10	0.38	57	
		3	3	brainstem	1,836	10	0.5	60	
		4	4	optic nerve (contralateral)	218	10	0.54	63	
		5	5	optic nerve (ipsilateral)	247	10	0.33	63	
				left eye	363			63	
				right eye	345			63	
				left lens	167			63	
				right lens	136			63	
				normal tissue ring 1 (0–1.5 cm from PTV)	6,723			62	
				normal tissue ring 2 (1.5–3 cm from PTV)	4,652			57	
				normal tissue ring 3 (> 3 cm from PTV)	13,037			45	
				TOTAL	34,258				
	prostate	1	2	PTV	3,586	-5	0.3	85.5	73
		2	1	rectum	8,766	8	0.4	78	
		3	3	bladder	5,373	2	0.85	78	
		4	4	penile bulb	294	1	1	85.5	
		5	5	femora	7,049	4	0.8	85.5	
				normal tissue ring 1 (0–1.5 cm from PTV)	2,700			83	
				normal tissue ring 2 (1.5–3 cm from PTV)	7,203			77	
				normal tissue ring 3 (> 3 cm from PTV)	9,419			65	
				TOTAL	44,390				

Note that, in addition to the tradeoff curves, in our experiments the treatment planners were also provided with summary dose distribution information for the different structures (such as minimum, maximum, and mean dose, as well as DVH endpoints) during the process.

2.3.2.1 Brain case

For the brain case we distinguish $L = 5$ major gEUD criteria, one for each of the first five structures listed in Table 2.1. We consider the two alternative priority scenarios A and

B as indicated in that table. Figures 2.2 and 2.3 show the $L - 1$ stages of Phase 2 of the SALO procedure for these two scenarios, while Figure 2.4 shows the final sets of DVH curves. These curves show how scenario A compares to (a) LO and (b) scenario B.

Consider scenario A for the brain case in Figure 2.2. When the treatment planner begins the process, Figure 2.2(a) is generated and presented to the planner. The treatment planner then uses this information to assess the relationship between the gEUDs delivered to the planning target volume (PTV) and the chiasm. In this instance, the treatment planner used this information to choose a lower bound on the gEUD to the PTV of 53.96 Gy, as indicated by the dot. From the graph we can then also conclude that this means that the gEUD to the chiasm will have to be greater than or equal to 47.94 Gy. The optimization model then adds the lower bound on the gEUD to the PTV to the set of constraints and generates Figure 2.2(b). As we can see, by slightly increasing gEUD to the chiasm it is possible to reduce the gEUD to the brainstem by a meaningful amount. However, as we allow more dose to the chiasm, the benefit to the brainstem lessens. Without this accurate information, the planner would not be able to identify the clinically most beneficial tradeoff between these two criteria. The treatment planner then follows this procedure for all other stages.

After optimizing the final criterion (i.e., minimizing the sum of all voxel doses and generating a new set of apertures given the chosen gEUD bounds) we obtain a treatment plan whose DVHs are shown in Figure 2.4(a). In addition, the DVHs obtained by using pure LO are shown as well. We conclude that, by accepting a minor reduction in PTV dose, all other high priority structures receive improved dose distributions, particularly the chiasm. Our treatment planners consider the plan generated by the SALO procedure to be superior to the one created using pure LO. This is consistent with expectations, for if the LO plan were more desirable, then the treatment planner would have selected the extreme points on the tradeoff curves (representing strict prioritization). Scenario B for the brain case provides an alternate prioritization for the criteria, and the choices made by the treatment planner are shown in Figure 2.3. In this scenario the chiasm is of higher importance than the PTV and is constrained before the gEUD to the PTV. Finally, Figure 2.4(b) shows that the difference between the two scenarios is relatively small, indicating a level of robustness of the procedure with respect to interchanging the priorities of PTV and chiasm.

2.3.2.2 Prostate case

For the prostate case we also distinguished $L = 5$ major gEUD criteria, again one for each of the first five structures listed in Table 2.1. We consider the two alternative priority scenarios A and B as indicated in that table. Figures 2.5 and 2.6 show the $L - 1$ stages of Phase 2 of the SALO procedure for these two scenarios, while Figure 2.7 shows the final set

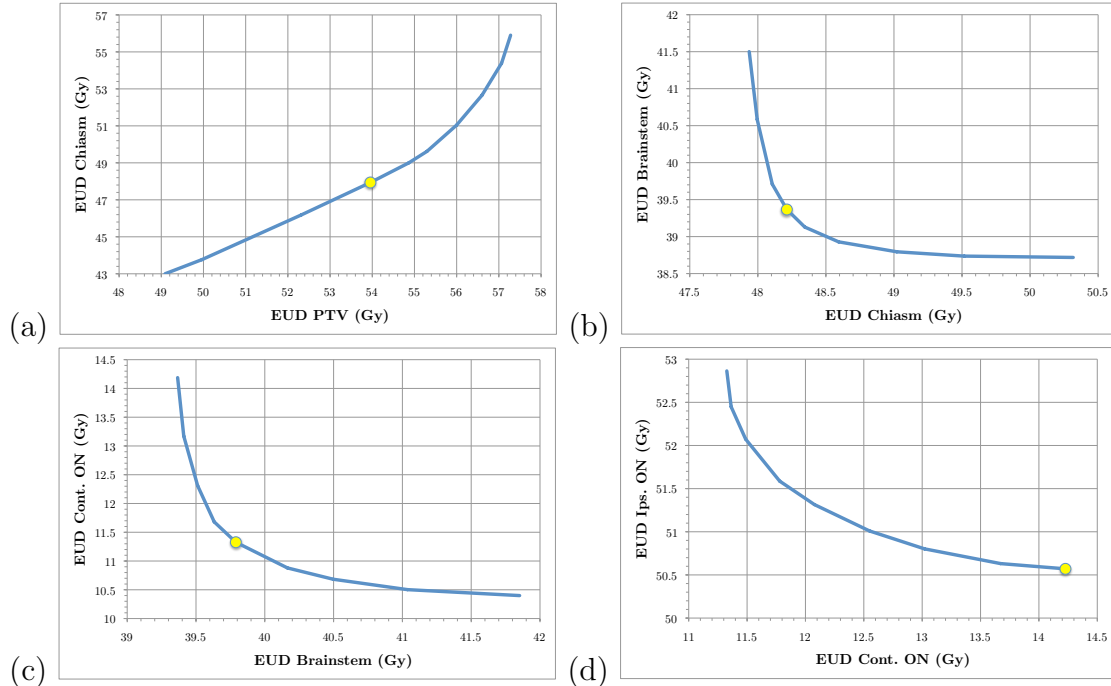


Figure 2.2: SALO progression for the brain case, scenario A: (a) stage 1, (b) stage 2, (c) stage 3, and (d) stage 4

of DVH curves. These curves show how scenario A compares to (a) LO and (b) scenario B.

For the prostate case, the procedure progresses in a similar fashion as for the brain case. However, in this case most of the clinically desired treatment planning goals were more easily satisfied. Therefore, instead of searching for a clinically feasible treatment plan, the SALO process as applied here primarily focused on finding the most desirable treatment plan. For scenario A the first tradeoff is presented in Figure 2.5(a). Instead of just focusing on meeting treatment goals, the treatment planner can decide how aggressively they wish to treat the PTV. For the next sets of tradeoffs, a similar line of reasoning is used, and Figure 2.7(a) allows a comparison of the DVHs obtained by the SALO procedure and pure LO. As in the brain case, a minor reduction in PTV dose allowed for significant reductions in dose to critical structures, especially the rectum.

When selecting desired tradeoffs on the plots, it is critical to pay close attention to the scales on the axes. For example, in Figure 2.5(c) it can be seen that small changes in the bladder yields relatively large improvements in the femora. In contrast, in Figure 2.5(d) the absolute differences are very small in magnitude and clinically insignificant due to the model being very tightly constrained at that stage in the optimization.

In scenario B for the prostate case the rectum received the highest priority. In this case, the treatment planner decided to be slightly more aggressive with respect to the PTV.

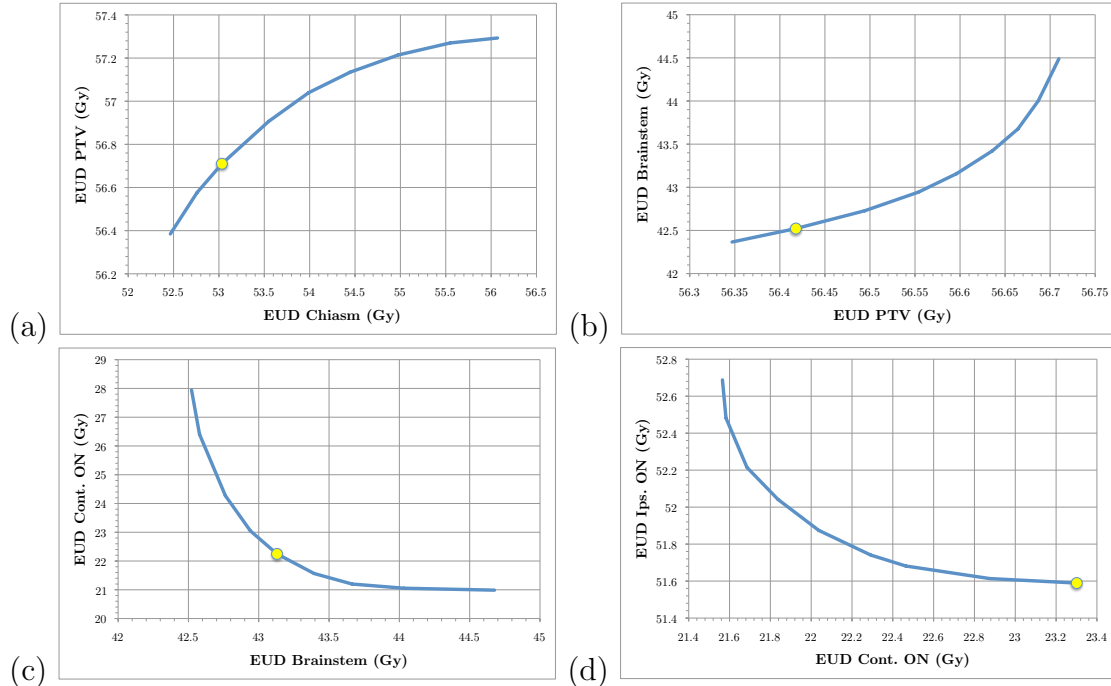


Figure 2.3: SALO progression for the brain case, scenario B: (a) stage 1, (b) stage 2, (c) stage 3, and (d) stage 4

As seen in Figure 2.7(b), the rectum receives more dose using the priorities in scenario B, while the femora receives less. However, these changes are less dramatic than those between scenario A and pure LO.

2.3.3 SALO Methodology Discussion

Through these processes, we can see how the prioritization aspect of LO has been integrated with the interactive nature of MCO. By combining these two characteristics, clinically desirable treatment plans were generated systematically and efficiently.

Without a treatment planning system with the flexibility to automate the SALO procedure, this analysis would not be clinically feasible. For treatment planning systems that allow plug-ins, implementing the SALO procedure is a straightforward process, and the clinical benefits could be realized quite easily. That is, the procedure can be implemented without changing a clinic's current treatment plan solver. Programmers need only to set up some background data structures and a coherent user interface.

The main downside to this type of implementation is that the usability heavily depends on efficiently approximating the tradeoff curve between criteria. If the solver is too slow, the treatment planner could be wasting time waiting for the tradeoff generation. One way




Figure 2.4: Brain case: DVHs for treatment plan for (a) scenario A versus LO and (b) scenario A versus scenario B

to speed up the solving process is use a solution methodology that benefits from previous solution information. For the linear program applied to the brain and prostate cases in

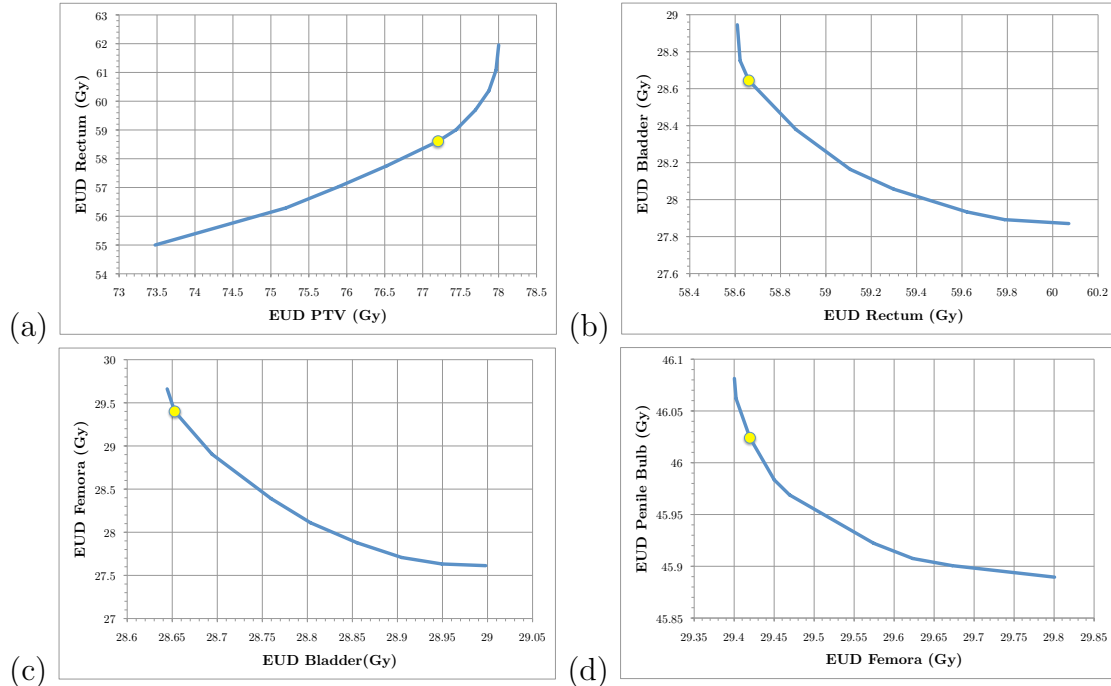


Figure 2.5: SALO progression for the prostate case, scenario A: (a) stage 1, (b) stage 2, (c) stage 3, and (d) stage 4

Section 2.3.2, the solution algorithm used the most recent iteration's solution to initialize the solver for the next point on the tradeoff curve. Another way to speed up the process is to design a solver to run on a graphics processing unit (GPU). When properly designed and coded, models solved using GPUs allow for significant increases in speed (see, e.g., *Men et al.*, 2010).

In practice, it might be beneficial to supplement the SALO procedure with other dose distribution information. Dose distribution statistics and DVHs for points along the tradeoff curves can be generated with little extra computational effort and would bolster the information presented to treatment planners. Because all calculations up to the first stage tradeoff assessment can be done without treatment planner interaction, different first stage scenarios can be generated to influence decisions on the full prioritization. That is, a treatment planner can first look at a few tradeoff curves before deciding the final importances of criteria. Finally, for the treatment planners interested in the final relative weights between the different criteria, these values can be recreated after the SALO procedure is completed (see *Breedveld et al.*, 2009 for this method).

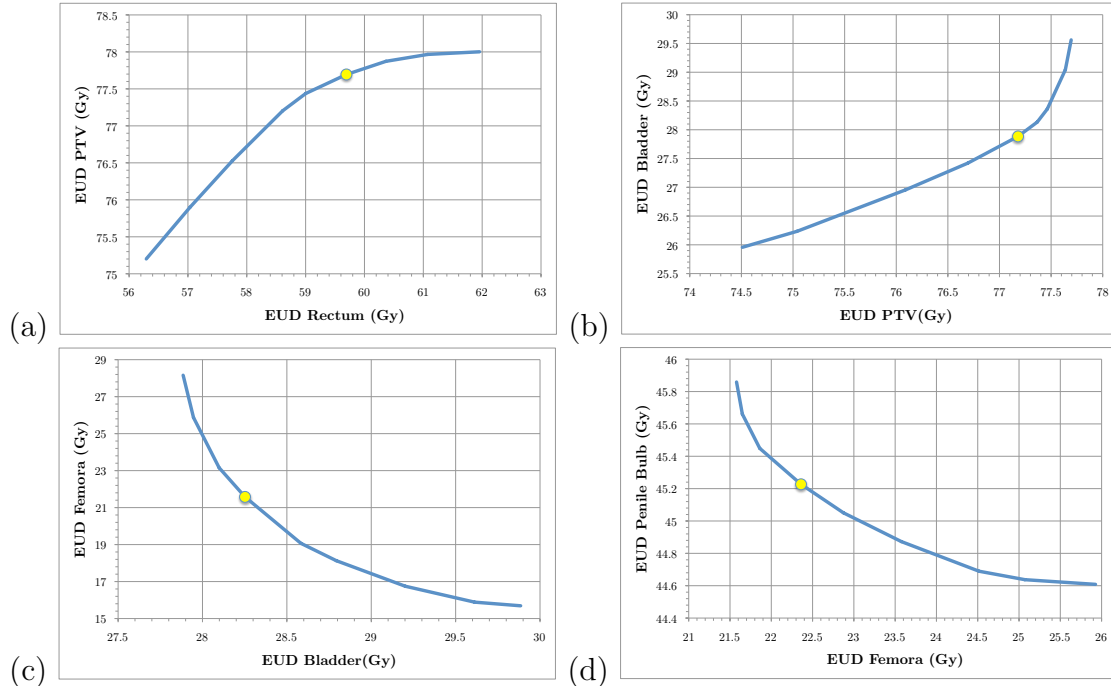


Figure 2.6: SALO progression for the prostate case, scenario B: (a) stage 1, (b) stage 2, (c) stage 3, and (d) stage 4

2.4 SALO Extensions and Improvements

The following section is an extension of the SALO methodology developed with assistance from two undergraduate researchers, Ruqing Ye and Christian Svetnicka. This section will also use the FMO formulation of the IMRT treatment planning problem rather than the DAO formulation presented earlier in the chapter. However, the technique can be extended easily for the DAO formulation.

2.4.1 Pros and Cons of SALO

As discussed earlier, SALO has many advantages that make it more attractive over traditional multicriteria optimization techniques. The SALO procedure is more computationally efficient than full-frontier methods of multicriteria optimization, mainly because it avoids the complex navigation of multidimensional Pareto frontiers. In addition, SALO is flexible in ways that lexicographic ordering methods are not. Lastly, SALO is interactive with the clinician, allowing him/her to rank criteria and maintain control over the entire process, rather than take a completely hands-off approach, like with lexicographic ordering.

However, SALO also has some disadvantages as well. For instance, the interactivity that was previously listed as an advantage can be time-consuming. Clinicians are busy

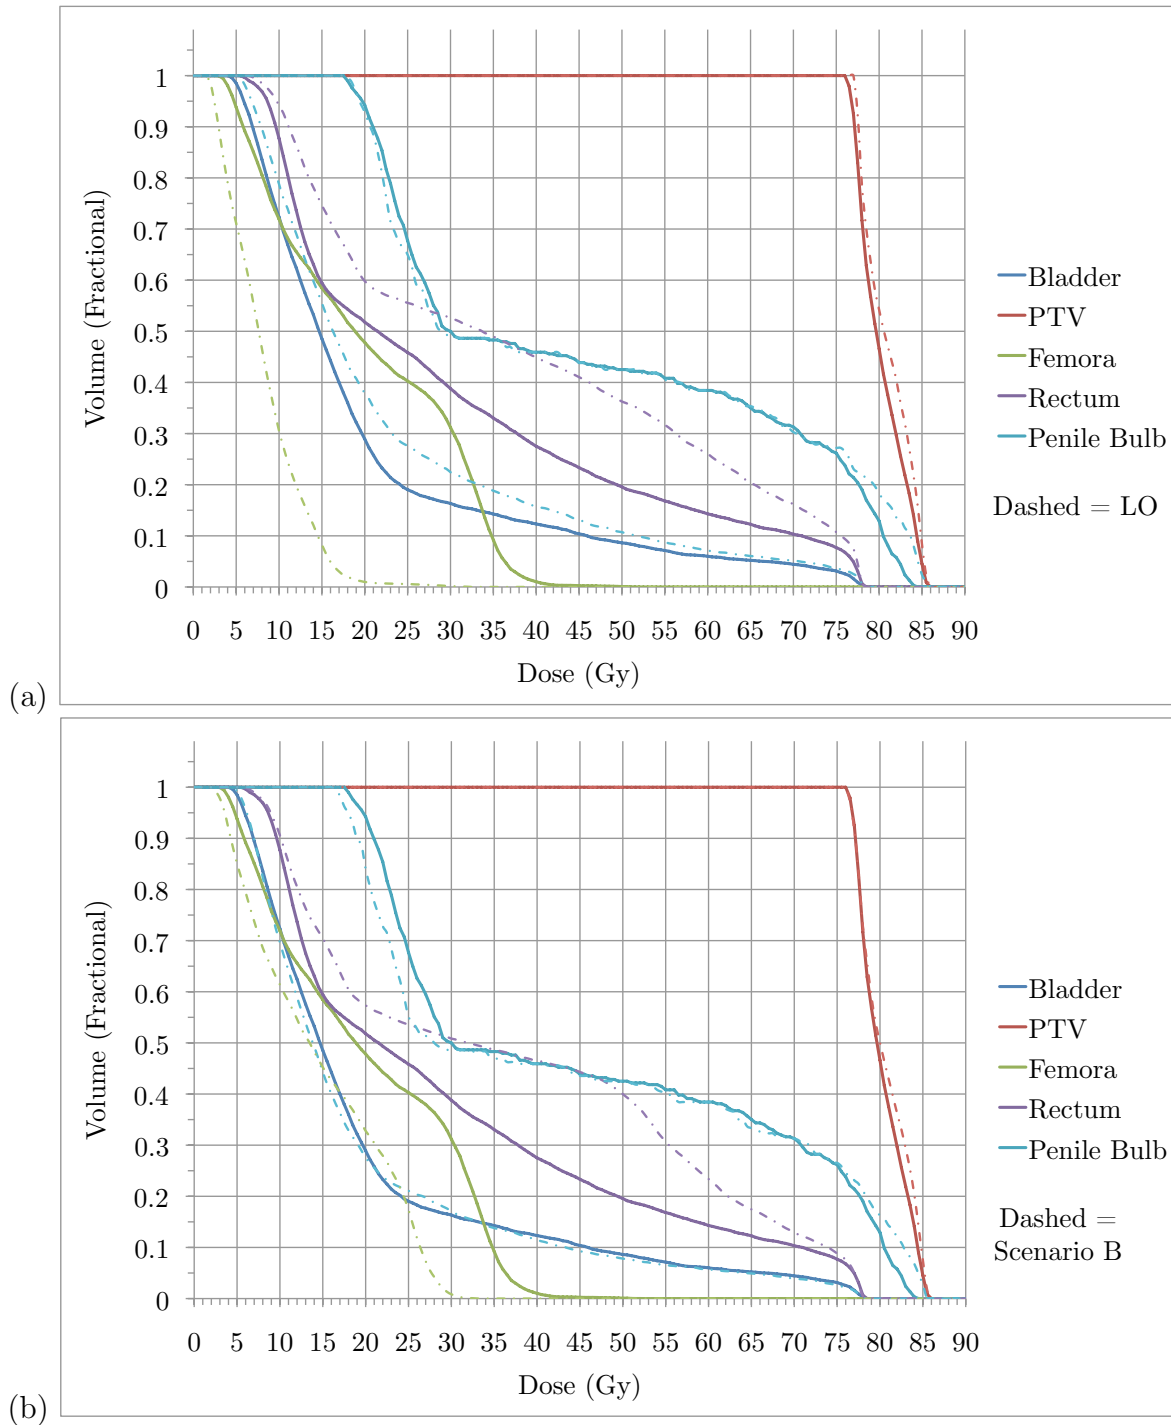


Figure 2.7: Prostate case: DVHs for treatment plan for (a) scenario A versus LO and (b) scenario A versus scenario B

medical professionals, and many don't have time to work with a fully interactive treatment planning method such as SALO. Along with that point, patients often need treatment as

soon as possible, and the slow interactivity can delay the treatment planning. Another disadvantage of SALO is that, while improving over multicriteria optimization, SALO is still prone to examining clinically unacceptable or irrelevant tradeoffs, which can again waste time. Lastly, SALO only deals with relative criteria values rather than absolute values. This is because SALO compares criteria against each other rather than against absolute values, which may be clinically significant. This can be problematic because certain tradeoffs may not make sense clinically (given absolute values), but SALO will still compare them because it is constructed to assume that there is still a relative tradeoff to be made.

2.4.2 SALO with Automation

It is not surprising that a procedure such as SALO has drawbacks. Two significant drawbacks are the slow interactivity of the procedure and the exploration of irrelevant or unnecessary tradeoffs. This is because these drawbacks contribute directly to wasted time. Therefore, we would like to improve the aspects of SALO that unnecessarily lengthen the interaction process. Trivial choices occur when the tradeoff calculated has no clinical relevance or produces an immediately clear choice, and the generation of unnecessary tradeoffs occurs when sections of the tradeoff curve are calculated that would never be considered in the final treatment plan. We seek to automate decisions regarding obvious and/or clinically irrelevant choices so as to improve the treatment planning process and save clinicians' valuable time.

In order to automate this process, we obtained more information from clinicians to incorporate into the treatment planning model. After speaking with our collaborators, we believe that clinicians are able to provide an upper and lower bounds on the clinically relevant ranges of our criteria. The upper bounds would signify a "worst case" limit that we will not allow the dose distribution to exceed. The lower bounds would signify a limit beyond which improvements in the treatment planning criteria no longer have clinical impact. That is, we gain no benefit for improving the criteria beyond that limit. These bounds will be used to identify a "critical region" where tradeoff selection is nontrivial.

2.4.3 SALOA Model

We propose a modification to SALO we're calling sensitivity analysis in lexicographic ordering with automation (SALOA) that improves upon the existing SALO methodology by incorporating new criteria bound information from clinicians. This information includes a lower limit, G_ℓ^L , and "worst case" upper limit, G_ℓ^U , for each criterion. Any tradeoffs that exist outside of the range determined by G_ℓ^L and G_ℓ^U are not clinically relevant and would not be calculated; it would be frivolous and inefficient to consider points that improve upon

the clinical lower limit, and it would also be clinically unacceptable to consider points worse than the worst case upper limit. Therefore, the added constraint $G_\ell(z) \leq G_\ell^U$ will ensure that unacceptable plans will not be considered. As mentioned before, we consider the FMO formulation of the IMRT treatment planning model.

The general model is similar to SALO, but with an additional constraint to bound the criteria by G_ℓ^U , and it is described as follows:

$$\begin{aligned} & \underset{(x,z)}{\text{minimize}} && \{G_\ell(z), G_{\ell+1}(z)\} \\ & \text{subject to} && \text{(P}_{\text{SALOA}}^{(\ell)}) \\ & z_j = \sum_{i \in N} D_{ij} x_i && \text{for } j \in \mathcal{V} && (2.21) \\ & G_{\ell'}(z) \leq \bar{G}_{\ell'} && \ell' = 1, \dots, \ell - 1 && (2.22) \\ & x_i \geq 0 && i \in N && (2.23) \\ & z \in Z && && (2.24) \\ & G_\ell(z) \leq G_\ell^U && \ell = 1, \dots, L && (2.25) \end{aligned}$$

where x_i represents the intensity of beamlet i in set of beamlets N and D_{ij} is the dose to voxel j from beamlet i at unit intensity.

The interactive process differs from SALO. In the original SALO procedure, endpoints, or points very close to the true endpoints, of the Pareto frontier are generated using an objective function of the form $\alpha G_\ell(z) + (1 - \alpha) G_{\ell+1}(z)$ for different values of α , using a high value of α to generate the endpoint $(G_\ell^{\bar{\alpha}}, G_{\ell+1}^{\bar{\alpha}})$ (for which G_ℓ will be lowest and $G_{\ell+1}$ will be highest) and a low value of α to generate an approximate endpoint $(G_\ell^\alpha, G_{\ell+1}^\alpha)$ (for which G_ℓ will be highest and $G_{\ell+1}$ will be lowest), high- and low- α being almost one and almost zero, respectively. With traditional SALO, the algorithm would then use the Sandwich Algorithm to approximate the Pareto frontier by constructing both an upper and a lower bound on the frontier. SALO thus produces a full or nearly full Pareto frontier for G_ℓ and $G_{\ell+1}$. However, this may result in calculating clinically irrelevant areas of the Pareto frontier. Instead, in SALOA, we compare these endpoints to the values of $G_\ell^L, G_{\ell+1}^L, G_\ell^U$ and $G_{\ell+1}^U$. If the endpoints are outside of the clinically relevant region, SALOA will automatically determine the next step that should be taken based on this information.

The following rules for determining the relevant region, in which the tradeoff curve intersects the critical region, help us generate a clinically relevant tradeoff curve for criteria G_ℓ and $G_{\ell+1}$:

$$\max\{G_\ell^L, G_\ell^{\bar{\alpha}}\} \leq G_\ell(z) \leq \max\{\min\{G_\ell^U, G_\ell^\alpha\}, \max\{G_\ell^L, G_\ell^\alpha\}\} \quad (2.26)$$

$$G_{\ell+1}^\alpha \leq G_{\ell+1}(z) \leq \min\{G_{\ell+1}^U, G_{\ell+1}^{\bar{\alpha}}\}. \quad (2.27)$$

Let us define the critical region as the rectangular region for which the tradeoff curve must be displayed to the physician if the tradeoff curve overlaps with it (see Figure 2.8). It is defined by the values in the rectangle generated by the corners (G_ℓ^L, G_ℓ^U) and $(G_{\ell+1}^L, G_{\ell+1}^U)$. Additionally, the endpoint-joint line is the line that connects the two endpoints determined by the high and low values of α .

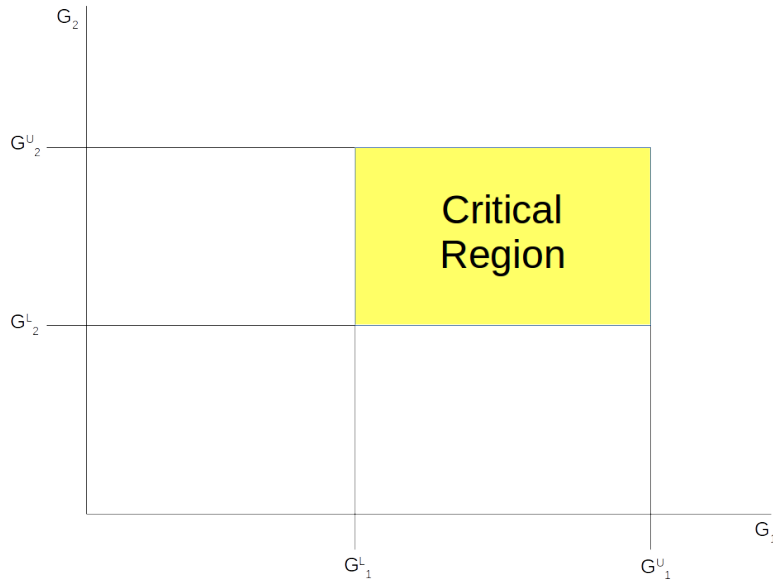


Figure 2.8: Critical region of clinical relevance for G_1 and G_2

2.4.4 Solution Methodology

First, we must determine whether there exists a tradeoff in the clinically significant tradeoff region for our criteria G_ℓ and $G_{\ell+1}$. If there is not a relevant tradeoff, we are then able to automatically select a value without having to engage the planner to choose a tradeoff. However, if there is a relevant tradeoff to be decided upon, then we must generate the tradeoff curve for the clinically relevant values of G_ℓ and $G_{\ell+1}$. To implement this procedure, we examine two major scenarios, with difference in the position of the high- α point: (I) $G_{\ell+1}^{\bar{\alpha}} \leq G_{\ell+1}^L$ and (II) $G_{\ell+1}^L < G_{\ell+1}^{\bar{\alpha}} \leq G_{\ell+1}^U$. Each major scenario will then have minor scenarios to consider and will be evaluated based on the constraints of the given stage.

The first major scenario, (I), to consider occurs when the tradeoff curve never crosses the threshold for the lower limit value of $G_{\ell+1}$, that is, $G_{\ell+1}^{\bar{\alpha}} \leq G_{\ell+1}^L$. Three minor scenarios are given under this major scenario:

$$(I_i) \quad G_{\ell}^{\bar{\alpha}} \leq G_{\ell}^L \text{ and } G_{\ell}^{\alpha} \leq G_{\ell}^L,$$

$$(I_{ii}) \quad G_{\ell}^{\bar{\alpha}} \leq G_{\ell}^L \text{ and } G_{\ell}^L < G_{\ell}^{\alpha} \leq G_{\ell}^U,$$

$$(I_{iii}) \quad G_{\ell}^L < G_{\ell}^{\bar{\alpha}} \leq G_{\ell}^U \text{ and } G_{\ell}^L < G_{\ell}^{\alpha} \leq G_{\ell}^U.$$

These three minor scenarios can be seen in Figure 2.9. In scenario (I_i), it is notable that the tradeoff curve never crosses the threshold for the lower criterion limit for the higher-weighted criterion G_{ℓ} . In other words, for all values of α , G_{ℓ} will always be outside of the clinically relevant region of tradeoffs. Hence, for scenario (i), we can choose the bound for G_{ℓ} to be G_{ℓ}^L , that is, $\overline{G}_{\ell} = G_{\ell}^L$. This will produce the most clinically desirable value for G_{ℓ} while also maintaining significant flexibility for the remaining tradeoffs between the subsequent criteria.

In scenario (I_{ii}), the tradeoff curve originates in the clinically irrelevant region for G_{ℓ} , but ends in the clinically relevant region for G_{ℓ} . In such a case, the bound for G_{ℓ} can again be selected automatically as G_{ℓ}^L . This is because all values between $G_{\ell}^{\bar{\alpha}}$ and G_{ℓ}^L are clinically irrelevant and need not be considered. Then, out of the remaining possible tradeoffs, G_{ℓ}^L achieves the most clinically desirable value for G_{ℓ} while also maintaining significant flexibility for the remaining tradeoffs between the subsequent criteria because of (I).

In scenario (I_{iii}), the tradeoff curve is contained entirely in the clinically relevant region for G_{ℓ} . In this case, it would be logical to choose the value of $G_{\ell}^{\bar{\alpha}}$ as \overline{G}_{ℓ} . This is the value that produces the most clinically desirable value for G_{ℓ} given the tradeoff curve, while still allowing flexibility for the remaining tradeoffs between the subsequent criteria.

Upon examination of the results of each of the three minor scenarios, it can be concluded that \overline{G}_{ℓ} is always the maximum of G_{ℓ}^L and $G_{\ell}^{\bar{\alpha}}$, when $G_{\ell+1}^{\bar{\alpha}} \leq G_{\ell+1}^L$. Therefore, $\overline{G}_{\ell} = \max\{G_{\ell}^L, G_{\ell}^{\bar{\alpha}}\}$.

The second major scenario, (II), to consider occurs when the tradeoff curve does cross the threshold for the lower limit of $G_{\ell+1}$, that is, $G_{\ell+1}^L < G_{\ell+1}^{\bar{\alpha}} \leq G_{\ell+1}^U$. Potential endpoint locations for major scenario (II) are shown in Figure 2.10. The different black lines are potential realizations of the tradeoff curve between the endpoints.

Two minor scenarios are given under this major scenario. For (II_i), the endpoint-joint line potentially intersects with the critical region with $G_{\ell}^{\alpha} \geq G_{\ell}^L$, and we would present to the physician the part of the curve which is included in the critical region to choose the upper bound \overline{G}_{ℓ} . It should be noted that a potential arrangement of endpoints not shown




Figure 2.9: Example endpoint locations and possible tradeoff curves for major scenario (I) with I_i : Red Square I_{ii} : Blue Diamond, I_{iii} : Green Octagon

on Figure 2.10 is that with one endpoint in the critical region and one not. The two shown in the figure are just to show that the endpoints themselves need not be in the critical region to be associated with this minor scenario. That that effect, the entire tradeoff curve does not need to be presented. Rather, only the curve contained within the critical region has to be presented to the physician. In Figure 2.11, the orange region represents the region where the curve could potentially exist. Only curves that intersect with the yellow critical region would be generated. Finding this intersection requires the calculation of new endpoints for the clinically relevant tradeoff curve (i.e., intersection points with the critical region). The new upper endpoint can be found by solving the model with the constraint $G_\ell \geq G_\ell^L$ and the new lower endpoint can be found by solving the model with the constraint $G_{\ell+1} \geq G_{\ell+1}^L$. The curve between the two points will then be generated and displayed to the physician for input.

For (II_{ii}) , the endpoint-joint line does not intersect with the critical region, thus there is no clinically relevant tradeoff to be made. The upper bound for G_ℓ can then be automatically selected with $\overline{G}_\ell = G_\ell^L$. This is the value that achieves the best, clinically significant value for G_ℓ while maintaining significant flexibility for the remaining tradeoffs with the subsequent criteria.

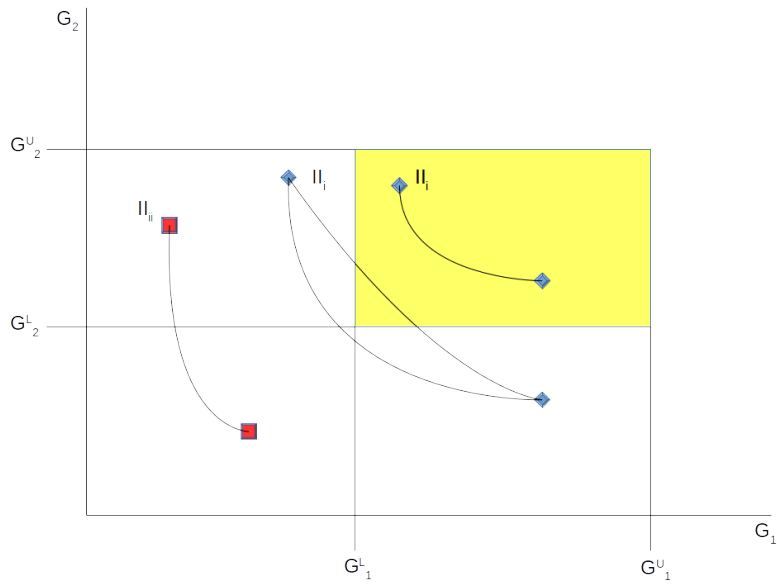


Figure 2.10: Example endpoint locations and possible tradeoff curves for major scenario (II) with II_i : Red Square, II_{ii} : Blue Diamond

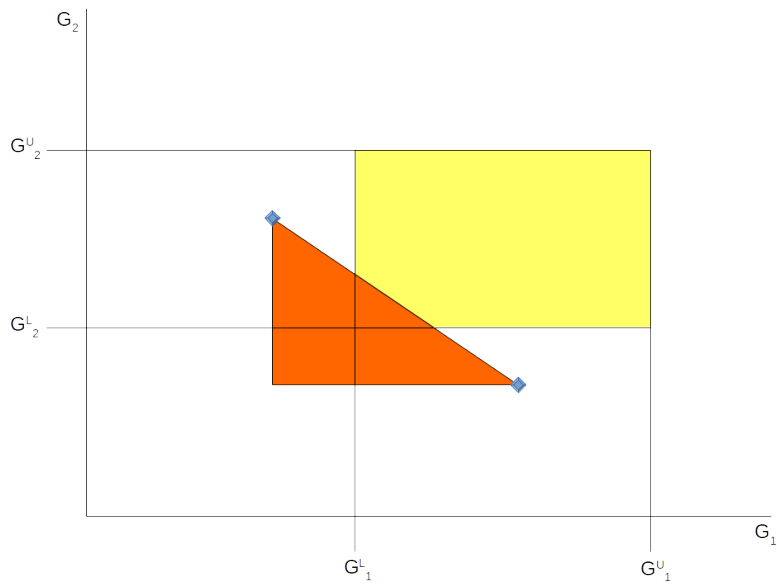


Figure 2.11: Feasible region for the convex tradeoff curve for minor scenario II_i given endpoint locations

2.4.5 SALOA Results

As a proof of concept, we looked at a clinical protocol ($L = 14$) for a head and neck case. This case used 11 beam directions. Instead of looking at the treatment plan quality of SALO and SALOA, we instead investigated the number of times the FMO solver, the most time-consuming subroutine, is called and the number of times physician input is required. Due to the influence of the first criterion selection’s aggressiveness, we consider different levels of constraining PTV coverage and see how these metrics are impacted versus SALO.

Table 2.2: Number of solver calls and physician input queries for SALO and SALOA

Technique: PTV coverage	# Solver Calls	# Physician Inputs
SALOA: low PTV	36	1
SALOA: med-low PTV	39	2
SALOA: med PTV	52	5
SALOA: med-high PTV	52	5
SALOA: high PTV	58	7
SALO	81	13

In Table 2.2, we see that the SALOA technique provides both fewer solver calls and fewer physician interactions while still presenting the treatment planning with relevant tradeoffs.

2.5 Conclusions

The SALO and SALOA procedures provide treatment planners with a directed, systematic process to treatment plan selection. By following a physician’s prioritization, the treatment planner can avoid wasting effort considering clinically inferior treatment plans. The planner is guided by criteria importance, but given the information necessary to accurately assess the tradeoff between criteria each stage. When applied to clinical cases, the SALO procedure efficiently generated desirable treatment plans. The SALOA procedure improved upon the computational efficiency of SALO. As treatment planning becomes more individualized and complex with new techniques and models, methods that efficiently guide the treatment planner towards desirable plans will be necessary to implement these advances at the clinical level.

Acknowledgments

This work was done in collaboration with Dr. Mary Feng, Dr. Benedick Fraass, Dr. Martha Matuszak, Dr. Edwin Romeijn, and Dr. Randall Ten Haken. We would like to

thank Dr. Christina Tsien, M.D. for her involvement in selecting tradeoffs in the SALO implementation. We also would like to thank Drs. James Balter, Ph.D., Marc Kessler, Ph.D., Dan McShan, Ph.D., and Jean Moran, Ph.D. of the Department of Radiation Oncology at the University of Michigan for their insights and contributions. We would like to thank Eduardo Acosta for his ongoing clinical implementation of the methodology. Lastly we would like to thank Christian Svetnicka and Ruqing Ye for their work on SALOA. This work was supported in part by NIH-P01-CA59827.

CHAPTER III

Beam Orientation Optimization

3.1 Introduction

The set of machine instructions necessary for defining a treatment plan depends on the treatment modality. In this chapter, we will focus on developing a model that integrates two levels of machine instructions for IMRT and an associated solution methodology. IMRT delivers radiation from a set of beam locations around the patient with the gantry stationary during delivery. The linear accelerator is equipped with an MLC that, using sliding tungsten leaves, dynamically produces shapes that partially block radiation (see Figure 1.2) called apertures. From each beam location, radiation from multiple apertures at different intensities enable a spatially complex dose distribution to be delivered within the patient. The machine instructions to be determined are (1) the orientation of the stationary beam locations around the patient and (2) the apertures and their corresponding intensities to be delivered at each of these beam locations. For (1), often referred to as beam orientation optimization (BOO), we would like a small number of stationary locations around the body (e.g., 5-9 beam locations) that, when used as the beam set \mathcal{B}' in an FMO problem (see 1.2.1), produces a high quality treatment plan. For (2), given beam set \mathcal{B}' , we would like to find the intensity profile using FMO followed by LS (there is potential for analogous development with DAO).

BOO can be approached in several ways: manual or pre-determined beam selection, local searches and metaheuristics, constructive methods, and combinatorial programming. An in-depth look into beam selection methodologies can be found in *Ehrgott et al., 2008*.

In clinical practice, beam selection has the potential to be an arduous, time consuming process. The potential gains of BOO in some more accessible cancer locations (e.g., prostate cancer) are believed by many physicians to be insignificant compared to the time investment. Therefore, treatment planners will often use a predetermined beam arrangement (e.g., 7 equispaced coplanar beams) to determine (1). Historically, these arrangements have been successful in allowing for clinically desirable dose distributions for the patient. In more

complicated cases, treatment planners will visually inspect the case and select beams that, in their experience, seem as if they will produce a high-quality treatment plan when solving for (2). While this method may produce quick results, the selected beams may not be optimal or even desirable when considering all possible arrangements. The added flexibility in allowing non-coplanar beam arrangements further exacerbates these concerns.

Most BOO literature presents algorithmic methods for determining beam locations. One class of methods for addressing (1) determines beam orientations by iteratively searching the set of potential beam arrangements. These can be local searches (see, e.g., *Aleman et al.*, 2008a; *Lee et al.*, 2011; *Craft*, 2007) or metaheuristic searches (see, e.g., *Bertsimas et al.*, 2013; *Bedford and Webb*, 2006, 2007; *Pugachev et al.*, 2000). Sometimes beam quality evaluation metrics are considered, such as beams-eye-view (BEV), pseudo-beams-eye-view (pBEV), and mean organ-at-risk (see *Goitein et al.*, 1983; *Pugachev and Xing*, 2001; *D'Souza et al.*, 2004, respectively). In these methods, beams are independently evaluated and ranked to be selected based on these criteria. Others consider a greedy framework and evaluate each iteration's beam selection with the knowledge of previous selected beams (see, e.g., *Breedveld et al.*, 2012). Lastly, some research considers integer programming formulations of (1) and (2) (see *Aleman et al.*, 2008b; *Lim et al.*, 2007; *Lim and Cao*, 2012; *Meedt et al.*, 2003; *Mišić et al.*, 2010; *Yarmand et al.*, 2013), but these often take too long to solve in a clinical setting when considering large numbers of non-coplanar potential beam locations.

We study the problem of determining (1) while explicitly incorporating treatment plan quality (i.e., the outputs of (2)). To this end, *we integrate the BOO and FMO treatment planning models* to simultaneously decide (1) and (2). The output of this model is a set of beam locations and the associated intensity profiles for each active location. LS will still need to be performed in order to obtain deliverable apertures, but we assume this occurs without incident. Section 3.2 details the BOO and FMO integrated model and our proposed solution methodology. We apply our methodology to several clinical cases and compare the results to existing methodologies in Section 3.4 and discuss the results and methods in Sections 3.4.2 and 3.4.3. Lastly, we present some theoretical work on objective function bounding in Section 3.5 and give some suggestions for clinical use and future research in Section 3.6.

3.2 Treatment Planning Models

Notation for these models is consistent with that presented in 1.2.1.

3.2.1 Traditional FMO Model

First we consider the traditional FMO model reproduced below from 1.2.1, given some fixed beam set $\mathcal{B}' \subset \mathcal{B}$ and treatment planning objective function $F(z)$. In P_{FMO} below, the fixed active beam set, \mathcal{B}' , is fixed before any other machine instruction decisions are made. P_{FMO} is solved with decision variables x and z .

$$\underset{(x,z)}{\text{minimize}} \quad F(z)$$

subject to (P_{FMO})

$$z_j = \sum_{b \in \mathcal{B}'} \sum_{i \in \mathcal{N}_b} D_{bij} x_{bi} \quad j \in \mathcal{V} \quad (3.1)$$

$$x_{bi} \geq 0 \quad b \in \mathcal{B}', i \in \mathcal{N}_b \quad (3.2)$$

3.2.2 Integrated BOO and FMO Model

For the integrated model, determining which $\mathcal{B}' \subset \mathcal{B}$ is a decision in the model. In BOO, a relatively small subset $\mathcal{B}' \subset \mathcal{B}$, such that $|\mathcal{B}'| \ll |\mathcal{B}|$, from the set of potential beam locations is selected. The chosen set of locations must allow for a high quality treatment plan to be delivered while being small enough to be deliverable in a clinically reasonable amount of time. The integrated BOO and FMO treatment planning model can be set up as follows. Let $F(z)$ be the treatment planning objective function. Model $P_{\text{FMO}}^{\text{BOO}}$ represents the full integrated model, where the decision variables are the contents of set $\mathcal{B}' \subset \mathcal{B}$ and variables x and z .

$$\underset{(\mathcal{B}', x, z)}{\text{minimize}} \quad F(z)$$

subject to ($P_{\text{FMO}}^{\text{BOO}}$)

$$z_j = \sum_{b \in \mathcal{B}} \sum_{i \in \mathcal{N}_b} D_{bij} x_{bi} \quad j \in \mathcal{V} \quad (3.3)$$

$$x_{bi} \geq 0 \quad b \in \mathcal{B}', i \in \mathcal{N}_b \quad (3.4)$$

$$x_{bi} = 0 \quad b \in \mathcal{B} \setminus \mathcal{B}', i \in \mathcal{N}_b \quad (3.5)$$

We would like to solve this integrated model to produce a desirable treatment plan while keeping $|\mathcal{B}'|$ small. Because treatment plan quality is dependent on patient geometry, a strict bound on $|\mathcal{B}'|$ should not be imposed. Instead, the tradeoff between $|\mathcal{B}'|$ and treatment plan quality is considered.

This problem can be expressed as a subset selection problem. Let $\mathcal{X}(\mathcal{B}')$ be the set of feasible solutions given active set \mathcal{B}' (i.e., $\mathcal{X}(\mathcal{B}') = \{(x, z) : (3.3), (3.4), (3.5)\}$) and $\mathcal{F}(\mathcal{B}')$ be a minimum $F(z)$ such that $(x, z) \in \mathcal{X}(\mathcal{B}')$. Therefore, the updated problem, $P_{\text{FMO}}^{\text{BOO}}$, can be written as follows keeping in mind that we want to control $|\mathcal{B}'|$:

$$\min_{\mathcal{B}' \subset \mathcal{B}} \mathcal{F}(\mathcal{B}'). \quad (3.6)$$

3.3 Solution Methodology for $P_{\text{FMO}}^{\text{BOO}}$

The goal of the integrated BOO and FMO model is to select beams while explicitly considering treatment plan quality in both coplanar and non-coplanar settings. However, solving this type of model to optimality can be prohibitively time consuming for large potential beam location sets (see *Yarmand et al.*, 2013), which occurs when considering high-resolution coplanar and most non-coplanar beam arrangements. Using heuristic methods is a common way of avoiding these computational pitfalls when selecting beam orientations provided that they produce high-quality arrangements. Because of the computational advantages and potentially high performance, we consider a greedy framework for subset selection. Let us define $\Delta\mathcal{F}(\mathcal{B}', b)$ as the difference between $\mathcal{F}(\mathcal{B}')$ and $\mathcal{F}(\mathcal{B}' \cup \{b\})$ for some $b \in \mathcal{B} \setminus \mathcal{B}'$ with $\mathcal{B}' \subset \mathcal{B}$:

$$\Delta\mathcal{F}(\mathcal{B}', b) = \mathcal{F}(\mathcal{B}' \cup \{b\}) - \mathcal{F}(\mathcal{B}'). \quad (3.7)$$

A greedy subset selection algorithm can now be presented in this context.

- Step 0: Set $\mathcal{B}' = \emptyset$
- Step 1: Find $b^* = \operatorname{argmin}_{b \in \mathcal{B} \setminus \mathcal{B}'} \Delta\mathcal{F}(\mathcal{B}', b)$
 - If $\Delta\mathcal{F}(\mathcal{B}', b^*) \geq 0$, stop
- Step 2: Set $\mathcal{B}' = \mathcal{B}' \cup \{b^*\}$
 - If $|\mathcal{B}'|$ reaches a pre-determined maximum size, stop
 - If $\mathcal{F}(\mathcal{B}')$ is of desirable quality, stop
 - Otherwise, go to Step 1

The output of this algorithm will be the set of beams to be delivered as well as the associated fluence maps for delivery (i.e., (\mathcal{B}', x, z)). A major computational drawback to

this general greedy algorithm is that Step 1 in each iteration involves $O(|\mathcal{B}|)$ evaluations of function $\mathcal{F}(\mathcal{B}')$, a large-scale optimization problem. When considering non-coplanar beam arrangements where $|\mathcal{B}|$ is large, this is clinically undesirable. Instead, we would like to replace the beam evaluation function, $\Delta\mathcal{F}(\mathcal{B}', b)$, with some approximation function, $\Delta\tilde{\mathcal{F}}(\mathcal{B}', b)$, in Step 1 of the greedy algorithm. We also want to define $\Delta\tilde{\mathcal{F}}(\mathcal{B}', b)$ such that the stopping condition in Step 1, $\Delta\tilde{\mathcal{F}}(\mathcal{B}', b) \geq 0$, holds if there is no benefit to adding any additional beams.

We would like the function $\Delta\tilde{\mathcal{F}}(\mathcal{B}', b)$ to accurately and efficiently predict the benefit of adding a beam without explicitly computing $\Delta\mathcal{F}(\mathcal{B}', b)$. Therefore, the remainder of the section will focus on methods of prediction. There has been work done in this area in signal processing. Matching pursuit (MP) is commonly used in subset selection problems (see *Mallat, 1993; Tropp et al., 2006*). In MP, the algorithm greedily selects the next element of the subset based on which element provides the best local improvement, i.e., using first-order information. While the integrated BOO and FMO model does not exactly fit this specific problem structure, the motivation behind these algorithms provides insight into quantifying $\Delta\tilde{\mathcal{F}}(\mathcal{B}', b)$, especially the idea of iteratively looking at first-order information. In Sections 3.3.1 and 3.3.2, we develop approaches that utilize first-order information to assess the quality of each beam. Section 3.3.3 explores using second-order information.

3.3.1 KKT Motivation

One of the methods for determining apertures and aperture intensities (decisions (2) from Section 3.1) in the literature is DAO (see, e.g., *Romeijn et al., 2005*). DAO utilizes an iterative, greedy approach to generating apertures and their associated intensities given a particular beam arrangement, explicitly solving for decision (2). The subproblem for quantifying aperture quality, which is evaluated for each beam, is derived from studying the Karush–Kuhn–Tucker (KKT) conditions for the treatment planning problem formulated for aperture optimization. KKT conditions are the necessary conditions for a solution to be optimal that use first order information (see *Boyd and Vandenberghe, 2004*). The general idea is to evaluate the KKT condition associated with beams in $\mathcal{B} \setminus \mathcal{B}'$ and add the beam that most violates that condition.

In order to have a single KKT condition associated with each beam, we present an artificially extended version of P_{FMO} . We introduce variable q_b for $b \in \mathcal{B}$ and objective function weight η_b . We show KKT multipliers π , ρ^+ , and ρ^- to the right of their associated constraints.

$$\underset{(x,z,q)}{\text{minimize}} \quad F(z) + \sum_{b \in \mathcal{B}} \eta_b q_b$$

subject to

$(P_{\text{FMO}}^{\text{KKT}})$

$$z_j = \sum_{b \in \mathcal{B}} \sum_{i \in \mathcal{N}_b} D_{bij} x_{bi} \quad j \in \mathcal{V} \quad (\pi_j)$$

$$x_{bi} \leq q_b \quad b \in \mathcal{B}; i \in \mathcal{N}_b \quad (\rho_{bi})$$

$$x_{bi} \geq 0 \quad b \in \mathcal{B}; i \in \mathcal{N}_b. \quad (\gamma_{bi})$$

$$q_b \geq 0 \quad b \in \mathcal{B} \quad (\xi_b)$$

The restricted version of this problem, show below, fixes set $\mathcal{B}' = \bar{\mathcal{B}}'$. This is the analogous to evaluating $\mathcal{F}(\bar{\mathcal{B}}')$ with the addition of variable q_b and objective weight η_b .

$$\underset{(x,z,q)}{\text{minimize}} \quad F(z) + \sum_{b \in \mathcal{B}} \eta_b q_b$$

subject to

$(P_{\text{FMO}}^{\text{KKT}}(\bar{\mathcal{B}}'))$

$$z_j = \sum_{b \in \mathcal{B}} \sum_{i \in \mathcal{N}_b} D_{bij} x_{bi} \quad j \in \mathcal{V}$$

$$x_{bi} \leq q_b \quad b \in \mathcal{B}; i \in \mathcal{N}_b$$

$$x_{bi} \geq 0 \quad b \in \bar{\mathcal{B}}', i \in \mathcal{N}_b$$

$$x_{bi} = 0 \quad b \in \mathcal{B} \setminus \bar{\mathcal{B}}', i \in \mathcal{N}_b$$

$$q_b \geq 0 \quad b \in \mathcal{B}$$

Each iteration, we solve $P_{\text{FMO}}^{\text{KKT}}(\bar{\mathcal{B}}')$ and check it's corresponding solution to the KKT conditions of $P_{\text{FMO}}^{\text{KKT}}$. The conditions read as follows:

$$\pi_j = [\nabla F(z)]_j \quad (3.8)$$

$$\sum_{j \in \mathcal{V}} D_{bij} \pi_j + \rho_{bi} - \gamma_{bi} = 0 \quad b \in \mathcal{B}; i \in \mathcal{N}_b \quad (3.9)$$

$$\eta_b - \sum_{i \in \mathcal{N}_b} \rho_{bi} - \xi_b = 0 \quad b \in \mathcal{B} \quad (3.10)$$

$$\rho_{bi} (x_{bi} - q_b) = 0 \quad b \in \mathcal{B}; i \in \mathcal{N}_b \quad (3.11)$$

$$-\gamma_{bi} x_{bi} = 0 \quad b \in \mathcal{B}; i \in \mathcal{N}_b \quad (3.12)$$

$$\xi_b q_b = 0 \quad b \in \mathcal{B} \quad (3.13)$$

$$\rho_{bi}, \gamma_{bi} \geq 0 \quad b \in \mathcal{B}; i \in \mathcal{N}_b \quad (3.14)$$

$$\xi_b \geq 0 \quad b \in \mathcal{B}. \quad (3.15)$$

For our problem, we assign η_b the value of 0 for all b to make model $P_{\text{FMO}}^{\text{KKT}}(\bar{\mathcal{B}}')$ and function $\mathcal{F}(\bar{\mathcal{B}}')$ achieve the same optimal dose distribution. With that replacement made, note that with a substitution, we must have the following:

$$\sum_{i \in \mathcal{N}_b} \rho_{bi} \leq 0. \quad (3.16)$$

Notice that a solution to $P_{\text{FMO}}^{\text{KKT}}(\bar{\mathcal{B}}')$ satisfies all conditions except the following:

$$\pi_j = [\nabla F(z)]_j \quad (3.17)$$

$$\rho_{bi} = \gamma_{bi} - \sum_{j \in \mathcal{V}} D_{bij} \pi_j \quad b \in \mathcal{B}; i \in \mathcal{N}_b \quad (3.18)$$

$$\sum_{i \in \mathcal{N}_b} \rho_{bi} \leq 0 \quad b \in \mathcal{B}. \quad (3.19)$$

These can be combined into a single constraint for each beam, including inactive beams, through substitution and a little logic (i.e., notice that $\gamma_{bi} \geq 0$). Let us define new notation \bar{v}_{bi} , such that,

$$\bar{v}_{bi} = \left(\sum_{j \in \mathcal{V}} D_{bij} [\nabla F(z)]_j \right)^- \quad b \in \mathcal{B}, i \in \mathcal{N}_b, \quad (3.20)$$

with the constraint in question written as the following:

$$\sum_{i \in \mathcal{N}_b} \bar{v}_{bi} \geq 0 \quad b \in \mathcal{B}. \quad (3.21)$$

By plugging in the optimal solution with active beam set \mathcal{B}' , these conditions may be checked. If they are satisfied, then \mathcal{B}' is an optimal beam arrangement. If they are not, then the beam for which the condition below has the largest violation is added. Because $\mathcal{F}(\mathcal{B}')$ was solved to optimality with active beam set \mathcal{B}' , the only constraints that might not be satisfied in the optimal solution to the restricted problem are those associated with inactive beams in $\mathcal{B} \setminus \mathcal{B}'$. Therefore, the largest violation is found solving the following subproblem through enumeration.

$$\max_{b \in \mathcal{B} \setminus \mathcal{B}'} \sum_{i \in \mathcal{N}_b} v_{bi} \quad (3.22)$$

It should be noted that the vector $\bar{v}_b = \{v_{bi} | i \in \mathcal{N}_b\}$ is precisely the projected gradient of the objective with respect to x_b for each beam. This can be seen by changing the objective to be based on x -values instead of z in the objective function through substitution of constraint (3.3). The expression for each beam can be rewritten as the ℓ_1 -norm of subvectors of these projected gradients.

$$\sum_{i \in \mathcal{N}_b} \bar{v}_{bi} = \|\bar{v}_b\|_1 \leq 0 \quad b \in \mathcal{B} \setminus \mathcal{B}' \quad (3.23)$$

Therefore, the KKT-motivated predictive measure can be defined as the following:

$$\Delta \tilde{\mathcal{F}}_{\text{KKT}}(\mathcal{B}', b) = -\|\bar{v}_b\|_1 \quad b \in \mathcal{B} \setminus \mathcal{B}'. \quad (3.24)$$

3.3.2 Steepest Descent Motivation

Another approach for selecting a beam utilizes the concept of steepest descent. Given an active beam subset \mathcal{B}' , the beam that provides the steepest descent direction among all potential “directions” that only involve adding a single additional beam is selected. To make the notation cleaner, let us first reformulate the objective function to be in terms of beamlet intensities by substituting in dose constraint (3.3).

$$G(x_{\mathcal{B}'}) = F \left(\left(\sum_{b \in \mathcal{B}'} \sum_{i \in \mathcal{N}_b} D_{bij} x_{bi} : j \in \mathcal{V} \right)^\top \right) \quad (3.25)$$

We can present an analogous function where a single additional beam is added to the set of active beams. Let beam $b \in \mathcal{B} \setminus \mathcal{B}'$ be the additional beam.

$$G(x_{\mathcal{B}'}, x_b) = F \left(\left(\sum_{b' \in \mathcal{B}'} \sum_{i \in \mathcal{N}_{b'}} D_{b'ij} x_{b'i} : j \in \mathcal{V} \right)^\top + \left(\sum_{i \in \mathcal{N}_b} D_{bij} x_{bi} : j \in \mathcal{V} \right)^\top \right) \quad (3.26)$$

Suppose that \bar{x} is an optimal solution to the integrated model using only a subset $\mathcal{B}' \subset \mathcal{B}$ of the beams. This would imply that the projected gradient, $[\nabla G(\bar{x}_{\mathcal{B}'})]^-$, would be zero. The steepest descent direction allowing an additional beam $b \in \mathcal{B} \setminus \mathcal{B}'$ with objective $G(x_{\mathcal{B}'}, x_b)$,

however, may be nonzero. Note that only subvector x_b will be involved with the projected gradient calculation at point $(\bar{x}_{\mathcal{B}'}, 0)$. Therefore, the following subproblem is considered with fixed $\bar{x}_{\mathcal{B}'}$.

$$\min_{x_b \geq 0} G(\bar{x}_{\mathcal{B}'}, x_b) \quad (3.27)$$

The steepest descent direction, i.e., the projected gradient corresponding to the additional beam, can be given by the following expression:

$$\bar{v}_b = [\nabla G(\bar{x}_{\mathcal{B}'}, x_b)]^- \Big|_{x_b=0} \quad b \in \mathcal{B}. \quad (3.28)$$

We then calculate the norm of this vector, $\|\bar{v}_b\|_2$, to get our metric of comparison for each beam. Therefore, the steepest-descent-motivated predictive measure can be defined as the following:

$$\Delta \tilde{\mathcal{F}}_{\text{sd}}(\mathcal{B}', b) = -\|\bar{v}_b\|_2 \quad b \in \mathcal{B} \setminus \mathcal{B}'. \quad (3.29)$$

3.3.3 Second-order Motivation

The previous two approaches both rely on the first-order rate of change in objective function value associated with individual beamlets in candidate beams. A perhaps more sophisticated method would use second-order information to determine a Newton-like approximation to the benefits of adding an additional beam's flexibility to the model.

Suppose again that \bar{x} is an optimal solution to the restricted problem only using $\mathcal{B}' \subset \mathcal{B}$. A Newton-like improving direction for $G(x_{\mathcal{B}'}, x_b)$ starting from $(\bar{x}_{\mathcal{B}'}, 0)$ may involve the subvector $x_{\mathcal{B}'}$, unlike the previous approaches. Therefore, for each beam $b \in \mathcal{B} \setminus \mathcal{B}'$, we consider the following subproblem.

$$\min_{x_{\mathcal{B}'}, x_b \geq 0} G(x_{\mathcal{B}'}, x_b) \quad (3.30)$$

The difficulty in this subproblem lies in the non-negativity constraints on the beamlet intensities $(x_{\mathcal{B}'}, x_b)$, as a projected Newton direction for this problem at $(\bar{x}_{\mathcal{B}'}, 0)$ is not guaranteed to be improving (see *Bertsekas*, 1982). Instead, beamlet indices that are not promising candidates for change are identified, i.e., beamlets with zero intensity and strictly positive partial derivative of the objective function. This would indicate that the improving direction for a beamlet with these characteristics would be to decrease in intensity, which is impossible. Let us denote the indices of all beamlets of concern that have this quality by set

$N^+(x_{\mathcal{B}'}, x_b)$.

$$N^+(x_{\mathcal{B}'}, x_b) = \bigcup_{b' \in \mathcal{B}' \cup \{b\}} \left\{ (b', i) : i \in \mathcal{N}_{b'}, x_{b'i} = 0, \left. \frac{\partial G(x_{\mathcal{B}'}, x_b)}{\partial x_{b'i}} \right|_{x=\bar{x}} > 0 \right\} \quad (3.31)$$

The Newton direction, w_b , at $x = (x_{\mathcal{B}'}, x_b)$, is generally in the form of

$$w_b = -\mathcal{M}(x)\nabla G(x), \quad (3.32)$$

where $\mathcal{M}(x)$ is a positive semi-definite matrix and $\nabla G(x)$ is the gradient of G at the current solution. Using set $N^+(x_{\mathcal{B}'}, x_b)$, we construct \mathcal{M} in the following way. This is an improving direction if one exists (see *Bertsekas*, 1982).

$$[\mathcal{M}(x)]_{ii'}^{-1} = \begin{cases} 0, & i \neq i', \text{ and either } i \in N^+(x_{\mathcal{B}'}, x_b) \text{ or } i' \in N^+(x_{\mathcal{B}'}, x_b) \\ \frac{\partial^2 G(x)}{\partial x_i \partial x_{i'}}, & \text{otherwise} \end{cases} \quad (3.33)$$

In order to quantify this value, we consider the Newton decrement. This value can be related to the quantity $G(\bar{x}_{\mathcal{B}'}, 0) - \inf_{x_{\mathcal{B}'}, x_b} \hat{G}(x_{\mathcal{B}'}, x_b)$, where $\hat{G}(x_{\mathcal{B}'}, x_b)$ is the second order approximation of G at $(\bar{x}_{\mathcal{B}'}, 0)$. This value is traditionally used in the unconstrained Newton's method, but we consider a different interpretation to adapt it for our use. In *Boyd and Vandenberghe* (2004), it is described as “the directional derivative of G at $(\bar{x}_{\mathcal{B}'}, 0)$ in the direction of the Newton step.” Therefore, the expression can be seen as the dot product between the projected gradient and the direction of the newton step. The expression (3.35) is derived from equation (3.34) using the direction \tilde{w}_b . We can then define $\Delta\tilde{\mathcal{F}}(\mathcal{B}', b)$ with $\bar{x} = (\bar{x}_{\mathcal{B}'}, 0)$.

$$\lambda(\bar{x}) = \sqrt{([\nabla G(\bar{x})]^+)^{\top} \nabla^2 G(\bar{x})^{-1} \nabla G(\bar{x})} \quad (3.34)$$

$$\Delta\tilde{\mathcal{F}}_{\text{so}}(\mathcal{B}', b) = -\sqrt{([\nabla G_r(\bar{x})]^+)^{\top} \mathcal{M}(\bar{x}) \nabla G_r(\bar{x})} \quad (3.35)$$

Numerical Issues Because the creation of set $N^+(\bar{x}_{\mathcal{B}'}, \bar{x}_b)$ involves identifying beamlet intensities that are zero, it may be beneficial to include beamlets with intensities that are very close to zero. Therefore, we consider ϵ -nearness when generating the set $N^+(\bar{x}_{\mathcal{B}'}, \bar{x}_b)$. That is, we can redefine set $N^+(\bar{x}_{\mathcal{B}'}, \bar{x}_b)$ as the following equation for implementation.

$$N^+(\bar{x}_{\mathcal{B}'}, \bar{x}_b) = \bigcup_{b' \in \mathcal{B}' \cup \{b\}} \left\{ (b', i) : i \in \mathcal{N}_{b'}, x_{b'i} < \epsilon, \left. \frac{\partial G(x_{\mathcal{B}'}, x_b)}{\partial x_{b'i}} \right|_{x=\bar{x}} > 0 \right\} \quad (3.36)$$

3.4 Application

3.4.1 Assessing the BOO Methodology

We would like to compare our methods to clinical practice and other beam-ranking strategies. To this end, we consider a coplanar equispaced strategy, BEV strategy, and pBEV strategy as methods with which to compare. The pBEV original methodology requires a lower PTV bound our FMO model lacks, so we modify pBEV to apply it to our FMO model. We also consider turning on all potential beams to see how far our solutions are from an “ideal” arrangement.

Another consideration when comparing plans is whether or not to change the objective function when beam arrangements change. Since we do not have a consistent and fair way to redefine objective function weights depending on beam arrangement or other factors, we keep the same model parameters and objective weights regardless of beam arrangement. While in practice parameter tweaking would occur once an arrangement is found, we do not consider that step of the planning process for our comparisons. We are assessing these models in terms of objective function value, which we assume to be a good indication of treatment plan quality relative to other plans optimized using the same objective function.

3.4.1.1 Equispaced Beam Strategy

An equispaced strategy is an arrangement of beams that are angularly equally-spaced in a coplanar arc around the patient. Clearly, given a desired number of active beams, there are multiple ways to select equispaced beams. In order to avoid selecting an equispaced arrangement that is subpar given the number of beams and resolution of beam sampling, we solve the FMO with \mathcal{B}' set to each potential equispaced arrangement given our discretization and take the best in terms of objective function value.

3.4.1.2 BEV Methodology

Goitein et al. (1983) provide a simple method for determining beam quality. This method was described in *Lee et al.* (2011), which is the source we reference when implementing the metric. For each beam, the BEV metric counts the number of PTV voxels that have a dose coefficient above some $\kappa > 0$. The metric is as follows:

$$\Delta \tilde{\mathcal{F}}_{\text{BEV}}(\mathcal{B}', b) = \sum_{s \in T} \sum_{j \in \mathcal{V}_s} \min \left\{ 1, \sum_{i \in \mathcal{N}_b} r_{ij} \right\}, \quad \text{where } r_{ij} = \begin{cases} 1, & D_{bij} \geq \kappa \\ 0, & \text{otherwise} \end{cases}, \quad (3.37)$$

with r_{ij} as the indicator variable used in the counting process. It should be noted that this metric does not take into account other active or inactive beam interactions or delivered dose.

In order to select κ , we first calculated the metric with a number of values between 0 Gy and the maximum D_{bij} value in target voxels. We then selected κ for comparison by selecting the κ that produced the lowest objective values when the arrangements associated with each tested κ were solved using an FMO.

3.4.1.3 pBEV Methodology

Pugachev et al. (2000) and *Pugachev and Xing (2001)* present pBEV. Let T_s be the PTV threshold value and \bar{D}_{bj} is the dose delivered to voxel j from beam b when considering the model with a single beam and some “maximum” intensity profile. In order to calculate this dose, the following steps are taken (these are reproduced here from their 2001 paper).

1. Find the voxels affected by the beamlet
2. Assign the beamlet an intensity that could deliver a dose equal to or higher than the prescription in every target voxel
3. For each organ-at-risk (OAR) or normal tissue voxel crossed by the beamlet, calculate the ratio by which the beamlet intensity has to be reduced to ensure that tolerance is not exceeded
4. Find the minimum ratio from the data in step 3
5. Reduce the beamlet Intensity assigned in Step 2 accordingly to the minimum ratio. This value represents the maximum usable intensity of the beamlet
6. Repeat steps 1-5 for all relevant beamlets to obtain the “maximum” beam intensity profile, in which none of the beamlet intensities can be further increased without violating the tolerance of some structure
7. Perform a forward dose calculation using the “maximum” beam intensity profile
8. Compute the overall score of the chosen beam direction according to an empiric score function as follows:

$$\Delta \tilde{\mathcal{F}}_{\text{pBEV}}(\mathcal{B}', b) = \sum_{s \in T} \frac{1}{|\mathcal{V}_s|} \sum_{j \in \mathcal{V}_s} \left(\frac{\bar{D}_{bj}}{T_s} \right)^2. \quad (3.38)$$

The difficulty in applying this approach is that we don't have threshold values for the OARs other than 0 Gy. Instead, we propose that we solve an FMO problem for each beam and then use the resulting dose distribution for \bar{D}_{bj} . Then, high quality beams should have higher values when evaluated with equation (3.38).

3.4.2 BOO Results

These cases were run using a C++ implementation of the pricing problem along with a custom graphical processing unit (GPU) FMO solver for the restricted master problem (RMP). Eigen basic linear algebra subprograms (BLAS) was used for the linear algebra subroutines on the CPU pricing problem. Some of these cases are available for optimizers in *Craft et al.* (2014). We considered a piece-wise quadratic voxel-based objective function (see *Romeijn et al.*, 2004). Let t_j be the target dose value for each voxel and α_j and β_j be the over- and under-dosing penalties, respectively.

$$F(z) = \sum_{j \in \mathcal{V}} \alpha_j ((z_j - t_j)^+)^2 + \sum_{j \in \mathcal{V}} \beta_j ((t_j - z_j)^+)^2 \quad (3.39)$$

First, we want to examine if our metrics do a good job in estimating $\Delta\mathcal{F}(\mathcal{B}', b)$. Let us consider the situation where we've added a few beams to the model and are considering adding an additional beam. In Figures 3.1 and 3.2, let the "Explicit" line be the outright calculation of $\Delta\mathcal{F}(\mathcal{B}', b)$, i.e., the optimal value of FMO for each potential beam. We present a first-order method and the second-order method estimations scaled using linear regression. We do this to show the relative differences in metrics across different beams for each metric. The asterisks represents the beam which would be chosen for each method. The beams are sorted by angular position around a coplanar axis. In Figure 3.1, 2 beams have already been selected. We see the second-order "Hess" methodology selects the beam physically closest to the "best." It should be noted that the "best" beam in the long run is not necessarily the one selected using explicit calculation of $\Delta\mathcal{F}(\mathcal{B}', b)$; that methodology is still heuristic in nature. In Figure 3.2, similar predictive ability can be seen with 4 beams already added for a prostate case. The large "humps" in the curve can be seen at places where beams are already delivering radiation, and thus would be poor candidates when adding beams to the model. This figure also includes the BEV and pBEV metrics. Since BEV and pBEV do not take beam interaction into account, their predictive abilities are quite poor. The non-coplanar results display similar prediction characteristics, but make for difficult-to-read plots and are not shown here.

Next, we investigate if we see similar improvements in objective function value when using our methods to the explicit methodology. We are also interested in the number of beams

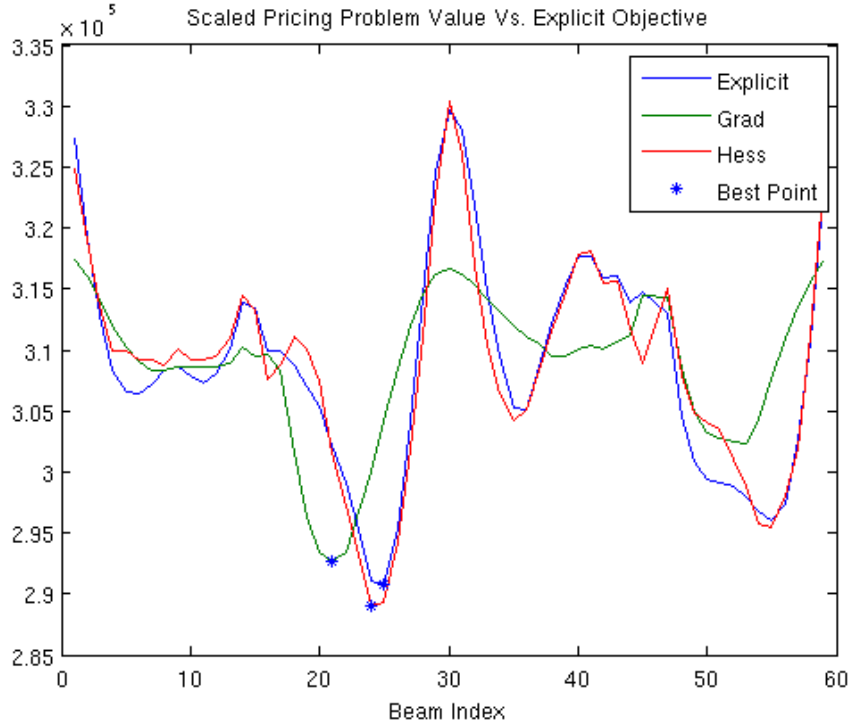


Figure 3.1: Predictive results for steepest descent and second-order methodologies scaled to explicit $\Delta\mathcal{F}(\mathcal{B}', b)$ calculation for $|\mathcal{B}'| = 2$ on a coplanar head-and-neck case case

needed to achieve the same quality as a coplanar arrangement. In Figure 3.3, we see that for an equispaced coplanar arrangement on a prostate case with 60 beams, we can achieve the objective of the best 5-beam equispaced plan with 4 beams using our methods. We also see that our methodologies cause decreases in the objective function almost identically to explicitly calculating $\Delta\mathcal{F}(\mathcal{B}', b)$. In Figure 3.4, a brain case, we look at the coplanar and non-coplanar versions of our methods. The non-coplanar arrangements get close to achieving the same objective as turning on all 60 coplanar beams. In this brain case, 394 non-coplanar beams are considered. Figure 3.5 compares our methods on a non-coplanar liver case with 56 beams. We outperform BEV and pBEV significantly.

Clinically, we are also interested in how using non-coplanar beam changes how dose is distributed in the body. Consider Figures 3.6 and 3.7. Figure 3.6 shows a dosewash of the patient treated with 7 equispaced beams and Figure 3.7 shows the same patient treated with 7 potentially non-coplanar beams selected using the second-order method. The redistribution of can be seen in a reduction in the normal tissue. Non-coplanar directions allow for the dose to spread out a little more around the patient. Another way of visualizing dose to a patient is in a dose-volume histogram (DVH). A DVH shows the fractional volume of a structure

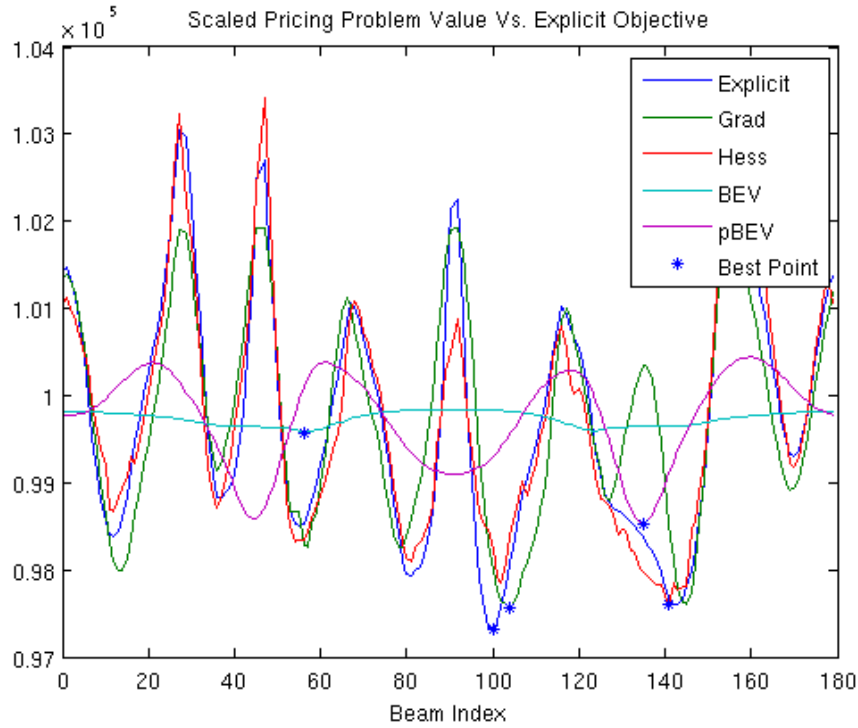


Figure 3.2: Predictive results for steepest descent and second-order methodologies scaled to explicit $\Delta\mathcal{F}(\mathcal{B}', b)$ calculation for $|\mathcal{B}'| = 4$ on a coplanar prostate case case

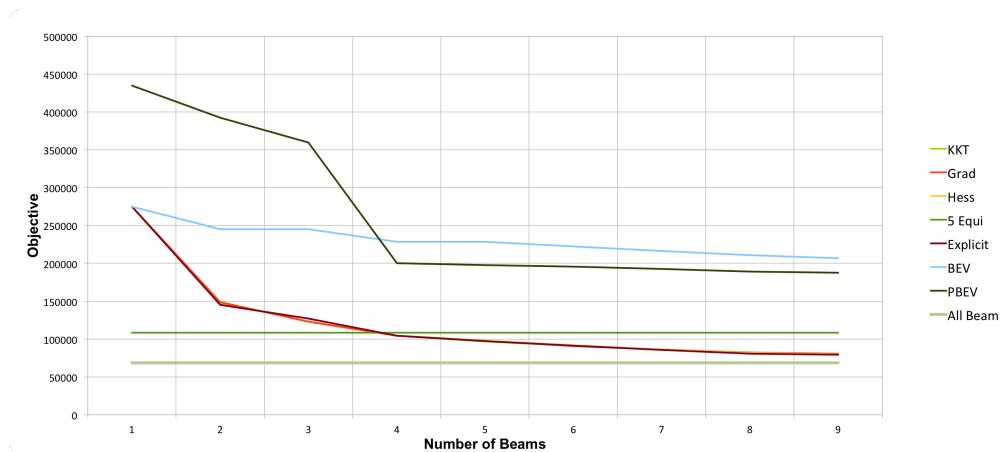


Figure 3.3: Objective function value vs. added number of beams for a coplanar prostate case with 60 beams

(y-axis) that receives some dose or higher(x-axis). The goal is to have a uniform high dose to the cancerous tissue, while delivering lower doses to organs at risk. On the DVH, this means that we want a step function that starts at 100% then steps down to 0% at the desired target dose and for the organs at risk to have curves that are near to the bottom left corner.

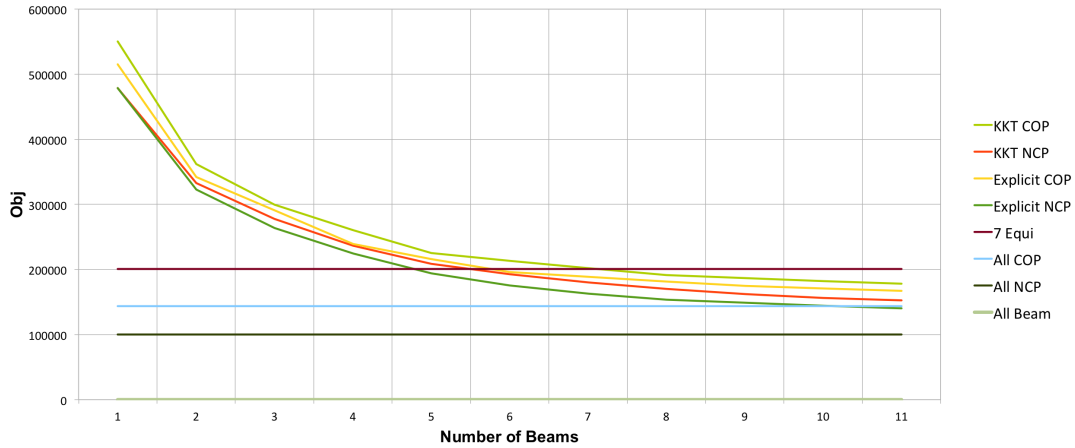


Figure 3.4: Objective function value vs. added number of beams for a brain case with 60 coplanar beams and 394 non-coplanar beams




Figure 3.5: Objective function value vs. added number of beams for a liver case with 56 non-coplanar beams

In the DVH in Figure 3.8, we see the 7-beam second-order plan (dashed) dominates the 7-beam equispaced plan almost everywhere. We observed this behavior in all non-coplanar arrangements with our methodology versus equispaced coplanar arrangements.

3.4.3 Methodology Discussion

For each of these methods, we see good performance that is better than what many are doing in the clinic. However, are these methodologies viable in a clinical environment? There are other considerations to take into account before answering this question.




Figure 3.6: Dosewash of 7 equispaced beams




Figure 3.7: Dosewash of 7 beams selected using the second order method

First, we assume a fixed objective function. In reality, this will be adjusted to the physician's and treatment planner's goals. New objective functions would call for a rerun of the integrated model, i.e., generate new beams, with the current framework. While the solution methodology might allow for that, generating a new beam arrangement can (dramatically) change the set of feasible dose distributions for the treatment plan. This




Figure 3.8: DVH for the patient from Figures 3.6 and 3.7

can lead to inefficient tweaking and too much guesswork involved in changing the objective function each iteration. Instead, it might be better to use some kind of pre-planning objective function weight estimation technique (see, e.g., *Chan et al., 2014*) to define the objective used to determine a beam arrangement, then tweak the plan with that fixed arrangement as is traditionally done.

Second, there is the issue of computational efficiency. Clinically, the act of finding beams needs to be fast enough not to hinder the treatment planner’s workflow. Consider the run times presented in Figure 3.9. We can see the computation time for the different methods. The KKT methodology and the first-order method have essentially the same computational footprint at about one second per iteration, so only one is shown (“Grad”). The second-order (“Hess”) and explicit methodology take much longer. Considering the performance of the first-order methods, we recommend that they be used clinically over the second-order methodology. It should be noted that this was implemented by a non-professional coder, but with efficient libraries. We believe these numbers are proportionately reflective of a high quality CPU implementation. However, these run times could be further improved with a cloud- or GPU-implementation of the pricing problem provided the large amount of coordinated data can be efficiently handled.

Finally, the greedy nature of the algorithm could be a problem. While we believe that

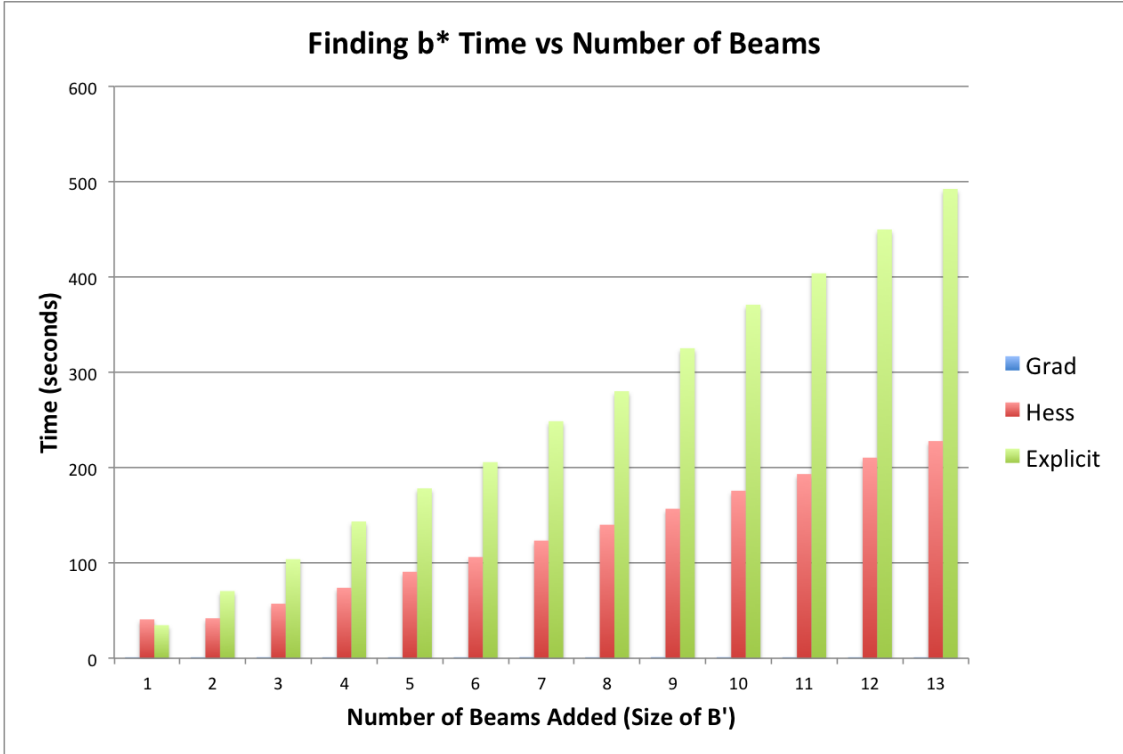


Figure 3.9: Computation time for calculating beam score, $|\mathcal{B}| = 60$

the inherent flexibility of beam selection problem (i.e., different “pretty good” arrangements will do almost as well as the optimal arrangement) will allow our methods to perform well, we did consider ways to control or reduce the impact from the greedy beam selection. First, we considered adding and removing beams. However, when we did some numerical experiments with this, dropping and adding new beams didn’t have a significant enough impact to merit further pursuits. We also considered post-processing heuristics, but others (e.g., *Aleman et al.*, 2008a) have studied neighborhood searches and other beam angle adjustment algorithms. Our methodology could produce a seed arrangement for these post-processing heuristics.

This last point about post processing led us to question how close we were to an optimal arrangement and if there were bounds we could derive to gather this information. While numerical results have not been generated on these bounds, the theoretical groundwork is presented in Section 3.5.

3.5 Bounding on the Optimal Objective Value

An aspect of this problem to consider is how we might bound the optimal objective value, both on $\mathcal{F}(\mathcal{B}')$ and $\mathcal{F}(\mathcal{B})$. We would like to assess the quality of our methodology against bounds in addition to the empirical case results presented earlier. We would also like to use these bounds to gain insight into the problem and inspire future work to be done on this topic. This section presents a series of proofs that lay the groundwork for bounding both the FMO and DAO versions of the integrated problem.

3.5.1 Bounding $\mathcal{F}(\mathcal{B})$ Using an FMO Model

First, our optimal objective value given a restriction on the number of active beams, \bar{B} , can be seen as the following:

$$\min_{\mathcal{B}' \subset \mathcal{B}, |\mathcal{B}'| \leq \bar{B}} \mathcal{F}(\mathcal{B}'). \quad (3.40)$$

Clearly, for $|\bar{B}'| < |\bar{B}''|$, we have the following relationship:

$$\min_{\mathcal{B}' \subset \mathcal{B}, |\mathcal{B}'| \leq \bar{B}'} \mathcal{F}(\mathcal{B}') \geq \min_{\mathcal{B}'' \subset \mathcal{B}, |\mathcal{B}''| \leq \bar{B}''} \mathcal{F}(\mathcal{B}''). \quad (3.41)$$

Therefore, a lower bound on the objective function value for any stage can be found evaluating $\mathcal{F}(\mathcal{B})$, i.e., allowing all beams. This value is the ideal plan objective that we would like to approach with a small subset of the beams. By comparing our solutions to this value, we can assess how close we get to a bound on the optimal solution.

However, evaluation of $\mathcal{F}(\mathcal{B})$ involves solving a $|\mathcal{B}|$ -beam FMO, which may be computationally undesirable or even infeasible. In the non-coplanar setting, the sizes of the $|\mathcal{B}|$ dose-to-point matrices prohibit explicit evaluation on our memory-constrained in-house GPU solver. While a custom CPU solver was developed for this purpose, the custom CPU solver still cannot handle more than around 300 beams (around 24GB of dose-to-point matrices stored in binary format). For larger beam sets, there are methods of approximating bounds on $\mathcal{F}(\mathcal{B})$.

3.5.1.1 Bounding $\mathcal{F}(\mathcal{B})$ given a feasible solution

In order to develop a bound on $\mathcal{F}(\mathcal{B})$, we assume that beamlet intensities are not allowed to go above some upper limit, m_{bi} , for each beam $b \in \mathcal{B}$ and beamlet $i \in \mathcal{N}_b$.

Consider an optimization problem of the following form:

$$\underset{(x,z)}{\text{minimize}} \quad F(z)$$

subject to (\bar{P}_{UB})

$$z_j = \sum_{b \in \mathcal{B}} \sum_{i \in \mathcal{N}_b} D_{bij} x_{bi} \quad j \in \mathcal{V} \quad (3.42)$$

$$x_{bi} \geq 0 \quad b \in \mathcal{B}, i \in \mathcal{N}_b \quad (3.43)$$

$$x_{bi} \leq m_{bi} \quad b \in \mathcal{B}, i \in \mathcal{N}_b. \quad (3.44)$$

The solution to this problem represents the “best” dose distribution possible given all beams can be modulated. Since we often cannot solve this problem due to the size of $|\mathcal{B}|$, we would like to develop a bound on the optimal solution. Let $x = \{x_{bi} : b \in \mathcal{B}, i \in \mathcal{N}_b\}$, $z = \{z_j : j \in \mathcal{V}\}$, and $D = [D_1 \dots D_{|\mathcal{B}|}]$ and let $G(x) = F(Dx)$.

Lemma III.1. *Let $F(z)$ be a convex function. Then $G(x) = F(Dx)$ is a convex function.*

Proof. $\mathbf{dom} G = \{x : Dx \in \mathbf{dom} F\}$. Therefore, the affine mapping preserves convexity, so G is convex (see *Boyd and Vandenberghe, 2004, Section 3.2.2*). \square

Using this knowledge, we can rewrite the problem in terms of x only for simplicity in notation.

$$\underset{x}{\text{minimize}} \quad G(x)$$

subject to (\bar{G}_{UB})

$$x_{bi} \geq 0 \quad b \in \mathcal{B}, i \in \mathcal{N}_b$$

$$x_{bi} \leq m_{bi} \quad b \in \mathcal{B}, i \in \mathcal{N}_b$$

Our next goal is to find a bound on the optimal value to the above problem using a feasible, but not optimal, solution. Let $\nabla G(x)$ be the gradient of $G(x)$ and $x \in \mathcal{X}_{UB}$ represent x satisfying the constraints of \bar{G}_{UB} .

Theorem III.2. *Let $\bar{x} \in \mathcal{X}_{UB}$ and $x^* = \operatorname{argmin}_{x \in \mathcal{X}_{UB}} G(x)$. Then,*

$$G(\bar{x}) + \nabla G(\bar{x})^\top \tau \leq G(x^*) \quad (3.45)$$

with $\tau \in \mathbb{R}^{|\mathcal{N}|}$ such that

$$\tau_{bi} = \begin{cases} m_{bi} - \bar{x}_{bi} & \text{if } [\nabla G(\bar{x})]_{bi} < 0 \\ -\bar{x}_{bi} & \text{if } [\nabla G(\bar{x})]_{bi} \geq 0 \end{cases} \quad \forall b \in \mathcal{B}, i \in \mathcal{N}_b. \quad (3.46)$$

Proof. By convexity of the function G we have the following relationship:

$$G(\bar{x}) + \nabla G(\bar{x})^\top (x^* - \bar{x}) \leq G(x^*). \quad (3.47)$$

Let $[\nabla G(x)]_{bi}$ be term $\{bi\}$ of the gradient of $G(x)$. Observe that we can calculate the largest improvement a change in particular variable can contribute to the objective. Every term $[\nabla G(\bar{x})]_{bi}(x_{bi}^* - \bar{x}_{bi})$ is bounded below by the following:

$$[\nabla G(\bar{x})]_{bi}(x_{bi}^* - \bar{x}_{bi}) \geq \begin{cases} [\nabla G(\bar{x})]_{bi}(m_{bi} - \bar{x}_{bi}) & \text{if } [\nabla G(\bar{x})]_{bi} < 0 \\ [\nabla G(\bar{x})]_{bi}(-\bar{x}_{bi}) & \text{if } [\nabla G(\bar{x})]_{bi} \geq 0 \end{cases} \quad \forall b \in \mathcal{B}, i \in \mathcal{N}_b. \quad (3.48)$$

With τ defined as above, we have the relationship:

$$\nabla G(\bar{x})^\top \tau \leq \nabla G(\bar{x})^\top (x^* - \bar{x}), \quad (3.49)$$

which when combined with previous statements shows that our initial claim holds.

$$G(\bar{x}) + \nabla G(\bar{x})^\top \tau \leq G(\bar{x}) + \nabla G(\bar{x})^\top (x^* - \bar{x}) \leq G(x^*) \quad (3.50)$$

□

Now that we have an expression for a bound, we would like to know if the information necessary can be calculated and the bound evaluated.

$$\begin{aligned} G(\bar{x}) &= \sum_{j \in \mathcal{V}} \alpha_j \left(\left(\sum_{b \in \mathcal{B}} \sum_{i \in \mathcal{N}_b} D_{bij} \bar{x}_{bi} - t_j \right)^+ \right)^2 \\ &+ \sum_{j \in \mathcal{V}} \beta_j \left(\left(t_j - \sum_{b \in \mathcal{B}} \sum_{i \in \mathcal{N}_b} D_{bij} \bar{x}_{bi} \right)^+ \right)^2 \end{aligned} \quad (3.51)$$

$$\begin{aligned} [\nabla G(\bar{x})]_{bi} &= 2 \sum_{j \in \mathcal{V}} \alpha_j D_{bij} \left(\sum_{b' \in \mathcal{B}} \sum_{i' \in \mathcal{N}_{b'}} D_{b'i'j} \bar{x}_{b'i'} - t_j \right)^+ \\ &- 2 \sum_{j \in \mathcal{V}} \beta_j D_{bij} \left(t_j - \sum_{b' \in \mathcal{B}} \sum_{i' \in \mathcal{N}_{b'}} D_{b'i'j} \bar{x}_{b'i'} \right)^+ \end{aligned} \quad (3.52)$$

Here α_j and β_j are the over-dosing and under-dosing penalties, respectively. t_j is the

threshold parameter. The relevant dimensions of the data are in the D matrix and in the parameters of the function G . For the FMO representation of the machine instructions, x_{bi} represents beamlet intensity, $D \in \mathbb{R}^{|\mathcal{N}| \times |\mathcal{V}|}$. Using the quadratic objective below, the other data exists in $\mathbb{R}_{|\mathcal{V}|}$. These sizes lend themselves to manageable calculation sizes and therefore the bound can be computed easily given any $\bar{x} \in \mathcal{X}_{UB}$.

3.5.1.2 Bounding $\mathcal{F}(\mathcal{B})$ given an optimal solution using a fixed subset of beams, \mathcal{B}'

Another piece of information we're interested in is a bound on the optimal objective function using only a subset of the beams, $|\mathcal{B}'| \ll \mathcal{B}$. Consider the following BOO and FMO integrated formulation.

$$\begin{aligned}
& \underset{(\mathcal{B}', x, z)}{\text{minimize}} && F(z) \\
& \text{subject to} && (\bar{P}_{\text{FMO}}^{\text{BOO}}) \\
& z_j = \sum_{b \in \mathcal{B}} \sum_{i \in \mathcal{N}_b} D_{bij} x_{bi} && j \in \mathcal{V} \\
& x_{bi} \geq 0 && b \in \mathcal{B}', i \in \mathcal{N}_b \\
& x_{bi} \leq m_{bi} && b \in \mathcal{B}', i \in \mathcal{N}_b \\
& x_{bi} = 0 && b \in \mathcal{B} \setminus \mathcal{B}', i \in \mathcal{N}_b \\
& |\mathcal{B}'| \leq \bar{B} &&
\end{aligned} \tag{3.53}$$

Theorem III.3. *Let \mathcal{B}'' be some arrangement $\mathcal{B}'' \subset \mathcal{B}$ with $|\mathcal{B}''| = \bar{B}$ with optimal solution $x_{\mathcal{B}''}^*$ (i.e., solving $\bar{P}_{\text{FMO}}^{\text{BOO}}$ with $\mathcal{B}' = \mathcal{B}''$). Suppose that we have an arrangement $\hat{\mathcal{B}}$ such that $\hat{\mathcal{B}} \subset \mathcal{B}''$ and $|\hat{\mathcal{B}}| < \bar{B}$. Let $\bar{x}_{\hat{\mathcal{B}}}$ be a feasible vector for $\bar{P}_{\text{FMO}}^{\text{BOO}}$. Then,*

$$G(\bar{x}_{\hat{\mathcal{B}}}) + \nabla G(\bar{x}_{\hat{\mathcal{B}}})^\top \phi \leq G(x_{\mathcal{B}''}^*), \tag{3.54}$$

with $\|\phi\|_0 = \bar{B}$ calculated as

$$[\phi_{\mathcal{B}''}]_{bi} = \begin{cases} m_{bi} - [\bar{x}_{\hat{\mathcal{B}}}]_{bi} & \text{if } [\nabla G(\bar{x}_{\hat{\mathcal{B}}})]_{bi} < 0, b \in \mathcal{B}'' \\ -[\bar{x}_{\hat{\mathcal{B}}}]_{bi} & \text{if } [\nabla G(\bar{x}_{\hat{\mathcal{B}}})]_{bi} \geq 0, b \in \mathcal{B}'' \\ 0 & \text{if } b \in \mathcal{B} \setminus \mathcal{B}'' \end{cases} \tag{3.55}$$

Proof. By convexity we have the following relationship:

$$G(\bar{x}_{\hat{\mathcal{B}}}) + \nabla G(\bar{x}_{\hat{\mathcal{B}}})^\top (x_{\mathcal{B}''}^* - \bar{x}_{\hat{\mathcal{B}}}) \leq G(x_{\mathcal{B}''}^*). \tag{3.56}$$

Note that each term $[x_{\mathcal{B}''}^*]_{bi} = \bar{x}_{\hat{\mathcal{B}}} = 0$ for all $b \in \mathcal{B} \setminus \mathcal{B}''$, $i \in \mathcal{N}_b$. Therefore, $\|(x_{\mathcal{B}''}^* - \bar{x}_{\hat{\mathcal{B}}})\|_0 \leq \bar{B}$ for any potential arrangement \mathcal{B}'' . By observing the maximum change in decision variables, each term $[\nabla G(\bar{x}_{\hat{\mathcal{B}}})]_{bi}[(x_{\mathcal{B}''}^* - \bar{x}_{\hat{\mathcal{B}}})]_{bi}$ is bounded below as follows:

$$[\nabla G(\bar{x}_{\hat{\mathcal{B}}})]_{bi}[(x_{\mathcal{B}''}^* - \bar{x}_{\hat{\mathcal{B}}})]_{bi} \geq \begin{cases} [\nabla G(\bar{x}_{\hat{\mathcal{B}}})]_{bi}(m_{bi} - [\bar{x}_{\hat{\mathcal{B}}}]_{bi}) & \text{if } [\nabla G(\bar{x}_{\hat{\mathcal{B}}})]_{bi} < 0, b \in \mathcal{B}'' \\ [\nabla G(\bar{x}_{\hat{\mathcal{B}}})]_{bi}(-[\bar{x}_{\hat{\mathcal{B}}}]_{bi}) & \text{if } [\nabla G(\bar{x}_{\hat{\mathcal{B}}})]_{bi} \geq 0, b \in \mathcal{B}'' \\ 0 & \text{if } , b \in \mathcal{B} \setminus \mathcal{B}'' . \end{cases} \quad (3.57)$$

With ϕ defined as above, the following relationship then holds:

$$G(\bar{x}_{\hat{\mathcal{B}}}) + \nabla G(\bar{x}_{\hat{\mathcal{B}}})^\top \phi_{\mathcal{B}''} \leq G(x_{\mathcal{B}''}^*). \quad (3.58)$$

□

We can extend this to apply to an optimal beam arrangement given that $\hat{\mathcal{B}}$ beams are included in the arrangement. Let $\mathcal{B}^*(\hat{\mathcal{B}})$ signify an optimal arrangement given $\hat{\mathcal{B}} \subset \mathcal{B}^*$. We have:

$$\min_{\mathcal{B}'' \subset \mathcal{B}} G(\bar{x}_{\hat{\mathcal{B}}}) + \nabla G(\bar{x}_{\hat{\mathcal{B}}})^\top \phi_{\mathcal{B}''} \leq G(x_{\mathcal{B}^*(\hat{\mathcal{B}})}^*). \quad (3.59)$$

Let $\phi = \phi_{\mathcal{B}''}$, such that,

$$\mathcal{B}'' = \operatorname{argmin}_{\mathcal{B}'' \subset \mathcal{B}} G(\bar{x}_{\hat{\mathcal{B}}}) + \nabla G(\bar{x}_{\hat{\mathcal{B}}})^\top \phi_{\mathcal{B}''}. \quad (3.60)$$

We can easily find the corresponding ϕ by greedily selecting the smallest $\bar{B} - |\hat{\mathcal{B}}|$ terms of $[\nabla G(\bar{x}_{\hat{\mathcal{B}}})]_{bi}[\phi]_{bi}$ as each term is independent of subset selection \mathcal{B}'' .

3.5.2 Bounding $\mathcal{F}(\mathcal{B})$ Using a DAO Model

Another method of determining a dose distribution given a fixed beam set to consider is DAO using potential beam locations as potential aperture locations. Column generation is the standard method of solving DAO, allowing us to avoid the previously mentioned memory issues of solving $\mathcal{F}(\mathcal{B})$. The resulting objective function value will be an approximation to $\mathcal{F}(\mathcal{B})$ if the column generation process is halted before all potentially beneficial columns are added. Generating 1000+ apertures can provide some insight into an “ideal” solution. However, this is still an approximation, and we would like to develop theoretical bounds on the objective function value. We consider a way to make the FMO presented earlier with upper bounds on the beamlet intensities equivalent to a DAO-like procedure, then use that

equivalence to develop new bounds on the optimal solution of $\mathcal{F}(\mathcal{B})$. Let us consider the DAO pricing problem with a beam-on-time bound.

3.5.2.1 DAO Pricing Problem with Beam-On-Time Bound

Let μ_b be the upper bound on total beam intensity, and let \mathcal{K}_b be the set of deliverable apertures for beam b . Let y_{bk} be the intensity of aperture k on beam b . The following is the DAO Master Problem:

$$\underset{(y)}{\text{minimize}} \quad H(y)$$

$$\text{subject to} \quad (\bar{H}_{BB})$$

$$y_{bk} \geq 0 \quad b \in \mathcal{B}, k \in \mathcal{K}_b \quad (3.61)$$

$$\sum_{k \in \mathcal{K}_b} y_{bk} \leq \mu_b \quad b \in \mathcal{B}. \quad (3.62)$$

Let γ_b be the dual variable associated with the beam-on-time constraint and ρ_{bk} be the dual variable associated with the non-negativity constraint. As before, the dose deposition constraint is contained in the objective function. Then, in addition to primal feasibility, the KKT conditions for this master problem would be the following:

$$\gamma_b \geq 0 \quad b \in \mathcal{B} \quad (3.63)$$

$$\rho_{bk} \geq 0 \quad b \in \mathcal{B}, k \in \mathcal{K}_b \quad (3.64)$$

$$\gamma_b \left(\sum_{k \in \mathcal{K}_b} y_{bk} - \mu_b \right) = 0 \quad b \in \mathcal{B} \quad (3.65)$$

$$-y_{bk} \rho_{bk} = 0 \quad b \in \mathcal{B}, k \in \mathcal{K}_b \quad (3.66)$$

$$[\nabla H(y)]_{bk} - \rho_{bk} + \gamma_b = 0 \quad b \in \mathcal{B}, k \in \mathcal{K}_b. \quad (3.67)$$

Now consider formulation \bar{H}_{BB}^{DAO} that only allows certain apertures $\mathcal{K}'_b \subset \mathcal{K}_b$ to have positive intensity. Let $y \in \mathcal{Y}_{BB}^{DAO}$ represent a feasible solution for \bar{H}_{BB}^{DAO} . This is the form of the DAO RMP:

$$\underset{(y)}{\text{minimize}} \quad H(y)$$

$$\text{subject to} \quad (\bar{H}_{BB}^{DAO})$$

$$y_{bk} \geq 0 \quad b \in \mathcal{B}, k \in \mathcal{K}_b$$

$$\begin{aligned} \sum_{k \in \mathcal{K}_b} y_{bk} &\leq \mu_b & b \in \mathcal{B} \\ y_{bk} &= 0 & b \in \mathcal{B}, k \in \mathcal{K}_b \setminus \mathcal{K}'_b. \end{aligned}$$

After obtaining optimal solution $y^* \in \mathcal{Y}_{BB}^{DAO}$ to \bar{H}_{BB}^{DAO} , we want to check if it violates any of the KKT conditions for the original master problem. First consider the KKT conditions to the DAO RMP below:

$$\begin{aligned} \gamma_b &\geq 0 & b \in \mathcal{B} \\ \rho_{bk} &\geq 0 & b \in \mathcal{B}, k \in \mathcal{K}_b \\ \gamma_b \left(\sum_{k \in \mathcal{K}_b} y_{bk}^* - \mu_b \right) &= 0 & b \in \mathcal{B} \\ -y_{bk}^* \rho_{bk} &= 0 & b \in \mathcal{B}, k \in \mathcal{K}_b \\ [\nabla H(y^*)]_{bk} - \rho_{bk} + \gamma_b &= 0 & b \in \mathcal{B}, k \in \mathcal{K}_b. \end{aligned}$$

Without loss of generality, we can assume that $y_{bk}^* > 0$ for $k \in \mathcal{K}'_b, b \in \mathcal{B}$. This can be easily realized by noticing that we can solve the model for some \mathcal{K}'_b , remove any zero-intensity apertures from \mathcal{K}'_b , and then resolve to obtain the same solution with all apertures having positive intensities. This means that for every $k \in \mathcal{K}'_b$ and $b \in \mathcal{B}$, we have $\rho_{bk} = 0$. For any beam b using the entire allowed intensity amount μ_b (i.e., allowing non-negative γ_b), we can calculate $\gamma_b = -[\nabla H(y^*)]_{bk}$ for any $k \in \mathcal{K}'_b$.

The conditions that the optimal solution to the RMP could potentially violate in the master problem are constraints (3.64) and (3.67). These can be summarized as the following:

$$\begin{aligned} \rho_{bk} &\geq 0 & b \in \mathcal{B}, k \in \mathcal{K}_b \\ [\nabla H(y)]_{bk} + \gamma_b &= \rho_{bk} & b \in \mathcal{B}, k \in \mathcal{K}_b. \end{aligned}$$

Therefore, we need to have the following condition satisfied.

$$[\nabla H(y)]_{bk} + \gamma_b \geq 0 \quad b \in \mathcal{B}, k \in \mathcal{K}_b \quad (3.68)$$

Then, for each beam $b \in \mathcal{B}$, we would like to find the following:

$$\min_{k \in \mathcal{K}_b} [\nabla H(y)]_{bk} + \gamma_b. \quad (3.69)$$

Note: $\gamma_b = 0$ for all beams that are not using all of their allowed μ_b intensity. Otherwise,

we use γ_b calculated as above.

This can be easily altered address the beam-on-time constraint in Theorem III.5.

3.5.2.2 Bounding $\mathcal{F}(\mathcal{B})$ Given an Optimal Solution Using a Subset of Apertures

Let $y \in \mathcal{Y}_{BB}$ represent a feasible solution to problem \bar{H}_{BB} and let $y \in \mathcal{Y}_{BB}^{DAO}$ represent a feasible solution for \bar{H}_{BB}^{DAO} .

Theorem III.4. *Let $\bar{y} = \operatorname{argmin}_{y \in \mathcal{Y}_{BB}^{DAO}} H(y)$ and $y^* = \operatorname{argmin}_{y \in \mathcal{Y}_{BB}} H(y)$. Then,*

$$H(\bar{y}) + \sum_{b \in \mathcal{B}} \mu_b \min_{k \in \mathcal{K}_b \setminus \mathcal{K}'_b} [\nabla H(\bar{y})]_{bk} \leq H(y^*). \quad (3.70)$$

Proof. Clearly, $\bar{y} \in \mathcal{Y}_{BB}$. By convexity,

$$H(\bar{y}) + \nabla H(\bar{y})^\top (y^* - \bar{y}) \leq H(y^*). \quad (3.71)$$

Now consider $[\nabla H(\bar{y})]_{bk}(y_{bk}^* - \bar{y}_{bk})$ for $k \in \mathcal{K}'_b, b \in \mathcal{B}$. Since no improving direction exists for problem \bar{H}_{BB}^{DAO} in those dimensions, we have

$$0 \leq \sum_{b \in \mathcal{B}} \sum_{k \in \mathcal{K}'_b} [\nabla H(\bar{y})]_{bk}(y_{bk}^* - \bar{y}_{bk}). \quad (3.72)$$

Thus, we only need to consider $[\nabla H(\bar{y})]_{bk}(y_{bk}^* - \bar{y}_{bk})$ for $k \in \mathcal{K}_b \setminus \mathcal{K}'_b, b \in \mathcal{B}$. Note that for each beam b , the total change in aperture dose is bounded above by μ_b by constraint (3.62) due to $\bar{y} \in \mathcal{Y}_{BB}^{DAO}$. Therefore, the following bound holds for each beam.

$$\mu_b \min_{k \in \mathcal{K}_b \setminus \mathcal{K}'_b} [\nabla H(\bar{y})]_{bk} \leq \sum_{k \in \mathcal{K}_b \setminus \mathcal{K}'_b} [\nabla H(\bar{y})]_{bk}(y_{bk}^* - \bar{y}_{bk}) \quad \forall b \in \mathcal{B}. \quad (3.73)$$

We then have a bound on $H(y^*)$.

$$H(\bar{y}) + \sum_{b \in \mathcal{B}} \mu_b \min_{k \in \mathcal{K}_b \setminus \mathcal{K}'_b} [\nabla H(\bar{y})]_{bk} \leq H(y^*) \quad (3.74)$$

□

In order to calculate this bound, we need to be able to find $\min_{k \in \mathcal{K}_b \setminus \mathcal{K}'_b} [\nabla H(\bar{y})]_{bk}$ for each beam $b \in \mathcal{B}$. The traditional DAO pricing problem finds $\min_{k \in \mathcal{K}_b} [\nabla H(\bar{y})]_{bk}$ and does not consider constraint (3.62). However, there is an updated pricing problem that incorporates constraint (3.62) (see Section 3.5.2.1).

This can be extended to problems with constraint (3.62) replaced by an overall beam-on-time constraint. Consider the following optimization model.

$$\begin{aligned}
& \underset{(y)}{\text{minimize}} && H(y) \\
& \text{subject to} && (\bar{H}_{OB}^{DAO}) \\
& && y_{bk} \geq 0 && b \in \mathcal{B}, k \in \mathcal{K}_b \\
& && \sum_{b \in \mathcal{B}} \sum_{k \in \mathcal{K}_b} y_{bk} \leq T \\
& && y_{bk} = 0 && b \in \mathcal{B}, k \in \mathcal{K}_b \setminus \mathcal{K}'_b
\end{aligned}$$

Theorem III.5. *Let $\bar{y} = \operatorname{argmin}_{y \in \mathcal{Y}_{OB}^{DAO}} H(y)$ and $y^* = \operatorname{argmin}_{y \in \mathcal{Y}_{OB}} H(y)$. Then,*

$$H(\bar{y}) + T \min_{k \in \mathcal{K}_b, b \in \mathcal{B}} [\nabla H(\bar{y})]_{bk} \leq H(y^*). \quad (3.75)$$

Proof. Clearly, this can be seen as all apertures having a single “beam” bound. Given this observation, this proof structure is identical to that in Theorem III.4.

3.5.2.3 Bounding $\mathcal{F}(\mathcal{B})$ Given a Feasible Solution Considering All Potential Apertures

We would like to consider \bar{y} to be any feasible solution to \bar{H}_{OB} below. This would mean that it is possible for all of the $|\mathcal{K}|$ apertures to have non-negative values.

$$\begin{aligned}
& \underset{(y)}{\text{minimize}} && H(y) \\
& \text{subject to} && (\bar{H}_{OB}) \\
& && y_{bk} \geq 0 && b \in \mathcal{B}, k \in \mathcal{K}_b \\
& && \sum_{b \in \mathcal{B}} \sum_{k \in \mathcal{K}_b} y_{bk} \leq T
\end{aligned} \quad (3.76)$$

Let the feasible region here be define as \mathcal{Y}_{OB} .

Theorem III.6. *Let $\bar{y} \in \mathcal{Y}_{OB}$ and $y^* = \operatorname{argmin}_{y \in \mathcal{Y}_{OB}} H(y)$. Then,*

$$H(\bar{y}) + T \min_{\substack{k \in \mathcal{K}_b \\ b \in \mathcal{B} \\ \bar{y}_{bk} < T}} [\nabla H(\bar{y})]_{bk} - T \max_{\substack{k \in \mathcal{K}_b \\ b \in \mathcal{B} \\ \bar{y}_{bk} > 0}} [\nabla H(\bar{y})]_{bk} \leq H(y^*). \quad (3.77)$$

Proof. Clearly, we have:

$$H(\bar{y}) + \nabla H(\bar{y})^\top (y^* - \bar{y}) \leq H(y^*). \quad (3.78)$$

For each $k \in \mathcal{K}_b, b \in \mathcal{B}$ we have:

$$[\nabla H(\bar{y})]_{bk}(y_{bk}^* - \bar{y}_{bk}) \geq \begin{cases} [\nabla H(\bar{y})]_{bk}(T) & \text{if } [\nabla H(\bar{y})]_{bk} < 0 \\ [\nabla H(\bar{y})]_{bk}(-T) & \text{if } [\nabla H(\bar{y})]_{bk} \geq 0 \end{cases} \quad \forall b \in \mathcal{B}, k \in \mathcal{K}'_b. \quad (3.79)$$

Because of constraint (3.76), we know that the total increase and total decrease in intensity are each bounded by T . That is,

$$\|(y^* - \bar{y})^+\|_1 \leq T \quad (3.80)$$

and

$$\|(\bar{y} - y^*)^+\|_1 \leq T. \quad (3.81)$$

Therefore, we have the following bounds:

$$T \min_{\substack{k \in \mathcal{K}_b \\ b \in \mathcal{B} \\ \bar{y}_{bk} < T}} [\nabla H(\bar{y})]_{bk} \leq \sum_{b \in \mathcal{B}} \sum_{\substack{k \in \mathcal{K}_b \\ \bar{y}_{bk} < T}} ([\nabla H(\bar{y})]_{bk})^- (y_{bk}^* - \bar{y}_{bk})^+ \quad (3.82)$$

and

$$-T \max_{\substack{k \in \mathcal{K}_b \\ b \in \mathcal{B} \\ \bar{y}_{bk} > 0}} [\nabla H(\bar{y})]_{bk} \leq - \sum_{b \in \mathcal{B}} \sum_{\substack{k \in \mathcal{K}_b \\ \bar{y}_{bk} > 0}} ([\nabla H(\bar{y})]_{bk})^+ (\bar{y}_{bk} - y_{bk}^*)^+. \quad (3.83)$$

Combining our previous statements, we have satisfy our theorem.

$$H(\bar{y}) + T \min_{\substack{k \in \mathcal{K}_b \\ b \in \mathcal{B} \\ \bar{y}_{bk} < T}} [\nabla H(\bar{y})]_{bk} - T \max_{\substack{k \in \mathcal{K}_b \\ b \in \mathcal{B} \\ \bar{y}_{bk} > 0}} [\nabla H(\bar{y})]_{bk} \leq H(y^*) \quad (3.84)$$

□

This may be difficult to calculate for the general case due to the calculation involving all potential apertures that could potentially occur.

3.5.2.4 Bounding $\mathcal{F}(\mathcal{B})$ Given a Feasible Solution Using a Subset of Apertures

Consider problem \bar{H}_{OB} and RMP \bar{H}_{OB}^{DAO} . We let \bar{y} be any feasible solution to \bar{H}_{OB}^{DAO} and then use that to calculate a bound on the optimal solution of \bar{H}_{OB}^{DAO} .

$$\begin{aligned} & \underset{(y)}{\text{minimize}} && H(y) \\ & \text{subject to} && (\bar{H}_{OB}^{DAO}) \end{aligned}$$

$$\begin{aligned} y_{bk} &\geq 0 && b \in \mathcal{B}, k \in \mathcal{K}_b \\ \sum_{b \in \mathcal{B}} \sum_{k \in \mathcal{K}_b} y_{bk} &\leq T \\ y_{bk} &= 0 && b \in \mathcal{B}, k \in \mathcal{K}_b \setminus \mathcal{K}'_b \end{aligned}$$

Theorem III.7. *Let $\bar{y} \in \mathcal{Y}_{OB}^{DAO}$ and $y^* = \operatorname{argmin}_{y \in \mathcal{Y}_{OB}^{DAO}} H(y)$. Then,*

$$H(\bar{y}) + T \min_{k \in \mathcal{K}_b \setminus \mathcal{K}'_b, b \in \mathcal{B}} [\nabla H(\bar{y})]_{bk} + T \min_{\substack{k \in \mathcal{K}'_b \\ b \in \mathcal{B} \\ \bar{y}_{bk} < T}} [\nabla H(\bar{y})]_{bk} - T \max_{\substack{k \in \mathcal{K}'_b \\ b \in \mathcal{B} \\ \bar{y}_{bk} > 0}} [\nabla H(\bar{y})]_{bk} \leq H(y^*) \quad (3.85)$$

Proof. By Theorem III.4, we have the following relationship:

$$T \min_{k \in \mathcal{K}_b \setminus \mathcal{K}'_b, b \in \mathcal{B}} [\nabla H(\bar{y})]_{bk} \leq \sum_{b \in \mathcal{B}} \sum_{k \in \mathcal{K}_b \setminus \mathcal{K}'_b} [\nabla H(\bar{y})]_{bk} (y_{bk}^* - \bar{y}_{bk}). \quad (3.86)$$

Now we must find a lower bound for the following:

$$\sum_{b \in \mathcal{B}} \sum_{k \in \mathcal{K}'_b} [\nabla H(\bar{y})]_{bk} (y_{bk}^* - \bar{y}_{bk}). \quad (3.87)$$

Note that these variables are in the structure as Theorem III.6. Therefore, we apply Theorem III.6 to get the following bound on $H(y^*)$.

$$H(\bar{y}) + T \min_{k \in \mathcal{K}_b \setminus \mathcal{K}'_b, b \in \mathcal{B}} [\nabla H(\bar{y})]_{bk} + T \min_{\substack{k \in \mathcal{K}'_b \\ b \in \mathcal{B} \\ \bar{y}_{bk} < T}} [\nabla H(\bar{y})]_{bk} - T \max_{\substack{k \in \mathcal{K}'_b \\ b \in \mathcal{B} \\ \bar{y}_{bk} > 0}} [\nabla H(\bar{y})]_{bk} \leq H(y^*) \quad (3.88)$$

□

3.5.3 Equivalence of Pricing Problem Constraints

Consider the bounded FMO model from Section 3.5.1.1 and the bounded DAO model from Section 3.5.2.2. The main difference is between constraints (3.44) in the FMO model and (3.62) in the DAO model. The physical differences in apertures and individual beamlets (i.e., deliverability requirements) contribute to this modeling discrepancy. The reason we consider these two models is that we want a bound on the BOO/FMO integrated model using a DAO bounding technique. In order to have the bound be relevant, we must have the FMO feasible region be contained in the DAO feasible region. Ideally, the regions would be equivalent.

$$x_{bi} \leq m_{bi} \quad b \in \mathcal{B}, i \in \mathcal{N}_b \quad (3.44)$$

and

$$\sum_{k \in \mathcal{K}_b} y_{bk} \leq \mu_b \quad b \in \mathcal{B} \quad (3.62)$$

Let us assume that $\mu_b = m_{bi}$ for all $b \in \mathcal{B}$ and $i \in \mathcal{N}_b$. First, we consider if (3.62) implies (3.44). That is, if we have a feasible solution to the DAO model using constraint (3.62), can we guaranty a feasible solution to the FMO model with constraint (3.44) with the same objective function. We do so by finding beamlets that produce the same dose distribution.

Lemma III.8. *Let $\bar{y} \in \mathcal{Y}_{BB}^{DAO}$. Then, $\exists \bar{x} \in \mathcal{X}_{UB}$ such that $z = Dx = \mathcal{D}y$.*

Proof. Let

$$\bar{x}_{bi} = \sum_{\substack{k \in \mathcal{K}_b \\ \text{s.t. } i \in \mathcal{N}_k}} \bar{y}_{bk} \quad b \in \mathcal{B}, i \in \mathcal{N}_b. \quad (3.89)$$

Then, by our our definition of $z = Dx$ and $z = \mathcal{D}y$, we have $z = D\bar{x} = \mathcal{D}\bar{y}$. \bar{x} satisfies the non-negativity constraints. We see that the total dose any \bar{x}_{bi} is assigned is less than μ_b , and thus constraint (3.44) holds. \square

Next, we consider if (3.44) implies (3.62). That is, if we have a feasible solution to the FMO model with constraint (3.44), can we guaranty a feasible solution to the DAO model using constraint (3.62) with the same objective function. We do so by finding apertures that produce the same dose distribution.

Unfortunately, this does not necessarily work. We can see using a simple example. Let us consider a single beam with a single row of three beamlets with intensities $(\mu_b, 0, \mu_b)$. We can see that there is not any deliverable aperture shape to satisfy this intensity mapping. Therefore more than μ_b intensity is necessary to deliver that map.

However, if we relax the MLC constraints on the DAO problem (i.e., allow undeliverable apertures to exist in \mathcal{K}_b), then we can deliver any beamlet intensity map with at most μ_b total aperture dose for each beam. Let $\tilde{\mathcal{K}}_b$ be the set of all potential apertures (disregarding deliverability) for beam $b \in \mathcal{B}$. Let $\tilde{\mathcal{Y}}_{BB}^{DAO}$ be the associated feasible region with \mathcal{K}_b replaced by $\tilde{\mathcal{K}}_b$.

Theorem III.9. $\bar{x} \in \mathcal{X}_{UB} \iff \bar{y} \in \tilde{\mathcal{Y}}_{BB}^{DAO}$ such that $D\bar{x} = \mathcal{D}\bar{y}$.

Proof. First, it follows from Lemma III.8 that if $y \in \tilde{\mathcal{Y}}_{BB}^{DAO}$, then we can find $\bar{x} \in \mathcal{X}_{UB}$ such that $D\bar{x} = \mathcal{D}y$ because $\mathcal{Y}_{BB}^{DAO} \subseteq \tilde{\mathcal{Y}}_{BB}^{DAO}$.

Second, we would like to show that if $\bar{x} \in \mathcal{X}_{UB}$, then $\exists \bar{y} \in \tilde{\mathcal{Y}}_{BB}^{DAO}$ such that $D\bar{x} = \mathcal{D}\bar{y}$. Since deliverability is no longer a concern, we use the following procedure to generate apertures and assign intensities.

- Step 1: Generate aperture $k' \in \tilde{\mathcal{K}}_b$ consisting of beamlet set $\mathcal{N}_{k'}$ such that $i \in \mathcal{N}_{k'}$ if $\bar{x}_{bi} > 0$
- Step 2: Set $\bar{y}_{bk'} = \min\{\bar{x}_{bi} : i \in \mathcal{N}_{k'}\}$
- Step 3: Set $\bar{x}_{bi} = \bar{x}_{bi} - \bar{y}_{bk'}$ for all $i \in \mathcal{N}_{k'}$
- Step 4: If any $\bar{x}_{bi} > 0$, then go to Step 1. Otherwise, end.

We can see that by constructing the aperture intensities in this manner, we have at most $|\mathcal{N}_b|$ iterations and thus at most $|\mathcal{N}_b|$ apertures. Each iteration, at least one beamlet will drop out from the set of beamlets with positive intensity. Because we are reducing all non-zero beamlet intensities uniformly, the highest total aperture dose we can have is the maximum beamlet intensity, μ_b . Therefore, constraint (3.62) holds. With all other non-generated apertures zero, the non-negativity constraints hold. Lastly, by our definition of $z = Dx$ and $z = \mathcal{D}y$, we have $z = D\bar{x} = \mathcal{D}\bar{y}$. Thus, our theorem holds. \square

It may be computationally beneficial to consider bounding the FMO model using the DAO model allowing undeliverable apertures with constraints (3.44) in the FMO model and (3.62) in the DAO model. The idea is that the fewer elements we use to undershoot the bound, the better the bound will be. The number of apertures used in a clinically desirable solution is traditionally fewer than the total number of beamlets.

3.5.4 Considerations and Potential Bounding Uses

While the current section has some methods for bounding $\mathcal{F}(\mathcal{B})$ and $\mathcal{F}(\mathcal{B}')$, there are some additional considerations to account for before finding numerical results. First, coherent

choices for m_{bi} and μ_b need to be determined. What these should be is not clear as they don't bound these explicitly and/or consistently clinically. Also, how these bounds should be incorporated into the algorithm's stopping criteria is still an open question. Lastly, the piecewise quadratic objectives used in our application may be ill-natured for these types of bounds. Voxels far away from the threshold, t_j , can produce high-valued gradients, which can, in turn, generate grossly loose bounds. Linear objectives might provide a tighter bound, but that has not yet been investigated

However, we believe that there are interesting uses for the theoretical work in Section 3.5, especially in an integrated BOO and DAO model. In each iteration of generating an aperture, the model will assess whether or not to add an aperture to an existing beam location or add an additional beam location. These types of bounding, which we believe will be more effective in the aperture model due to the number of apertures being less than the number of beamlets, can be integrated into the beam selection and stopping conditions of this type of model.

3.6 Clinical Applicability and Future Research

3.6.1 Potential Clinical Implementation

Clinically, our integrated BOO and FMO model can be added into a treatment planning system. The major downside is that the data required for the integrated model are the D_{bij} values for each beam, which can be slow to calculate and expensive to store. There are several ways to utilize the methods while sidestepping this issue.

One way to use this methodology is to use the integrated BOO and FMO model only for the beam selection aspects of treatment planning and not for the end-of-planning FMO. This would mean that the beam locations chosen by the integrated model, \mathcal{B}' , would be the fixed beam set used by treatment planners and physicians to tweak before delivery. The presented methodology, in this respect, could be automated. It could also use downsampled data. We did not investigate the effects of downsampling, but believe that this methodology would continue to perform well with lower resolution voxels.

Another useful way to utilize the greedy framework is in considering when to add an additional beam. Consider a treatment in the midst of the planning process with 8 active beams that is unable to achieve a clinically acceptable dose distribution. Adding an additional beam will add flexibility to the model. The developed methodology can be used to select the next beam to add to the model, even though the beam selection algorithm might not have been used for the first 8. Estimates of the potential benefits of adding the additional beam can also be observed.

Lastly, this work was developed when also working on FusionArc (see *Matuszak et al.*, 2013), a hybrid delivery method that combines the continuous arc-based delivery of VMAT with the stationary delivery of IMRT. A modified BOO-like algorithm was used to determine where along an existing VMAT arc the gantry should stop and deliver an IMRT beam.

3.6.2 Future Work

Much of the groundwork for the BOO and DAO integrated model has been developed, and we hope to study this model in the future. This type of constructive methodology is also relevant to methods where beam locations of importance are iteratively selected (e.g., coplanar VMAT, FusionArc, Non-coplanar VMAT), which will be thematically evident in chapter IV.

3.7 Acknowledgements

This work was done in collaboration with Dr. Marina Epelman, Dr. Martha Matuszak, and Dr. Edwin Romeijn. We would like to thank Dr. Dan Ruan, Dr. Ke Sheng, and Dr. Peng Dong for their contributions in the infancy of this project as well as their help in collaborative publications (see *Dong et al.*, 2013a,b, 2014; *Nguyen et al.*, 2014). We would also like to thank Dr. Martha Matuszak for her help in generating data in the University of Michigan treatment planning system.

CHAPTER IV

Non-coplanar Volumetric Modulated Arc Therapy

4.1 Introduction

The dose distribution a patient receives is dependent upon the machine constraints of the delivery modality. Beams are fitted with an MLC, a device that, using sliding tungsten leaves, can cast the beam's output into shapes called apertures. The combination of aperture modulation, dose rate, and beam orientation and/or movement define the different methods of radiation delivery. Recall that there are several different treatment modalities that are commonly used in modern radiation therapy: IMRT and VMAT. Both of these techniques utilize a beam of radiation that is mounted on a moveable gantry that can rotate about the body. With the ability to move the device to which the patient is fixed, called the couch, non-coplanar beam directions can be achieved. In VMAT, the beam moves along a coplanar arc (i.e., fixed couch position with a moving gantry) while delivering radiation through actively changing apertures. The decisions that define a VMAT treatment plan are the coplanar arc path, the movements of the MLC along the arc, and the dynamic dose rate of the beam. When making these decisions, we must consider machine constraints such as MLC leaf movement restrictions, dose rate limits, gantry speed constraints, and gantry acceleration bounds. VMAT treatment planning is currently a hot topic among optimizers studying treatment planning due to the difficult nature of the problem and the lack of a standard treatment planning method (see, e.g., *Peng et al.*, 2012; *Craft et al.*, 2012; *Papp and Unkelbach*, 2014; *Unkelbach et al.*, 2015).

Coplanar VMAT takes advantage of the majority of the capabilities of the treatment machine and thus is seen as one of the most state-of-the-art treatment modalities. However, as treatment machines become more versatile in gantry and couch movement, the assumption that VMAT arcs must be coplanar can be relaxed. Nearly all current VMAT modeling techniques assume a pre-determined coplanar arc as an input to the inverse optimization for the other decisions. We would like to consider a treatment modality where gantry and couch

movement, dose rate, and MLC leaf movement are all dynamic *without the coplanar restriction applied to the beam's arc*. That is, determining the non-coplanar path of the arc through the $4\text{-}\pi$ space around the patient would be combined with the previously mentioned VMAT modeling restrictions (i.e., integrating a routing problem with an already-difficult treatment planning problem). This creates a mathematically challenging problem where solving to optimality is clinically infeasible. At a workshop at Massachusetts General Hospital where the top researchers in VMAT treatment plan optimization discussed the top methodologies and considered the future of VMAT, non-coplanar VMAT emerged as one of the most promising next-steps for cutting-edge algorithm research. An overview of modern optimization algorithms for VMAT can be found in *Unkelbach et al. (2015)*. Other modalities may approach delivery angles in a more flexible manner, such as a robotic arm, but these modalities are not nearly as widespread.

4.2 Treatment Planning Models

The VMAT optimization model requires us to introduce new notation and a few new delivery concepts. We first look at the coplanar VMAT treatment planning model, then study the non-coplanar counterpart.

In coplanar VMAT, the beam arc is predefined and modeled as a series of ordered control points. The set of decisions to make are the MLC leaf positions modeled as apertures at each control point, the dose rate through each control point, and the gantry speed through each control point. The gantry speed can also be looked at as the time spent traveling between control points and will be referred to as such for the remainder of the chapter. MLC leaf movement constraints must also be applied to allow adequate time between control points for MLC leaves to shift. Lastly, we assume that the gantry and couch can accelerate and decelerate instantaneously. This is not a realistic assumption when dealing with the standard delivery machinery. It should be noted that the delivery machines used for IMRT are also used for VMAT. However, this assumption is more realistic for flexible delivery modalities like the CyberKnife M6 FIM System. In the future work of the project, we would like to incorporate acceleration/deceleration constraints in the model, but for our study they are omitted.

The notation in this chapter follows that presented in Section 1.2.2 with a few additions. In VMAT, the path of the beam is often modeled as passing through a series of locations around the patient called control points. Let \mathcal{K} be the set of all control points and let k_ℓ be the ℓ^{th} control point along the arc. For decision variables, let A_{k_ℓ} be the aperture delivered at k_ℓ , r_{k_ℓ} be the dose rate at k_ℓ , t_{k_ℓ} be the time spent per degree traveling through control

point k_ℓ (i.e., inverse of the gantry speed), and y_{k_ℓ} be the intensity per degree at control point k_ℓ . Let m be the total number of control points in the delivery arc.

For parameters, let $R_{k_\ell}^U$ be the upper bound on dose rate for control point k_ℓ in dose per time, $T_{k_\ell}^L$ and $T_{k_\ell}^U$ be the lower and upper bounds, respectively, on time per degree at each control point k_ℓ , δ_{k_ℓ} be the distance in degrees at control point k_ℓ , and $\mathcal{D}_{k_\ell j}(A_{k_\ell})$ be dose received by voxel $j \in \mathcal{V}$ from aperture A_{k_ℓ} at unit intensity. Lastly, let $T_{k_\ell, k_{\ell+1}}^L(A_{k_\ell}, A_{k_{\ell+1}})$ be minimum time the gantry can take traveling from k_ℓ to $k_{\ell+1}$ and still allow for the MLC leaf changes necessary to change from A_{k_ℓ} to $A_{k_{\ell+1}}$ between control points k_ℓ and $k_{\ell+1}$.

4.2.1 Traditional VMAT Model

With the above definitions, we present the traditional VMAT model with a fixed path, (COP-VMAT-FP). In the clinical, this uses a coplanar arc.

$$\text{(COP-VMAT-FP)} \quad \underset{(A, y, t, r, z)}{\text{minimize}} \quad F(\mathbf{z}) \quad (4.1)$$

$$\text{subject to } z_j = \sum_{\ell=1}^m \mathcal{D}_{k_\ell j}(A_{k_\ell}) \delta_{k_\ell} y_{k_\ell} \quad j \in \mathcal{V} \quad (4.2)$$

$$y_{k_\ell} = r_{k_\ell} t_{k_\ell} \quad \ell = 1, \dots, m \quad (4.3)$$

$$\delta_{k_\ell} t_{k_\ell} \in [T_{k_\ell}^L, T_{k_\ell}^U] \quad \ell = 1, \dots, m \quad (4.4)$$

$$r_{k_\ell} \in [0, R_{k_\ell}^U] \quad \ell = 1, \dots, m \quad (4.5)$$

$$t_{k_\ell} \geq T_{k_\ell, k_{\ell+1}}^L(A_{k_\ell}, A_{k_{\ell+1}}) \quad \ell = 1, \dots, m \quad (4.6)$$

$$A_{k_\ell} \in \mathcal{A} \quad \ell = 1, \dots, m \quad (4.7)$$

Constraint (4.2) relates aperture shape and angular fluence rate to dose received. Constraint (4.3) combines time per degree at a control point and dose rate to create fluence. Constraints (4.4) and (4.5) bound variables t and r . Constraint (4.7) ensures apertures are deliverable and constraint (4.6) controls the leaf movement restrictions based on gantry movement and adjacent apertures.

4.2.2 Non-coplanar VMAT Full Problem

We extend the traditional VMAT model by relaxing the assumptions that the arc is predetermined and that the arc is restricted to a coplanar path. Because of these relaxations, k_ℓ is now a decision variable representing the ℓ^{th} control point in the arc. Let \mathcal{K} be the set of all potential control points in the 4π space around the patient. In order to construct a path,

we add a constraint to the model requiring $k_{\ell+1}$ to be adjacent to k_ℓ . Let \mathcal{K}_{k_ℓ} be the set of control points adjacent to k_ℓ , with adjacency defined as being within one unit of angular discretization in parameter space.

We present the overall non-coplanar (NCP) VMAT full problem (FP) below. We would like to select a subset of the control points that form a path to minimize some function of the dose while adhering to deliverability constraints. While this formulation of the model may be unsolvable in a realistic amount of time, it may help when finding a heuristic solution.

$$(FP) \quad \underset{k, A, y, t, r, z}{\text{minimize}} \quad F(\mathbf{z}) \quad (4.8)$$

$$\text{subject to } z_j = \sum_{\ell=1}^m \mathcal{D}_{k_\ell j}(A_{k_\ell}) \delta_{k_\ell} y_{k_\ell} \quad j \in \mathcal{V} \quad (4.9)$$

$$y_{k_\ell} = r_{k_\ell} t_{k_\ell} \quad \ell = 1, \dots, m \quad (4.10)$$

$$\delta_{k_\ell} t_{k_\ell} \in [T_{k_\ell}^L, T_{k_\ell}^U] \quad \ell = 1, \dots, m \quad (4.11)$$

$$r_{k_\ell} \in [0, R_{k_\ell}^U] \quad \ell = 1, \dots, m \quad (4.12)$$

$$t_{k_\ell} \geq T_{k_\ell, k_{\ell+1}}^L(A_{k_\ell}, A_{k_{\ell+1}}) \quad \ell = 1, \dots, m \quad (4.13)$$

$$A_{k_\ell} \in \mathcal{A} \quad \ell = 1, \dots, m \quad (4.14)$$

$$k_{\ell+1} \in \mathcal{K}_{k_\ell} \quad \ell = 1, \dots, m \quad (4.15)$$

As stated earlier, k_ℓ represents the ℓ th control point along the beam path, with the path of the gantry passing through exactly m control points. Constraint (4.9) represents the dose delivered to voxel j from apertures at active control points k_1 to k_m . Constraint (4.10) calculates the aperture intensity at control point k_ℓ . Constraints (4.11) and (4.12) apply upper and lower bounds on absolute (non-directional) gantry travel time and dose rate, respectively. Constraint (4.13) ensures sufficient gantry travel time to maintain deliverability of adjacent apertures A_{k_ℓ} and $A_{k_{\ell+1}}$. Function $T_{k_\ell, k_{\ell+1}}^L(A_{k_\ell}, A_{k_{\ell+1}})$ represents the minimum time the gantry can take traveling from k_ℓ to $k_{\ell+1}$ and still allow for the MLC leaf changes necessary to change from A_{k_ℓ} to $A_{k_{\ell+1}}$ between control points k_ℓ and $k_{\ell+1}$. Constraint (4.14) ensures apertures are deliverable shapes. Constraint (4.15) requires adjacent apertures in the path. Dummy apertures are added to the beginning and end of the path $\ell = 0, m + 1$. These dummy apertures have no MLC movement restrictions or dose rate and serve only as placeholders for the pricing problem.

Using a representation of travel time (rather than speed) will help us out later when constructing a feasible solution. Travel time can be converted back to speed after the model has been solved. For simplicity, we assume that the same speed restrictions apply to gantry

and couch movement, but separate speed restricts can be applied to the gantry and couch in an extension of our model.

4.2.3 Master Problem

Solving the FP outright is computationally intractable, so we investigate ways to simplify the solution process. We start by assuming the beam travels at its slowest speed to remove a decision from the model. A post-processing step can be performed to raise speed as much as possible given a feasible solution. The travel time is fixed to the control-point maximum (i.e., movement speed to the control-point minimum), $T_{k_\ell}^U$ (see Section 2.1 in *Peng et al.*, 2012 for the speed-centered version of this model). The resulting upper bound on fluence rate of aperture y_{k_ℓ} is $\delta_{k_\ell} R_{k_\ell}^U T_{k_\ell}^U$, which we denote with $Y_{k_\ell}^U$. We get the following Master Problem (MP). A feasible solution to this problem implies a feasible solution to the FP.

$$\begin{aligned}
 \text{(MP)} \quad & \underset{k, A, y, z}{\text{minimize}} && F(\mathbf{z}) \\
 & \text{subject to} && z_j = \sum_{\ell=1}^m \mathcal{D}_{k_\ell j}(A_{k_\ell}) \delta_{k_\ell} y_{k_\ell} && j \in \mathcal{V} \\
 & && y_{k_\ell} \in [0, Y_{k_\ell}^U] && \ell = 1, \dots, m && (4.16) \\
 & && \delta_{k_\ell} T_{k_\ell}^U \geq T_{k_\ell, k_{\ell+1}}^L(A_{k_\ell}, A_{k_{\ell+1}}) && \ell = 1, \dots, m && (4.17) \\
 & && A_{k_\ell} \in \mathcal{A} && \ell = 1, \dots, m \\
 & && k_{\ell+1} \in \mathcal{K}_{k_\ell} && \ell = 1, \dots, m
 \end{aligned}$$

4.3 Heuristic Methodology for NCP VMAT

What we would like to do with our methodology is to generate a feasible, high-quality solution to the FP. Based on the success of the column-generation framework developed in *Peng et al.*, 2012, we consider a similar framework for this problem. The general idea is to start with all k_ℓ control point locations unassigned and $y_{k_\ell} = 0$. Setting these decision variables forms a RMP (see Section 4.3.1). We then solve the RMP to generate a deliverable, but incomplete, treatment plan. Iteratively, we attempt to improve the current plan by generating an aperture at an inactive control point that guarantees that we will have a feasible solution when the algorithm terminates. That is, when m control points and apertures have been assigned, the solution is feasible for (MP) with a connected path of control points with apertures that are deliverable and reachable given the gantry travel time. The process by which we select a beneficial aperture is the pricing problem (PP), and the PP will be discussed in Section 4.3.2. After adding an aperture to the RMP, we solve the updated RMP.

The solution to the RMP feeds into the next iteration's PP. We continue this procedure until m apertures are added.

4.3.1 Restricted Master Problem

Let us define the RMP given some control point set $\mathcal{K}' \subset \mathcal{K}$. In $\text{RMP}(\mathcal{K}')$, we fix a set of $|\mathcal{K}'| = m' < m$ control points $k_\ell \in \mathcal{K}'$ for $\ell = 1, \dots, m'$ and their associated apertures, \bar{A}_{k_ℓ} . The following program solves for the associated aperture intensities with fixed control points and aperture shapes.

$$\begin{aligned}
(\text{RMP}(\mathcal{K}')) \quad & \underset{y, z}{\text{minimize}} && F(\mathbf{z}) \\
& \text{subject to} && z_j = \sum_{\ell=1}^{m'} \mathcal{D}_{k_\ell j}(\bar{A}_{k_\ell}) \delta_{k_\ell} y_{k_\ell} && j \in \mathcal{V} \\
& && y_{k_\ell} \in [0, Y_{k_\ell}^U] && \ell = 1, \dots, m'
\end{aligned}$$

This program is an optimization problem with linear constraints, and thus it is a convex optimization problem if F is a convex function. Using convex solvers, we can efficiently solve this problem for manageable m' .

4.3.2 Pricing Problem

The overall idea for the PP is to select a control point and aperture that provide the largest immediate benefit to the model while maintaining feasibility of MP (i.e. satisfies constraints (4.14), (4.15), and (4.17)).

Suppose an aperture $A_{\bar{k}}$ is inserted between control points k_ℓ and $k_{\ell+1}$. Let us define $\pi_{k_\ell, k_{\ell+1}}^{\bar{k}}(A_{\bar{k}})$ as the rate of improvement in the objective function with the addition of this aperture based on first-order optimality conditions. $A_{\bar{k}} \in \mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$, where $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}} \subseteq \mathcal{A}$ is the set of deliverable apertures that can feasibly be added at control point \bar{k} to the plan given preceding aperture A_{k_ℓ} and succeeding aperture $A_{k_{\ell+1}}$ at control points k_ℓ and $k_{\ell+1}$, respectively. This value can be calculated as in equation (4.18) given optimal solution $(\bar{\mathbf{y}}, \bar{\mathbf{z}})$ to $\text{RMP}(\mathcal{K}')$.

$$\pi_{k_\ell, k_{\ell+1}}^{\bar{k}}(A_{\bar{k}}) = \sum_{j \in \mathcal{V}} [-\nabla F(\bar{\mathbf{z}})(\bar{\mathbf{z}})]_j \mathcal{D}_{\bar{k}j}(A_{\bar{k}}) \delta_{\bar{k}} \quad (4.18)$$

Our strategy is to select a control point \bar{k} and aperture $A_{\bar{k}}$ to fit in the path between preceding aperture k_ℓ and succeeding aperture $k_{\ell+1}$. Let $P(k_\ell, k_{\ell+1})$ be the set of potential control points to consider between existing control points k_ℓ and $k_{\ell+1}$ (see Section 4.3.2.1 for

more detail). PP can be broadly described as below in equation (4.19).

$$(PP) \quad \max_{\ell=0,\dots,m'} \max_{\bar{k} \in P(k_\ell, k_{\ell+1})} \max_{A_{\bar{k}} \in \mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}} \pi_{k_\ell, k_{\ell+1}}^{\bar{k}}(A_{\bar{k}}) \quad (4.19)$$

At the ends, a search is done between the end control point and some dummy control point that has no MLC leaf restriction. Section 4.3.2.1 goes into defining the search neighborhood $P(k_\ell, k_{\ell+1})$. Section 4.3.2.2 details the formulation and effect of $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$. Section 4.3.3 describes the PP solution procedure. Lastly, Section 4.3.3.1 discusses tie-breaking strategies.

4.3.2.1 Determining $P(k_\ell, k_{\ell+1})$

$P(k_\ell, k_{\ell+1})$ is the control point search space between given control points k_ℓ and $k_{\ell+1}$. The overall idea in specifying $P(k_\ell, k_{\ell+1})$ is to do two things: (i) ensure that there are always enough remaining control points to complete the m -length path and (ii) restrict the search space in between existing path control points to hedge against being overly greedy.

First, we consider (i). Since control points k_ℓ and $k_{\ell+1}$ may not be adjacent, the gantry would need to travel through a certain number of intermediate control points to reach $k_{\ell+1}$ from k_ℓ . Let $E(k_\ell, k_{\ell+1})$ be the number of control points necessary to travel between control points k_ℓ and $k_{\ell+1}$ along a minimum-time path (defined below) consisting of adjacent control points. (If k_ℓ and $k_{\ell+1}$ are adjacent, then $E(k_\ell, k_{\ell+1}) = 0$.)

Suppose \mathcal{K}' , with $|\mathcal{K}'| = m'$ has been specified at the current step of the constructive algorithm. Let $J(\mathcal{K}')$ be the number of “not immediately necessary” control points, i.e., control points that are not “used up” to connect points in \mathcal{K}' with a sequence of adjacent control points. We can compute $J(\mathcal{K}')$ as the following:

$$J(\mathcal{K}') = m - m' - \sum_{\ell=1}^{m'-1} E(k_\ell, k_{\ell+1}). \quad (4.20)$$

Given this metric and aperture set $k_\ell \in \mathcal{K}'$ for $\ell = 1, \dots, m'$, we can make the following claims:

- The longest feasible path between two control points k_ℓ and $k_{\ell+1}$ is bounded by $E(k_\ell, k_{\ell+1}) + J(\mathcal{K}')$.
- The longest feasible path beyond an endpoint of the path is bounded by $J(\mathcal{K}')$.

Next, we must incorporate (ii) to restrict our search space in between path control points. One method of doing this is to consider the ratio in travel time between the path $\ell \rightarrow \bar{k} \rightarrow \ell + 1$ and the path $\ell \rightarrow \ell + 1$. Let us define $T_t(k_\ell, k_{\ell+1})$ as the shortest time it takes the beam

to travel from k_ℓ to $k_{\ell+1}$ given fixed angular travel time t and control points k_ℓ and $k_{\ell+1}$. Then, let \bar{T} be some upper bound on this ratio. We would like to consider control points \bar{k} that satisfy the following ratio condition.

$$\frac{T_t(k_\ell, \bar{k}) + T_t(\bar{k}, k_{\ell+1})}{T_t(k_\ell, k_{\ell+1})} \leq \bar{T} \quad (4.21)$$

We define our set of control points to consider, $P(k_\ell, k_{\ell+1})$, given a t , k_ℓ and $k_{\ell+1}$.

$$P(k_\ell, k_{\ell+1}) = \left\{ \bar{k} : \frac{T_t(k_\ell, \bar{k}) + T_t(\bar{k}, k_{\ell+1})}{T_t(k_\ell, k_{\ell+1})} \leq \bar{T}, E(k_\ell, \bar{k}) + E(\bar{k}, k_{\ell+1}) \leq J(\mathcal{K}') + 1 \right\} \quad (4.22)$$

In order to calculate function T_t , let us make some additional assumptions. Assume that the set of potential control point locations is equally spaced in parameter space in a grid of gantry and couch orientations. We assume this for the proof-of-concept nature of the project, but this assumption can be relaxed with updated T_t formulas. We can calculate this time metric given an assumption on beam movement.

Let g and c be the number of control points apart k_ℓ and $k_{\ell+1}$ are on the gantry axis and couch axis, respectively. For example, see Figure 4.1. Let δ be the number of degrees in spacing at each control point. Lastly, let us assume that the couch and gantry can move simultaneously. We use a path heuristic in which the beam travels along a diagonal until it can move in a parameter-space-orthogonal path to the destination.

$$T_t(k_\ell, k_{\ell+1}) = \min(g, c)\delta t + (\max(g, c) - \min(g, c))\delta t \quad (4.23)$$

With these assumptions, calculating this expressing reduces to finding the max between g and c .

One questions that still remains is how to deal with the endpoints of the path. Clearly, we cannot extend past the endpoints further than $J(\mathcal{K}')$ control points. However, we may want to further restrict the search space from the endpoints. A simple bound, E_{end} , only considering points within $\min(J(\mathcal{K}'), E_{\text{end}})$ control points from the endpoint may be all that is necessary. This bound will affect how the algorithm grows the path. We could also look at the time, T_t , to the control points past the end and restrict the search using that value. Lastly, it may be logical to consider adding a condition that the beam cannot go back on itself sharply from the endpoints to promote path movement around the body. While we do not add this bound, a more in depth study of how the algorithm behaves will need to be done before any definitive “endpoint rules” can be finalized.




Figure 4.1: Illustration of $g = 2$ and $c = 1$ values between two red control points

4.3.2.2 Determining $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$ and Its Implications

We need to specify set $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$, which is our aperture search space, in a particular way to satisfy constraint 4.17 in MP. This constraint says that MLC leaves must be able to feasibly move in the time between control points k_ℓ , \bar{k} , and $k_{\ell+1}$. Note that \bar{k} need not be adjacent to the other two.

$$\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}} \subseteq \{A \in \mathcal{A} : \delta_{k_\ell} T_{k_\ell}^U \geq T_{k_\ell, \bar{k}}^L(A_{k_\ell}, A), \delta_{k_\ell} T_{\bar{k}}^U \geq T_{\bar{k}, k_{\ell+1}}^L(A, A_{k_{\ell+1}})\} \equiv \mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}U} \quad (4.24)$$

How we choose $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$ will define the MLC capabilities at some control point \bar{k} . As we can see in equation (4.24), the definition of $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$ is dependent upon the preceding and following apertures and the upper bound on the time spent at the control point. With the preceding and following apertures fixed, our concern is with the upper bound on the time spent between control points. Keeping the time at the machine upper bound could allow for a larger $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$, i.e., a larger search space for deliverable apertures, but may reduce the flexibility of the model in future iterations.

One thing to note is that the choice of $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$ will influence the behavior of the algorithm, especially in later stages. If we choose a smaller set $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$, that will result in less flexibility in the current set of apertures, but may allow for greater flexibility in later stages. We could consider a family of potential choices for $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$ (see Appendix A). However, let us continue with our assumption of a fixed travel time. This information will be used when determining the potential MLC leaf arrangements when solving the pricing problem.

Lastly, we want to know explicitly how our selected max angular travel time, $T_{\bar{k}}^U$, and preceding and following apertures A_{k_ℓ} and $A_{k_{\ell+1}}$ control the feasible aperture region. That is, for each row, we would like to determine the feasible MLC leaf positions given $T_{\bar{k}}^U$, k_ℓ , and $k_{\ell+1}$ for potential control point \bar{k} .

Let us assume that $T_{\bar{k}}^U$, k_ℓ , and $k_{\ell+1}$ are fixed along with apertures A_{k_ℓ} , and $A_{k_{\ell+1}}$. Let $A_{\bar{k}r}$ be the set of feasible rows for the r th row of control point \bar{k} of aperture $A_{\bar{k}}$.

$$A_{\bar{k}r} = \left\{ a : T_{\bar{k}}^U \delta_{k_\ell} \geq T_{k_\ell, \bar{k}}^L(A_{k_\ell r}, A_{\bar{k}r}), T_{\bar{k}}^U \delta_{\bar{k}} \geq T_{\bar{k}, k_{\ell+1}}^L(A_{\bar{k}r}, A_{k_{\ell+1}r}) \right\} \quad (4.25)$$

Let v be the maximum speed an MLC leaf can move and N be the rightmost leaf setting. The total distance a leaf can cover is v times the time it is moving. We can see the following relationship.

$$\left\{ a : T_{\bar{k}}^U \delta_{k_\ell} \geq T_{k_\ell, \bar{k}}^L(A_{k_\ell r}, A_{\bar{k}r}) \right\} = \{(L, R) : 0 \leq L \leq R \leq N; \\ |L - L_{k_\ell}|, |R - R_{k_\ell}| \leq vT_t(k_\ell, \bar{k})\}, \quad (4.26)$$

and we can define the leaf positions as the following:

$$A_{\bar{k}r} = \{(L, R) : 0 \leq L \leq R \leq N; |L - L_{k_\ell}|, |R - R_{k_\ell}| \leq vT_t(k_\ell, \bar{k}); \\ |L - L_{k_{\ell+1}}|, |R - R_{k_{\ell+1}}| \leq vT_t(\bar{k}, k_{\ell+1})\}. \quad (4.27)$$

4.3.3 Solving the PP

Recall the pricing problem from equation (4.19).

$$(PP) \quad \max_{\ell=0, \dots, m'} \max_{\bar{k} \in P(k_\ell, k_{\ell+1})} \max_{A_{\bar{k}} \in \mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}} \pi_{k_\ell, k_{\ell+1}}^{\bar{k}}(A_{\bar{k}})$$

For each ℓ, \bar{k} , we want to solve the following:

$$(PP(\ell, \bar{k})) \quad \max_{A_{\bar{k}} \in \mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}} \pi_{k_\ell, k_{\ell+1}}^{\bar{k}}(A_{\bar{k}}). \quad (4.28)$$

For a given ℓ and \bar{k} , we can calculate (4.28) in the following manner for a beam with M rows (see *Peng et al.*, 2012).

$$\max_{A_{\bar{k}} \in \mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}} \pi_{k_\ell, k_{\ell+1}}^{\bar{k}}(A_{\bar{k}}) = \sum_{r=1}^M \max_{(L, R) \in A_{\bar{k}r}} \int_L^R \pi_{k_\ell, k_{\ell+1}}^{\bar{k}}(x) dx \quad (4.29)$$

We integrate across the leaf positions in order to allow fractional (i.e., non-discrete) leaf positions. Details can be found in *Peng et al.* (2012), Section 2.6.1.

The overall pricing problem becomes the following problem, over which we can enumerate due to the fact that leaf positions can only be at integer points or the leaf limit.

$$(PP) \quad \max_{\ell=0, \dots, m'} \max_{\bar{k} \in P(k_\ell, k_{\ell+1})} \sum_{r=1}^M \max_{(L, R) \in A_{\bar{k}r}} \int_L^R \pi_{k_\ell, k_{\ell+1}}^{\bar{k}}(x) dx \quad (4.30)$$

We represent aperture dose in the same manner as in *Peng et al.* (2012), and thus will not be described in detail here. In short, the continuous row of beamlet positions is discretized into N beamlets. Dose is the summation of the open beamlets plus the fraction of the partially-blocked beamlets at the locations occupied by the MLC leaves. The function we are integrating over is a step function.

4.3.3.1 Breaking Ties in the PP

It may be possible that, for a particular control point \bar{k} , the pricing problem produces the same $(PP(\ell, \bar{k}))$ values for different ℓ . In this situation, we need rules to determine where to insert control point \bar{k} and associated aperture in the path. The list below shows some methods of how to break ties in order of execution. In our implementation, we go by the pricing problem score (4.30), then the lowest ratio (i), then the most central (iii).

- i Insert the control point where the $\frac{T_t(k_\ell, \bar{k}) + T_t(\bar{k}, k_{\ell+1})}{T_t(k_\ell, k_{\ell+1})}$ ratio is the smallest
- ii Insert the control point such that $T_t(k_\ell, \bar{k}) + T_t(\bar{k}, k_{\ell+1})$ is the smallest
- iii Insert the control point so that it is most central (i.e., $\frac{T_t(k_\ell, \bar{k})}{T_t(k_\ell, k_{\ell+1})}$ is closest to 1)
- iv Insert the control point into the path with the smallest ℓ (to definitively break all ties)

4.4 Application

4.4.1 Brain Case Details

This model was applied to a brain case as a proof-of-concept. The data is from common optimization for radiation therapy (CORT) dataset in *Craft et al. (2014)*. Potential control point set, \mathcal{K} , consists of control points spaced 10 degrees apart along both the gantry and couch axes totaling 648 control points. Each control point contains 160 beamlets, and the patient was discretized into 481,915 voxels. The prescription dose for the case is 60 Gy. The dose-to-points data were generated in the open source Computational Environment for Radiotherapy Research (CERR) (see *Deasy et al., 2003*).

This case was chosen primarily for the availability of data. For this type of algorithm, dose-to-points matrices, D , must be generated for each potential control point. While this may be avoided in a clinical implementation of this algorithm with additional path restrictions, considering the full \mathcal{K} is necessary in the presented form of the algorithm. The algorithm was implemented in C++ using Gurobi (see *Bixby et al., 2010*) to solve the RMP.

4.4.2 Model Assumptions

Several assumptions were made for this initial implementation. $P(k_\ell, k_{\ell+1})$, the set of control points to search between k_ℓ and $k_{\ell+1}$, has an additional restriction of not allowing more than \bar{E} control points between k_ℓ , selected point c , and $k_{\ell+1}$. We also restricted the search at the endpoints of the path to at most \bar{E}_{end} control points. The values we tested can be seen in Table 4.1. We found that \bar{E} was too influential and opted for finding a value in future test. $P(k_\ell, k_{\ell+1})$ was constructed in the following way:

$$P(k_\ell, k_{\ell+1}) = \left\{ \bar{k} : \frac{T_t(k_\ell, \bar{k}) + T_t(\bar{k}, k_{\ell+1})}{T_t(k_\ell, k_{\ell+1})} \leq \bar{T}, \right. \\ \left. E(k_\ell, \bar{k}) + E(\bar{k}, k_{\ell+1}) \leq \min(J(\mathcal{K}'), \bar{E}) + 1 \right\}. \quad (4.31)$$

A number of assumptions were made on how the gantry and couch move. As mentioned before, these assumptions may not hold with conventional IMRT delivery systems, but are more realistic for some other systems, e.g., CyberKnife M6 FIM Systems. We assume that the beam can feasibly deliver from all 4π angles. Ideally, we want to restrict the path to only visit feasible control points (i.e., avoid collisions between the gantry and the couch/patient), and that change fits within our framework. However, the collision information was not available at the time of this project. We assume that the gantry and couch move in equispaced parameter space at the same speed. This speed is not decoupled in the current formulation

Table 4.1: Parameters tested over for the initial set of runs

m	$\overline{E}_{\text{end}}$	\overline{E}	\overline{T}
36	4	4	2.5
36	4	8	2.5
36	4	100	2.5
36	4	100	3.5
36	8	4	2.5
36	8	8	2.5
36	8	100	2.5
36	8	100	3.5
36	10	100	2.5
36	10	100	4

of the model. Further studies must be done to determine these changes.

We assume that the couch and gantry can move simultaneously. Movement is assumed to begin and end instantaneously (i.e., perfect acceleration/deceleration). We also assume that the beam has no upper fluence bound.

The leaf movement speed upper bound is 4.5 cm per second and the gantry and couch speed upper bound is 6 degrees per second. The parameters $\overline{E}_{\text{end}}$, \overline{E} , and \overline{T} were tested to get ballpark estimates of what good values of these should be. A single parameter set, bolded in Table 4.1 showing those tested, was used for the runs in the results section.

4.4.3 Results

With this proof-of-concept, we will examine treatment plan quality through dose-volume histograms (DVHs, see 3.4.2 for DVH details) and dosewashes. These metrics are commonly used for clinical assessment of the quality of treatment plans.

First, consider a 36 control point solution. We are interested in how well our algorithm preforms compared with the state-of-the-art coplanar VMAT algorithm presented in *Peng et al.* (2012). In Figure 4.2, we see the DVH comparing the 36 control point NCP VMAT (solid) solution to the 36 control point coplanar VMAT (dashed) solution. The NCP VMAT solution dominates the coplanar VMAT solution almost everywhere. This shows us that considering non-coplanar directions for VMAT treatment can potentially yield higher-quality treatment plans than using only the traditional coplanar arcs. However, only using 36 control points heavily restricts our ability to take advantage of the flexibility of VMAT treatment plant (i.e., dynamically adjusting the MLC leaf positions over the arc). We can see in Figure 4.3 that we fail to get close to the “ideal” IMRT solution that uses 36 non-coplanar control points discovered by the NCP VMAT algorithm as beam directions. Therefore, we ask the

question “Can we do better?”




Figure 4.2: 36 control point NCP VMAT (solid) vs. coplanar VMAT (dashed)




Figure 4.3: 36 control point NCP VMAT (dashed) vs. IMRT (solid)

To improve our solution, we interpolate along the 36-point NCP VMAT path and add in four additional, equispaced control points in between the control points on the arc and four more on the end to get up to 180 control points. This is a finer discretization than we initially planned. An example of this can be seen in Figure 4.4. We then reoptimize the

treatment plan (i.e., determine aperture shapes and aperture intensities) using the algorithm in *Peng et al.* (2012) (again, without gantry acceleration/deceleration constraints) along a predetermined non-coplanar arc.

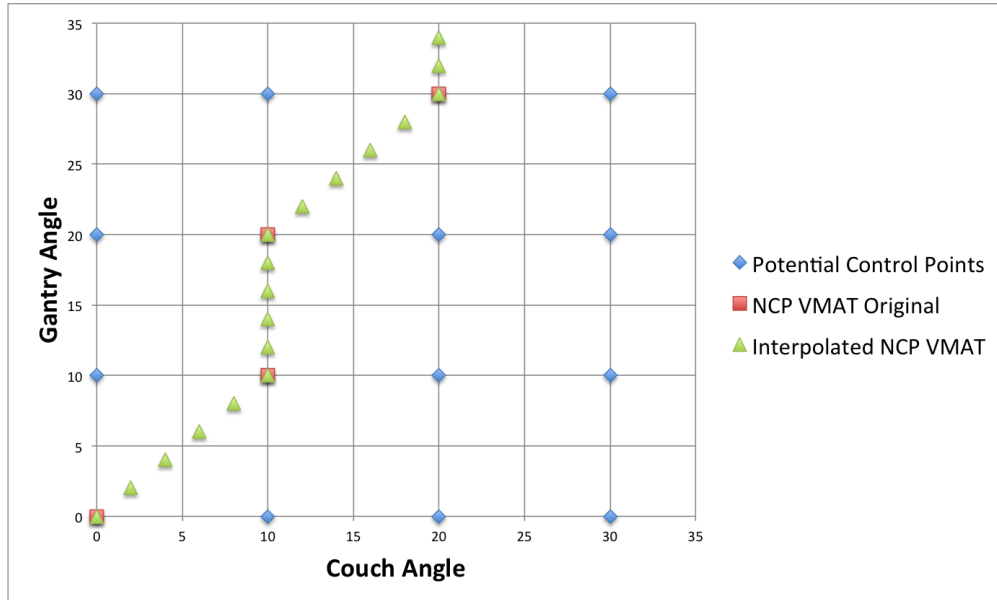


Figure 4.4: Illustration of potential control points (blue), 36-point NCP VMAT trajectory (red), and interpolated 180-point NCP VMAT trajectory (green)

The DVH showing this comparison can be seen in Figure 4.5. We see that a large improvement is made by allowing greater control over the MLC leaves. We are also interested in how far from an ideal plan this solution is. Consider optimizing an IMRT plan along the same 180 control points. The DVH for this solution and the 180 control point NCP VMAT along the 36-point NCP VMAT path can be seen in Figure 4.6. The 180 control point NCP VMAT is much closer to the ideal solution than the 36 control point NCP VMAT is to the 36 control point IMRT solution.

Lastly, we would like to compare our 180 control point NCP VMAT taken from extrapolating the 36 control point NCP VMAT path to a 180 control point coplanar VMAT treatment plan. These results can be seen in the DVH in Figure 4.7. While the improvements are not as significant as when considering 36 control point plans, the NCP VMAT plan dominates the coplanar VMAT plan nearly everywhere. It should also be noted that many of the parameter effects of the NCP VMAT model have not been studied, and the NCP VMAT solutions presented here should not be considered to be the “best” the technique can offer. The effects of allowing non-coplanar beam directions in VMAT treatment are more evident when considering dosewashes. In Figure 4.8, we can see how some of the dose in hotspots (areas of high dose in undesirable areas) is lowered by using NCP VMAT.




Figure 4.5: 180 control point VMAT along 36 control point NCP VMAT path (solid) vs. 36 control point NCP VMAT (dashed)




Figure 4.6: 180 control point VMAT along 36 control point NCP VMAT path (dashed) vs. 180 control point IMRT (solid)

4.5 Conclusions and Future Work

Preliminary results suggest that superior treatments can be obtained without significant increase in treatment time when using non-coplanar beams, but more extensive testing is needed. Acceleration and deceleration constraints, as they pertain to specific VMAT delivery

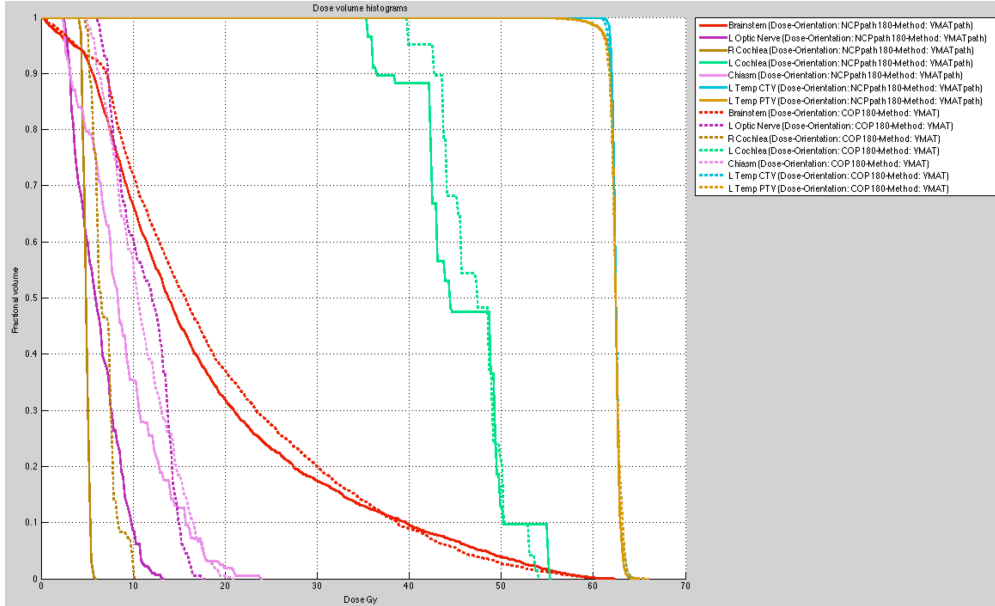


Figure 4.7: 180 control point VMAT along 36 control point NCP VMAT path (solid) vs. 180 control point coplanar VMAT (dashed)

modalities, should be implemented in the model. Gantry and couch speeds should also be decoupled, although the gantry/couch parameter space may not be the ideal way of describing beam location for other delivery modalities like the CyberKnife M6 FIM System.

On the clinical side, more cases should be tested with this planning framework. A full, 2-degree control point spacing should be optimized to see the effects of the algorithm on a high-resolution search space. However, the method of interpolating between control points on paths found using a coarse resolution is probably more clinically feasible. Fewer dose-to-points matrices would need to be generated and the solution time for the algorithm would be significantly faster. Lastly, a study in treatment time versus treatment plan quality could be done using this type of algorithm for non-coplanar VMAT treatment plans.

A final thought for future work is in using NCP VMAT to improve existing plans or partial plans. An incomplete path of control points can be used as a seed for this algorithmic framework provided they control points can be connect in the end. A treatment planner might identify a few “high priority” locations they want the beam to pass through, and the NCP VMAT algorithm could figure out the connecting path and associated apertures along the path. Techniques like BOO (see III) could also be used to identify these “high priority” control points. This is an exciting research area, and we hope to further study NCP VMAT models in the future.

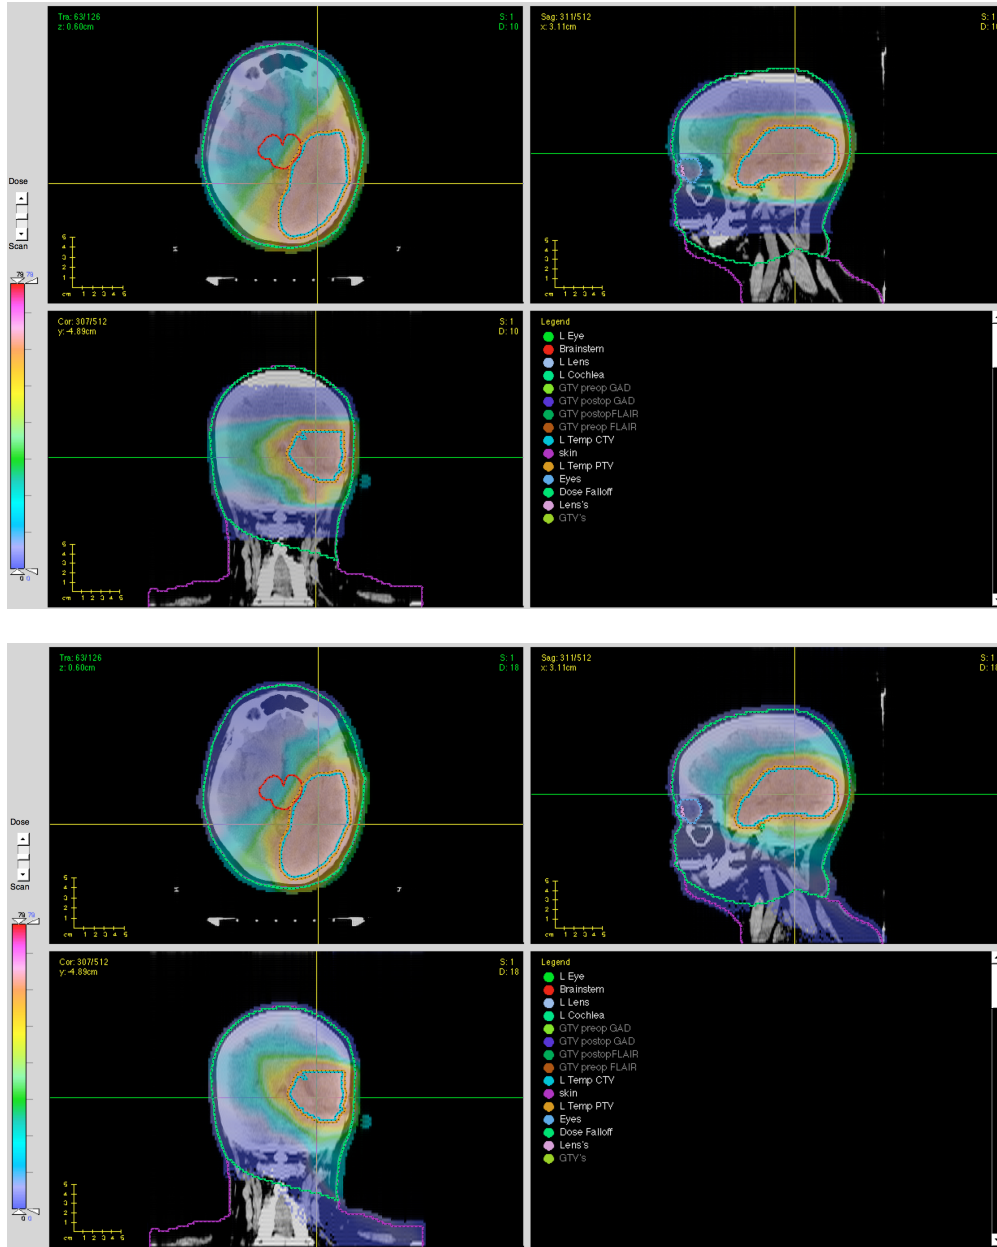


Figure 4.8: (top) 180 control point coplanar VMAT dosewash, (bottom) 180 control point VMAT along 36 control point NCP VMAT path dosewash

4.6 Acknowledgements

This work was done in collaboration with Dr. Thomas Bortfeld, Dr. David Craft, Dr. Edwin Romeijn, Dr. David Papp, and Dr. Jan Unkelbach. We would also like to thank Dr. Bram Gorissen for his insight in handling large data sets and Dr. Mark Bangert for his insight into non-coplanar pathfinding.

CHAPTER V

Adaptive Treatment Planning for Lung Cancer

5.1 Introduction

When planning a radiation therapy treatment, known patient information is used to generate a desirable treatment plan. Geometric information along with other pre-treatment characteristics influence the resulting plan. Over the course of treatment, these patient characteristics may change, and the planner’s initial understanding of patient information may evolve, causing inaccuracies in the original planning model. Due to these inaccuracies, the pre-treatment plan is often adapted to the new data by replanning and reoptimizing the treatment plan. By incorporating the anticipation and knowledge of future changes into the pre-treatment planning model, or “future-proofing” the treatment plan, planners may have greater control of patient outcomes and improve the effectiveness of the treatment.

5.1.1 Adaptation Motivation

In traditional radiation therapy treatment planning, the general practice is to develop a treatment plan and use that plan for the course of treatment. This plan is delivered over n daily fractions, where fluence $\frac{x}{n}$ (and thus dose $\frac{z}{n}$) is delivered in each fraction. This has several advantages. First, it is an established method and straightforward in practice. Second, it puts a relatively small strain on treatment planning resources (i.e., planner time, computation time, physician time, etc.). However, not all patient information is known a priori. The information describing a patient changes over the course of treatment, both geometrically and physiologically. Treatments are often adapted (i.e., replanned given the already delivered dose) when significant new information is realized.

Without considering adaptation when plans are initially generated pre-treatment, the dose distribution already delivered at the time of adaptation may force the adapted treatment to be inflexible in the amount it can adapt to the new information and still stay within the

treatment planning protocol. Therefore, more sophisticated methods of adaptive planning are sometimes considered (see *Ding et al.*, 2007; *Wu et al.*, 2008; *Nishi et al.*, 2013; *Hurkmans et al.*, 2012; *de la Zerda et al.*, 2007 for examples of adaptive treatment planning).

There are several reasons in addition to raising treatment quality that adaptation-conscious planning should be considered. With treatment planning, imaging, and biomarker analysis getting faster and more efficient (see *Breedveld et al.*, 2012; *Men et al.*, 2010; *Long et al.*, 2012, etc.), there are additional resources freed up for replanning. Dosemetrists are also starting to plan with replanning in mind because they are realizing that considering future-proofing initial plans can improve treatment. The quality of the treatment plan improves as information is effectively incorporated into the model.

5.2 Lung Cancer Setting

In this chapter, we discuss possible modeling approaches to adaptive treatment planning in radiation oncology. To make our discussion concrete, we use a specific cancer site (lung cancer) and a specific form of new information (a particular biomarker which, as indicated by preliminary studies, has the potential to be a good predictor of the likelihood of a certain radiation side effect). It should be noted that further studies, including data analysis and clinical validation of functions and probability estimates used in our models in this chapter, are ongoing. However, the types of optimization models we present for the adaptive planning framework would be applicable in any treatment environment that has similar stages, namely: make a pre-treatment plan, begin treatment, observe change in the patient or obtain additional patient-specific information, re-plan, and finish treatment. We expect that these treatment planning approaches will become increasingly applicable as mid-treatment biomarker measurements and other monitoring of each patient’s status during treatment become more prevalent.

The two main goals in lung cancer treatment are to eradicate the tumor and spare healthy organs. Towards the first goal, we consider maximizing the chance of no tumor progression after 2 years, which we refer to as the probability of local tumor control, P_{LTC} . There are numerous healthy organs to consider (e.g., spinal cord, esophagus, heart), but the lung itself necessarily receives radiation. Therefore, in our models we particularly focus on the adverse effect of radiation-induced lung toxicity (RILT), namely, the probability a patient will develop RILT of grade 2 or higher, which we denote by P_{RILT} .

Table 5.1: Parameter estimates for calculating the probability of local tumor control

Parameter	Estimate
$S_0(2)$	0.298
β_0	5.928
β_1	0.0988
a	0.813

5.2.1 Local Tumor Control

The relationship between dose distribution vector z and P_{LTC} can be characterized as follows: let $F(z)$ be D95 to the PTV (that is, 95% of the PTV receives dose of at least D95); then

$$P_{\text{LTC}}(z) = S_0(2)e^{\beta_0 - \beta_1(F(z))}, \quad (5.1)$$

with parameters $S_0(2)$, β_0 , and β_1 (see Table 5.1 for estimates of values of these parameters and *Cox*, 1972 for more information on a fitted Cox regression model used to obtain these estimates). However, due to the computational difficulties presented by incorporating the functional form of D95 into optimization problems, we will use the linearized EUD function as a proxy for D95 instead. That is, in our models we will use

$$F(z) = z_{\text{PTV}}^{\text{EUD}} \equiv az_{\text{PTV}}^{\min} + (1 - a)\bar{z}_{\text{PTV}}, \quad (5.2)$$

where z_{PTV}^{\min} and \bar{z}_{PTV} are the minimum and mean doses to the PTV, respectively, and a is a structure-specific parameter used in the definition of linearized EUD (*Thieke et al.*, 2002; see Table 5.1 for the value of a used). Note that, since P_{LTC} is increasing in $F(z)$, $z_{\text{PTV}}^{\text{EUD}}$ can be maximized instead.

5.2.2 Radiation Induced Lung Toxicity

Higher probability of RILT is the result of higher lung dose, specifically, higher mean lung dose. Analysis of data collected for a large population of patients suggests the relationship between mean lung dose and P_{RILT} depicted in Figure 5.1, which is traditionally used to estimate patient's P_{RILT} pre-treatment (in view of the following discussion, the relationship can be interpreted as the *expected* value of P_{RILT} , as a function of mean lung dose, for a patient sampled from the population uniformly at random).

Gathering additional information part-way into the treatment can give planners a patient-specific estimate of the relationship between P_{RILT} and mean lung dose. In particular, our clinical collaborators have considered the ratio of Transforming Growth Factor β_1 (TGF β_1) level measured 2 weeks into the treatment, and its pre-treatment level (see *Kong et al.*, 2008




Figure 5.1: Population based dependence of P_{RILT} on mean lung dose (or expected P_{RILT}).

for background information). Let u be the value of 2-week TGF β 1 ratio measured for a particular patient:

$$u = \frac{\text{TGF}\beta 1_{2\text{wk}}}{\text{TGF}\beta 1_{0\text{wk}}}. \quad (5.3)$$

According to the biostatistical studies by our collaborators, the estimate of P_{RILT} as a function of dose distribution z (specifically, mean lung dose \bar{z}_{lung}) for this patient can be expressed as

$$P_{\text{RILT}}(z; u) = \frac{e^{X(\bar{z}_{\text{lung}}; u)}}{1 + e^{X(\bar{z}_{\text{lung}}; u)}}, \quad (5.4)$$

where

$$X(\bar{z}_{\text{lung}}; u) = \gamma_0^2 + (\gamma_1^2 + \gamma_2^2 u) \bar{z}_{\text{lung}} \quad (5.5)$$

and

$$\bar{z}_{\text{lung}} = \frac{1}{|V_{\text{lung}}|} \sum_{j \in V_{\text{lung}}} z_j. \quad (5.6)$$

(Table 5.2 shows estimates of relevant parameter values.)

Table 5.2: Parameter estimates for P_{RILT}

Parameter	Estimate
γ_0^2	-5.3731
γ_1^2	0.1697
γ_2^2	0.0315

Figure 5.2 depicts the population-based (i.e., expected) P_{RILT} function, as well as two patient-specific functions, one for a patient whose lung tissue is more resistant (i.e., having a lower value of u), and another for a patient whose lung tissue is more sensitive (i.e., having a higher value of u).

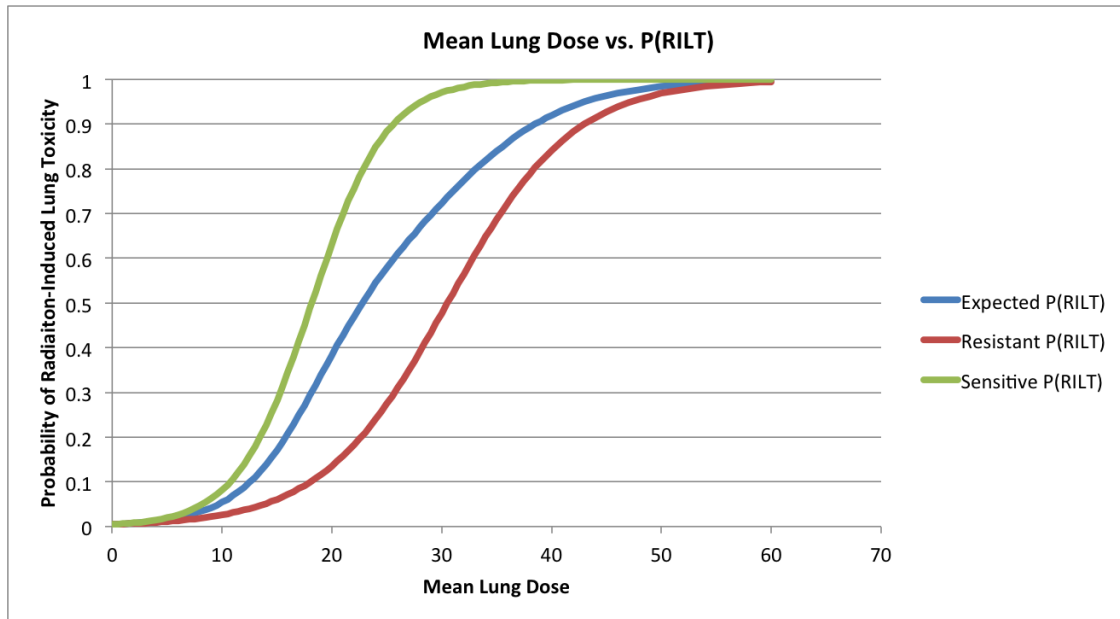


Figure 5.2: P_{RILT} as a function of mean lung dose for the population (expected P_{RILT}) and for sensitive and resistant patients.

5.2.3 Bounding P_{RILT} in Pre-Treatment Planning Models

Since RILT is a highly undesirable side effect of radiation therapy in lung cancer cases, it is common to include an upper bound on P_{RILT} in optimization models used for treatment planning. In light of the above discussion of patient-specific functional forms of P_{RILT} , we will discuss the meaning of such bounds, and two possible forms they can take. In this subsection, we discuss the pre-treatment planning setting (i.e., planning done before the value of u specific to the patient can be observed), but similar ideas are utilized in the discussion of adaptive models later in this chapter.

Consider planning a treatment for a new patient, before a measurement of $\text{TGF}\beta 1$ ratio specific to him can be made. We can view this patient as being selected randomly (uniformly) from the entire population of patients; thus we can view his $\text{TGF}\beta 1$ ratio as a random variable U . We assume that U is a continuous nonnegative random variable, with CDF $H(\cdot)$ reflecting the distribution of values of this biomarker in the patient population. We consider two possible approaches to constraining P_{RILT} in our treatment planning models:

Option 1 (“pop”): population-based, or expected value constraint, stating that *expected* P_{RILT} for a given patient does not exceed $1 - \alpha$ (for a given $\alpha \in (0, 1)$). This constraint is commonly used in the traditional, non-adaptive approach to treatment planning. Essentially, this constraint is a bound on the height of the blue curve in Figures 5.1 and 5.2; we can express it as

$$E [P_{\text{RILT}}(z; U)] \equiv \int_{\mathbb{R}} P_{\text{RILT}}(z; u) dH(u) \leq 1 - \alpha. \quad (5.7)$$

(5.7) provides an important bound because a clinic is interested in their expected treatment outcomes for a population of patients. However, due to the nature of an expectation, probability of RILT in some patients will be higher, in some patients — lower, than the target value of $1 - \alpha$, thus, this planning approach is not ideal on a patient-by-patient basis.

Although U is a continuous random variables, for computational tractability we use a scenario-based discrete approximation. We discretize the range of U into S scenarios, ordered so that $u^1 < u^2 < \dots < u^{S-1} < u^S$ (recall that higher values of u indicate higher levels of patient’s sensitivity to radiation, so u^1 corresponds to the most resistant patients, and u^S — to the most sensitive ones), and let p^s be the probability of scenario s occurring. Details on how the scenario set is constructed can be found in Section 5.4.1. With this discretization, constraint (5.7) can be rewritten as

$$\sum_{s=1}^S p^s P_{\text{RILT}}(z; u^s) \leq 1 - \alpha. \quad (5.8)$$

Option 2 (“rob”): robust constraint, stating that P_{RILT} for the given patient does not exceed $1 - \alpha$. The attractive feature of this constraint is that it provides a bound on the complication probability for the *specific* patient, which is attractive to clinicians and planners, as well as patients. However, in pre-treatment planning, without any patient-specific information available to the planner, this constraint is expressed as

$$P_{\text{RILT}}(z; u) \leq 1 - \alpha \quad \forall u \in \mathfrak{R}, \quad (5.9)$$

or, using our scenario approximation,

$$P_{\text{RILT}}(z; u^s) \leq 1 - \alpha, \quad s = 1, \dots, S. \quad (5.10)$$

In other words, without patient-specific information available, this is a worst-case bound on P_{RILT} for *any* patient, which leads to rather conservative treatment plans.

5.3 Planning Strategies and Models

In this section, we present three treatment planning strategies, formulate corresponding treatment planning optimization models, and discuss solution approaches.

The first treatment strategy we describe is non-adaptive (NA), which does not consider future adaptation explicitly and will serve as a baseline for comparison the other two strategies. Currently, the NA strategy is used in most clinics.

The other two treatment strategies incorporate new information (namely, observation of biomarker value u part-way into the treatment) to allow adaptation, i.e., changing the remaining portion of the treatment after the observation is made. The stochastic adaptive (SA) model uses two-stage stochastic optimization to fully incorporate adaptation of the treatment plan into the pre-treatment treatment planning model. However, as we will discuss, the SA model can be expensive to solve. Therefore, modeling strategy approximate adaptive (AA) was developed to generate high quality “future-conscious” treatment plans without the heavy computation overhead of the SA model.

We will use the IMRT FMO treatment planning framework in our models, with the same notation as in Section 1.2.1. To keep our formulations concise, we let $Z \subset \mathbb{R}^{|V|}$ denote the set of feasible dose distributions, i.e., ones that satisfy any upper- and lower-bound constraints and other clinical constraints on the dose distributions to structures other than the lung and the PTV. We will assume that the set Z is convex, which is commonly true in clinical settings; in our computational experiments set Z was a convex polyhedron.

5.3.1 Non-adaptive Treatment Planning Models

We begin by presenting the non-adaptive treatment planning models with two options for constraining P_{RILT} , and exploring relevant mathematical structure of each formulation. The basic structure of these models is:

$$\begin{aligned}
 \text{(NA)} \quad & \underset{x,z}{\text{maximize}} && P_{\text{LTC}}(z) \\
 & \text{s.t.} && z \in Z \\
 & && z_j = \sum_{i \in N} D_{ij} x_i && j \in V \\
 & && x_i \geq 0 && i \in N \\
 & && \text{constraint on } P_{\text{RILT}},
 \end{aligned}$$

where, as discussed in Section 5.2.3, we have two options for constraining P_{RILT} .

Option 1: Population-based bound on expected P_{RILT} . In the model (NA_{pop}) we use constraint (5.7) to limit P_{RILT} :

$$\text{(NA}_{\text{pop}}) \quad \underset{x,z}{\text{maximize}} \quad P_{\text{LTC}}(z) \quad (5.11)$$

$$\text{s.t.} \quad z \in Z \quad (5.12)$$

$$z_j = \sum_{i \in N} D_{ij} x_i \quad j \in V \quad (5.13)$$

$$x_i \geq 0 \quad i \in N \quad (5.14)$$

$$E[P_{\text{RILT}}(z, U)] \leq 1 - \alpha \quad (5.15)$$

Referencing expressions (5.4) and (5.5), we can express the left-hand side of (5.15) as follows:

$$\begin{aligned} E [P_{\text{RILT}}(z; U)] &\equiv \int_{\mathbb{R}} P_{\text{RILT}}(z, U) dH(u) \\ &= \int_{\mathbb{R}} \frac{e^{X(\bar{z}_{\text{lung}}, u)}}{1 + e^{X(\bar{z}_{\text{lung}}, u)}} dH(u) \\ &= \int_{\mathbb{R}} 1 - \frac{1}{1 + e^{X(\bar{z}_{\text{lung}}, u)}} dH(u) \\ &= 1 - \int_{\mathbb{R}} \frac{1}{1 + e^{X(\bar{z}_{\text{lung}}, u)}} dH(u) \\ &= 1 - \int_{\mathbb{R}} \frac{1}{1 + e^{\gamma_0^2 + (\gamma_1^2 + \gamma_2^2 u) \bar{z}_{\text{lung}}}} dH(u), \end{aligned}$$

and, using an S -scenario discrete approximation, we can express constraint (5.15) as

$$\approx 1 - \sum_{s=1}^S \frac{p_s}{1 + e^{\gamma_0^2 + (\gamma_1^2 + \gamma_2^2 u) \bar{z}_{\text{lung}}}} \leq 1 - \alpha.$$

Notice that the function on the left-hand side is monotone decreasing in \bar{z}_{lung} , and so we can reduce this constraint to a simple upper bound constraint on \bar{z}_{lung} :

$$\bar{z}_{\text{lung}} \leq \overline{\text{MLD}}^{(1)}, \quad (5.16)$$

where $\overline{\text{MLD}}^{(1)}$ is defined as the solution to

$$\alpha = \sum_{s=1}^S \frac{p_s}{1 + e^{\gamma_0^2 + (\gamma_1^2 + \gamma_2^2 u^s) \text{MLD}^{(1)}}}, \quad (5.17)$$

which can be obtained numerically. Moreover, P_{LTC} is a monotone increasing function of \bar{z}_{lung} , and thus objective function (5.11) can be replaced with \bar{z}_{lung} . Therefore, if Z is a polyhedral set, (NA_{pop}) is a linear programming problem.

Option 2: Robust bound on P_{RILT} . Model (NA_{rob}) uses constraint (5.9), or its S -scenario discretization (5.10) for controlling P_{RILT} :

$$(\text{NA}_{\text{rob}}) \quad \underset{x, z}{\text{maximize}} \quad P_{\text{LTC}}(z) \quad (5.18)$$

$$\text{s.t.} \quad z \in Z \quad (5.19)$$

$$z_j = \sum_{i \in N} D_{ij} x_i \quad j \in V \quad (5.20)$$

$$x_i \geq 0 \quad i \in N \quad (5.21)$$

$$P_{\text{RILT}}(z, u^S) \leq 1 - \alpha. \quad (5.22)$$

Again, objective function of (NA_{rob}) can be replaced with \bar{z}_{lung} , and constraint (5.22) can be simplified by incorporating expressions (5.4) and (5.5) and rewriting the inequality as

$$\gamma_0^2 + (\gamma_1^2 + \gamma_2^2 u^S) \bar{z}_{\text{lung}} \leq \ln \left(\frac{1 - \alpha}{\alpha} \right),$$

further simplified to

$$\bar{z}_{\text{lung}} \leq \frac{\ln \left(\frac{1 - \alpha}{\alpha} \right) - \gamma_0^2}{\gamma_1^2 + \gamma_2^2 u^S}. \quad (5.23)$$

5.3.2 Framework for Adaptive Treatment Planning Models

In this subsection we present several adaptive treatment planning models. To recap, the general framework for these models is as follows: an initial treatment plan is developed prior to observation of any patient-specific information captured by the TGF β 1 ratio. Treatment delivery proceeds based on this plan, using equal fractions, until patient-specific information (i.e., the value of the biomarker u) is observed (we refer to this as “stage 1” of the treatment). At this point, the treatment is adapted, based on the dose delivered so far and the observed value of u , and the remainder (“stage 2”) of the treatment is performed using the updated plan.

We will use the following additional notation: we define superscripts (1) and (2) to designate the two treatment stages discussed above. For example, we will denote by $z^{(1)}$ and $z^{(2)}$ the dose distributions of the treatment plans obtained for stages 1 and 2 of the treatment, respectively (with $x^{(1)}$ and $x^{(2)}$ denoting the vectors of corresponding beamlet intensities). Let $w \in (0, 1)$ denote the length of the first stage, specifically, the percentage of treatment that elapses before u is observed. Then, the patient receives a total dose distribution of $wz^{(1)}$ in stage 1 of the treatment, $(1 - w)z^{(2)}$ — in stage 2, and a total distribution of $wz^{(1)} + (1 - w)z^{(2)}$.

5.3.3 Stochastic Adaptive Treatment Planning Models

This section presents model SA for adaptive treatment planning (with two options for constraining P_{LTC}) based on the standard two-stage stochastic programming approach. The advantage of this approach is that treatment planning for the first stage incorporates second-stage treatment adaptation into the model. In other words, the first stage treatment plan produced by SA is calculated in an *informed* way, which should enable it to take better advantage of the re-planning opportunities than any other model. (See *Shapiro et al. (2014)* for more information on stochastic programming.)

As before, we use a scenario-based approximation of the distribution of U in these models. The basic structure of SA models is:

(SA_{pop})

$$\begin{aligned}
 & \underset{\substack{x^{(1)}, z^{(1)}, x^{(2)}(s), \\ z^{(2)}(s), z(s), s=1, \dots, S}}{\text{maximize}} && \sum_{s=1}^S p^s P_{\text{LTC}}(z(s)) \\
 \text{s.t.} && z_j^{(1)} = \sum_{i \in N} D_{ij} x_i^{(1)} && j \in V \\
 && z_j^{(2)}(s) = \sum_{i \in N} D_{ij} x_i^{(2)}(s) && j \in V; s = 1, \dots, S \\
 && z_j(s) = wz_j^{(1)} + (1 - w)z_j^{(2)}(s) && j \in V; s = 1, \dots, S \\
 && \text{constraint on } P_{\text{RILT}} \\
 && z^{(1)} \in Z^{(1)} \\
 && z^{(2)}(s) \in Z^{(2)} && s = 1, \dots, S \\
 && z \in Z \\
 && x_i^{(1)} \geq 0 && i \in N \\
 && x_i^{(2)}(s) \geq 0 && i \in N, s = 1, \dots, S.
 \end{aligned}$$

The additional notation in this model includes $x^{(2)}(s)$ and $z^{(2)}(s)$ for $s = 1, \dots, S$ — the beamlet intensities and dose distribution, respectively, to be used in the second stage of the treatment if the biomarker value u^s is realized, as well as the corresponding total dose distribution $z(s)$, for $s = 1, \dots, S$. Also, because of the explicit incorporation of biomarker realizations, we alter our objective function to be the expectation of P_{LTC} . This objective function is separable in $F(z(s))$ and concave when the probability of complication exceeds $1 - e^{-1} \approx 63.2\%$ (see Appendix B for derivation).

As in the (NA) models, we have two options for constraining P_{RILT} in (SA).

Option 1: Population-based bound on expected P_{RILT} . In the model (SA_{pop}) we use a constraint similar to (5.7) to limit P_{RILT} , with the modification reflecting scenario-based values $\bar{z}_{\text{lung}}(s)$ incorporated into the expectation:

(SA_{pop})

$$\begin{aligned} & \underset{\substack{x^{(1)}, z^{(1)}, x^{(2)}(s), \\ z^{(2)}(s), z(s), s=1, \dots, S}}{\text{maximize}} && \sum_{s=1}^S p^s P_{\text{LTC}}(z(s)) && (5.24) \end{aligned}$$

$$\text{s.t.} \quad z_j^{(1)} = \sum_{i \in N} D_{ij} x_i^{(1)} \quad j \in V \quad (5.25)$$

$$z_j^{(2)}(s) = \sum_{i \in N} D_{ij} x_i^{(2)}(s) \quad j \in V; s = 1, \dots, S \quad (5.26)$$

$$z_j(s) = w z_j^{(1)} + (1 - w) z_j^{(2)}(s) \quad j \in V; s = 1, \dots, S \quad (5.27)$$

$$\sum_{s=1, \dots, S} P_{\text{RILT}}(z(s); u^s) p^s \leq 1 - \alpha \quad (5.28)$$

$$z^{(1)} \in Z^{(1)} \quad (5.29)$$

$$z^{(2)}(s) \in Z^{(2)} \quad s = 1, \dots, S \quad (5.30)$$

$$z \in Z \quad (5.31)$$

$$x_i^{(1)} \geq 0 \quad i \in N \quad (5.32)$$

$$x_i^{(2)}(s) \geq 0 \quad i \in N, s = 1, \dots, S. \quad (5.33)$$

Constraint (5.28) can be written as follows:

$$\sum_{s=1}^S \frac{p_s}{1 + e^{\gamma_0^2 + (\gamma_1^2 + \gamma_2^2 u^s) \bar{z}_{\text{lung}}(s)}} \geq \alpha. \quad (5.34)$$

The left hand side of this constraint is decreasing in $\bar{z}_{\text{lung}}(s)$ and concave as long as $\bar{z}_{\text{lung}}(s)$ corresponds to a P_{RILT} less than 50% (see appendix B for derivation).

Option 2: Robust bound on P_{RILT} . Model (SA_{rob}) imposes the robust version of the constraint on P_{RILT} :

(SA_{rob})

$$\begin{aligned} & \text{maximize} && \sum_{s=1}^S p^s P_{\text{LTC}}(z(s)) && (5.35) \\ & x^{(1)}, z^{(1)}, x^{(2)}(s), && && \\ & z^{(2)}(s), z(s), s=1, \dots, S && && \end{aligned}$$

$$\text{s.t.} \quad z_j^{(1)} = \sum_{i \in N} D_{ij} x_i^{(1)} \quad j \in V \quad (5.36)$$

$$z_j^{(2)}(s) = \sum_{i \in N} D_{ij} x_i^{(2)}(s) \quad j \in V; s = 1, \dots, S \quad (5.37)$$

$$z_j(s) = w z_j^{(1)} + (1 - w) z_j^{(2)}(s) \quad j \in V; s = 1, \dots, S \quad (5.38)$$

$$P_{\text{RILT}}(z; u^s) \leq 1 - \alpha \quad s = 1, \dots, S \quad (5.39)$$

$$z^{(1)} \in Z^{(1)} \quad (5.40)$$

$$z^{(2)}(s) \in Z^{(2)} \quad s = 1, \dots, S \quad (5.41)$$

$$z \in Z \quad (5.42)$$

$$x_i^{(1)} \geq 0 \quad i \in N \quad (5.43)$$

$$x_i^{(2)}(s) \geq 0 \quad i \in N, s = 1, \dots, S. \quad (5.44)$$

Constraint (5.39) can be expressed as simple upper bounds on $\bar{z}_{\text{lung}}(s)$:

$$\bar{z}_{\text{lung}}(s) \leq \frac{\ln\left(\frac{1-\alpha}{\alpha}\right) - \gamma_0^2}{\gamma_1^2 + \gamma_2^2 u^s} \quad s = 1, \dots, S. \quad (5.45)$$

5.3.3.1 Second Stage Planning Strategies

With the SA model, second-stage dose distributions, $z^{(2)}(s)$ are found for S specific biomarker ratio observations. However, the observed biomarker ratio u may, and is most likely, not one of these values. Therefore, the second stage dose distributions found in (SA) may not correspond precisely to a patient's estimated predisposition to P_{RILT} . Physicians must find a suitable second stage dose distribution for the patient. Two common strategies exist for selecting a second stage treatment plan. Physicians can use the pre-computed second stage treatment plan $(x^{(2)}(s), z^{(2)}(s))$ with scenario s that corresponds to the observed value of u (see Section 5.4.1 for more details on matching u with s), or they can use a patient-

specific plan found by reoptimizing the second stage dose specific to the observed ratio u . This replanning model is presented below.

Let \bar{z} denote the dose distribution used in the first stage of the treatment (i.e., at the time of re-planning, $w\bar{z}$ has been already delivered to the patient), and u denote the observed value of the biomarker. The corresponding instance of the re-planning problem $(\text{RP}(\bar{z}, u))$ calculates the second-stage dose $z^{(2)}$ and total dose $z = w\bar{z} + (1 - w)z^{(2)}$:

$$(\text{RP}(\bar{z}, u)) \quad \underset{\bar{x}^{(2)}, z^{(2)}, z}{\text{maximize}} \quad P_{\text{LTC}}(z) \quad (5.46)$$

$$\text{s.t.} \quad z \in Z \quad (5.47)$$

$$z_j^{(2)} = \sum_{i \in N} D_{ij} x_i^{(2)} \quad j \in V \quad (5.48)$$

$$x_i^{(2)} \geq 0 \quad i \in N \quad (5.49)$$

$$z_j = w\bar{z}_j + (1 - w)z_j^{(2)} \quad j \in V \quad (5.50)$$

$$P_{\text{RILT}}(z, u) \leq 1 - \alpha \quad (5.51)$$

Due to the patient-specific nature of replanning model, constraint (5.51) of (RP) bounds P_{RILT} for the particular observed value of u . This constraint's algebraic form is similar to (5.23).

It should be noted that the (SA_{pop}) model may allow $P_{\text{RILT}}(z(s); u^s)$ to exceed $1 - \alpha$ for some scenarios. Using the pre-computed second stage dose distributions may result in patients with higher than $1 - \alpha$ probabilities of P_{RILT} . This characteristic can be seen in the results Section 5.4.3.

5.3.4 Approximate Adaptive Treatment Planning Model

While the SA model offers the most control over P_{RILT} and explicitly incorporates multiple potential stage 2 biomarker realizations into the pre-treatment planning, its solution time may prove clinically prohibitive in practice (see Table 5.4 for run time and problem size details for our application). As S increases, the stochastic model solve time increases to multiple hours, a time too long for practical clinical use.

In light of this, we also consider a simpler model for first-stage treatment planning in the adaptive framework. This model approximates the stochastic programming model (we refer to it as the approximate model, or AA). In particular, in this model we deliberately avoid the creation of a second-stage dose distribution $z^{(2)}(s)$ for each scenario (which makes the (SA) models large and hard to solve). Instead, we use a single variable $z^{(2)}$ to approximate eventual second-stage treatment decisions, imposing a number of constraints to encourage a

good approximation.

The approximate adaptive model is as follows:

$$(AA_{\text{rob}}) \quad \underset{\tilde{x}^{(1)}, \tilde{z}^{(1)}, x^{(2)}, z^{(2)}, z}{\text{maximize}} \quad P_{\text{LTC}}(z) \quad (5.52)$$

$$\text{s.t.} \quad z \in Z \quad (5.53)$$

$$\tilde{z}_j^{(1)} = \sum_{i \in N} D_{ij} \tilde{x}_i^{(1)} \quad j \in \mathcal{V} \quad (5.54)$$

$$\tilde{x}_i^{(1)} \geq 0 \quad i \in N \quad (5.55)$$

$$z_j^{(2)} = \sum_{i \in N} D_{ij} x_i^{(2)} \quad j \in \mathcal{V} \quad (5.56)$$

$$x_i^{(2)} \geq 0 \quad i \in N \quad (5.57)$$

$$z_j = w \tilde{z}_j^{(1)} + (1 - w) z_j^{(2)} \quad j \in \mathcal{V} \quad (5.58)$$

$$P_{\text{RILT}}(\tilde{z}^{(1)}, u^S) \leq 1 - \alpha \quad (5.59)$$

$$P_{\text{LTC}}(\tilde{z}^{(1)}) \geq P_{\text{LTC}}(z^*) - \epsilon \quad (5.60)$$

$$P_{\text{RILT}}(z, u^1) \leq 1 - \alpha. \quad (5.61)$$

Variables $\tilde{x}^{(1)}$ and $\tilde{z}^{(1)}$ represent the first stage beamlet intensities and dose distribution, respectively, produced by this model, and variables $x^{(2)}$ and $z^{(2)}$ are approximations of second stage treatment decisions. It should be noted that the size of this problem is independent of S .

Constraints (5.59) and (5.60) ensure that the first stage dose distribution is no worse than the one obtained by solving (NA_{rob}) (here, z^* denotes the optimal value of (NA_{rob})). Since the model (AA_{rob}) has limited foresight into second-stage treatment, we utilized the robust option for bounding P_{RILT} , to ensure that there will exist a second-stage treatment plan for each patient that is safe, i.e., has P_{RILT} that does not exceed $1 - \alpha$ regardless of this patient's u value. We allow a small deviation, ϵ , from the optimal value z^* to avoid numerical issues without sacrificing plan quality by any significant amount. Constraint (5.61) allows the second stage dose distribution to push more dose into the lung given the initial stage one dose distribution, leading to a first-stage treatment plan $(z^{(1)}, x^{(1)})$ that will allow for a more flexible plan adaptation.

Because only the extremes of the biomarker ratio range are considered, we would solve $(RP(\tilde{z}^{(1)}, u))$ to get the second stage treatment plan specific to this patient after delivering dose $w \tilde{z}^{(1)}$ to the patient and observing u .

5.4 Results and Discussion

In the previous section, we introduced 5 models for treatment planning: two non-adaptive models, (NA_{pop}) and (NA_{rob}) , and three models for computing treatment plans in an adaptive environment, (SA_{pop}) , (SA_{rob}) , and (AA_{rob}) .

These models were tested on four clinical lung cancer cases. The size of the data for each case can be seen in Table 5.3. These cases were selected out of a pool of lung cancer patients for their flexibility in allowing higher doses to the lung while staying within other organs' dosimetric limits. We are interested in the tradeoff between P_{RILT} and P_{LTC} . In this section, we would like to show that the different models generate different first-stage dose distributions, and that these distributions have an impact on the tradeoff between P_{RILT} and P_{LTC} in the resulting full treatment.

Table 5.3: The sizes of clinical cases: number of voxels and number of beamlets. The clinical number is the identifier for the case in UMPlan, the University of Michigan's in-house treatment planning system, included here for figure references and in case data needs to be retrieved.

Case	Clinical number	V	N
1	38	20221	1184
2	44	14965	1388
3	45	15155	1411
4	49	19473	1270

Table 5.4: The number of variables, constraints, and non-zeros as well as run time ranges in minutes for case 3.

Model	Variables	Constraints	Non-zeros	Time min	Time max
(NA_{rob})	16572	17651	17871794	2	5
(SA_{pop}) , S= 2	80029	83664	53702943	85	130
(SA_{rob}) , S= 2	80027	83574	53702763	40	110
$(AA_{\text{rob}}^{(1)})$	48299	50457	35789053	8	50
(SA_{pop}) , S= 4	143487	149695	89538301	130	170
(SA_{rob}) , S= 4	143483	149498	89537905	120	180

Throughout Section 5.3, we discussed the specific mathematical structure of P_{LTC} and P_{RILT} -related functions used in the objectives and constraints of our models. Although in many cases we could simplify these functions and constraints to be linear, a few remaining functions are still nonlinear, notably, the objective function of both (SA) models and the population-based P_{RILT} constraint in the model (SA_{pop}) . In our computational work, we

approximated these nonlinear functions using inner piecewise linear approximations within their convex regions at 1 Gy spacing on the x-axis. These approximation techniques are discussed in Appendix C, which contains the full linear programming formulation of the stochastic adaptive models from Section 5.3.3. This enabled us to use an off-the-shelf linear program solver for this project. Due to the proof-of-concept nature of this chapter, we believe it is sufficient to use a general purpose solver (in particular, clinical solution times are not directly considered in this project).

5.4.1 Discretizing the Biomarker Ratio Distribution

As mentioned in Section 5.2.2, the random variable U is a continuous random variable, whose distribution can be estimated using a (finite) set of already-treated patients. However, the full set of patient outcomes cannot be integrated into the stochastic model in a computationally tractable way; thus we used a partitioning of the domain of U into potential scenarios $s = 1, \dots, S$, with a single value, u^s , representing each scenario. We partitioned the range of observed values of TGF β 1 ratios into S intervals, and chose the largest value in each interval to be the “representative” value. Recall that larger values of u are observed in patients who are more sensitive to radiation; this choice of representative values allows us to provide robust guarantees on P_{RILT} even using the scenario-based modeling approach. This guarantee holds when we select a pre-computed second stage dose distribution from (SA) to deliver to a patient after u is observed.

When it came to the choice of intervals to cover the domain of U , we considered a quantile partition strategy and a uniform partition strategy.

Quantile Strategy Given the desired cardinality S , each scenario s corresponds to the appropriate quantile of the underlying distribution of patient TGF β 1 ratios. Since the distribution is divided by quantiles, $p_s = 1/n$ for each $s = 1, \dots, S$.

Uniform Strategy The domain of TGF β 1 ratios is split into S intervals of equal length. The probability of a realization falling into the interval S is then:

$$p_s = \frac{|\mathcal{S}(s)|}{\sum_{s'=1}^S |\mathcal{S}(s')|}, \quad (5.62)$$

where $|\mathcal{S}(s)|$ is the number of ratio realizations in interval s in the underlying data set.

After speaking with clinical collaborators and seeing little difference in the performance of these partitioning strategies, we continue the chapter using *only the quantile strategy*. Table 5.5 shows the u^s values for the n -quantile strategy. We artificially imposed 4.5 as

the maximum value in order to keep a few outliers from the true distribution from overly influencing the model.

Table 5.5: Scenario-specific biomarker ratio realizations for different numbers of quantiles

	1-quantile	2-quantile	3-quantile	4-quantile	5-quantile
u^0	4.5	0.795323131	0.576232926	0.503458878	0.4182977484
u^1		4.5	1.069878213	0.795323131	0.6559276511
u^2			4.5	1.4	0.8919406165
u^3				4.5	1.5589825295
u^4					4.5

5.4.2 Sanity Check: Flexibility in stage-one dose distribution

First, we determine if different models generate different $z^{(1)}$ dose distributions for stage 1 of the treatment (for non-adaptive models, obviously, stage 1 and stage 2 dose distributions are the same). Dropping the superscript (1) for the purposes of this comparison, let us consider $z_{\text{rob}}^{\text{NA}} - z_{\text{rob}}^{\text{AA}}$ and $z_{\text{rob}}^{\text{NA}} - z_{\text{rob}}^{\text{SA}}$. The elements of these difference vectors are sorted in increasing order and displayed for the PTV and lung structures in Figure 5.3. We can see that there is a significant amount of dose redistribution.

5.4.3 Analysis for a Two-scenario Patient Population

Suppose that a patient can only have two biomarker ratio realizations that correspond to the 2-quantile discretization presented in Section 5.4.1. We will relax this assumption later on and consider the entire biomarker ratio distribution, but as a proof of concept, let us consider this simple case. Because these biomarker ratio realizations are far apart, they can be seen as a patient either being resistant to radiation ($u = 0.795323131$), or sensitive to radiation ($u = 4.5$).

In this two-scenario environment, we do not solve model $(\text{RP}(\bar{z}, u))$ to get a stage 2 dose distribution with a robust bound (i.e., a stage 2 dose distribution with P_{RILT} bounded by constraint (5.51)). Instead, we consider the stage 2 dose distributions as determined by the pre-treatment models to highlight the differences in P_{RILT} control. Clinically, provided there are available resources, model $(\text{RP}(\bar{z}, u))$ could be solved to determine the patient-specific stage 2 dose distribution. With only two scenarios, it is sufficient for the AA model to use the overall dose distribution $z = wz^{(1)} + (1 - w)z^{(2)}$ for the resistant patient and stage 1 dose distribution $\tilde{z}^{(1)}$ for the sensitive patient (from model (AA_{rob})).




Figure 5.3: Sorted differences in first-stage dose distribution vectors (a) $z_{\text{rob}}^{\text{NA}} - z_{\text{rob}}^{\text{AA}}$ lung, (b) $z_{\text{rob}}^{\text{NA}} - z_{\text{rob}}^{\text{AA}}$ PTV, (c) $z_{\text{rob}}^{\text{NA}} - z_{\text{rob}}^{\text{SA}}$ lung, and (d) $z_{\text{rob}}^{\text{NA}} - z_{\text{rob}}^{\text{SA}}$ PTV.

It is interesting to analyze what is considered to be a feasible range of mean lung doses for sensitive and resistant patients in different pre-treatment models. For example, in (NA_{rob}) the treatment is planned assuming the patient is sensitive. On the other hand, the (SA) models would optimize for both the initial dose distribution and the distributions corresponding to the potential biomarker ratio realizations, with population (or expected value) based model allowing for more flexibility than the robust (or worst-case) model. The different overall mean lung dose feasible regions can be found in Figure 5.4.

Another key insight we want to look at is the tradeoff between P_{RILT} and P_{LTC} for the patient population, as well as for sensitive and resistant patients, for each value of α . Consider the tradeoff curves for case 2 in Figures 5.5, 5.6, and 5.7 that compare the realized P_{RILT} and P_{LTC} for each patient type (as well as the population averages) if treated with the plans obtained from models (NA) and (SA). (Recall that scenario 1 represents resistant patients, and scenario 2 — sensitive patients.) Connecting black lines show the points on the curve corresponding to a particular α . These connecting lines are vertical for the stochastic

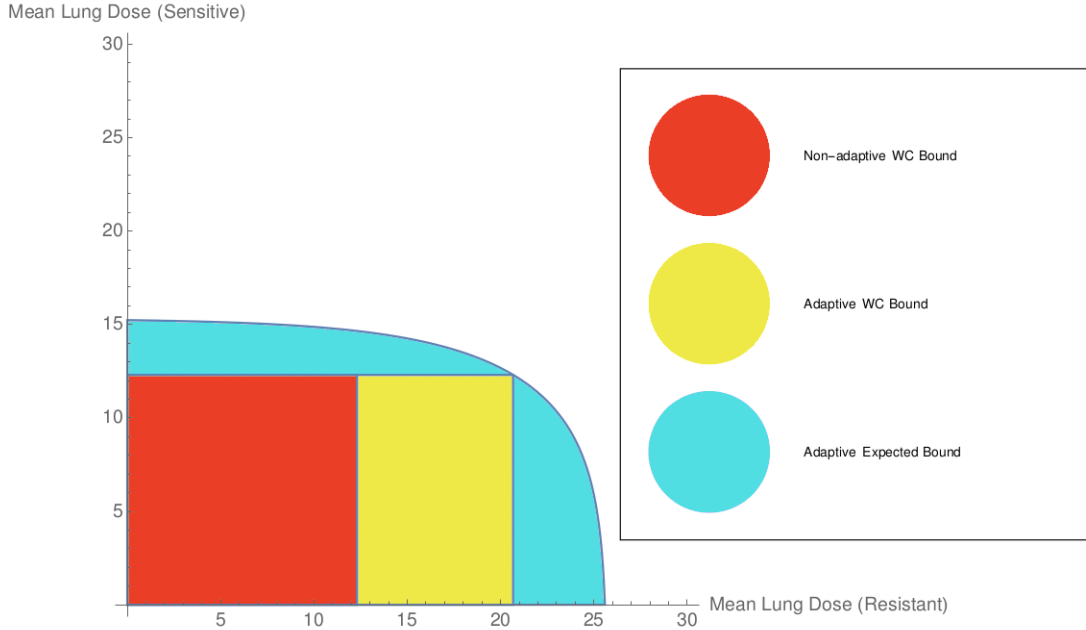


Figure 5.4: Feasible region with respect to mean lung dose for the second stage dose distribution for the different methods (WC Bound = “rob” strategies)

adaptive robust method (i.e., the P_{RILT} bound is kept and additional dose is added to raise the P_{LTC}), and these lines are horizontal for the non-adaptive method (i.e., the overall dose stays the same and only the P_{RILT} changes based on patient sensitivity). We see that both techniques produce realizations that are on similar curves, but how these are related differ greatly. Consider P_{RILT} being bounded at 15%. Table 5.6 shows the realized values for each of the techniques (recall, without solving model (RP(\bar{z}, u))). We see that the added flexibility in the stochastic adaptive model with population bound gives us the best population-wide average P_{LTC} , but it comes at the expense of the sensitive patient receiving more dose. We also see that the approximate adaptive methodology performs almost as well as the stochastic adaptive robust model, with much less computational overhead. The approximate adaptive method performs similarly to the stochastic adaptive robust model in Figure 5.7, and thus is not shown separately. Figures 5.5, 5.6, and 5.7 show the same plots for case 4.

5.4.4 The Impact of Treatment Adaptation to Exact Biomarker Value, and of Adaptation Timing

Having demonstrated the potential structural impact of treatment re-planning on patient outcomes in the previous subsection, let us consider more realistic measurements of biomarker values: after w percent of the treatment has been delivered, an exact biomarker value is observed; in particular, we no longer crudely partition the population into “sensitive” and




Figure 5.5: P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 2, NA - pop and rob




Figure 5.6: P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 2, SA - pop

“resistant” patients, but consider the full continuous range of potential biomarker values.

Out of necessity, the formulations of (SA) models are still scenario-based, and second-stage treatment decisions $(x^{(2)}(s), z^{(2)}(s))$ for biomarker values u^s , $s = 1, \dots, S$ included in the models are produced in the process of obtaining first-stage solutions. As mentioned in Section 5.3.3.1, when the precise value of u is observed, the treatment planner faces a choice: to invest the time and create a second-stage treatment plan from scratch by solving




Figure 5.7: P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 2, SA - rob

Table 5.6: Realized P_{RILT} and P_{LTC} for $\alpha = 0.85$

	P_{RILT}			P_{LTC}		
	Resis.	Sens.	Average	Resis.	Sens.	Average
Non-adaptive - rob	0.0432	0.15	0.0966	0.6225	0.6225	0.6225
Non-adaptive - pop	0.0608	0.2392	0.15	0.7202	0.7202	0.7202
Approx Adaptive - rob	0.15	0.15	0.15	0.8378	0.6225	0.7301
Stochastic - rob	0.15	0.15	0.15	0.8384	0.6225	0.7304
Stochastic - pop	0.1034	0.1963	0.15	0.8079	0.6827	0.7453

an instance of problem (RP), or to utilize the pre-computed second stage treatment plan $(x^{(2)}(s), z^{(2)}(s))$ with scenario s that corresponds to the observed value of u (see Section 5.4.1).

Suppose the treatment planner decides to use the latter approach, i.e., not to fine-tune the second stage of the treatment based on the exact observed value of u . Figure 5.11 shows resulting treatment outcomes (P_{RILT} and P_{LTC}) as a function of the patient's biomarker value. (The (SA) models were solved with $S = 4$ quantile-based scenarios). The figure shows that there are benefits to be gained over non-adaptive treatments even when using this limited version of treatment adaptation: more dose is given to the patients who are identified as being more resistant, leading to improved P_{LTC} .

Next, let us study the treatments resulting from full re-planning to the exact observed biomarker value. We will focus our attention on the robust (option 2) treatment planning models. The results can be seen in Figure 5.12. These lines are essentially on top of each other. We see that the flexibility of the treatment planning optimization model, clinical pro-




Figure 5.8: P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 4, NA - pop and rob




Figure 5.9: P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 4, SA - pop

tol, specific case geometry, and choice of w allow us to perform well regardless of the dose distribution. As mentioned before, these cases were chosen for their flexibility in planning (i.e., tumor is position relatively far from organs at risk). While this can be considered wor-




Figure 5.10: P_{RILT} vs. P_{LTC} for a two-scenario patient population, case 4, SA - rob

risome for these techniques (i.e., little benefit gained from solving more complicated models when model $(RP(z, u))$ is used), not all treatment plans will be as forgiving.

Sometimes not all of these flexibility-generating elements occur and an informed initial stage distribution offers significant benefits. One example of when an informed approach leading into reoptimization is desirable is when w is large (i.e., reoptimization occurs closer to the end of treatment). Consider the non-adaptive method and the approximate adaptive method. As a proof of concept, the initial treatment plan was developed for $u = 4.5$ and then adapted to a realized biomarker ratio of $u = 0.795323131$ with $w = 0.33, 0.67, 0.90, 0.95$. We assume that the accuracy of the biomarker-informed P_{RILT} relationship does not change based on w for this example.

As we see in Figure 5.13, as w is increased, the impact of the initial dose distribution increases as well. There are two main takeaways from these plots. First, it is better to test earlier than later. This intuitively makes sense as we have more room for adapting the treatment. Second, the approximated adaptive methodology's initial dose distribution is better for replanning and allows for more flexibility as w increases.

5.4.5 When to Use These Techniques

The need for these techniques occurs frequently in the clinic. Treatment planners should use a technique when adaptation is likely, as the choice of initial dose distribution will likely have some kind of an impact on the flexibility of adapting the plan. For example, esophagitis often occurs when treating lung cancer patients. If the patient's symptoms get




Figure 5.11: (top) Treatment outcomes for non-adaptive plans, and plans with limited (scenario-based only) adaptation, vs. patient’s biomarker value; Realized P_{RILT} (top) and P_{LTC} (bottom) realized.

too extreme during the treatment, the patient is replanned to spare the esophagus. Without proper preparation in the treatment planning process, the patient may experience even more adverse effects or subpar tumor coverage. If adaptation is not likely to be forced to occur due to complications, then it might be useful to plan for “boosting” the dose near the end of treatment if the patient is responding well to the radiation. In most cases, and especially when w is not clearly known, we recommend the approximate adaptive methodology.

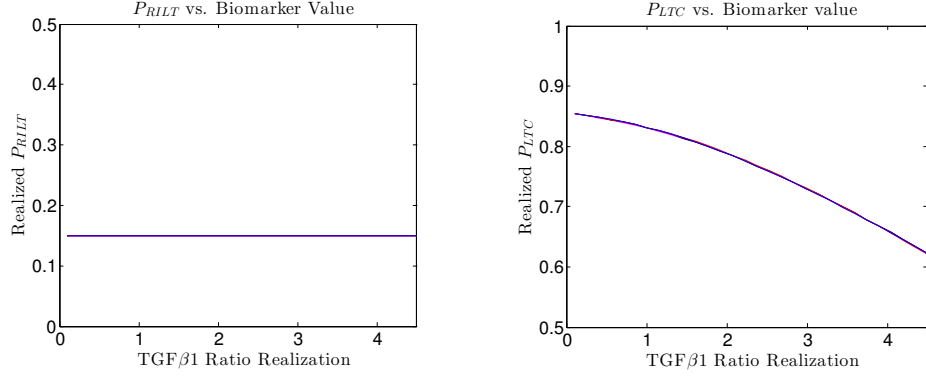


Figure 5.12: For $w = 0.33$, (left) realized P_{RILT} vs. biomarker ratio realization for reoptimized plans, (right) realized P_{LTC} vs. biomarker ratio realization for reoptimized plans; all models on nearly the same lines

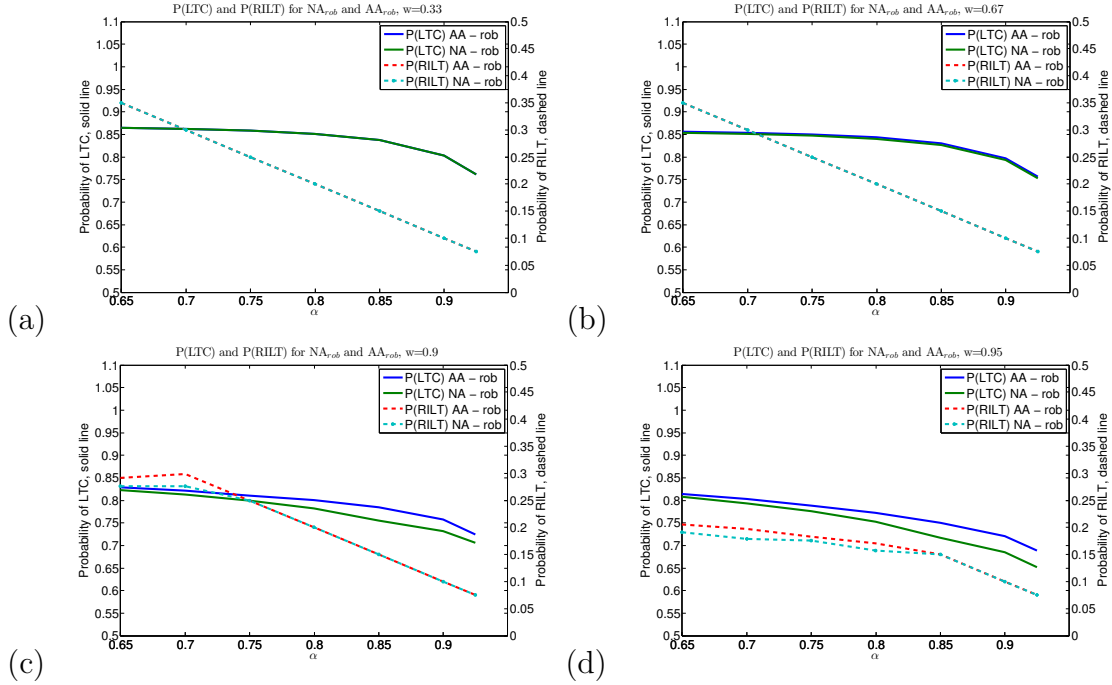


Figure 5.13: P_{RILT} vs. P_{LTC} for $w =$ (a) 0.33, (b) 0.67, (c) 0.90, and (d) 0.95

5.5 Conclusions and Future Research

In conclusion, the initial stage dose distribution has an important effect on the quality of the adapted dose distribution. Without some kind of methodology in place to address this issue, treatment planners could produce subpar treatment plans when they adapt to changes in patient information. This methodology can be extended to handle multiple potential adaptations and “future proof,” to a degree, a plan that is likely to be adapted.

One assumption that was made in this work was that the functions representing P_{RILT}

were certain and that the biomarker ratio realizations were exactly representative the patient's true sensitivity to radiation. That is, there was no uncertainty around the probability functions or biomarker ratio realizations incorporated into the models. A natural extension to this work is to consider the probability functions and/or the biomarker ratio realization as estimates and incorporate the associated uncertainty to the treatment planning model.

Another extension to the model could be made when considering the timing of replanning, i.e., the value of w . Suppose that there are only enough clinical resources to adapt a patient's treatment to new information exactly once. Also suppose that the information you obtain from testing the patient (e.g., a biomarker test) increases in accuracy as w increases. Waiting to get better information competes with the notion from Section 5.4.4 that more adaptive flexibility is gained by obtaining information earlier. In these models, treatment planners would need to decide pre-treatment when is the best time gather new information and adapt, w , and how to plan knowing that adaptation will occur.

5.6 Acknowledgements

This work was done in collaboration with Dr. Marina Epelman, Dr. Martha Matuszak, Dr. Edwin Romeijn, Dr. Matthew Schipper, and Dr. Randall Ten Haken.

APPENDICES

APPENDIX A

Extending the Feasible Aperture Region by Restricting Gantry Travel Time

We can consider a family of potential choices for $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$. Let $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}(t)$ be defined as the following equation with fixed angular speed t .

$$\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}(t) = \left\{ A \in \mathcal{A} : t\delta_{k_\ell} \geq T_{k_\ell, \bar{k}}^L(A_{k_\ell}, A), t\delta_{\bar{k}} \geq T_{\bar{k}, k_{\ell+1}}^L(A, A_{k_{\ell+1}}) \right\} \\ \text{for } t\delta_{\bar{k}} \leq T_{\bar{k}}^U, t\delta_{k_\ell} \leq T_{k_\ell}^U \quad (\text{A.1})$$

Clearly, $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}(t) \subseteq \mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}(t')$ for $t \leq t'$ and $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}(T_{k_\ell}^U) = \mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}U}$ is the largest set. This information will be used when determining the potential MLC leaf arrangements when solving the pricing problem. One thing to note is that the choice of $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}} \subseteq \mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}U}$ will influence the behavior of the algorithm in later stages. If we choose a smaller set $\mathcal{A}_{k_\ell, k_{\ell+1}}^{\bar{k}}$, that will result in less flexibility in the current set of apertures, but may allow for greater flexibility in later stages.

We want to know explicitly how our selected angular travel time t and preceding and following apertures k_ℓ and $k_{\ell+1}$ control the feasible aperture region. That is, for each row, we would like to determine the feasible MLC leaf positions given s , k_ℓ , and $k_{\ell+1}$ for potential control point \bar{k} .

Let us assume that t , k_ℓ , and $k_{\ell+1}$ are fixed along with apertures A_{k_ℓ} , and $A_{k_{\ell+1}}$. We would like to know the feasible MLC leaf positions for apertures at control point \bar{k} . First, let $A_{\bar{k}r}$ be the r th row of aperture $A_{\bar{k}}$. Then, we can define set $A_{\bar{k}r}(t)$ as the set of feasible rows for the r th row of control point \bar{k} with some abuse of notation for in our time constraint.

$$A_{\bar{k}r}(t) = \left\{ a : t\delta_{k_\ell} \geq T_{k_\ell, \bar{k}}^L(A_{k_\ell r}, A_{\bar{k}r}), t\delta_{\bar{k}} \geq T_{\bar{k}, k_{\ell+1}}^L(A_{\bar{k}r}, A_{k_{\ell+1}r}) \right\} \quad (\text{A.2})$$

Let v be the maximum speed an MLC leaf can move and N be the rightmost leaf setting. The total distance a leaf can cover is v times the time it is moving. Therefore, we can see the following relationship.

$$\left\{ a : t\delta_{k_\ell} \geq T_{k_\ell, \bar{k}}^L(A_{k_\ell r}, A_{\bar{k}r}) \right\} = \{(L, R) : 0 \leq L \leq R \leq N; \\ |L - L_{k_\ell}|, |R - R_{k_\ell}| \leq vT_t(k_\ell, \bar{k})\} \quad (\text{A.3})$$

We can define the leaf positions in the following manner.

$$A_{\bar{k}r}(t) = \{(L, R) : 0 \leq L \leq R \leq N; |L - L_{k_\ell}|, |R - R_{k_\ell}| \leq vT_t(k_\ell, \bar{k}); \\ |L - L_{k_{\ell+1}}|, |R - R_{k_{\ell+1}}| \leq vT_t(\bar{k}, k_{\ell+1})\} \quad (\text{A.4})$$

This family of apertures is an easy extension to combat greediness when solving the pricing problem.

APPENDIX B

Derivations for P_{RILT} and P_{LTC}

P_{RILT} derivation

Consider constraint (5.34) reproduced below:

$$\sum_{s=1}^S \frac{p_s}{1 + e^{\gamma_0^2 + (\gamma_1^2 + \gamma_2^2 u^s) \bar{z}_{\text{lung}}(s)}} \geq \alpha.$$

We will show that the left hand side of (5.34) is decreasing in $\bar{z}_{\text{lung}}(s)$ and concave as long as $\bar{z}_{\text{lung}}(s)$ corresponds to a P_{RILT} less than 50%.

We can rewrite this complicating constraint as the following:

$$\sum_{s=1}^S \frac{p_s}{1 + e^{\delta_1^s z^s - \delta_0}} \geq \alpha, \tag{B.1}$$

where, for convenience, we have

$$\delta_0 = -\gamma_0^2 \tag{B.2}$$

$$\delta_1^s = \gamma_1^2 + \gamma_2^2 u^s \tag{B.3}$$

$$z^s = \bar{z}_{\text{lung}}(s). \tag{B.4}$$

We expect that $\delta_0, \delta_1 > 0$ (the initial parameter estimates show have $\gamma_0^2 < 0, \gamma_2^1, \gamma_2^1 > 0$).

Now consider the following function:

$$f^s(z^s) = \frac{1}{1 + e^{\delta_1^s z^s - \delta_0}}. \quad (\text{B.5})$$

It is easy to see that this function is decreasing in z^s . Moreover,

$$\lim_{z^s \downarrow -\infty} = 1 \quad (\text{B.6})$$

$$f^s(0) = \frac{1}{1 + e^{-\delta_0}} \quad (\text{B.7})$$

$$f^s\left(\frac{\delta_0}{\delta_1^s}\right) = \frac{1}{2} \quad (\text{B.8})$$

$$\lim_{z^s \uparrow \infty} = 0, \quad (\text{B.9})$$

where the first is clinically irrelevant but helps understand the shape of the function. Consider the solution to the following:

$$\frac{1}{1 + e^{\delta_1^s z^s - \delta_0}} = \alpha, \quad (\text{B.10})$$

which is a feasible solution to the patient-specific strategy. In particular, this yields:

$$\bar{z}_{\text{lung}}(s) = \frac{\ln\left(\frac{1-\alpha}{\alpha}\right) + \delta_0}{\delta_1^s}. \quad (\text{B.11})$$

Suppose $\alpha = 0.85$; then for $u^s = 0$ this is at 21.4 Gy, and for $u^s = 5$ this is at 11.1 Gy. These would correspond to bounds on the mean lung dose that correspond to a particular P_{RILT} for a given biomarker realization. Now, consider the point $\hat{z}^s \equiv \delta_0/\delta_1^s = \frac{5.37}{0.17+0.031u^s}$ Gy. This point is an inflection point. For $u^s = 0$ this is at 31.7 Gy, and for $u^s = 4.5$ this is at 17.4 Gy. The constraint function is concave whenever $z^s \leq \hat{z}^s$ for all $s = 1, \dots, S$. This means that the function is concave on the feasible region as long as we assume that no patient gets a MLD that would increase their probability of RILT to $\frac{1}{2}$ or more. Restricting P_{RILT} to 50% is a clinically reasonable hard bound for RILT. Because the complicating constraint function is not linear or quadratic, linear program (LP) solvers (such as Gurobi or CPLEX) cannot handle it directly. However, note that it is separable in z^s . Therefore, it is easy to approximate by a piecewise-linear function as long as we limit ourselves to a part of the function where it is concave. We do this by having a hard bound on P_{RILT} in all models where this might be violated (i.e., $P_{\text{RILT}}(z, u) \leq \frac{1}{2}$ for all dose distributions) included in the construction of dose distribution set Z .

P_{LTC} derivation

Consider presented stochastic adaptive objective function presented in Section 5.3.3 with P_{LTC} function from 5.2.1 reproduced below. We will show that this function is separable in $F(z(s))$ and concave when the probability of complication exceeds $1 - e^{-1} \approx 63.2\%$.

$$\sum_{s=1}^S p_s P_{\text{LTC}}(z(s)) = \sum_{s=1}^S p_s S_0(2)^{e^{\beta_0 - \beta_1 F(z(s))}} \quad (\text{B.12})$$

While we would like $F(z)$ to be D95 (a dose-volume function), we instead use some kind of proxy function to facilitate the solving of the optimization model. We let $F(z)$ be a linearized EUD:

$$F(z) = a z_{\text{PTV}}^{\min} + (1 - a) \bar{z}_{\text{PTV}} \quad (\text{B.13})$$

$$z_{\text{PTV}}^{\min} = \min_{j \in V_{\text{PTV}}} z_j \quad (\text{B.14})$$

$$\bar{z}_{\text{PTV}} = \frac{1}{|V_{\text{PTV}}|} \sum_{j \in V_{\text{PTV}}} z_j, \quad (\text{B.15})$$

with function parameter estimates $S_0(2) = 0.298$, $\beta_0 = 5.928$, and $\beta_1 = 0.0988$. Note that the objective function is separable in $F(z(s))$. Now let

$$f(d) = 1 - S_0(2)^{e^{\beta_0 - \beta_1 d}} = \exp(\ln S_0(2) \cdot e^{\beta_0 - \beta_1 d}), \quad (\text{B.16})$$

(i.e., $f(d)$ represents probability of local progression (which we want to minimize) and $1 - f(d)$ is probability of no local progression) where $d = F(z)$. Then its derivative is the following:

$$f'(d) = \exp(\ln S_0(2) \cdot e^{\beta_0 - \beta_1 d}) \cdot \ln S_0(2) \cdot \beta_1 \cdot e^{\beta_0 - \beta_1 d} < 0. \quad (\text{B.17})$$

Since $0 < S_0(2) < 1$, f is decreasing in d . Moreover, its second order derivative is the following:

$$f''(d) = -\beta_1^2 \cdot \ln S_0(2) \cdot e^{\beta_0 - \beta_1 d} \cdot \exp(\ln S_0(2) \cdot e^{\beta_0 - \beta_1 d}) \cdot [\ln S_0(2) \cdot e^{\beta_0 - \beta_1 d} + 1]. \quad (\text{B.18})$$

This means that $f''(d) > 0$ and hence f is convex in w whenever satisfies the following:

$$\ln S_0(2) \cdot e^{\beta_0 - \beta_1 d} + 1 > 0, \quad (\text{B.19})$$

or, equivalently,

$$F(z) = d > \frac{\beta_0 - \ln \left[\frac{1}{\ln(1/S_0(2))} \right]}{\beta_1}, \quad (\text{B.20})$$

and concave in d otherwise. This lower bound on $F(z)$ is added to our models as a constraint in feasible dose distribution set Z presented in Section 5.3. Clinically, plans would not be allowed to have this low of $F(z)$, so adding a clinically reasonable bound to maintain a desirable model structure will not impact treatment plan quality.

APPENDIX C

Full Linear Models for Stochastic Adaptive Lung Formulations

Option 1: Population-based bound on expected P_{RLT} . The following model is the LP formulation of the (SA_{pop}) model including the piecewise linear approximations.

$$(\text{SA}_{\text{pop}}^{\text{LP}}) \quad \begin{array}{l} \text{maximize} \\ x^{(1)}, z^{(1)}, x^{(2)}(s), z^{(2)}(s), \\ z(s), h_s^1, h_s^2, s=1, \dots, S \end{array} \sum_{s=1}^S p_s h_s^1 \quad (\text{C.1})$$

subject to

$$z_j^{(1)} = \sum_{i \in N} D_{ij} x_i^{(1)} \quad j \in V \quad (\text{C.2})$$

$$z_j^{(2)}(s) = \sum_{i \in N} D_{ij} x_i^{(2)}(s) \quad j \in V; s = 1, \dots, S \quad (\text{C.3})$$

$$z_j(s) = w z_j^{(1)} + (1 - w) z_j^{(2)}(s) \quad j \in V; s = 1, \dots, S \quad (\text{C.4})$$

$$z_{\text{PTV}}^{\min}(s) \leq z_j(s) \quad j \in V_{\text{PTV}}; s = 1, \dots, S \quad (\text{C.5})$$

$$\bar{z}_{\text{PTV}}(s) = \frac{1}{|V_{\text{PTV}}|} \sum_{j \in V_{\text{PTV}}} z_j(s) \quad s = 1, \dots, S \quad (\text{C.6})$$

$$z_{\text{PTV}}^{\text{EUD}}(s) = a z_{\text{PTV}}^{\min}(s) + (1 - a) \bar{z}_{\text{PTV}}(s) \quad j \in V_{\text{PTV}}; s = 1, \dots, S \quad (\text{C.7})$$

$$\bar{z}_{\text{lung}}(s) = \frac{1}{|V_{\text{lung}}|} \sum_{j \in V_{\text{lung}}} z_j(s) \quad s = 1, \dots, S \quad (\text{C.8})$$

$$z_{\text{PTV}}^{\text{EUD}}(s) > \frac{\beta_0 - \ln \left[\frac{1}{\ln(1/S_0(2))} \right]}{\beta_1} \quad s = 1, \dots, S \quad (\text{C.9})$$

$$\sum_{s=1}^S p_s h_s^2 \geq \alpha \quad (\text{C.10})$$

$$h_s^1 \leq H^1(z_{\text{PTV}}^{\text{EUD}}(s)) \quad s = 1, \dots, S \quad (\text{C.11})$$

$$h_s^2 \leq H_s^2(\bar{z}_{\text{lung}}(s)) \quad s = 1, \dots, S \quad (\text{C.12})$$

$$z^{(1)} \in Z^{(1)} \quad (\text{C.13})$$

$$z^{(2)}(s) \in Z^{(2)} \quad s = 1, \dots, S \quad (\text{C.14})$$

$$z(s) \in Z \quad s = 1, \dots, S \quad (\text{C.15})$$

$$x_i^{(1)} \geq 0 \quad i \in N \quad (\text{C.16})$$

$$x_i^{(2)}(s) \geq 0 \quad i \in N; s = 1, \dots, S, \quad (\text{C.17})$$

where $H^1(z)$ is a piecewise-linear and concave function and $H_s^2(\bar{z}_{\text{lung}}(s))$, for $s = 1, \dots, S$, is a piecewise-linear and concave function that satisfy the following:

$$H^1(z) \approx S_0(2)e^{\beta_0 - \beta_1 z} \quad (\text{C.18})$$

$$H_s^2(\bar{z}_{\text{lung}}(s)) \approx \frac{1}{1 + e^{\gamma_0^2 + (\gamma_1^2 + \gamma_2^2 u_1^s) \bar{z}_{\text{lung}}(s)}} \quad s = 1, \dots, S. \quad (\text{C.19})$$

Option 1: Robust bound on P_{RILT} . The following model is the LP formulation of the (SA_{rob}) model including the piecewise linear approximations.

$$(\text{SA}_{\text{rob}}^{\text{LP}}) \quad \begin{array}{l} \text{maximize} \\ x^{(1)}, z^{(1)}, x^{(2)}(s), z^{(2)}(s), \\ z(s), h_s^1, h_s^2, s=1, \dots, S \end{array} \sum_{s=1}^S p_s h_s^1 \quad (\text{C.20})$$

subject to

$$z_j^{(1)} = \sum_{i \in N} D_{ij} x_i^{(1)} \quad j \in V \quad (\text{C.21})$$

$$z_j^{(2)}(s) = \sum_{i \in N} D_{ij} x_i^{(2)}(s) \quad j \in V; s = 1, \dots, S \quad (\text{C.22})$$

$$z_j(s) = w z_j^{(1)} + (1 - w) z_j^{(2)}(s) \quad j \in V; s = 1, \dots, S \quad (\text{C.23})$$

$$z_{\text{PTV}}^{\min}(s) \leq z_j(s) \quad j \in V_{\text{PTV}}; s = 1, \dots, S \quad (\text{C.24})$$

$$\bar{z}_{\text{PTV}}(s) = \frac{1}{|V_{\text{PTV}}|} \sum_{j \in V_{\text{PTV}}} z_j(s) \quad s = 1, \dots, S \quad (\text{C.25})$$

$$z_{\text{PTV}}^{\text{EUD}}(s) = a z_{\text{PTV}}^{\min}(s) + (1 - a) \bar{z}_{\text{PTV}}(s) \quad j \in V_{\text{PTV}}; s = 1, \dots, S \quad (\text{C.26})$$

$$\bar{z}_{\text{lung}}(s) = \frac{1}{|V_{\text{lung}}|} \sum_{j \in V_{\text{lung}}} z_j(s) \quad s = 1, \dots, S \quad (\text{C.27})$$

$$z_{\text{PTV}}^{\text{EUD}}(s) > \frac{\beta_0 - \ln \left[\frac{1}{\ln(1/S_0(2))} \right]}{\beta_1} \quad s = 1, \dots, S \quad (\text{C.28})$$

$$h_s^1 \leq H^1(z_{\text{PTV}}^{\text{EUD}}(s)) \quad s = 1, \dots, S \quad (\text{C.29})$$

$$\bar{z}_{\text{lung}}(s) \leq \frac{\ln \left(\frac{1-\alpha}{\alpha} \right) - \gamma_0^2}{\gamma_1^2 + \gamma_2^2 u_1^s} \quad s = 1, \dots, S \quad (\text{C.30})$$

$$z^{(1)} \in Z^{(1)} \quad (\text{C.31})$$

$$z^{(2)}(s) \in Z^{(2)} \quad s = 1, \dots, S \quad (\text{C.32})$$

$$z(s) \in Z \quad s = 1, \dots, S \quad (\text{C.33})$$

$$x_i^{(1)} \geq 0 \quad i \in N \quad (\text{C.34})$$

$$x_i^{(2)}(s) \geq 0 \quad i \in N; s = 1, \dots, S, \quad (\text{C.35})$$

where $H^1(z)$ is a piecewise-linear and concave function that satisfies the following:

$$H^1(z) \approx S_0(2)^{e^{\beta_0 - \beta_1 z}}. \quad (\text{C.36})$$

Note: As long as we limit ourselves to values $\alpha > \frac{1}{2}$ then, by construction, the bounds in (C.10) are uniformly looser than the bounds in (C.30), so that option 2 is indeed more conservative. We should also make sure to restrict ourselves to values of α that ensure that the upper bounds (C.30) are nonnegative, i.e.,

$$\alpha < \frac{1}{1 + e^{\gamma_0^2}},$$

which, for our parameter estimates, means $\alpha < 0.99538$.

BIBLIOGRAPHY

BIBLIOGRAPHY

- Aleman, D., A. Kumar, R. Ahuja, H. Romeijn, and J. Dempsey (2008a), Neighborhood search approaches to beam orientation optimization in intensity modulated radiation therapy treatment planning, *Journal of Global Optimization*, 42(4), 587–607.
- Aleman, D., H. E. Romeijn, and J. Dempsey (2008b), A Response Surface Approach to Beam Orientation Optimization in Intensity-Modulated Radiation Therapy Treatment Planning, *INFORMS Journal on Computing*, 21(1), 62–76.
- Bedford, J. L., and S. Webb (2006), Constrained segment shapes in direct-aperture optimization for step-and-shoot IMRT, *Medical Physics*, 33(4), 944.
- Bedford, J. L., and S. Webb (2007), Direct-aperture optimization applied to selection of beam orientations in intensity-modulated radiation therapy., *Physics in Medicine and Biology*, 52(2), 479–98.
- Bertsekas, D. (1982), Projected newton methods for optimization problems with simple constraints, *SIAM Journal on Control and Optimization*, 20, 221–246.
- Bertsimas, D., V. Cacchiani, D. Craft, and O. Nohadani (2013), A hybrid approach to beam angle optimization in intensity-modulated radiation therapy, *Computers & Operations Research*, 40(9), 2187–2197.
- Bixby, R., Z. Gu, and E. Rothberg (2010), Gurobi optimization.
- Boyd, S., and L. Vandenberghe (2004), *Convex optimization*, Cambridge university press.
- Breedveld, S., P. Storchi, and B. Heijmen (2009), The equivalence of multi-criteria methods for radiotherapy plan optimization, *Physics in Medicine and Biology*, 54, 123–135.
- Breedveld, S., P. Storchi, P. Voet, and B. Heijmen (2012), iCycle: Integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans., *Medical Physics*, 39(2), 951–63.
- Burman, C., G. Kutcher, B. Emami, and M. Goitein (1991), Fitting of normal tissue tolerance data to analytical function, *International Journal of Radiation Oncology Biology Physics*, 21(1), 7199–7209.
- Chan, T., T. Craig, T. Lee, and M. Sharpe (2014), Generalized inverse multiobjective optimization with application to cancer therapy, *Operations Research*, 62(3), 680–695.

- Choi, B., and J. Deasy (2002), The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning, *Physics in Medicine and Biology*, *47*, 3579–3589.
- Clark, V., Y. Chen, J. Wilkens, J. Alaly, K. Zakaryan, and J. Deasy (2008), IMRT treatment planning for prostate cancer using prioritized prescription optimization and mean-tail-dose functions, *Linear Algebra and its Applications*, *428*(5–6), 1345–1364.
- Cox, D. R. (1972), Regression models and life-tables, *Journal of the Royal Statistical Society. Series B (Methodological)*, pp. 187–220.
- Craft, D. (2007), Local beam angle optimization with linear programming and gradient search., *Physics in Medicine and Biology*, *52*(7), N127–35.
- Craft, D., and T. Bortfeld (2008), How many plans are needed in an IMRT multi-objective plan database?, *Physics in Medicine and Biology*, *53*, 2785–2796.
- Craft, D., and M. Monz (2010), Simultaneous navigation of multiple Pareto surfaces, with an application to multicriteria IMRT planning with multiple beam angle configurations, *Medical Physics*, *37*(2), 736–741.
- Craft, D., T. Halabi, and T. Bortfeld (2005), Exploration of tradeoffs in intensity-modulated radiotherapy, *Physics in Medicine and Biology*, *50*, 5857–5868.
- Craft, D., T. Halabi, H. Shih, and T. Bortfeld (2006), Approximating convex pareto surfaces in multiobjective radiotherapy planning, *Mathematical Programming*, *33*(9), 3399–3407.
- Craft, D., D. McQuaid, J. Wala, W. Chen, E. Salari, and T. Bortfeld (2012), Multicriteria VMAT optimization, *Medical Physics*, *39*(2), 686–696.
- Craft, D., M. Bangert, T. Long, D. Papp, and J. Unkelbach (2014), Shared data for intensity modulated radiation therapy (IMRT) optimization research: the cort dataset, *GigaScience*, *3*(1), 37.
- de la Zerda, A., B. Armbruster, and L. Xing (2007), Formulating adaptive radiation therapy (art) treatment planning into a closed-loop control framework, *Physics in Medicine and Biology*, *52*(14), 4137.
- Deasy, J., A. Blanco, and V. Clark (2003), Cerr: a computational environment for radiotherapy research, *Medical Physics*, *30*(5), 979–985.
- Ding, G. X., D. M. Duggan, C. W. Coffey, M. Deeley, D. E. Hallahan, A. Cmelak, and A. Malcolm (2007), A study on adaptive IMRT treatment planning using kv cone-beam ct, *Radiotherapy and Oncology*, *85*(1), 116–125.
- Dong, P., P. Lee, D. Ruan, T. Long, E. Romeijn, Y. Yang, D. Low, P. Kupelian, and K. Sheng (2013a), 4π non-coplanar liver sbrt: A novel delivery technique, *International Journal of Radiation Oncology Biology Physics*, *85*(5), 1360 – 1366.

- Dong, P., P. Lee, D. Ruan, T. Long, E. Romeijn, D. A. Low, P. Kupelian, J. Abraham, Y. Yang, and K. Sheng (2013b), 4π noncoplanar stereotactic body radiation therapy for centrally located or larger lung tumors, *International Journal of Radiation Oncology Biology Physics*, 86(3), 407 – 413.
- Dong, P., D. Nguyen, D. Ruan, C. King, T. Long, E. Romeijn, D. A. Low, P. Kupelian, M. Steinberg, Y. Yang, and K. Sheng (2014), Feasibility of prostate robotic radiation therapy on conventional c-arm linacs, *Practical Radiation Oncology*, 4(4), 254 – 260.
- D’Souza, W., R. Meyer, and L. Shi (2004), Selection of beam orientations in intensity-modulated radiation therapy using single-beam indices and integer programming, *Physics in Medicine and Biology*, 49(15), 3465–3481.
- Ehrgott, M., A. Holder, and J. Reese (2008), Beam selection in radiotherapy design, *Linear Algebra and its Applications*, 428(5-6), 1272–1312.
- Ehrgott, M., Ç. Güler, H. Hamacher, and L. Shao (2010), Mathematical optimization in intensity modulated radiation therapy, *Annals of Operations Research*, 175, 309–365.
- Goitein, M., M. Abrams, D. Rowell, H. Pollari, and J. Wiles (1983), Multi-dimensional treatment planning: II. beam’s eye-view, back projection, and projection through ct sections, *International Journal of Radiation Oncology Biology Physics*, 9(6), 789–797.
- Hoffmann, A., A. Siem, D. den Hertog, J. Kaanders, and H. Huizenga (2006), Derivative-free generation and interpolation of convex Pareto optimal IMRT plans, *Physics in Medicine and Biology*, 51, 6349–6369.
- Hong, T., D. Craft, F. Carlsson, and T. Bortfeld (2008), Multicriteria optimization in intensity-modulated radiation therapy treatment planning for locally advanced cancer of the pancreatic head, *International Journal of Radiation Oncology Biology Physics*, 72(5), 1208–1214.
- Hurkmans, C., I. Dijckmans, M. Reijnen, J. van der Leer, C. van Vliet-Vroegindewij, and M. van der Sangen (2012), Adaptive radiation therapy for breast IMRT-simultaneously integrated boost: Three-year clinical experience, *Radiotherapy and Oncology*, 103(2), 183–187.
- Jee, K., D. McShan, and B. Fraass (2007), Lexicographic ordering: intuitive multicriteria optimization for IMRT, *Physics in Medicine and Biology*, 52, 1845–1861.
- Kong, F., X. Ao, L. Wang, and T. Lawrence (2008), The use of blood biomarkers to predict radiation lung toxicity: a potential strategy to individualize thoracic radiation therapy., *Cancer control: journal of the Moffitt Cancer Center*, 15(2), 140–150.
- Küfer, K.-H., M. Monz, A. Scherrer, P. Süß, F. Alonso, A. Sultan, T. Bortfeld, and C. Thieke (2009), Multicriteria optimization in intensity modulated radiotherapy planning, in *Handbook of Optimization in Medicine*, edited by P. Pardalos and H. Romeijn, chap. 5, pp. 123–167, Springer.

- Lawrence, Y., X. Li, I. El Naqa, C. Hahn, L. Marks, T. Merchant, and A. Dicker (2010), Radiation dose-volume effects in the brain, *International Journal of Radiation Oncology Biology Physics*, 76(3, Supplement), S20–S27.
- Lee, C.-H., D. Aleman, and M. Sharpe (2011), A set cover approach to fast beam orientation optimization in intensity modulated radiation therapy for total marrow irradiation, *Physics in Medicine and Biology*, 56(17), 5679.
- Lim, G., and W. Cao (2012), A two-phase method for selecting IMRT treatment beam angles: Branch-and-Prune and local neighborhood search, *European Journal of Operational Research*, 217(3), 609–618.
- Lim, G., J. Choi, and R. Mohan (2007), Iterative solution methods for beam angle and fluence map optimization in intensity modulated radiation therapy planning, *OR Spectrum*, 30(2), 289–309.
- Long, T., M. Matuszak, M. Feng, B. Fraass, R. T. Haken, and H. Romeijn (2012), Sensitivity analysis for lexicographic ordering in radiation therapy treatment planning, *Medical Physics*, 39(6), 3445–3455.
- Mallat, S. (1993), Matching pursuits with time-frequency dictionaries, *IEEE Transactions on Signal Processing*, 41(12), 3397–3415.
- Marks, L., E. Yorke, A. Jackson, R. T. Haken, L. Constine, A. Eisbruch, S. Bentzen, J. Nam, and J. Deasy (2010), Use of normal tissue complication probability models in the clinic, *International Journal of Radiation Oncology Biology Physics*, 76(3, Supplement), S10–S19.
- Matuszak, M., J. Steers, T. Long, D. McShan, B. Fraass, H. Romeijn, and R. T. Haken (2013), Fusionarc optimization: A hybrid volumetric modulated arc therapy (VMAT) and intensity modulated radiation therapy (IMRT) planning strategy, *Medical Physics*, 40(7), 071,713.
- Mayo, C., E. Yorke, and T. Merchant (2010), Radiation associated brainstem injury, *International Journal of Radiation Oncology Biology Physics*, 76(3, Supplement), S36–S41.
- Meedt, G., M. Alber, and F. Nüsslin (2003), Non-coplanar beam direction optimization for intensity-modulated radiotherapy., *Physics in Medicine and Biology*, 48(18), 2999–3019.
- Men, C., H. Romeijn, Z. Taşkın, and J. Dempsey (2007), An exact approach to direct aperture optimization in IMRT treatment planning, *Physics in Medicine and Biology*, 52(24), 7333–7352.
- Men, C., X. Jia, and S. Jiang (2010), GPU-based ultra-fast direct aperture optimization for online adaptive radiation therapy, *Physics in Medicine and Biology*, 55, 4309–4319.
- Michalski, J., H. Gay, A. Jackson, S. Tucker, and J. Deasy (2010), Radiation dose-volume effects in radiation-induced rectal injury, *International Journal of Radiation Oncology Biology Physics*, 76(3, Supplement), S123–S129.

- Mišić, V., D. Aleman, and M. Sharpe (2010), Neighborhood search approaches to non-coplanar beam orientation optimization for total marrow irradiation using IMRT, *European Journal of Operational Research*, 205(3), 522–527.
- Nguyen, D., P. Dong, T. Long, D. Ruan, D. A. Low, E. Romeijn, and K. Sheng (2014), Integral dose investigation of non-coplanar treatment beam geometries in radiotherapy, *Medical Physics*, 41(1).
- Niemierko, A. (1999), A generalized concept of equivalent uniform dose, *Medical Physics*, 26(6), 1100.
- Nishi, T., Y. Nishimura, T. Shibata, M. Tamura, N. Nishigaito, and M. Okumura (2013), Volume and dosimetric changes and initial clinical experience of a two-step adaptive intensity modulated radiation therapy (IMRT) scheme for head and neck cancer, *Radiotherapy and Oncology*, 106(1), 85–89.
- Papp, D., and J. Unkelbach (2014), Direct leaf trajectory optimization for volumetric modulated arc therapy planning with sliding window delivery, *Medical Physics*, 41(1), 011,701.
- Peng, F., X. Jia, X. Gu, M. Epelman, H. E. Romeijn, and S. Jiang (2012), A new column-generation-based algorithm for VMAT treatment plan optimization, *Physics in Medicine and Biology*, 57(14), 4569–4588.
- Preciado-Walters, F., R. Rardin, M. Langer, and V. Thai (2004), A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer, *Mathematical Programming*, 101(2), 319–338.
- Pugachev, A., and L. Xing (2001), Pseudo beam’s-eye-view as applied to beam orientation selection in intensity-modulated radiation therapy, *International Journal of Radiation Oncology Biology Physics*, 51(5), 1361–1370.
- Pugachev, A., A. Boyer, and L. Xing (2000), Beam orientation optimization in intensity-modulated radiation treatment planning, *Medical Physics*, 27(6), 1238–1245.
- Roach, III, M., J. Nam, G. Gagliardi, I. El Naqa, J. Deasy, and L. Marks (2010), Radiation dose-volume effects and the penile bulb, *International Journal of Radiation Oncology Biology Physics*, 76(3, Supplement), S130–S134.
- Romeijn, H., R. Ahuja, J. Dempsey, A. Kumar, and J. Li (2003), A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning, *Physics in Medicine and Biology*, 48(21), 3521–3542.
- Romeijn, H., J. Dempsey, and J. Li (2004), A unifying framework for multi-criteria fluence map optimization models, *Physics in Medicine and Biology*, 49, 1991–2013.
- Romeijn, H., R. Ahuja, J. Dempsey, and A. Kumar (2005), A column generation approach to radiation therapy treatment planning using aperture modulation, *SIAM Journal on Optimization*, 15(3), 838–862.

- Ruzika, S., and M. Wiecek (2005), Approximation methods in multiobjective programming, *Journal of Optimization Theory and Applications*, 126(3), 473–501.
- Shapiro, A., D. Dentcheva, et al. (2014), *Lectures on stochastic programming: modeling and theory*, vol. 16, SIAM.
- Shepard, D., M. Ferris, G. Olivera, and T. Mackie (1999), Optimizing the delivery of radiation therapy to cancer patients, *Siam Review*, 41(4), 721–744.
- Shepard, D., M. Earl, X. Li, S. Naqvi, and C. Yu (2002), Direct aperture optimization: a turnkey solution for step-and-shoot IMRT, *Medical Physics*, 29(6), 1007–1018.
- Thieke, C., T. Bortfeld, and K.-H. Küfer (2002), Characterization of dose distributions through the max and mean dose concept, *Acta Oncologica*, 41(2), 158–161.
- Tropp, J., A. Gilbert, and M. Strauss (2006), Algorithms for simultaneous sparse approximation. part i: greedy pursuit, *Signal Processing*, 86, 572–588.
- Unkelbach, J., T. Bortfeld, D. Craft, M. Alber, M. Bangert, R. Bokrantz, D. Chen, R. Li, L. Xing, C. Men, S. Nill, D. Papp, E. Romeijn, and E. Salari (2015), Optimization approaches to volumetric modulated arc therapy planning, *Medical Physics*, 42(3), 1367–1377.
- Viswanathan, A., E. Yorke, L. Marks, P. Eifel, and W. Shipley (2010), Radiation dose-volume effects of the urinary bladder, *International Journal of Radiation Oncology Biology Physics*, 76(3, Supplement), S116–S120.
- Waley, A. (1961), *Monkey*, Penguin UK.
- Warkentin, B., P. Stavrev, N. Stavreva, C. Field, and B. Fallone (2004), A tcp-ntcp estimation module using dvhs and known radiobiological models and parameter sets, *Journal of Applied Clinical Medical Physics*, 5(1).
- Webb, S., and A. Nahum (1993), A model for calculating tumour control probability in radiotherapy including the effects of inhomogeneous distributions of dose and clonogenic cell density, *Physics in Medicine and Biology*, 38(6), 653.
- Wilkins, J., J. Alaly, K. Zakarian, W. Thorstad, and J. Deasy (2007), IMRT treatment planning based on prioritizing prescription goals, *Physics in Medicine and Biology*, 52, 1675–1692.
- Wu, Q., D. Thongphiew, Z. Wang, B. Mathayomchan, V. Chankong, S. Yoo, W. Lee, and F. Yin (2008), On-line re-optimization of prostate IMRT plans for adaptive radiation therapy, *Physics in Medicine and Biology*, 53(3), 673.
- Xia, P., and L. Verhey (1998), Multileaf collimator leaf sequencing algorithm for intensity modulated beams with multiple static segments, *Medical Physics*, 25(8), 1424–1434.
- Yarmand, H., B. Winey, and D. Craft (2013), Guaranteed epsilon-optimal treatment plans with the minimum number of beams for stereotactic body radiation therapy, *Physics in Medicine and Biology*, 58(17), 5931.