

Optimization Methods for Volumetric Modulated Arc Therapy and Radiation Therapy Under Uncertainty

by

Fei Peng

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
2013

Doctoral Committee:

Professor H. Edwin Romeijn, Co-Chair
Associate Professor Marina A. Epelman, Co-Chair
Associate Professor Amy M. Cohn
Professor Jeffrey A. Fessler
Clinical Instructor Martha M. Matuszak

© Fei Peng 2013

All Rights Reserved

ACKNOWLEDGEMENTS

This thesis would not have been possible without the advice from my advisors, Dr. Edwin Romeijn and Dr. Marina Epelman. I am deeply grateful for Dr. Romeijn's support and guidance throughout my time in graduate school. He encourages me to challenge and not be content with myself, and his attention to the big picture has always led me to think more about the "why" below the surface of a problem. I want to give my sincerest thanks to him for being patient, supportive and inspiring. I am truly lucky to have him as an advisor. I would also like to thank Dr. Epelman for her advice every step along the way. Her attention to detail has helped me to never lose sight of the practical and important issues and motivations, and she can always give me suggestions to improve my research in ways I had not thought of. She is a great mentor, and I am grateful for the sustained help she provided me over the past few years.

I appreciate Dr. Amy Cohn, Dr. Jeff Fessler, and Dr. Martha Matuszak for serving on my committee. I have had the pleasure to work with Dr. Cohn on a research project during my graduate studies. Her enthusiasm and wholeheartedness have taught me a great deal. Her dedication to teaching also helped me tremendously as a new GSI. I thank Dr. Fessler for his feedback on my research, and for his help in my medical imaging class. Being on a topic outside of IOE, his class was extremely challenging for me, yet his teaching was one of the those I enjoyed the most at U of M. I am grateful to Dr. Matuszak for providing her experience and perspective. Her suggestions and advice helped me understand and focus more attention on the clinical side of the problems, which are of great importance in radiation therapy research. It has been a breeze working with all of my committee members over the

past couple of years. Their advice and perspectives have helped a lot in shaping this thesis. They are the best committee I can ask for.

I want to also thank Dr. Steve Jiang's group at UC San Diego for giving me the opportunity to spend two summers there for research, and introducing me to many practical issues in radiation therapy. Dr. Jiang is a great person to learn from, and to sit around with. The time I spent in San Diego was both productive and enjoyable, and I am grateful for that.

I wish to thank the staff members at IOE for their support during my time here. Tina has never failed to brighten my day with a big smile on her face. She is the most positive person I have ever met.

I have made some wonderful friends at IOE. The company of Troy Long, Ilbin Lee, Robert Riggs, and Majid Al-Gwaiz has made the time in the office feel much shorter. They are the ones I share the happiness over a good joke, or a bad day with little progress.

I am blessed to have parents like mine. Their love and encouragement have carried me over many obstacles in my life, and they have always trusted me. I am forever grateful to them for what they have done for me. Finally, I would like to thank my wife Sujie Liu. Graduate school is tough, and I could not have made it through without her. She is always there when I am down, and her belief in me has helped me put faith in myself in the moments of doubt. She is the best thing to happen to me in my life.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	ii
LIST OF FIGURES	vii
LIST OF TABLES	ix
LIST OF APPENDICES	xi
CHAPTER	
I. Introduction and Preliminaries	1
1.1 Background	1
1.2 Intensity modulated radiation therapy	2
1.2.1 Treatment planning for IMRT	4
1.2.2 Fluence map optimization	6
1.2.3 Evaluating a treatment plan — Dose Volume Histogram	9
1.3 Volumetric modulated arc therapy	10
1.3.1 VMAT with constant gantry speed and dose rate	12
1.4 Uncertainty and adaptive radiation therapy	13
1.5 Contributions and outline of the thesis	14
II. VMAT Treatment Plan Optimization	17
2.1 Introduction	17
2.2 The VMAT optimization model	20
2.3 An algorithm for solving (MP)	23
2.4 Restricted master problem	25
2.5 Intermediate stage master problem	27
2.6 Pricing problem	29
2.6.1 Pricing problem derivation	29
2.6.2 Solving the pricing problem	31
2.7 From (MP) to (FP): gantry speeds and dose rates	33
2.8 Other considerations	34

2.8.1	Upper bound on fluence rate	34
2.8.2	Lower bounds on fluence and dose rates	35
2.8.3	Interdigitation and other MLC constraints	35
2.8.4	Transmission effects	37
2.9	Data and implementation	38
2.10	Benchmark	41
2.11	Performance of VMAT plans	41
2.11.1	The most greedy heuristic	42
2.11.2	Effect of the value of s on treatment quality and time	43
2.11.3	Least greedy heuristic	44
2.11.4	VMAT treatment plan quality compared to benchmark	45
2.11.5	Rate of change in gantry speed	47
2.11.6	Transmission dose	49
2.12	Conclusions	51
 III. VMAT with Constant Gantry Speed and Dose Rate		52
3.1	Introduction	52
3.2	VMATC optimization problem formulation	55
3.3	Decomposition approaches to (VC)	57
3.3.1	Decomposition approach 1: gantry speed and dose rate selection as a sub-problem	57
3.3.2	Decomposition approach 2: aperture selection as a sub-problem	59
3.4	Methods for (approximately) solving $(LS^{(S,R)})$	60
3.4.1	Determining apertures at all control points given set \mathcal{C}	61
3.4.2	Methods for selecting/refining apertures	62
3.5	Two frameworks for finding (approximate) solutions to (VC)	66
3.5.1	Framework A: approximately solving (SR)	66
3.5.2	Framework B: alternating optimization	70
3.6	Experiments	72
3.6.1	Exact vs direct search methods for refining apertures	74
3.6.2	Overall strategy under frameworks A and B	75
3.6.3	Comparison with VMAT plans	80
3.7	Conclusions	82
 IV. Adaptive Radiation Therapy		85
4.1	Introduction	85
4.2	The adaptive IMRT optimization model	90
4.3	Setting m a-priori	92
4.4	Monte-Carlo bounding techniques	94
4.4.1	Lower bound	94
4.4.2	Upper bound	95
4.4.3	Calculating the confidence intervals	96

4.4.4	Applying the bounding technique to problem (P)	97
4.5	Comparison with a model in the literature	99
4.6	Experiments	102
4.6.1	Comparing the stochastic model with the conventional model	105
4.6.2	Re-optimization and adaptive radiation therapy	107
4.6.3	Comparison with model (CH)	111
4.7	Conclusions	115
V. Conclusions and Future Research		120
5.1	VMAT treatment plan optimization problem	121
5.2	Treatment planning for VMATC	122
5.3	IMRT plan optimization under uncertainty	123
APPENDICES		124
BIBLIOGRAPHY		135

LIST OF FIGURES

Figure

1.1	A Varian Truebeam radiation therapy treatment system. Image courtesy of Varian Medical Systems of Palo Alto, California. Copyright 2013, Varian Medical Systems. All rights reserved.	2
1.2	(a): leaves in the MLC block the radiation from the source, creating an irregular aperture; (b): view of the source from the MLC. Images courtesy of Varian Medical Systems of Palo Alto, California. Copyright 2013, Varian Medical Systems. All rights reserved.	4
1.3	(a): 3D geometry of a prostate cancer patient; (b): one CT slice of the same patient	5
1.4	Illustration of a fluence map consisting of intensities of beamlets in seven different beam angles	9
1.5	A typical DVH plot for prostate cancer treatment plan	10
1.6	The set of control points along the treatment arc. The circle in the middle represents the patient on the treatment couch.	11
2.1	Flow chart for the column generation based greedy heuristic for (MP) . . .	26
2.2	Distribution of 177 control points around the arc for (a) Cases 1 — 4, and (b) Case 5.	39
2.3	DVH comparison for Cases 1 — 5. Solid: 177-beam IMRT; dashed: VMAT	46
2.4	Gantry speeds (left column) and dose rates (right column) vs. control points for Case 2: (a) $s = 6$ deg/sec, ignoring rate of change constraint; (b) $s = 6$ deg/sec, considering rate of change constraints; (c) $s = 4$ deg/sec, ignoring rate of change constraint; (d) $s = 4$ deg/sec, considering rate of change constraints.	48
2.5	Case 2 $s = 6$, (a): final dose with (dashed) and without (solid) transmission; (b): final dose with (dashed) and without (solid) transmission, where transmission dose normalized to 95% target coverage; (c): normalized final dose with transmission, plan optimization with (dashed) and without (solid) considering transmission	50
3.1	Possible outcomes of one Nelder-Mead iteration, starting from initial simplex $\{y^0, y^1, y^2\}$	70
3.2	Optimization scheme for VMATC treatment plans	71

3.3	Objective function value in exact and direct search methods for refining apertures	75
3.4	Objective function value in exact and direct search methods for refining apertures, with the direct search method run until convergence	76
3.5	(a): 15 initial triangles tested in the Nelder-Mead algorithm; (b): (S, R) pairs (stars) corresponding to the best 80% solutions	78
3.6	5 additional triangles tested in the Nelder-Mead algorithm	78
3.7	Objective, treatment time and total MU comparison for all 20 triangles tested in the Nelder-Mead algorithm. 1 — 9: small; 10 — 13: medium; 14 — 15: large; 16 — 20: additional five triangles	79
3.8	(S, R) pairs (star) corresponding to the best 80% solutions from 10 initial runs under Framework B	81
3.9	Objective, treatment time and total MU of 10 plans under framework B (star) relative to the best plan from framework A	83
3.10	Comparison of average aperture size for VMAT (orange) vs VMATC (green) plans	83
3.11	DVH of VMATC (dashed) and VMAT plans (solid) for Cases 1-5	84
4.1	DVH clouds for the stochastic plans (solid) and the conventional plans (dashed) for 5 sample treatments	106
4.2	Probability of covering a % volume of the ITV in one fraction for conventional vs. stochastic plans	107
4.3	Coverage probability for conventional vs. adjusted stochastic plan, Case 3 .	108
4.4	DVH clouds for the initial stochastic plans (solid) and adaptive plans (dashed) for Case 1 in five sample treatments. Left: reoptimizing once, right: weekly reoptimization	116
4.5	Coverage probability for initial stochastic and adaptive plans for Cases 1. (a): reoptimizing once at fraction 20; (b-1) — (b-5): weekly reoptimization at fractions 8, 15, 22, 29, and 36, respectively	116
4.6	DVH clouds for the stochastic plans (solid) and the plans from model (CH) (dashed) without adaptive reoptimization for 5 sample treatments	117
4.7	Probability of covering a certain volume of the ITV in one fraction for stochastic plans vs. plans obtained by solving model (CH)	118
4.8	DVH clouds for the adaptive (SAA) plans (solid) and the adaptive plans from model (CH) (dashed) of 5 sample treatments for Case 1. Left: reoptimization once; right: weekly reoptimization	118
4.9	Probability of covering a certain volume of the ITV in one fraction for adaptive stochastic plans vs. adaptive plans from model (CH) for Case 1. (a): reoptimizing once at fraction 20; (b-1) — (b-5): weekly reoptimization at fractions 8, 15, 22, 29, and 36, respectively	119
A.1	A schematic diagram of the source, MLC leaf, and isocenter plane	127
A.2	A ray passing the MLC tip and a second ray tangent to the leaf end	128
B.1	Schematic of a ray tangent to the leaf end and a ray intersecting the leaf .	130

LIST OF TABLES

Table

2.1	Problem dimensions of different cases (after downsampling).	39
2.2	Physical machine parameters.	40
2.3	Performance of 177-beam IMRT treatment plans.	42
2.4	Performance of VMAT treatment plans with $s = S^L = 0.83$ deg/sec.	43
2.5	Performance of VMAT treatment plans for Case 2 obtained with different s values.	44
2.6	Performance of VMAT treatment plans with $s = S^U$	45
2.7	Impact of disallowing interdigitation on algorithm run time.	47
2.8	Impact of the presence of the bound ΔS on the rate of change in gantry speed on treatment time.	49
2.9	Impact of considering transmission on optimization runtime, treatment time and total MU	50
3.1	Clinical criteria for the prostate cancer cases.	73
3.2	Problem dimensions of the downsampled cases.	73
3.3	Machine parameters used in VMATC experiments.	74
3.4	Average runtime of the aperture selection problem from 30 random (S, R) pairs	77
3.5	VMATC and full VMAT treatment plans for cases 1-5	82
4.1	Problem dimensions of our test cases.	103
4.2	DVH criteria for critical structures.	104
4.3	Sample sizes and bounds in the initial (SAA) problem.	105
4.4	Sample sizes and bounds on the objective function in the adaptive reoptimization problems for Case 1.	109
4.5	Effect on cumulative dose of increasing the intensity by 10% at different stages of the treatment course.	110
4.6	Comparison of (SAA) solution \hat{x} with solution from model (CH) for the initial stochastic optimization problem	112
4.7	Comparison of (SAA) solutions with those from model (CH) for reoptimization problems under one-time reoptimization for Case 1	112
4.8	Comparison of (SAA) solutions with those from model (CH) for weekly reoptimization problems for Case 1	113

4.9	Relative difference between actual and expected dose over the undelivered fractions at different stages of the treatment course	114
4.10	Relative difference between actual and expected dose over the entire treatment at different stages of the treatment course	114

LIST OF APPENDICES

Appendix

A.	Relationship between Moving Speed of Radiation Field Edge and Moving Speed of MLC Leaf	125
B.	Derivation of Transmission Coefficient for a Beamlet Outside the Radiation Field	129
C.	Convergence of the SAA Approach	132

CHAPTER I

Introduction and Preliminaries

1.1 Background

Cancer is the second leading cause of death in the United States, responsible for approximately 23% of all deaths according to the National Institute of Health. Of all cancer patients, approximately 50% undergo radiation therapy as part of their treatment.

Radiation therapy is one of the cancer treatment methods. It uses high-energy radiation to control and eliminate the tumor by damaging the DNA of the malignant cancer cells. As radiation is deposited to the cancerous regions, it also affects the cells in the normal tissues that receive radiation dose in this process, leading to various side effects. Therefore, radiation therapy treatments must be carefully designed and delivered in order to minimize the side effects and achieve the desired treatment outcome.

There are two common types of radiation therapy: internal radiation therapy (also known as brachytherapy) and external beam radiation therapy. Brachytherapy, sometimes called internal radiation therapy, is a procedure that involves placing radioactive seeds or sources in or near the tumor, thereby irradiating a very localized area and reducing the radiation delivered to healthy tissues away from the radioactive material. Throughout the rest of this thesis we will focus on external beam radiation therapy, for which radiation is delivered from a treatment machine outside of the patient body. The treatment machine is equipped with a linear accelerator that accelerates electrons, which can be used to treat the patient directly

or to create photons which are used to treat the patient. The radiation is delivered through a gantry-mounted two-dimensional beam that is aimed at precise areas of the patient body. Figure 1.1 is a Truebeam™ commercial radiation therapy treatment system produced by Varian Medical Systems, Inc.



Figure 1.1: A Varian Truebeam radiation therapy treatment system. Image courtesy of Varian Medical Systems of Palo Alto, California. Copyright 2013, Varian Medical Systems. All rights reserved.

The amount of radiation used in radiation therapy is measured in Gray (Gy), which defines the absorption of one joule of radiation energy by one kilogram of matter. While the total amount prescribed in the treatment depends on many factors including the type and progression of the cancer, the treatment is usually delivered in a series of daily sessions spread over several weeks. Each of these sessions is referred to as a *fraction*. Delivering the treatment in small daily doses exploits the difference in radiation responses between the tumor and normal cells, and allows the normal cells to recover between fractions (*Halperin et al.* (2008)).

1.2 Intensity modulated radiation therapy


Intensity modulated radiation therapy (IMRT) is an advanced radiation therapy treatment method, for which equipment and treatment planning algorithms have been contin-

uously developed for more than a decade. With the help of computer-controlled linear accelerator, the IMRT machine is capable of delivering a high-precision dose distribution that conforms to the three-dimensional shape of the tumor, and creates sharp dose gradients (measured by how quickly the dose changes between two adjacent areas with different dose levels) to effectively avoid the tissues surrounding the tumor (see *Intensity Modulated Radiation Therapy Collaborative Working Group* (2001) for an overview of IMRT).


IMRT treatments usually consist of sequential radiation delivery from a (usually small) number of pre-defined beam angles around the patient body. The use of different beams allows treatment planners to effectively avoid placing critical organs directly in the path of radiation targeted at the tumor, and deliver the desired dose distribution collectively from all beam angles. Selecting the number and orientation of beam angles can be considered as an optimization branch in and of itself and has spurred the interest of many researchers (see *Stein et al.* (1997); *Pugachev et al.* (2001); *Jia et al.* (2011) for example). We will focus on cases for which the selection of the beam angles is performed by experienced planners and done before the treatment plan optimization.

The IMRT treatment machines are equipped with a device called the *Multileaf Collimator* (MLC). The MLC is comprised of dozens of leaf pairs that can stay stationary or move during the treatment. Individual leaves in the MLC work together to dynamically change the shape of the exposed beam (also known as *aperture*). Figure 1.2 (a) illustrates the location of the MLC inside the gantry and the aperture created with the MLC, and (b) is a view of the radiation source through the leaves. Note that the fact that the radiation source is a point leads to the size of the radiation field being different from the size of the MLC aperture, and the movement speed of the radiation field edge different from the physical speed of the leaves. We explain the relationship between these quantities in detail in Appendix A. Throughout the rest of this thesis, we assume that the apertures are MLC openings projected onto the isocenter plane, which is the plane passing the center of the tumor and parallel to the beam, and the leaf speed is the projected leaf speed for individual leaves. The planner can easily

convert these quantities back to those used by the MLC system before the plan is delivered on the treatment machine. In IMRT treatments, a beam angle usually includes multiple apertures, each with a distinctive intensity level. The apertures at the same beam angle are delivered sequentially, and together contribute to an intensity “landscape” for the beam, which is usually referred to as the *fluence map*. The ability of IMRT machines to position the MLC leaves with high precision allows planners to control the intensity of very small areas of the beams, e.g., $5\text{mm}\times 2\text{mm}$ in size, called *beamlets*. Determining the intensities for individual beamlets is one of the most important components of the treatment planning process, which we will describe below.



(a)



(b)

Figure 1.2: (a): leaves in the MLC block the radiation from the source, creating an irregular aperture; (b): view of the source from the MLC. Images courtesy of Varian Medical Systems of Palo Alto, California. Copyright 2013, Varian Medical Systems. All rights reserved.

1.2.1 Treatment planning for IMRT

The treatment planning process for radiation therapy starts with a visualization of the patient’s internal geometry through the use of Computed Tomography (CT) images. By

examining the CT images, physicians identify regions that contain cancer tumor that need to be irradiated, and regions that belong to critical organs which need to be spared in order to avoid side effects. Figure 1.3 (a) shows the 3D geometry of a prostate cancer patient, and (b) shows one CT slice that contains the internal structures of the same patient.




Figure 1.3: (a): 3D geometry of a prostate cancer patient; (b): one CT slice of the same patient

Based on this anatomical information and the type of cancer, a physician then defines the appropriate prescription dose for the tumor, as well as a set of criteria used for evaluating the dose distribution to the tumor, such as “at least 95% of the target volume must receive the prescription dose”. Usually it is not possible to entirely spare the normal organs of any radiation dose and deliver the prescribed dose to the tumor due to the proximity of critical organs, especially when the cancerous cells have invaded into such organs. The physician will, in addition, prescribe restrictions on the dose to normal organs, such as “no more than 10% of the brainstem volume can receive more than 10 Gy of dose”. Often times the restrictions for the target and the critical structures are conflicting with each other, and the process of determining their specifics usually depends on a combination of clinical experience and guidelines provided by existing treatment protocols.

The next step in the treatment planning process is to perform dose calculation. A large number of *voxels*, obtained by dividing the patient body (usually evenly) into a three-dimensional grid, as sample points to measure the dose absorbed. The dose calculation

process determines the dose contribution to different voxels if a beamlet is exposed for certain amount of time (called the *dose deposition coefficient*). The IMRT treatment machine is configured so that the amount of radiation output per unit of time, which we refer to as *monitor unit* (MU), is fixed. However, because different tissues have different radiation absorption, the amount of radiation that reaches each point in the patient body also depends on factors including, among others, the type of tissues the radiation beam passes, and the sequence they are passed. Therefore the dose calculation process needs to be done while taking the patient-specific anatomy into account. Common dose calculation methods include pencil beam (*Ahnesjö et al. (1992); Jeleń et al. (2005)*), model-based methods (*Mackie et al. (1985); Papanikolaou et al. (1993); Van Esch et al. (2006)*), and Monte Carlo simulation or Monte-Carlo-like algorithms (*Rogers and Bielajew (1990); Wang et al. (1998); Jia et al. (2010); Fogliata et al. (2011)*). Once the dose calculation process is completed, the dose deposition coefficients are used as input parameters in the treatment plan optimization problem.

1.2.2 Fluence map optimization

The treatment plan optimization problem, of which the basic type is called *fluence map optimization*, determines the optimal fluence map from all beam angles that collectively deliver the treatment plan that can best achieve the prescription dose for the target, and meet or exceed the normal tissue requirements as much as possible. The fluence map optimization is done by solving an “inverse” planning problem, which starts with a desired dose distribution and ends with the beamlet intensities that can best achieve this distribution:

- \mathcal{V} : the set of all voxels
- \mathcal{I} : the set of all beamlets
- d_{ij} : the dose deposition coefficient corresponding to beamlet $i \in \mathcal{I}$ and voxel $j \in \mathcal{V}$
- x_i : decision for the intensity of beamlet $i \in \mathcal{I}$

- z_j : decision for the dose delivered to voxel $j \in \mathcal{V}$
- $F(\cdot)$: voxel based penalty function

$$\begin{aligned}
 & \text{(FMO) minimize}_{x,z} && F(\mathbf{z}) \\
 & \text{subject to:} && \\
 & && z_j = \sum_{i \in \mathcal{I}} d_{ij} x_i && \forall j \in \mathcal{V} \\
 & && x_i \geq 0 && \forall i \in \mathcal{I}.
 \end{aligned}$$

Here we make the reasonable assumption that the dose delivered from different beamlets is additive. The objective function used in the optimization problem needs to be able to properly capture the different requirements we have for different structures, as well as to address the trade-off between structures. Researchers have proposed various functional forms for the objective function for this purpose, including:

- tumor control probability (TCP) and normal tissue complication probability (NTCP) (*Wolbarst (1984); Lyman (1985); Zaider and Minerbo (1999)*). These methods first fit the dose response data of a single cell to a mathematical function (which usually has a sigmoidal shape), then derive the probability of having no clonogenic cells in the target, and the probability of not causing complications to the normal organs, after receiving a (homogeneous or inhomogeneous) dose. For the target we want the probability to be as close to 1 as possible, and for the normal organs as close to 0 as possible.
- equivalent uniform dose (EUD) and generalized EUD (gEUD) (*Niemierko (1997, 1999)*). The EUD method establishes an equivalent homogeneous dose for an inhomogeneous dose distribution: if the EUD is uniformly distributed to the entire target volume, the number of surviving clonogens will be the same as that under the corresponding inhomogeneous dose distribution. This concept was later extended to make it applicable to critical organs, and its name changed to gEUD. The gEUD has often been a method

of choice and used by many researchers since its introduction (*Choi and Deasy (2002); Thieke et al. (2003); Wu et al. (2005)*).

- voxel-based convex function. Rather than trying to model the biological effects, this family of approaches measures the quality of the treatment plan by the value of a convex function at the dose delivered to each voxel, and use the average of such function values as the objective function. The resulting function can thus be separated for each voxel, making it simple to work with mathematically in the optimization problem. Typical choices for the convex function include absolute value function and least square functions.

Romeijn et al. (2004) showed that optimizing with most of the above objective functions or a combination of them will lead to the same treatment plan as optimizing with a corresponding voxel-based convex objective function. Therefore, we will use structure-specific, voxel-based functions as the objective function in our optimization models. It provides simplicity in the optimization problem, and can generate high quality treatment plans (see, for example, *Shepard et al. (1999); Romeijn et al. (2003)*).

Besides the simple constraints used in (FMO) above, investigators have utilized constraints of different types to enforce restrictions on the intensities and/or voxel dose, or to achieve a desired dose distribution. Common examples of such constraints include minimum, maximum, or average voxel dose in a certain structure, constraints based on biological criteria such as TCP and EUD, and excess or shortfall criteria (see *Romeijn and Dempsey (2008)* for an overview of related topics).

The solution to problem (FMO) provides us with a vector of beamlet intensities, as well as the dose delivered to each voxel given this intensity profile. Figure 1.4 shows a fluence map consisting of intensities of beamlets in seven different beam angles, where brighter colors corresponds to higher intensities. The beamlet intensities are subsequently fed into a post-processing step (also called the leaf sequencing step). The leaf sequencing step ensures that the fluence map with irregular intensities can be delivered by the MLC by converting it into

a number of apertures, each with uniform intensity. It is desirable that the beam-on time (the time the source is turned on in the treatment) and/or the number of total apertures is minimized as a result of the leaf sequencing step (see *Crooks et al. (2002)*; *Taşkın et al. (2010)*).




Figure 1.4: Illustration of a fluence map consisting of intensities of beamlets in seven different beam angles

1.2.3 Evaluating a treatment plan — Dose Volume Histogram

The most common way of evaluating the quality of a treatment plan, or equivalently the corresponding dose distribution, is by examining the associated Dose Volume Histogram (DVH). DVH is a method for summarizing the 3-D dose distribution in different structures with a set of 2-D curves, and it has been one of the most important tools in the process of evaluating IMRT treatment plans. Figure 1.5 is a typical DVH plot for prostate cancer. Each curve on the plot represents an individual structure, and a point on a curve represents a dose value (horizontal axis) and the corresponding percentage of volume (vertical axis) in the corresponding structure that receive dose above that value. We can visualize on the DVH many common constraints and criteria that impose upper or lower bounds on the percentage of volume at certain dose value (diamond markers) as well.

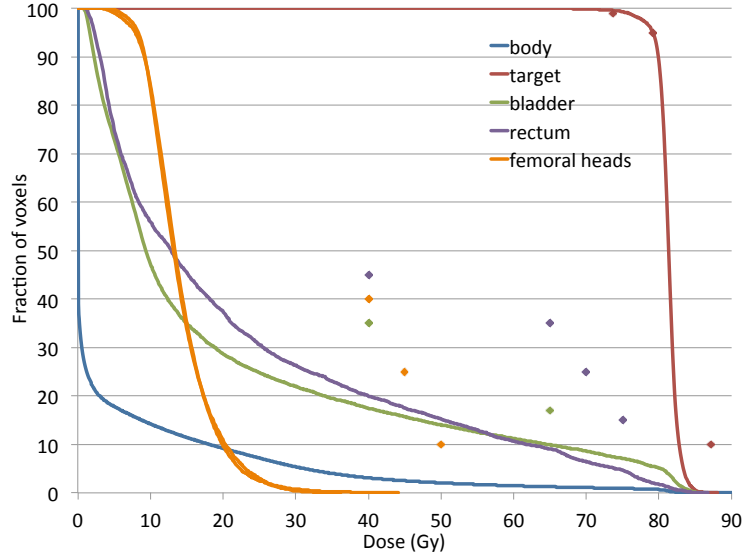


Figure 1.5: A typical DVH plot for prostate cancer treatment plan

1.3 Volumetric modulated arc therapy

Volumetric modulated arc therapy (VMAT) is a new radiation therapy treatment modality that delivers radiation while the gantry and the attached radiation source are in continuous motion. VMAT allows the source output (*dose rate*), gantry speed and the aperture shape to simultaneously vary. As opposed to IMRT treatments which require setup time in between consecutive beam angles, VMAT treatments are usually delivered in one or two continuous rotations (also referred to as *arcs*), therefore drastically reducing the treatment time. Since *Yu* (1995) introduced the intensity modulated arc therapy, which later led to the full capacity VMAT systems, VMAT has attracted the attentions of both the medical and the operations research communities. While VMAT delivery equipments and treatment planning algorithms are still being actively developed, clinical studies (see, for example, *Cozzi et al.* (2008); *Bertelsen et al.* (2010); *Matuszak et al.* (2010)) that compare VMAT to traditional IMRT and other types of treatment modalities with existing commercial systems have found that VMAT is capable of both substantially reducing the overall treatment time and providing comparable or superior treatment plan quality. As a result of the shorter treatment delivery time, clinics can benefit from increased throughput and alleviated patient

discomfort during the treatment. Moreover, reduced delivery time also means the treatments are less susceptible to intra-fraction motions, i.e., motions that happen during the treatment, such as breathing and organ fill-up/depletion, therefore further improving the quality of the actual treatments.

The concept of a *control point* is introduced to help treatment planners to control and describe a VMAT treatment. Although decisions involved in a treatment, namely gantry speed, dose rate and aperture shape, need to be specified at every angle along the gantry trajectory, doing this will quickly make the planning problem intractable. Instead, it is a common practice in research as well as in commercial planning systems to restrict these decisions to be made at only a pre-determined set of angles, namely the control points. The treatment machine is configured so that in between consecutive control points, the gantry speed, dose rate as well as aperture shape transition smoothly. Figure 2.2 shows a sample set of control points along a sample treatment arc. Note that the treatment arc can utilize the entire 360° circle, or can have a small gap under the treatment couch for reasons such as to avoid delivery from under the couch, and to avoid placing a critical organ directly in between the radiation source and the target.

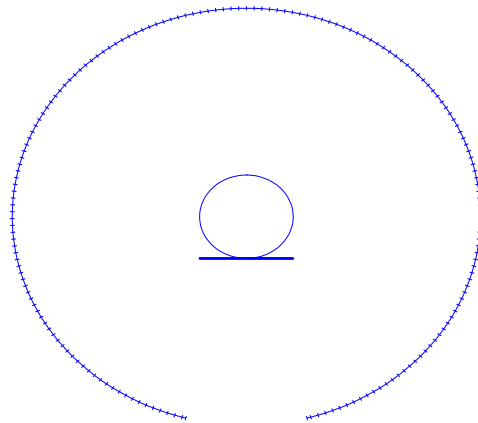


Figure 1.6: The set of control points along the treatment arc. The circle in the middle represents the patient on the treatment couch.

Despite the advantages of the VMAT technology, treatment planning for VMAT is much

more challenging than for IMRT. Because VMAT utilizes a much larger set of angles, the amount of data involved in the treatment planning process, as well as the efforts required in formulating and solving the plan optimization problem are much larger. Moreover, the continuous gantry motion adds restrictions on the relationship between the machine parameters used at different control points. In particular, the change in gantry speed from one control point to a subsequent one must be compatible with the treatment machine's physical acceleration/deceleration limits, and the change in aperture shapes cannot be too drastic between two control points. Therefore, special care must be taken when designing treatment plans for VMAT. We will discuss an optimization algorithm for VMAT in Chapter II.

1.3.1 VMAT with constant gantry speed and dose rate

IMRT machines equipped with conventional linear accelerator and MLC are able to deliver rotational arc therapy that resemble VMAT treatments. This type of treatment has less flexibility compared to full VMAT, allowing only constant gantry speed and dose rate during the treatment. However, like VMAT, it has the freedom to dynamically change the aperture shapes through the utilization of the MLC. It was first referred to as the Intensity Modulated Arc Therapy (IMAT) or Dynamic Arc Conformal (DAC) radiotherapy. Research efforts and clinical experiments on IMAT/DAC were the earliest of those on VMAT before the invention of dedicated VMAT treatment systems (*Yu (1995); Ma et al. (2001); Verellen et al. (2002); Crooks et al. (2003); Cao et al. (2007)*). In fact, some have used the term IMAT to refer to treatments methods in which dose rate and/or gantry speed can change dynamically. To avoid confusion, we refer to VMAT treatments that only allow constant gantry speed and dose rate as VMATC in the rest of this thesis.

Even though VMATC treatments lack some of the most important capabilities of VMAT, especially those that allow the gantry speed and dose rate to dynamically change during the treatment, fully realizing its potentials can still be beneficial. Being able to deliver the treatment in one rotation potentially allows the delivery time to be reduced, thus allowing

clinics with conventional IMRT systems to enjoy many of the benefits of VMAT treatments and avoid the large capital expenditure required to upgrade to the full VMAT systems.

In spite of the similarities between VMAT and VMATC treatments, VMATC's lack of the ability to change gantry speed and dose rate substantially limits our choice of apertures in the planning process, and poses difficulties that need to be addressed with substantially different methods from those used for VMAT. In Chapter III we propose and study optimization algorithms for the VMATC treatment planning problem.

1.4 Uncertainty and adaptive radiation therapy

The presence of uncertainty in the delivery of radiation therapy treatments often causes the treatment quality to be compromised and the actual dose distribution to deviate from the planned one. Setup errors in the positioning of the patient, changes in patient anatomy due to the changed progression of the disease, which happen in between fractions, as well as intra-fraction uncertainty that are typically caused by, e.g., breathing and bowel motions during the treatment all contribute to uncertainties in the treatment.

The use of immobilization devices have been explored in clinics to control and reduce uncertainty (*Bentel (1998)*). However, these aids cannot fully eliminate small errors and motion that affect the position of the patient relative to the radiation beams. The traditional approach to dealing with this uncertainty is by expanding the Clinical Target Volume (CTV), which is delineated from the CT images of a stationary patient, by an internal margin to produce the Internal Target Volume (ITV), and a margin to form the Planning Target Volume (PTV). In the treatment plan optimization the PTV is used as the target to account for uncertainties during the treatment. However, many researchers have shown that the use of margin may not achieve the desired target coverage, and may introduce overdosing of the organs surrounding the PTV (for example, *Olafsson and Wright (2006)*; *Men et al. (2011)*).

Alternatively, researchers have applied various techniques in stochastic optimization, dynamic programming and robust optimization in order to explicitly incorporate uncertainty

into the plan optimization problem (see, for example, *Löf et al. (1999)*; *Baum et al. (2006)*). However, it is important to ensure that in solving the optimization models, the solution does not only depend on extreme cases but also takes the majority of possible scenarios into account, and that the optimization models are not oversimplified in the solution process. In Chapter IV we propose a stochastic optimization based model, and a solution procedure that can guarantee high quality solutions.

Recent advancements in imaging technology has made it possible to monitor and record the information associated with the treatment delivery in each fraction, such as fraction-to-fraction setup variation and organ deformation, as well as to reconstruct in vivo the 3D dose distribution for a given fraction (*Partridge et al. (2002)*; *Yang et al. (2007)*). This additional information allows researchers and practitioners to not only verify the correct delivery of the planned treatment, but also review and correct any discrepancies in the actual dose distribution. We extend our stochastic optimization based model to adaptive treatments, where plan optimization is performed during the treatment course to account for the actual delivered dose.

1.5 Contributions and outline of the thesis

The rest of this thesis is organized as follows:

- Chapter II discusses a method for solving the treatment plan optimization problem for VMAT. We propose a new column generation based algorithm that takes into account bounds on the gantry speed and dose rate, as well as an upper bound on the rate of change of the gantry speed, in addition to MLC constraints. To our knowledge, the constraints have not all been considered in any existing research. The algorithm iteratively adds one aperture at each control point along the treatment arc. In each iteration, a restricted problem optimizing intensities at previously selected apertures is solved, and its solution is used to formulate a pricing problem, which selects an

aperture at another control point that is compatible with previously selected apertures and leads to the largest rate of improvement in the objective function value of the restricted problem. Once a complete set of apertures is obtained, their intensities are optimized and the gantry speeds and dose rates are adjusted to minimize treatment time while satisfying all machine restrictions. Comparisons of treatment plans obtained by our algorithm to idealized IMRT plans of 177 beams on 5 clinical prostate cancer cases demonstrate high quality with respect to clinical dose-volume criteria. For all cases our algorithm yields treatment plans that can be delivered in around 2 minutes. Our approach can be easily modified to accommodate a range of other settings of the VMAT treatment machine.

- Chapter III considers the VMATC treatment plan optimization problem. We specifically consider the simultaneous optimization of the constant gantry speed and dose rate, in addition to the optimization of leaf positions. We propose two algorithmic frameworks for (approximately) solving the optimization problem. One framework separates leaf setting optimization, which chooses the leaf positions given a dose rate and gantry speed combination, from the optimization of the constant dose rate and gantry speed. The framework searches for a dose rate and gantry speed combination that leads to the best plan qualities. The alternative framework applies alternating optimization to dose rate and gantry speed and the leaf settings. We compare both optimization frameworks on clinical patient cases, and compare the resulting VMATC plans to the VMAT plans. Our results show that VMATC is capable of producing very high quality treatment plans compared to VMAT, albeit at the expense of long computation time and higher total radiation output from the source.
- Chapter IV focuses on the optimization of IMRT treatment plans under inter-fraction uncertainty. We propose a stochastic model that incorporates the uncertainty in dose delivered in one fraction, as well as over the treatment course. Instead of solving

this problem directly, which easily becomes intractable for real-world cases, we draw samples from the error distribution and solve the sample average approximation problem. We apply a dynamic sampling procedure that establishes bounds on the objective function, therefore enabling us to find verifiably high quality, approximate solutions to the optimal solution. We show with clinical prostate cancer cases that using the stochastic model can improve the treatment plan quality compared to the conventional approach. If delivered dose information is available during the treatment course, we can extend our solution framework to adaptive radiation therapy optimization problems. However, in the adaptive optimization problems the solution puts more emphasis on the per-fraction target coverage, and usually results in higher dose to the surrounding structures. This effect needs to be taken into account when using this model in adaptive treatment planning.

- In Chapter V we conclude the thesis, and propose directions for future research.

CHAPTER II

VMAT Treatment Plan Optimization

2.1 Introduction

Rotational delivery of Intensity Modulated Radiation Therapy (IMRT) was first proposed by *Yu* (1995) as Intensity Modulated Arc Therapy (IMAT). This radiation therapy treatment modality delivers radiation while the gantry rotates around the patient in one or more arcs while, at the same time, field shapes (apertures) are formed to modulate fluence using a multi-leaf collimator (MLC) system. Recently, this modality has gained popularity and interest due to the introduction of several commercially available implementations, and is now usually referred to as Volumetric Modulated Arc Therapy (VMAT). These advances have also made it possible to dynamically change the dose rate and gantry speed. As compared to more standard IMRT treatments, which use a relatively small number of fixed beam directions, VMAT has the potential to significantly reduce treatment time. This has the additional benefits of decreasing patient discomfort, mitigating intrafraction motion uncertainties, and increasing the utilization rate of the equipment.

Since the conception of VMAT, researchers have investigated inverse planning techniques as well as methods to make VMAT plans deliverable in one single arc. Various studies utilized techniques including simulated annealing (*Cameron* (2005); *Earl et al.* (2003)) and transformation of an IMRT treatment plan (*Crooks et al.* (2003)) to design VMAT treatments. These studies showed that, with constant gantry speed and dose rate, either multiple arcs or

an excessively long time for a single arc are required to deliver a high-quality treatment. In a theoretical study, *Bortfeld and Webb (2009)* compared single-arc VMAT treatment to traditional IMRT on a phantom and argued that single-arc VMAT may not be able to replicate the optimal intensity profile without intensity modulation at individual angles – a claim that was denied in a response by *Otto (2009)*. Today the potential quality of VMAT treatment plans is generally believed to be comparable to IMRT treatment plans (*Palma et al. (2008)*; *Verbakel et al. (2009)*).

The nature of VMAT treatments makes the treatment plan optimization problem much more complex than an IMRT treatment plan optimization problem due to the continuous nature of the gantry motion, gantry speed constraints, dose rate constraints, and MLC leaf speed constraints. Even when the set of angles is discretized, exact methods quickly become intractable. Hence the literature has exclusively focused on simplifications of the model and/or heuristic solution approaches. Many approaches to the VMAT treatment plan optimization problem rely on converting an IMRT plan in some fashion. *Luan et al. (2008)*, *Shepard et al. (2007)*, and *Wang et al. (2008)* proposed algorithms that rely on (i) first identifying an optimal IMRT treatment plan using 10° -spaced beam directions and then (ii) producing a deliverable VMAT plan by minimizing the difference between the sequenced and the ideal IMRT fluence maps. *Cao et al. (2009)* started the treatment planning process by performing IMRT planning with direct machine parameter optimization, which directly optimizes aperture shapes and weights. They then used a simulated annealing based arc sequencer to convert the resulting fluence maps into deliverable VMAT plans while again minimizing the difference between the IMRT and VMAT fluence maps. *Craft et al. (2011)* performed IMRT optimization on 2° -spaced beam angles and, noting that adjacent fluence maps were similar, used a combination of beam merging and sliding window leaf sequences to obtain a VMAT deliverable plan whose quality is close to the “ideal” IMRT plan. However, their observation that adjacent fluence maps in the high-resolution IMRT treatment plan are similar appears to depend strongly on the particular treatment plan optimization algorithm

employed, so the approach may not be robust to changes in treatment plan evaluation criteria that would require using a different algorithm.

Several approaches to VMAT treatment plan optimization use the concept of *control points*, which represent beam directions along the arc where machine parameters such as gantry speed, aperture shape, and dose rate are controlled. Between control points, the machine parameters are then either kept constant or interpolated in some fashion. *Otto* (2008) proposed a simulated annealing algorithm that adds control points sequentially and, in each iteration, randomly samples aperture weights and shapes. *Bedford* (2009) discretized an IMRT fluence map into multiple intensity levels and assigned each level to a control point close to the corresponding angle in the IMRT plan. A subsequent local search procedure attempts to improve leaf positions and segment weights. *Bzdusek et al.* (2009) started with an IMRT plan on a coarse set of angles, then generated two apertures to approximate each of the fluence maps and placed these at predetermined control points. Additional control points were assigned interpolations of these apertures, and the remaining machine parameters were calculated in a subsequent optimization step. *Gözbasi* (2010) formulated the problem as a large integer programming model, but for tractability reasons resorted to grouping beamlets and solved a series of approximate problems. Unfortunately, the required computation time of the algorithm are still impractical and the algorithm also failed to consistently provide clinically acceptable results on their test cases.

Finally, *Men et al.* (2010b) proposed a column generation based approach which sequentially adds aperture shapes at control points while ensuring that each newly generated aperture is compatible with previously chosen ones. However, no bounds on the dose rate were taken into account, and gantry speed was considered to be constant. In this chapter we formalize and extend this approach to explicitly take the following physical restrictions into account during the course of the algorithm:

- apertures at consecutive control points must be compatible, i.e., the leaves in individual rows of the MLC must be able to finish their transition from the previous aperture to

the next aperture before the gantry reaches the next control point;

- restrictions on machine parameters, such as upper and lower bounds on the gantry speed and dose rate, as well as the upper bound on the rate of change of the gantry speed must be obeyed.

Although some of the algorithmic approaches proposed in the literature and discussed above account for some of these restrictions by either explicitly incorporating them into the treatment plan optimization or by modifying the optimized plan into one that satisfies them in a post-processing step, to the best of our knowledge none incorporates all of the restrictions explicitly in the treatment plan optimization process. Our algorithm is implemented using the Compute Unified Device Architecture (CUDA) to take advantage of the parallel processing power of the Graphic Processing Unit (GPU).

2.2 The VMAT optimization model

Let K denote the total number of (not necessarily equispaced) control points along one or more treatment arc(s) (in case of multiple arcs, K represents the total number of control points along all arcs, which will be ordered so that the first control point on the second arc succeeds the last control point on the first arc, etc.). We associate with each control point k an aperture A_k , a dose rate r_k (in MU/second), and a gantry speed s_k (in degrees/second). In addition, we can express the fluence rate (in MU/degree) at control point k as $y_k = r_k/s_k$. Note that each control point represents a snapshot of the continuous gantry rotation, i.e., r_k , s_k , y_k , and A_k represent the state of the treatment machine as the gantry passes through control point k . For convenience, we will add a dummy control point 0 in the beginning of the arc. For tractability we will then calculate the dose delivered to each voxel by making the approximation that the aperture, dose rate, and gantry speed (and hence fluence rate) are constant throughout the arc spanning from one control point to the next. Note in particular that this means that we do not need to specify any of the variables at the dummy control

point 0. If the angular distances δ_k between pairs of control points k and $k - 1$ are small, then this approximation will be sufficiently accurate (*Otto* (2008)).

The gantry speeds and apertures specified at consecutive control points need to be compatible with each other. This compatibility requirement derives from the characteristics of the MLC system used to form the apertures during the treatment. As the gantry travels from control point $k - 1$ to control point k , the leaves of the MLC system need to shift from positions prescribed by aperture A_{k-1} to those prescribed by aperture A_k . Since the speed of such leaf movement is bounded, the time spent by the gantry moving between the control points needs to be sufficiently large (and hence its speed sufficiently small) to allow for the required leaf movement to take place. We will denote the maximum gantry speed that will allow aperture A' at control point k to be reached from aperture A at control point $k - 1$ by $S_{k-1,k}^U(A, A')$ for $k = 1, \dots, K$, where, for convenience, $S_{0,1}^U(A, A') = \infty$ for all pairs of apertures A and A' .

In addition to the discretization parameters δ_k ($k = 1, \dots, K$), there are several other machine parameters that need to be taken into account. First, a VMAT delivery machine may have upper and lower bounds on gantry speed (denoted by S^U and S^L), and an upper bound on dose rate (denoted by R^U). Furthermore, there typically is an upper bound on the rate of change in speed that the gantry can sustain, which may be given per degree or per control point. We will denote the upper bound on the change in speed between control points $k - 1$ and k by ΔS_k . Finally, we let \mathcal{A} denote the set of all apertures that are deliverable by the MLC system. In particular, we will assume that left and right leaves within each pair can be positioned at any non-overlapping continuous points within their range, and that motion is independent between leaf pairs. In particular, we assume that there are no interdigitation constraints. After formulating our base model and developing our algorithm, we will discuss how additional constraints, such as a lower bound on dose rate, independent bounds on the fluence rate, and interdigitation constraints, can be incorporated into our solution approach when necessary.

Finally, let the discretized set of voxels be \mathcal{V} , and denote the delivered dose distribution by $\mathbf{z} = (z_j : j \in \mathcal{V})^\top$, where z_j is the total dose delivered to voxel $j \in \mathcal{V}$. The quality of the dose distribution \mathbf{z} is evaluated using an objective function $F : \mathbb{R}^{|\mathcal{V}|} \rightarrow \mathbb{R}$. Under our approximation z_j is equal to the sum of the doses delivered to the voxel from all control points. Moreover, since the delivered dose is linear in fluence, it is convenient to denote the dose received by voxel $j \in \mathcal{V}$ from aperture $A \in \mathcal{A}$ at control point k at unit fluence by $D_{kj}(A)$ ($k = 1, \dots, K$).

The complete VMAT optimization model, which we refer to as the full problem (FP), then reads:

$$\begin{aligned}
\text{(FP) minimize} \quad & F(\mathbf{z}) \\
\text{subject to} \quad & \\
z_j = \sum_{k=1}^K D_{kj}(A_k) \delta_k y_k & \quad j \in \mathcal{V} \\
y_k = \frac{r_k}{s_k} & \quad k = 1, \dots, K \\
|s_k - s_{k-1}| \leq \Delta S_k & \quad k = 2, \dots, K \\
s_k \in [S^L, S^U] & \quad k = 1, \dots, K \\
r_k \in [0, R^U] & \quad k = 1, \dots, K \\
s_k \leq S_{k-1,k}^U(A_{k-1}, A_k) & \quad k = 1, \dots, K \quad (2.1) \\
A_k \in \mathcal{A} & \quad k = 1, \dots, K,
\end{aligned}$$

where the terminal aperture A_0 in (2.1) can be chosen arbitrarily and is added for convenience only.

We can simplify this model by first noting that for any feasible solution $\mathbf{s}, \mathbf{r}, \mathbf{y}, \mathbf{z}$ and $(A_k : k = 1, \dots, K)$ of (FP) we can obtain another feasible solution with the same apertures and the same values of \mathbf{y} and \mathbf{z} (and hence the same objective function value), but with all the gantry speeds at their lower bound value S^L by scaling the dose rates by a factor of $S^L/s_k \leq 1$, $k = 1, \dots, K$, to compensate. As a result, we can eliminate the gantry speed

variables from the model. In addition, we can then replace the upper bound constraints on the dose rates by upper bounds on fluence rate and eliminate the dose rate variables as well. The problem thus reduces to the following *master problem* (MP):

$$\text{(MP) minimize} \quad F(\mathbf{z})$$

subject to

$$z_j = \sum_{k=1}^K D_{kj}(A_k) \delta_k y_k \quad j \in \mathcal{V} \quad (2.2)$$

$$y_k \in [0, Y^U] \quad k = 1, \dots, K \quad (2.3)$$

$$S^L \leq S_{k-1,k}^U(A_{k-1}, A_k) \quad k = 1, \dots, K \quad (2.4)$$

$$A_k \in \mathcal{A} \quad k = 1, \dots, K, \quad (2.5)$$

where $Y^U \equiv R^U/S^L$. We will refer to (MP) as the master problem, for reasons made clear in the following section. By construction, for any feasible solution to (MP) we can construct an equivalent feasible solution to (FP), in the sense that it has the same apertures and values of \mathbf{y} and \mathbf{z} , by setting $s_k = S^L$ and $r_k = y_k \cdot s_k$ for $k = 1, \dots, K$. Moreover, there may be several combinations of gantry speeds and dose rates that lead to equivalent solutions to (FP), and some may be more desirable than others, e.g., due to shorter treatment times. We will discuss this in more detail in Section 2.7.

Note that (MP) is not convex because of the nonlinearities associated with the apertures in constraints (2.2) — (2.5). Therefore, instead of solving problem (MP) exactly, we propose a column generation based algorithm.

2.3 An algorithm for solving (MP)

Our proposed algorithm starts with a dose distribution $\mathbf{z} = \mathbf{0}$ and without an aperture specified at any of the control points. The algorithm then proceeds by, in each iteration, attempting to improve the current treatment plan and dose distribution by selecting an aperture at one of the control points for which none has been specified yet (that is, *generating*

a column in constraint (2.2) of (MP)), and terminates when a completely specified treatment plan is obtained. In any given iteration, which is characterized by a set $\mathcal{C} \subseteq \{1, \dots, K\}$ of control points and the corresponding collection of apertures $\{\bar{A}_k, k \in \mathcal{C}\}$, we optimize fluence rates of these apertures by solving the so-called *restricted master problem* (RMP). The RMP is constructed so that a feasible solution to this problem implicitly provides a deliverable plan that can be fully specified by feasibly completing the sequence of apertures while retaining the current aperture fluence rates (in particular, $y_k = 0$ for $k \notin \mathcal{C}$). The optimal solution to the RMP then defines the so-called *pricing problem* (PP) which is used to select the next control point and aperture to be added to the plan. The intuition behind the PP is that, relative to the current solution, each candidate aperture at each control point $k \notin \mathcal{C}$ has an associated *price* which is defined as the rate of improvement in the objective function value if the fluence rate of that aperture is increased (from its current value of 0). The pricing problem then finds the control point/aperture combination that has the best price. In contrast with direct aperture optimization (DAO) for IMRT treatment planning (*Preciado-Walters et al. (2004)*; *Romeijn et al. (2005)*; *Shepard et al. (2002)*), only a single aperture is added for each control point. This makes our algorithm heuristic in nature and, in particular, a greedy heuristic.

More formally, our algorithm can be described as follows:

Column generation based greedy heuristic for (MP)

Step 0. Set $\mathcal{C} = \emptyset$ and $\bar{\mathbf{z}} = \mathbf{0}$.

Step 1. Use the information on the current treatment plan (control points \mathcal{C} , apertures \bar{A}_k for $k \in \mathcal{C}$, and dose distribution $\bar{\mathbf{z}}$) to formulate and solve an instance of the PP.

Step 2. If the optimal value of the PP is nonpositive, go to Step 5. Otherwise, denote the optimal solution to the PP by \bar{c} and $\bar{A}_{\bar{c}}$, and replace \mathcal{C} by $\mathcal{C} \cup \{\bar{c}\}$.

Step 3. Solve the instance of the RMP associated with \mathcal{C} and $A_k = \bar{A}_k, k \in \mathcal{C}$.

Step 4. Remove apertures in the set $\mathcal{C}=\{k \in \mathcal{C} : y_k = 0\}$, i.e., set $\mathcal{C} \leftarrow \mathcal{C} \setminus \mathcal{C}$.

Step 5. If $|\mathcal{C}| < K$, return to Step 1.

Step 6. If necessary, complete the treatment plan by identifying feasible apertures (which will have fluence 0) at control points $c \notin \mathcal{C}$, and denote the final set of fluence rates by \bar{y}_k ($k = 1, \dots, K$).

Figure 2.1 provides a graphical representation of the structure of this algorithm.

In the remainder of this section, we will provide additional details on Steps 1, 2 and 3 of this algorithm in Sections 2.4 and 2.6. Section 2.7 discusses obtaining a solution to (FP) based on the solution to (MP) returned by our algorithm. Section 2.8 discusses how additional constraints on the setting of the VMAT delivery machine can be incorporated into (FP) formulation and our solution. Finally, we present experiments and results in Section 2.9.

2.4 Restricted master problem

At each iteration of the algorithm we obtain the RMP from (MP) by fixing the apertures $A_k = \bar{A}_k$ for all control points in the current set $\mathcal{C} \subseteq \{1, \dots, K\}$ and setting the fluence rates $y_k = 0$ for all control points $k \notin \mathcal{C}$. The RMP then reads:

$$\begin{aligned}
 \text{(RMP}^{(\mathcal{C})}) \quad & \text{minimize} && F(\mathbf{z}) \\
 & \text{subject to} && \\
 & && z_j = \sum_{k \in \mathcal{C}} D_{kj}(\bar{A}_k) \delta_k y_k && j \in \mathcal{V} \\
 & && y_k \in [0, Y^U] && k \in \mathcal{C}.
 \end{aligned} \tag{2.6}$$

(RMP^(C)) can be thought of as a restriction of (MP) in that it considers a subset of control points, and considers the apertures \bar{A}_k , $k \in \mathcal{C}$, to be fixed and given. Note that (RMP^(C))

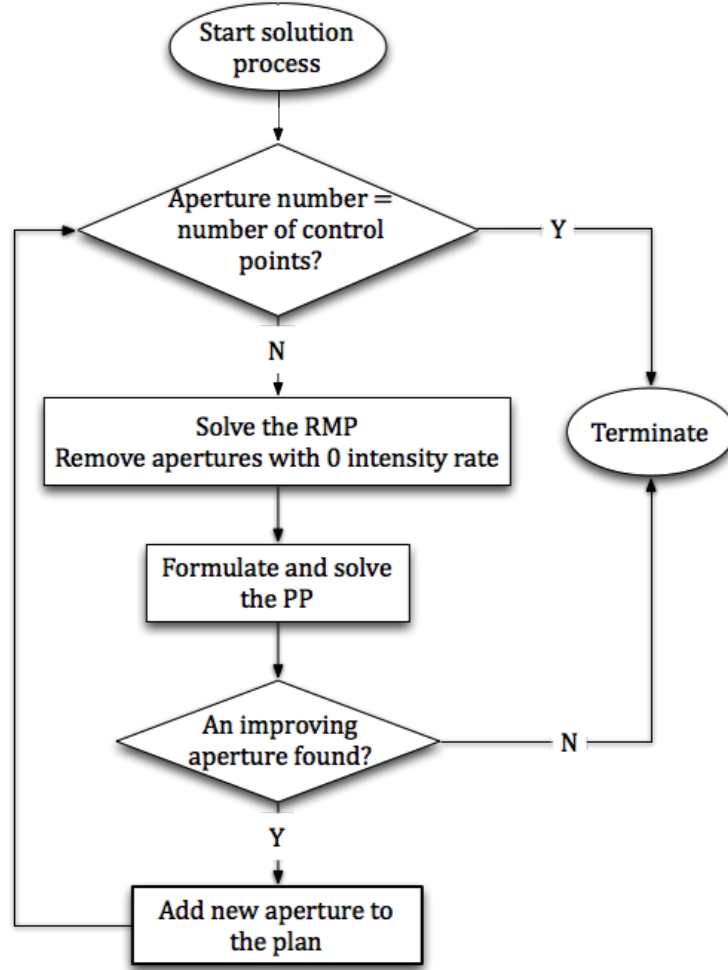


Figure 2.1: Flow chart for the column generation based greedy heuristic for (MP)

is a nonlinear optimization problem with linear constraints, and it is a convex optimization problem provided that F is a convex function.

Recall that any feasible solution to $(RMP^{(\mathcal{C})})$ should correspond to a deliverable plan that can be fully specified by completing the sequence of apertures at control points $k \notin \mathcal{C}$ in a way that is compatible with apertures \bar{A}_k , $k \in \mathcal{C}$, and setting $y_k = 0$, $k \notin \mathcal{C}$. To ensure that this is indeed the case, the set of apertures \bar{A}_k , $k \in \mathcal{C}$ in the restricted problem should satisfy the following condition. For all $k \in \mathcal{C}$, let k^- be the predecessor of k in \mathcal{C} , i.e., $k^- = \max\{k' \in \mathcal{C} : k' < k\}$, where $k^- \equiv 0$ if $k = \min\{k' : k' \in \mathcal{C}\}$. Any feasible solution to $(RMP^{(\mathcal{C})})$ for any $\mathcal{C} \subseteq \{1, \dots, K\}$ can be extended to a feasible solution to (MP), provided

that

$$S_{k-k}^U(\bar{A}_{k-}, \bar{A}_k) \geq S^L, \quad k \in \mathcal{C}. \quad (2.7)$$

Here we have generalized the notation of (FP) and (MP) by denoting the maximum gantry speed that will allow aperture A' at control point k' to be reached from aperture A at control point $k < k'$ by $S_{kk'}^U(A, A')$ (where again, for convenience, we will let $S_{0,k}^U(A, A') = \infty$ for all pairs of apertures A and A' and all $k = 1, \dots, K$). In Section 2.6 we will discuss in detail the pricing problem designed to ensure that the apertures added at each iteration of the column generation algorithm satisfy (2.7).

If, from solving (RMP^(C)), any aperture \bar{A}_k ($k \in \mathcal{C}$) has intensity rate 0, we can remove these apertures without affecting the plan quality. Apertures added early in the algorithm, when the total delivered dose is relatively low, are likely to have large openings and/or high fluence rates. As the algorithm progresses, these apertures are no longer effective, their fluence rates are set to 0 in the (RMP^(C)) and they are subsequently removed. Since only one aperture can be used at every control point, removing these redundant apertures allows better ones to be added later on in the column generation process. Let $\mathcal{Q} = \{k \in \mathcal{C} : y_k = 0\}$ be the set of control points that have redundant apertures, after solving (RMP^(C)), we set $\mathcal{C} \leftarrow \mathcal{C} \setminus \mathcal{Q}$.

2.5 Intermediate stage master problem

The RMP is solved at every iteration to find the fluence rate for apertures $k \in \mathcal{C}$. However, our algorithm is not complete before we have a procedure to search for the next improving aperture. We examine the master problem at an intermediate stage, where only control points $k \in \mathcal{C}$ have apertures specified, and all possible apertures are included for control points $k \notin \mathcal{C}$. This problem is not solved during the optimization process, rather we use its associated first order optimality conditions to find the rate of improvement in objective function value per unit fluence rate for apertures at control points $k \notin \mathcal{C}$. This information will help us identify the best aperture to add to the treatment plan.

Note that for control points $k \notin \mathcal{C}$, we associate an fluence rate value with every possible aperture, therefore there will be several fluence rates associated with one control point, and thus constraints (2.2) and (2.3) will have to be modified to accommodate this. Let

$$y_{kA} \quad k \notin \mathcal{C}, A \in \mathcal{A}.$$

represent the fluence rate for aperture $A \in \mathcal{A}$ at control point k , we can use

$$\sum_{A \in \mathcal{A}} y_{kA} \quad k \notin \mathcal{C}$$

instead of y_k in (2.2) and (2.3). The intermediate stage master problem can then be formulated by modifying formulation (MP) as:

$$(MI) \underset{z, y}{\text{minimize}} \quad F(\mathbf{z})$$

subject to

$$z_j = \sum_{k \in \mathcal{C}} D_{kj}(\bar{A}_k) \delta_k y_k + \sum_{k \notin \mathcal{C}} \sum_{A \in \mathcal{A}} D_{kj}(A) \delta_k y_{kA} \quad j \in \mathcal{V} \quad (\pi_j)$$

$$y_k \geq 0 \quad k \in \mathcal{C} \quad (\rho_k)$$

$$y_k \leq Y^U \quad k \in \mathcal{C} \quad (\gamma_k)$$

$$\sum_{A \in \mathcal{A}} y_{kA} \leq Y^U \quad k \notin \mathcal{C} \quad (\gamma_k)$$

$$y_{kA} \geq 0 \quad k \notin \mathcal{C}, A \in \mathcal{A} \quad (\beta_k(A))$$

We associate the set of constraints in (MI) with dual variables π_j ($j \in \mathcal{V}$), ρ_k ($k \in \mathcal{C}$), γ_k ($k \in \{1, \dots, K\}$), and $\beta_k(A)$ ($k \notin \mathcal{C}, A \in \mathcal{A}$). The first order optimality conditions, which are necessary and sufficient for optimality for this problem (see, for example, *Bazaraa et al.* (2006) for reference) read:

$$-\pi_j \in \frac{\partial F(\mathbf{z})}{\partial z_j} \quad j \in \mathcal{V} \quad (2.8)$$

$$\rho_k - \gamma_k = - \sum_{j \in \mathcal{V}} \pi_j D_{kj}(\bar{A}_k) \delta_k \quad k \in \mathcal{C} \quad (2.9)$$

$$\gamma_k (y_k - Y^U) = 0 \quad k \in \mathcal{C} \quad (2.10)$$

$$\rho_k \cdot y_k = 0 \quad k \in \mathcal{C} \quad (2.11)$$

$$-\gamma_k + \beta_k(A) = -\sum_{j \in \mathcal{V}} \pi_j D_{kj}(A) \delta_k \quad k \notin \mathcal{C}, A \in \mathcal{A} \quad (2.12)$$

$$\gamma_k \left(\sum_{A \in \mathcal{A}} y_{kA} - Y^U \right) = 0 \quad k \notin \mathcal{C} \quad (2.13)$$

$$\beta_k(A) \cdot y_{kA} = 0 \quad k \notin \mathcal{C}, A \in \mathcal{A}$$

$$\gamma_k \geq 0 \quad k \in \{1, \dots, K\}$$

$$\rho_k \geq 0 \quad k \in \mathcal{C}$$

$$\beta_k(A) \geq 0 \quad k \notin \mathcal{C}, A \in \mathcal{A} \quad (2.14)$$

Observe that the optimal solution to problem (RMP^(C)), $(\bar{\mathbf{y}}, \bar{\mathbf{z}})$, and the corresponding dual variables $\pi_j(\bar{\mathbf{z}})(j \in \mathcal{V})$, $\rho_k, \gamma_k(k \in \mathcal{C})$ satisfy conditions (2.8) — (2.11). We can set variables $y_{kA} = 0$ ($k \notin \mathcal{C}, A \in \mathcal{A}$), and set $\gamma_k = 0$ ($k \notin \mathcal{C}$), and verify if (2.12) and (2.14) are satisfied. If so, the current solution is already optimal for problem (MP). Otherwise, (2.12) and (2.14) provide us with the *price* for aperture $A \in \mathcal{A}$:

$$\beta_k(A) = -\sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) D_{kj}(A_k) \delta_k \quad k \notin \mathcal{C}, A \in \mathcal{A}. \quad (2.15)$$

Care needs to be taken to ensure that any aperture we select for $k \notin \mathcal{C}$ is compatible with the apertures in control points $k \in \mathcal{C}$. Next we explain in detail how improving apertures are chosen so that compatibility constraint (2.4) is satisfied.

2.6 Pricing problem

2.6.1 Pricing problem derivation

Suppose we were to add aperture $A \in \mathcal{A}_k$ at a control point $k \notin \mathcal{C}$, where $\mathcal{A}_k \subseteq \mathcal{A}$ is a set of deliverable apertures that can feasibly be added to the current treatment plan. From the discussion in Section 2.5, the rate of *improvement* (i.e., decrease) in objective function

value per unit fluence rate in this aperture can be calculated as

$$\beta_k(A) \equiv - \sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) D_{kj}(A_k) \delta_k,$$

where $(\pi_j(\bar{\mathbf{z}}) : j \in \mathcal{V})^\top = \boldsymbol{\pi}(\bar{\mathbf{z}}) \equiv -\nabla F(\bar{\mathbf{z}})$. Our strategy will be to select a control point k and aperture A that solve the following optimization problem, which we will refer to as the *pricing problem*:

$$\text{(PP)} \quad \max_{k \notin \mathcal{C}} \max_{A \in \mathcal{A}_k} \beta_k(A).$$

In order to ensure that solutions to (PP) satisfy condition (2.7), we need to specify the set \mathcal{A}_k in such a way that

$$\mathcal{A}_k \subseteq \{A \in \mathcal{A} : S_{k-k}^U(\bar{A}_{k-}, A) \geq S^L, S_{kk+}^U(A, \bar{A}_{k+}) \geq S^L\} \equiv \mathcal{A}_k^U.$$

The choice of $\mathcal{A}_k \subseteq \mathcal{A}_k^U$ in the PP dictates a tradeoff in the behavior of our heuristic algorithm. On the one hand, choosing a smaller set \mathcal{A}_k at the current iteration results in lesser flexibility in the current selection of apertures. On the other hand, this can potentially allow for more flexibility in later iterations of the heuristic. We will consider the following family of potential choices for \mathcal{A}_k in (PP), parameterized by a lower bound s on the speed between two control points:

$$\mathcal{A}_k(s) = \{A \in \mathcal{A} : S_{k-k}^U(\bar{A}_{k-}, A) \geq s, S_{kk+}^U(A, \bar{A}_{k+}) \geq s\} \text{ for } s \geq S^L. \quad (2.16)$$

It is easy to see that $\mathcal{A}_k(s') \subseteq \mathcal{A}_k(s)$ whenever $s' \geq s$, with $\mathcal{A}_k(S^L) = \mathcal{A}_k^U$. Note also that, although the choice of s does not explicitly determine the final gantry speed used at different control points, because it affects the selection of apertures $\mathcal{A}_k(s)$, it implicitly affects the final gantry speed and thus the delivery time. We will explore the impact of the choice of parameter s on the overall performance of the algorithm in our experiments.

If the optimal value of (PP) is non-positive, we conclude that the current solution cannot be improved by adding any aperture at any control point (at least in the framework of our algorithm). In this case, the algorithm will be terminated (see Step 2), and fluence rates at control points $k \notin \mathcal{C}$ will remain at zero. However, if an improving control point/aperture

are found, the algorithm will proceed to Step 3.

2.6.2 Solving the pricing problem

We will now describe a solution approach for the pricing problem (PP). First, it is easy to see that this problem decomposes into a collection of independent problems for the different candidate control points k , so that we can focus on the problem for a given $k \notin \mathcal{C}$:

$$(\text{PP}^{(k)}) \max_{A \in \mathcal{A}_k} \sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) D_{kj}(A),$$

where we have also eliminated the scaling constant δ_k .

Next, in the absence of interdigitation constraints, we can decompose $(\text{PP}^{(k)})$ by MLC row. In particular, if the MLC system consists of M leaf pairs, we can represent any aperture $A \in \mathcal{A}$ as a collection of leaf settings for its individual rows. Let us denote an aperture by $A = (a_1, \dots, a_M)$, where $a_m \in \alpha$ describes the leaf settings of MLC row m of aperture A , and α is the set of deliverable ‘‘row apertures.’’ The set \mathcal{A}_k can be written as $\mathcal{A}_k = \times_{m=1}^M \mathcal{A}_{km}$, with

$$\mathcal{A}_{km} \subseteq \{a : S_{k-k}^U(\bar{a}_{k-m}, a) \geq S^L, S_{kk+}^U(a, \bar{a}_{k+m}) \geq S^L\}, \quad (2.17)$$

where \bar{a}_{cm} is the m th row aperture of \bar{A}_c at control point $c \in \mathcal{C}$ and, with a slight abuse of notation, $S_{k'k}^U(a, a')$ denotes the maximum gantry speed that will allow row aperture a' of aperture A' at control point c to be reached from row aperture a at control point $k' < k$. Let $D_{kmj}(a)$ denote the dose received by voxel $j \in \mathcal{V}$ from row aperture $a \in \alpha$ in row m and control point k at unit fluence rate, with $D_{kj}(A) = \sum_{m=1}^M D_{kmj}(a_m)$. The pricing problem $(\text{PP}^{(k)})$ can then be decomposed by row:

$$(\text{PP}^{(km)}) \max_{a \in \mathcal{A}_{km}} \sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) D_{kmj}(a).$$

Recall that we would like to consider the family of feasible regions $\mathcal{A}_k = \mathcal{A}_k(s)$ defined by (2.16) for $s \geq S^L$. The corresponding family $\mathcal{A}_{km}(s)$ of feasible regions for $(\text{PP}^{(km)})$ is

$$\mathcal{A}_{km}(s) = \{a : S_{k-k}^U(\bar{a}_{k-m}, a) \geq s, S_{kk+}^U(a, \bar{a}_{k+m}) \geq s\} \text{ for } s \geq S^L. \quad (2.18)$$

Let us analyze the structure of these sets in more detail. Any row aperture a can be characterized as a pair of leaf settings (ℓ, r) , where $0 \leq \ell \leq r \leq N$ are the positions of the left and right leaf, respectively (N , which we assume to be integer, is the range of the MLC row). Furthermore, let v denote the maximum leaf speed (in distance/second, with the unit of distance depending on the choice of N). Then, for $k < c$,

$$\{a \in \alpha : S_{kk}^U(a, a') \geq s\} = \left\{ (\ell, r) : 0 \leq \ell \leq r \leq N; |\ell - \ell'|, |r - r'| \leq \frac{v\delta_{ck}}{s} \right\},$$

where $\delta_{ck} = \sum_{k'=c}^{k-1} \delta_{k'}$. Therefore,

$$\mathcal{A}_{km}(s) = \left\{ (\ell, r) : 0 \leq \ell \leq r \leq N; |\ell - \bar{\ell}_{k-m}|, |r - \bar{r}_{k-m}| \leq \frac{v\delta_{k-k}}{s}; \right. \\ \left. |\ell - \bar{\ell}_{k+m}|, |r - \bar{r}_{k+m}| \leq \frac{v\delta_{k+k}}{s} \right\}, \quad (2.19)$$

with the appropriate and obvious definitions of $(\bar{\ell}_{k-m}, \bar{r}_{k-m})$ and $(\bar{\ell}_{k+m}, \bar{r}_{k+m})$. Thus, the set $\mathcal{A}_{km}(s)$ is simply the set of all (ℓ, r) such that $0 \leq \ell \leq r \leq N$ and with upper and lower bounds on ℓ and r which get tighter as s increases.

Finally, in order to fully specify $(PP^{(km)})$, we need to characterize $D_{kmj}(a) = D_{kmj}(\ell, r)$. Specifying $D_{kmj}(\ell, r)$ as a function of continuous variables (ℓ, r) will require an immense amount of data. Instead, we use a common approximation obtained by discretizing each MLC row into N beamlets (thus specifying an appropriate MLC row range), where beamlet n represents the interval $[n-1, n]$. We then precompute traditional beamlet-based dose deposition coefficients D_{kmnj} , i.e., the dose rate to voxel $j \in \mathcal{V}$ from beamlet n in MLC row m at control point k ($n = 1, \dots, N$, $m = 1, \dots, M$, $k = 1, \dots, K$) at unit fluence rate and make the following approximation:

$$D_{kmj}(\ell, r) = \int_{\ell}^r \phi_{kmj}(x) dx,$$

where $\phi_{kmj} : [0, N] \rightarrow \mathbb{R}^+$ is the following step function:

$$\phi_{kmj}(x) = D_{kmnj}, \quad n-1 < x \leq n; \quad n = 1, \dots, N.$$

Note that the coefficient D_{kmnj} is similar, except in indexing, to the dose deposition coefficient

d_{ij} defined in Section 1.2.2.

To summarize, the pricing problem (PP^{km}) with $\mathcal{A}_{km} = \mathcal{A}_{km}(s)$ can be written as:

$$\max_{(\ell, r) \in \mathcal{A}_{km}(s)} \int_{\ell}^r \boldsymbol{\pi}_{km}(x; \bar{\mathbf{z}}) dx, \quad (2.20)$$

where

$$\boldsymbol{\pi}_{km}(x; \bar{\mathbf{z}}) = \sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) \phi_{kmj}(x).$$

The problem (2.20) can be solved efficiently by noting that the only candidate values for ℓ and r that have to be considered are (i) integers, and (ii) the bounds derived from (2.19). This is because function $\boldsymbol{\pi}_{km}(x; \bar{\mathbf{z}})$ is monotonically increasing or decreasing for $n-1 < x \leq n$; $n = 1, \dots, N$, due to the structure of function $\phi_{kmj}(x)$.

2.7 From (MP) to (FP): gantry speeds and dose rates

Our column generation algorithm will return a feasible solution $\bar{\mathbf{y}}, \bar{\mathbf{z}}$ and \bar{A}_k , $k = 1, \dots, K$ to the problem (MP). As a final step of our solution procedure, we need to compute gantry speed and dose rate vectors \mathbf{s} and \mathbf{r} consistent with this solution. Substituting the apertures and fluence rates into constraints of (FP), we need to find a feasible solution to the following system:

$$\bar{y}_k = \frac{r_k}{s_k} \quad k = 1, \dots, K \quad (2.21)$$

$$|s_k - s_{k-1}| \leq \Delta S_k \quad k = 1, \dots, K-1 \quad (2.22)$$

$$s_k \in [S^L, S^U] \quad k = 1, \dots, K \quad (2.23)$$

$$r_k \in [0, R^U] \quad k = 1, \dots, K \quad (2.24)$$

$$s_k \leq S_{k-1, k}^U(\bar{A}_{k-1}, \bar{A}_k) \quad k = 1, \dots, K. \quad (2.25)$$

This system has a feasible solution since the apertures and fluence rates computed by the algorithm are guaranteed to be compatible with the gantry moving at speed S^L . However, this solution would lead to a treatment with undesirably long delivery time. Instead, we

would like to identify a solution to (2.21)–(2.25) that can be delivered in the shortest time possible. Adding this objective function and rewriting constraints (2.21)–(2.25) in terms of the gantry speeds only, we obtain the following convex optimization problem:

$$\begin{aligned}
\text{(SP) minimize} \quad & \sum_{k=1}^K \frac{\delta_k}{s_k} \\
\text{subject to} \quad & \\
& s_k \in [S^L, S_k^U] \quad k = 1, \dots, K \\
& |s_k - s_{k-1}| \leq \Delta S_k \quad k = 1, \dots, K - 1, \quad (2.26)
\end{aligned}$$

where

$$S_k^U = \min \left\{ S^U, S_{k-1,k}^U(\bar{A}_{k-1}, \bar{A}_k), \frac{R^U}{\bar{y}_k} \right\}, \quad k = 1, \dots, K.$$

If \bar{s}_k is an optimal solution to this problem, then the corresponding dose rates are $\bar{r}_k = \bar{y}_k \bar{s}_k$, $k = 1, \dots, K$.

In order to assess the impact on delivery time of the bound on the change in speed we will also consider a relaxation of (SP) in which constraints (2.26) are removed. It is easy to see that the corresponding solution will be to choose $s_k = S_k^U$ and $r_k = \bar{y}_k S_k^U$ (for $k = 1, \dots, K$).

2.8 Other considerations

In this section we discuss how additional constraints on the setting of the VMAT delivery machine can be incorporated into the model (FP), its reformulation (MP), and our solution heuristic.

2.8.1 Upper bound on fluence rate

Our basic model contained no bounds on the fluence rate, other than those implied by bounds on the gantry speed and dose rate. An independent upper bound on the fluence rate can be easily incorporated by including it among the constraints of (FP), and by calculating Y^U in (MP) and (RMP^(C)) as the minimum of this upper bound and the ratio R^U/S^L .

2.8.2 Lower bounds on fluence and dose rates

Our basic model did not include lower bounds on fluence and dose rates, other than nonnegativity constraints. If such lower bounds Y^L and R^U , respectively, are dictated by the VMAT delivery specification, they can be easily included among the constraints of (FP). However, inclusion of one or both such bounds invalidates the reformulation steps that lead to the equivalent problem (MP). We could then, instead, consider a generalization (MP) in which we add lower bounds on the fluence which are the larger of Y^L and R^L/S^U and apply our heuristic solution procedure without modification (aside from the obvious addition of lower bounds on fluence rates to (RMP^(C))). The resulting master problem, however, is no longer equivalent to (FP), but rather is a relaxation. The consequences of using this relaxation are:

- (i) An intermediate solution obtained by solving (RMP^(C)) with $|\mathcal{C}| < K$ may not be extendable to a feasible solution to (FP). This means, in particular, that we have to modify Step 2 of the greedy heuristic and add an aperture even if the optimal solution value to the PP is nonpositive.
- (ii) Postprocessing problem (SP) may not have a feasible solution. Therefore, we may have to slightly modify the final treatment plan to be able to satisfy the constraints on the rate of change in gantry speed.

2.8.3 Interdigitation and other MLC constraints

Lastly, consider a VMAT delivery system that does not allow interdigitation of MLC leaves in adjacent rows. The main impact of this constraint is a modification of the solution process of the pricing problem discussed in Section 2.6.2.

Before discussing the required modification, we will argue that if interdigitation constraints are satisfied in each iteration of the algorithm (i.e., for all apertures \bar{A}_k , $k \in \mathcal{C}$), and we obtain a VMAT plan by linearly interpolating these apertures between control points

in \mathcal{C} , then the interdigitation constraints will also hold for MLC apertures throughout the treatment. To obtain the interpolation, we assume that both the gantry and each of the leaves move at constant speeds between two consecutive control points. Let us examine two adjacent MLC rows m and $m + 1$ at two control points k and k^+ . We assume that

$$\bar{\ell}_{km} \leq \bar{r}_{k,m+1} \text{ and } \bar{\ell}_{k^+m} \leq \bar{r}_{k^+,m+1}, \quad (2.27)$$

i.e., the left leaf in row m and the right leaf in row $m + 1$ do not overlap at either control point, and denote the time it takes the gantry to rotate from k to k^+ by Δt . At an intermediate time $t \in (0, \Delta t)$, with the gantry at location $c : k < c < k^+$, the leaf positions are computed by linear interpolation as follows:

$$\ell_m = \bar{\ell}_{km} + (\bar{\ell}_{k^+m} - \bar{\ell}_{km}) \cdot \frac{t}{\Delta t} = \bar{\ell}_{km} \cdot \frac{\Delta t - t}{\Delta t} + \bar{\ell}_{k^+m} \cdot \frac{t}{\Delta t}$$

and

$$r_{m+1} = \bar{r}_{k,m+1} + (\bar{r}_{k^+,m+1} - \bar{r}_{k,m+1}) \cdot \frac{t}{\Delta t} = \bar{r}_{k,m+1} \cdot \frac{\Delta t - t}{\Delta t} + \bar{r}_{k^+,m+1} \cdot \frac{t}{\Delta t}.$$

It is easy to see that because of (2.27), we have $\ell_m \leq r_{m+1}$; using the same technique we can show that $\ell_{m+1} \leq r_m$, assuming same is true at control points k and $k + 1$. Thus, at any location between the control points the interdigitation constraints are satisfied.

Now we study the impact of the interdigitation constraints on the pricing problem. Because of the dependency between individual MLC rows created by these constraints, we cannot no longer decompose the problem (PP^(k)) by row. As before, we characterize row aperture $a_m \in \alpha$ at MLC row m of aperture A by (ℓ_m, r_m) where $0 \leq \ell_m \leq r_m \leq N$ and the (integer) N represents the range of the MLC row, and denote by $S_{ck}^U(a'_m, a_m)$ the maximum gantry speed that will allow row m of aperture $A = (a_1, \dots, a_M)$ at control point k to be reached from row m of aperture $A' = (a'_1, \dots, a'_M)$ at control point $c < k$. To incorporate the interdigitation constraints, we add the inequalities

$$\ell_m \leq \min\{r_{m+1}, r_{m-1}\}, \quad m = 1, \dots, M \quad (2.28)$$

to the description of the set $A_c(s)$ in (2.16), where we let $r_0 = r_{M+1} = \infty$. By construction,

$A_k(s) \neq \emptyset$ throughout the algorithm.

The pricing problem for $k \notin \mathcal{C}$ under interdigtation constraints can then be written as:

$$\text{(PPI}^{(k)}) \quad \text{maximize} \quad \sum_{j \in \mathcal{V}} \pi_j(\bar{\mathbf{z}}) \sum_{m=1}^M D_{kmj}(\ell_m, r_m)$$

subject to

$$\begin{aligned} 0 \leq \ell_m \leq r_m \leq N & \quad m = 1, \dots, M \\ |\ell_m - \bar{\ell}_{k-m}|, |r_m - \bar{r}_{k-m}| \leq \frac{v\delta_{k-k}}{s} & \quad m = 1, \dots, M \end{aligned} \quad (2.29)$$

$$|\ell_m - \bar{\ell}_{k+m}|, |r_m - \bar{r}_{k+m}| \leq \frac{v\delta_{kk^+}}{s} \quad m = 1, \dots, M \quad (2.30)$$

$$\ell_m \leq \min\{r_{m+1}, r_{m-1}\} \quad m = 1, \dots, M,$$

where $D_{kmj}(\ell_m, r_m)$, as before, is approximated by an integral of the beamlet-based step function $\phi_{kmj} : [0, N] \rightarrow \mathbb{R}^+$. Given this discretization of each MLC row into beamlets, this version of the pricing problem can still be solved efficiently using a Dynamic Programming approach described in *Romeijn et al.* (2005), with potential left and right leaf positions including integers and bounds in (2.29) and (2.30).

2.8.4 Transmission effects

The leaves in the MLC system are designed to block all radiation directed at them. However, since these leaves have finite height, they cannot completely prevent the radiation from passing. *Transmission* refers to the radiation dose leaked through individual leaves and in between adjacent leaves in the MLC. Moreover, many MLC systems utilize leaves that have round ends (*Galvin et al.* (1992)). Transmission effects are more serious in areas near the end of leaves because of the smaller leaf height. *Boyer et al.* (2001) provided recommendations for MLC leaf design that require the leaves attenuate the beam to less than 5%.

Without explicitly taking transmission into account, the optimization dose will differ from the actual delivered dose. In Appendix B we present a detailed derivation of the transmission coefficient, $\alpha(\tau_1, \tau_2)$, for a beamlet that covers a segment that lies between τ_1 and τ_2 , with

$\tau_1 < \tau_2$, from the edge of the radiation field. With this information we can incorporate transmission effects into the dose deposition coefficient for any aperture $A \in \mathcal{A}$. Recall that we defined the coefficient

$$D_{kj}(A) = \sum_{m=1}^M D_{kmj}(\ell_{km}, r_{km}) \quad \forall k = 1, \dots, K, j \in \mathcal{V}.$$

We can expand this definition to include the transmission dose. Define:

$$\begin{aligned} \mathcal{D}_{kj}(A) = & D_{kj}(A) \\ & + \sum_{m=1}^M \left[\sum_{n=0}^{\lfloor \ell_{km} \rfloor - 1} \alpha(\ell_{km} - n - 1, \ell_{km} - n) D_{kmnj} + \alpha(0, \ell - \lfloor \ell_{km} \rfloor) (\ell_{km} - \lfloor \ell_{km} \rfloor) D_{km\lfloor \ell_{km} \rfloor j} \right. \\ & \left. + \alpha(0, \lceil r_{km} \rceil - r_{km}) (\lceil r_{km} \rceil - r_{km}) D_{km\lceil r_{km} \rceil j} + \sum_{n=\lceil r_{km} \rceil}^{N-1} \alpha(n - r_{km}, n - r_{km} + 1) D_{kmnj} \right], \end{aligned}$$

where the un-attenuated dose deposition coefficients for partial beamlets $[\lfloor \ell_{km} \rfloor, \ell_{km}]$ and $[r_{km}, \lceil r_{km} \rceil]$ are approximated by the proportion of those of whole beamlets. We can use coefficient $\mathcal{D}_{kj}(A)$ in the place of $D_{kj}(A)$ to obtain a better approximation of the true coefficient. This can be done in both the restricted master problem and the pricing problem. However, doing this in the pricing problem means that the solution is no longer guaranteed to be either integers or the bounds derived from (2.19), and we may have to modify the solution algorithm and find approximate solutions to the pricing problem. Therefore we only take the transmission effects into account in (RMP^(C)).

2.9 Data and implementation

Our algorithm was evaluated on a clinical dataset of prostate cancer patients treated at UCSD Moores Cancer Center. The preprocessing steps, including image processing and calculation of the dose deposition coefficients D_{kmnj} are performed using our in-house treatment planning system. The dose calculations in this system are based on a finite-size pencil-beam algorithm with 3-D density correction (*Gu et al. (2009, 2011)*). For each case we used beamlets of size $1 \times 1 \text{ cm}^2$ and voxels of size $4 \times 4 \times 2.5 \text{ mm}^3$. The full voxel grid will be used

in evaluating the dose distribution. However, in order to reduce the problem size, the optimization problem used a downsampled voxel grid that selected 1 of every 2 voxels along each of the three dimensions in critical structures and 1 of every 4 in unspecified tissue. The dimensions (after downsampling) and other characteristics of the 5 cases are summarized in Table 2.1. Figure 2.2 illustrates the distribution of control points around the arc. For all cases we used $K = 177$ control points, and the angular distances between control points are shown in Table 2.1. Finally, the physical machine parameters we used are based on the Varian VMAT machines (Varian Medical Systems, Inc., Palo Alto, CA, USA) and are shown in Table 2.2 (where $\Delta S_k = \Delta S$ for all $k = 1, \dots, K$).

Case	# voxels	# beamlets	# $D_{kmnj} \neq 0$	δ_k ($k = 3, \dots, 176$)	δ_2, δ_{177}	δ_1
1	18,375	14,337	37,224,301	$\frac{330}{174}$	$\frac{1}{2} \cdot \frac{330}{174}$	0
2	14,392	11,328	26,751,546	$\frac{330}{174}$	$\frac{1}{2} \cdot \frac{330}{174}$	0
3	18,469	21,417	52,850,633	$\frac{330}{174}$	$\frac{1}{2} \cdot \frac{330}{174}$	0
4	9,747	14,337	47,707,351	$\frac{330}{174}$	$\frac{1}{2} \cdot \frac{330}{174}$	0
5	17,342	17,700	61,606,811	$\frac{340}{174}$	$\frac{1}{2} \cdot \frac{340}{174}$	0

Table 2.1: Problem dimensions of different cases (after downsampling).




Figure 2.2: Distribution of 177 control points around the arc for (a) Cases 1 — 4, and (b) Case 5.

R^L (MU/sec)	R^U (MU/sec)	S^L (deg/sec)	S^U (deg/sec)	ΔS (deg/sec)	v (cm/sec)	SAD (cm)	SCD (cm)
0	10	0.83	6.0	0.75	2.25	100	53.9

Table 2.2: Physical machine parameters.

The prescription dose to the PTV was set to 79.2 Gy, corresponding to a 44-fraction treatment with 1.8 Gy delivered in each fraction. All Dose-Volume Histogram (DVH) based criteria for the PTV and critical structures, in particular rectum, bladder, and femoral heads, that were used to evaluate all treatment plans are based on clinical protocols at UCSD as well as RTOG protocol H-0126 (*Radiation Therapy Oncology Group (2004)*) (see the tables in the remainder of this section for details). In our optimization model, we chose F to be a voxel-based penalty function that penalizes the dose received by a voxel quadratically but asymmetrically with respect to a threshold dose that depends on the structure containing the voxel; consequently, F is a smooth piece-wise quadratic convex function. The structure-dependent relative weights of these penalty functions were determined using a manual procedure that iteratively modified these weights until satisfactory treatment plans were obtained for all 5 cases. This led to a common set of weights for all cases, thereby eliminating variation between cases caused by individual parameter tuning. In order to facilitate the comparison between the treatment plans we normalized each treatment plan by scaling the dose distribution so that 95% of the PTV receives the prescription dose.

Finally, we implemented our algorithm using the Compute Unified Device Architecture (CUDA) to take advantage of the processing power of the Graphic Processing Unit (GPU). GPU is a platform for implementation of large-scale parallelization of computer code, and has been shown to be very suitable for solving radiation therapy treatment plan optimization problem (*Men et al. (2009, 2010a,b)*). The RMP is solved using our in-house gradient based solver. On an Nvidia Tesla C1060 GPU card, our algorithm takes 2 — 6.5 minutes to complete. The high efficiency of our VMAT optimization engine provides a significant step towards a clinical application of VMAT for adaptive radiotherapy, although of course

a treatment planning system for adaptive radiotherapy would contain several additional components.

2.10 Benchmark

Before evaluating our VMAT treatment plans, we provide a benchmark for comparison by solving an IMRT treatment plan optimization problem with the same objective function as (FP) and using the $K = 177$ beam angles that correspond to control points in the VMAT treatment plan optimization model. These IMRT plans are used to assess the quality of our VMAT plans and are not suitable for clinical use due to the unreasonably long time required to deliver an IMRT plan with 177 beam angles. Since the IMRT treatment plan optimization model allows intensity modulation at each beam angle and contains no constraints on aperture compatibility nor gantry speed, they can be viewed as an ideal, likely clinically unattainable, solution. Therefore, any VMAT treatment plan that is close to this ideal IMRT treatment plan will also be close to the optimal VMAT treatment plan.

Table 2.3 includes the clinical dose-volume criteria and the 177-beam IMRT plans evaluated with these criteria for all 5 cases. All plans are able to avoid significant hot- and cold-spots in the PTV, and are able to satisfy most of the DVH criteria for the critical structures. For Case 2 14% of the bladder volume and 18% of the rectum volume overlap with the PTV, and for Case 4 almost 30% of the bladder volume overlaps with the PTV. For these patients, even additional patient-dependent parameter tuning did not allow us to resolve the violations highlighted in the table without introducing more serious ones.

2.11 Performance of VMAT plans

As discussed in Section 2.6.1, the column generation heuristic is parameterized by the parameter s , which represents how “greedy” the heuristic is. Moreover, since the choice of s affects the choice of \bar{A}_k , which affects S_k^U in the time-minimization problem (SP), it also

Structure	Threshold Dose (Gy)	Volume Criterion (%)	Case				
			1	2	3	4	5
PTV	73.7	≥ 99	100	99	100	99	100
	79.2	≥ 95	95	95	95	95	95
	87.1	≤ 10	0	0	0	6	0
Rectum	75	≤ 15	1	11	8	5	5
	70	≤ 25	2	13	11	8	6
	65	≤ 35	3	16	13	10	9
	40	≤ 45	12	26	24	28	20
Bladder	65	≤ 17	5	20	13	42	10
	40	≤ 35	10	38	20	70	20
Femoral heads (L/R)	50	≤ 10	0/0	0/0	0/0	0/1	0/0
	45	≤ 25	0/0	0/0	0/0	0/3	0/0
	40	≤ 40	0/0	0/0	0/0	1/6	0/0

Table 2.3: Performance of 177-beam IMRT treatment plans.

impacts the delivery time of the final treatment plans. In this section we test our heuristic under various settings of s , and study the performance of the resulting treatment plans in terms of quality and delivery time.

2.11.1 The most greedy heuristic

Recall that a smaller value of s means that there are more apertures to choose from in each iteration. We will start by choosing the smallest value of s , i.e., we will start by investigating the most greedy variant of our heuristic. Note that we use the same objective function as for the 177-beam IMRT optimization without any additional tuning, so that the results only reflect the differences between the two modalities and treatment planning techniques.

The results in Table 2.4 show that our VMAT treatment plans exhibit high quality in general, closely resembling the IMRT treatment plans when evaluated with the same set of clinical criteria. The major differences between the IMRT and VMAT plans are with respect to the dose distribution in the femoral heads, but the VMAT plans still easily meet all clinical criteria.

Structure	Threshold Dose (Gy)	Volume Criterion (%)	Case				
			1	2	3	4	5
PTV	73.7	≥ 99	100	100	100	99	100
	79.2	≥ 95	95	95	95	95	95
	87.1	≤ 10	0	0	0	1	1
Rectum	75	≤ 15	2	11	8	4	5
	70	≤ 25	3	15	11	6	7
	65	≤ 35	4	17	14	9	10
	40	≤ 45	13	29	25	20	23
Bladder	65	≤ 17	4	21	12	41	10
	40	≤ 35	11	39	21	66	19
Femoral heads (L/R)	50	≤ 10	0/0	0/0	0/0	0/0	0/0
	45	≤ 25	0/0	0/0	0/0	0/0	0/0
	40	≤ 40	0/0	0/0	0/0	1/0	0/0
Runtime (s)			284.9	323.8	271.4	223.9	299.0
Treatment time (s)			225.4	197.5	242.8	217.5	241.0

Table 2.4: Performance of VMAT treatment plans with $s = S^L = 0.83$ deg/sec.

The second to last row of Table 2.4 contains the plan optimization times, in seconds, using the algorithm for each case, while the last line shows the treatment delivery time for each case. These treatment times range between 3.29 to just over 4 minutes, versus the fastest allowable arc delivery time of 1 minute.

2.11.2 Effect of the value of s on treatment quality and time

Unfortunately, the treatment plans in Table 2.4 require relatively long treatment times, which negates some of the benefits of using VMAT treatments. We therefore next study whether we can choose a larger value of s , which can be expected to decrease treatment time, without significant changes in treatment plan quality. In these tests we focus on Case 2, which is one of the more challenging cases.

Table 2.5 shows the performance of the VMAT treatment plans obtained for Case 2 with various values of s ranging from S^L to S^U . We can conclude that the treatment plan quality is not sensitive to the choice of s . However, the overall treatment time is reduced dramatically, by a factor of more than 1.3 from 3.3 minutes to just over 2 minutes, when using $s = S^U = 6$.

Finally, we provide an estimate of the total number of MUs required for treatment given by $\sum_{k=1}^K \delta_k \bar{y}_k$ where the \bar{y}_k is the fluence rate at control point k in the final solution (see also the definition of the fluence rates y_k in Section 2.2). Note that even though the treatment times for $s = 4$ and $s = 6$ are similar (slightly shorter with $s = 4$), a comparison of total MU shows that using $s = 6$ is more desirable as the total MU required is much lower.

Structure	Threshold Dose (Gy)	Volume Criterion (%)	s (deg/sec)			
			0.83	2	4	6
PTV	73.7	≥ 99	100	100	100	100
	79.2	≥ 95	95	95	95	95
	87.1	≤ 10	0	0	0	0
Rectum	75	≤ 15	11	11	12	12
	70	≤ 25	15	14	15	15
	65	≤ 35	17	17	18	18
	40	≤ 45	29	29	32	32
Bladder	65	≤ 17	21	21	22	21
	40	≤ 35	39	40	42	42
Femoral heads (L/R)	50	≤ 10	0/0	0/0	0/0	0/0
	45	≤ 25	0/0	0/0	0/0	0/0
	40	≤ 40	0/0	0/0	0/2	0/0
Runtime (s)			323.8	260.1	116.1	71.5
Treatment time (s)			197.5	171.8	145.8	149.3
Total MU			873.5	859.1	737.2	667.7

Table 2.5: Performance of VMAT treatment plans for Case 2 obtained with different s values.

2.11.3 Least greedy heuristic

Based on the results in the previous section we focus our further analysis on the least greedy heuristic. Table 2.6 shows the performance of the VMAT treatment plans obtained with $s = S^U = 6$. Similarly to Case 2 we conclude that there is no significant changes clinically in treatment plan quality as compared to the plans in Table 2.4, but the treatment time is greatly reduced for all cases. In particular, the treatment time is consistently just over 2 minutes for all 5 cases.

Structure	Threshold Dose (Gy)	Volume Criterion (%)	Case				
			1	2	3	4	5
PTV	73.7	≥ 99	100	100	100	99	100
	79.2	≥ 95	95	95	95	95	95
	87.1	≤ 10	0	0	0	2	0
Rectum	75	≤ 15	1	12	9	4	5
	70	≤ 25	3	15	13	6	8
	65	≤ 35	4	18	16	9	10
	40	≤ 45	18	32	31	27	28
Bladder	65	≤ 17	5	21	13	43	11
	40	≤ 35	12	42	24	71	22
Femoral heads (L/R)	50	≤ 10	0/0	0/0	0/0	0/1	0/0
	45	≤ 25	0/1	0/0	1/2	1/4	0/0
	40	≤ 40	0/4	0/0	5/5	4/6	3/1
Runtime (s)			78.8	71.5	59.2	54.2	67.6
Treatment time (s)			149.3	149.3	141.1	142.6	155.5
Total MU			611.5	667.7	539.0	557.6	610.4

Table 2.6: Performance of VMAT treatment plans with $s = S^U$.

2.11.4 VMAT treatment plan quality compared to benchmark

Figure 2.3 shows the DVHs of the 177-beam IMRT and VMAT (with $s = S^U = 6$) treatment plans for Cases 1 — 5. As can be gleaned from the corresponding tables, the largest differences in these plans are with respect to the dose distribution in the femoral heads, but these are still well within the clinical protocol. Given the restrictions that VMAT delivery imposes on a treatment plan, we feel that the quality of the VMAT plan is not only clinically acceptable, but also high since the degradation from the idealized 177-beam IMRT plan is relatively small.

2.11.4.1 Disallowing interdigitation

In order to study the effect of using an MLC system that does not allow interdigitation, we implemented the modified pricing problem (PPI) described in Section 2.8.3 and applied our resulting algorithm on all 5 cases. Our experiments revealed no significant change in solution quality, while the solution time increased substantially for all cases, due to the

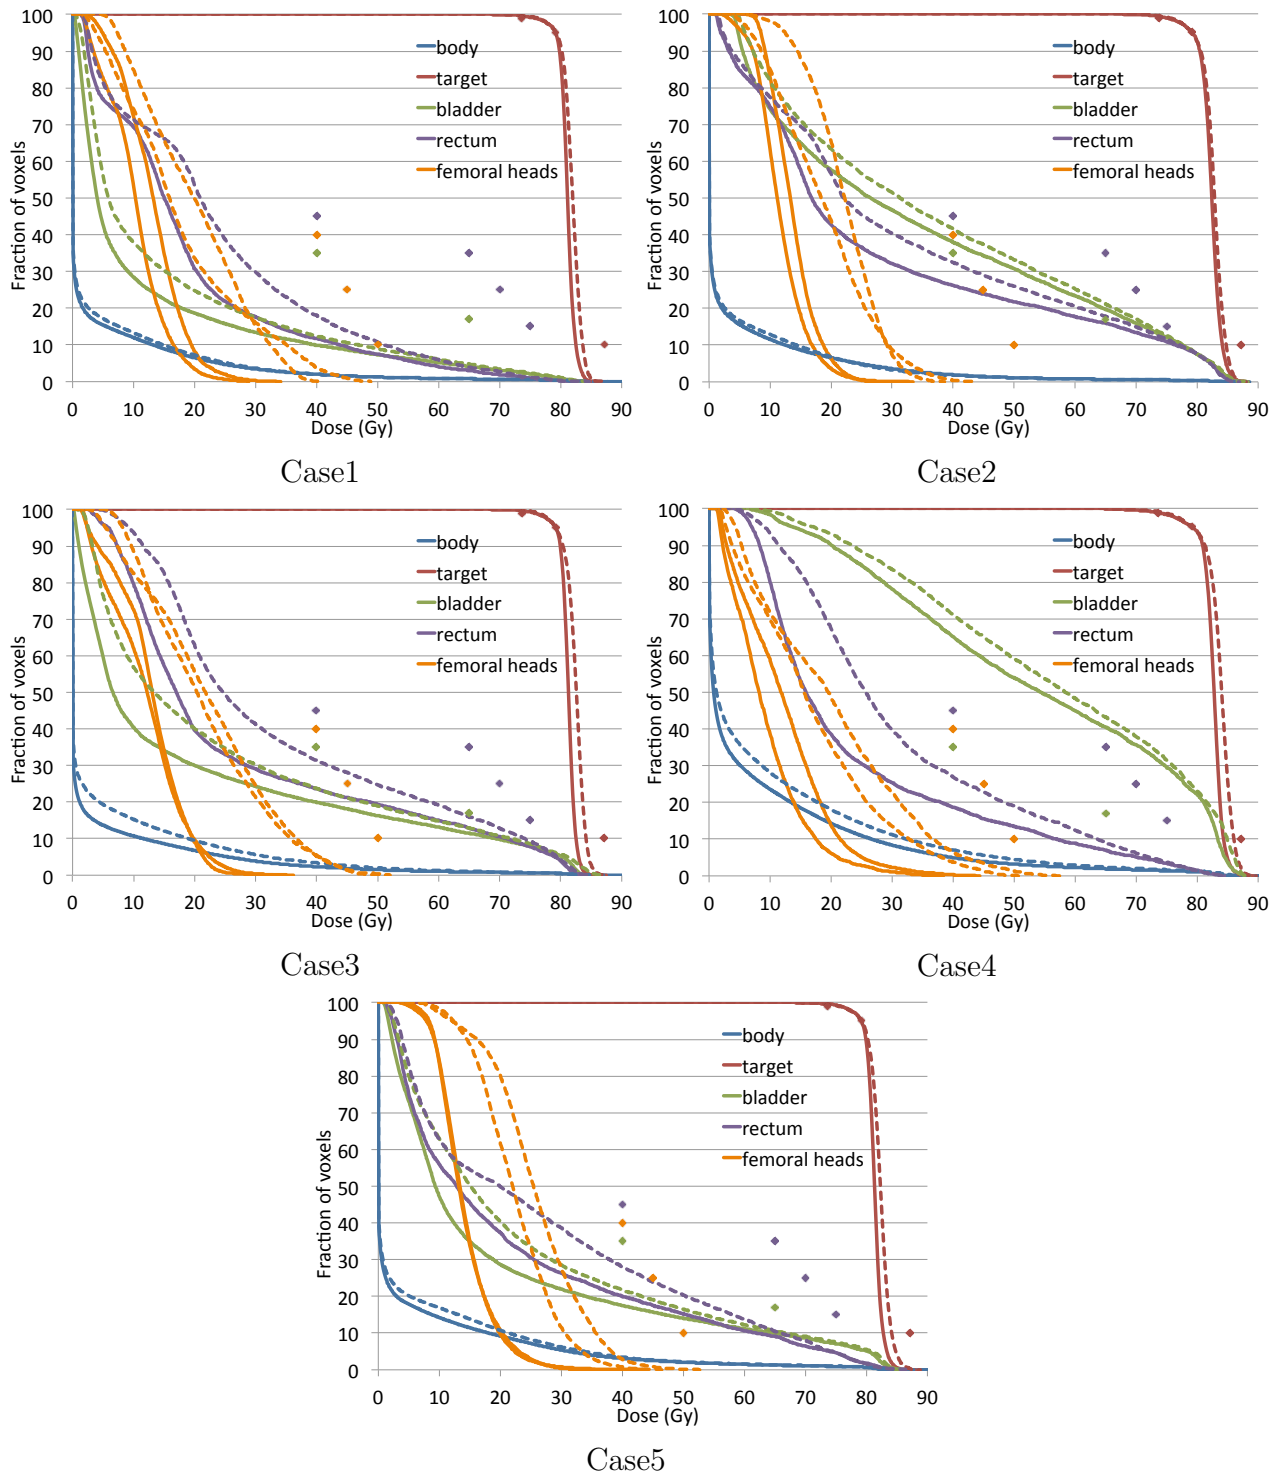


Figure 2.3: DVH comparison for Cases 1 — 5. Solid: 177-beam IMRT; dashed: VMAT

increased complexity of the pricing problem (see Table 2.7).

Solution time (s)	Case				
	1	2	3	4	5
Allowing interdigitation	78.8	71.5	59.2	54.2	67.6
Disallowing interdigitation	511.5	351.7	685.3	300.1	660.5

Table 2.7: Impact of disallowing interdigitation on algorithm run time.

2.11.5 Rate of change in gantry speed

Now we investigate the effect of the constraint on the rate of change in gantry speed on treatment time. As discussed in Section 2.7, we can assess the impact of the rate of change restriction by removing the corresponding constraint (2.26) from problem (SP). As before, we focus this study on Case 2.

Figure 2.4 plots the gantry speeds and dose rates for control points 1 through 176 for Case 2. Graphs (a-1) and (b-1) show gantry speeds, and (a-2) and (b-2) show dose rates, obtained by setting $s = S^U = 6$ in the column-generation algorithm, with rate of change constraint (2.26) omitted for graph (a-1) and (a-2), but included for graph (b-1) and (b-2). Graphs (c-1) — (d-2) contain analogous results with $s = 4$. It is clear that without a constraint on the rate of change, the gantry speed changes rapidly during the treatment in order to extend the time spent at apparently favorable control points while limiting the time spent at less favorable ones. Notice as well that at most control points the speed remains equal to the value s used in the column generation algorithm even after solving (SP).

When the rate of change constraint is enforced, as in parts (a-1) and (c-1), the sequence of speeds is smoothed so that the change in speed between control points is no more than ΔS_k . Note that, as discussed in Section 2.7, the rate of change constraint does not have any impact on treatment plan quality. However, if it were technologically feasible to increase or eliminate the bound on the rate of change in gantry speed, the treatment time could be shortened by about 35% to 1.7 minutes, as shown in Table 2.8.

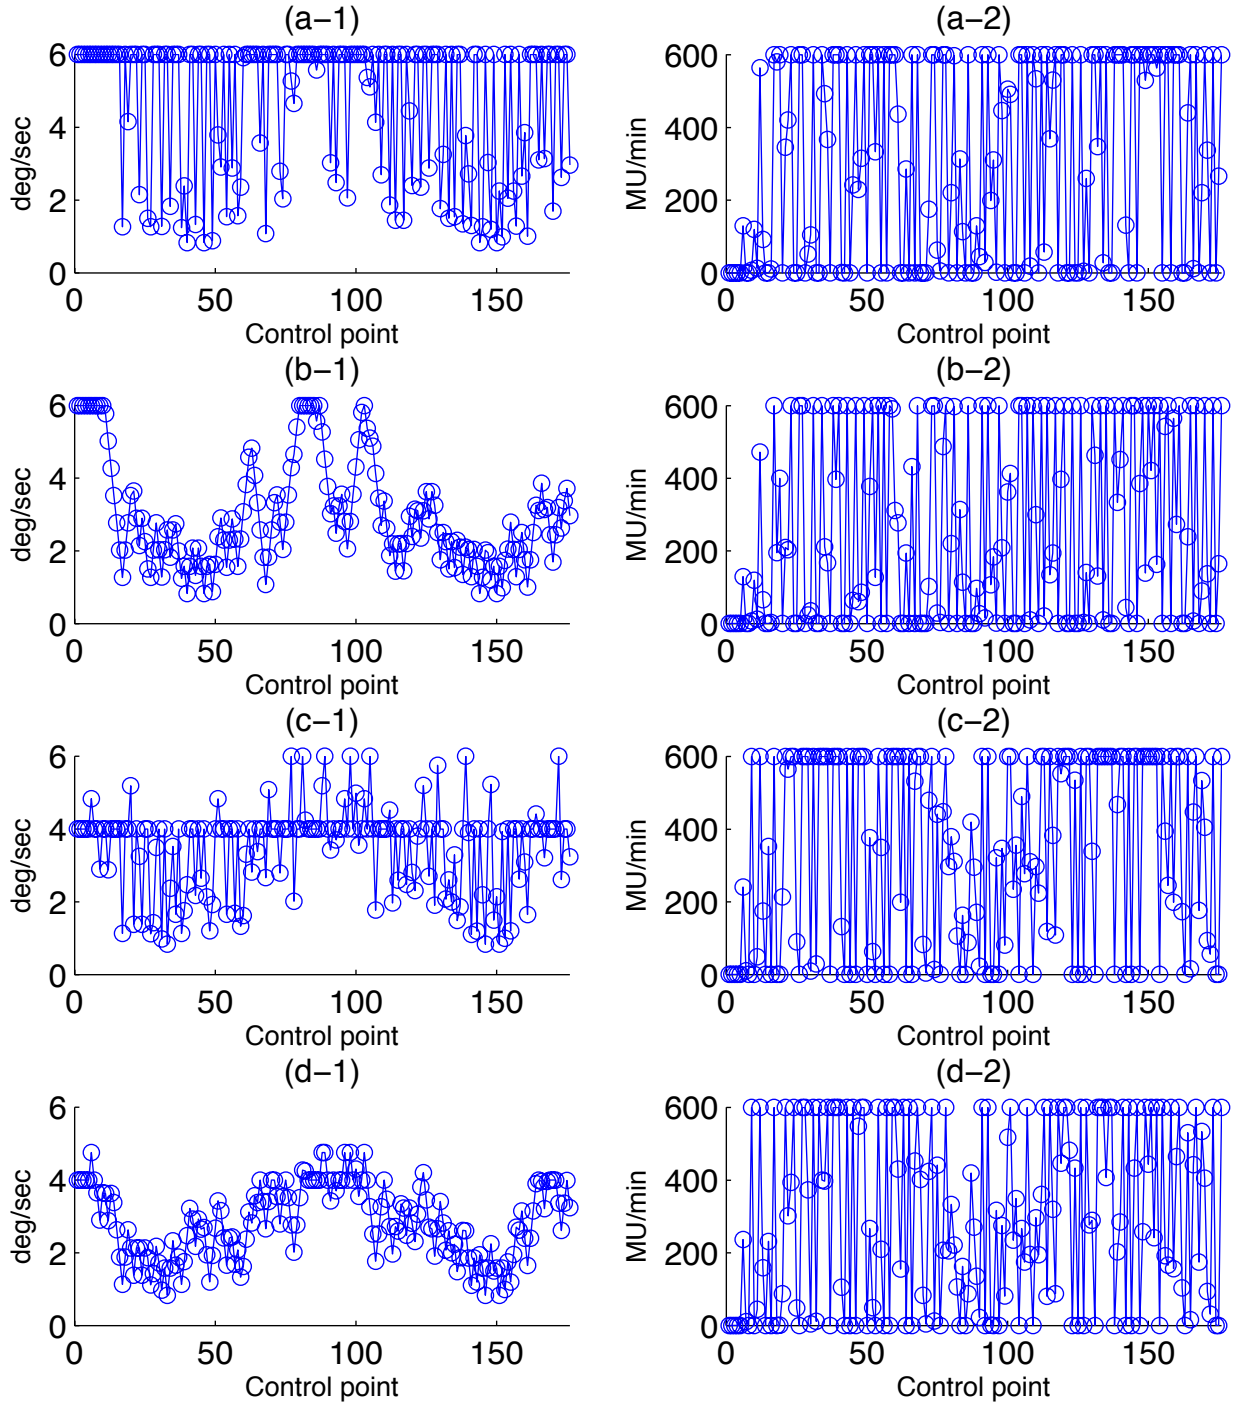


Figure 2.4: Gantry speeds (left column) and dose rates (right column) vs. control points for Case 2: (a) $s = 6$ deg/sec, ignoring rate of change constraint; (b) $s = 6$ deg/sec, considering rate of change constraints; (c) $s = 4$ deg/sec, ignoring rate of change constraint; (d) $s = 4$ deg/sec, considering rate of change constraints.

Treatment time (s)	Case				
	1	2	3	4	5
Without bound	97.8	99.4	91.6	94.9	98.9
With bound	149.3	149.3	141.1	142.6	155.5

Table 2.8: Impact of the presence of the bound ΔS on the rate of change in gantry speed on treatment time.

2.11.6 Transmission dose

Finally we study the effect of incorporating the transmission dose into the dose calculation during the optimization process. Again we incorporate the transmission dose by using the modified coefficient $\mathcal{D}_{kj}(A)$, explained in Section 2.8.4 in the (RMP^c). We first examine the impact of transmission on the dose distribution by applying the plan generated for Case 2 with $S = S^U$, and calculating the delivered dose with and without transmission. Figure 2.5 (a) shows a comparison of the DVH curves obtained in these two different settings. Because of transmission dose, the delivered dose is higher for all of the structures. However, the shapes of the corresponding curves remain similar. Normalizing the final dose distributions, for which the dose is scaled down 6.5%, reveals that given the same target coverage, transmission dose adds slightly to the dose to critical structures and normal tissues.

As discussed in Section 2.8.4, considering transmission helps us find better solutions to the (RMP^c), which lead to more accurate aperture prices in (PP^(k)) and better apertures being added in the column generation process. We compare two solutions, one obtained with transmission taken into account in the optimization process, and one without. For each plan we include transmission dose in calculating the the final dose distribution. The results can be seen in Figure 2.5 (c). Being able to incorporate the transmission dose in the optimization problem results in slightly lower doses to the critical structures and more homogeneous dose in the target. Moreover, the delivery time, MU and optimization time are all similar to the plan obtained without taking transmission into account, as shown in Table 2.9.

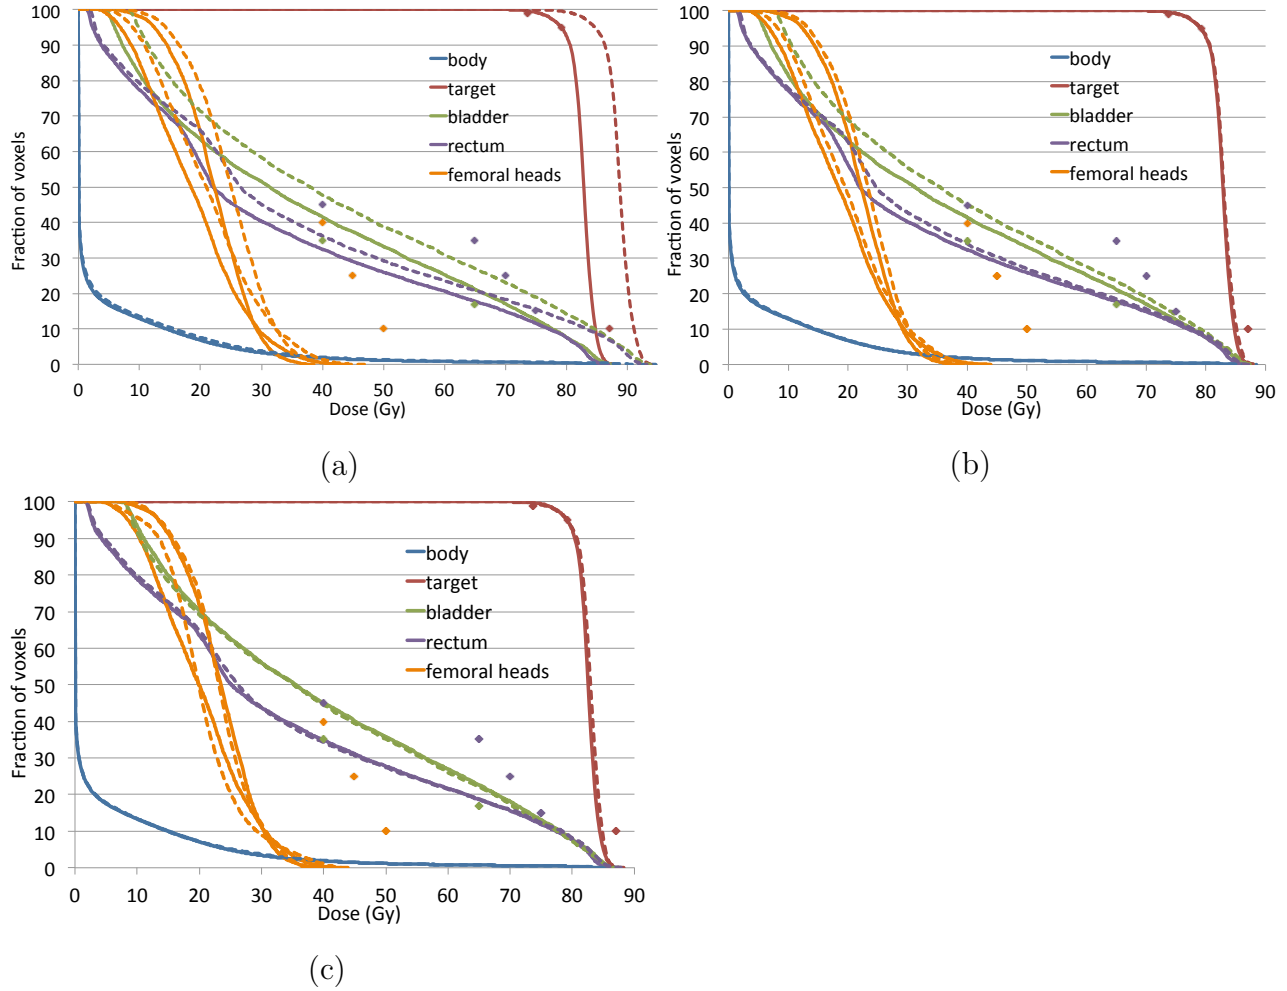


Figure 2.5: Case 2 $s = 6$, (a): final dose with (dashed) and without (solid) transmission; (b): final dose with (dashed) and without (solid) transmission, where transmission dose normalized to 95% target coverage; (c): normalized final dose with transmission, plan optimization with (dashed) and without (solid) considering transmission

Optimization scheme	runtime (s)	treatment time (s)	total MU
With transmission	85.8	143.1	665.4
Without transmission	71.5	149.3	667.7

Table 2.9: Impact of considering transmission on optimization runtime, treatment time and total MU

2.12 Conclusions

We have developed an efficient algorithm for generating VMAT treatment plans. Our approach incorporates various physical constraints, including bounds on dose rate, gantry speed, and rate of change in the gantry speed during treatment. Moreover, the approach is flexible enough to allow for additional constraints, including MLC delivery constraints. We have shown that our algorithm is capable of generating treatment plans of high quality on 5 prostate cancer cases. The VMAT plans produced by our algorithm can be delivered in around 2 minutes. Moreover, a GPU implementation allows us to complete the optimization within a couple minutes. Along with additional 15 seconds for computing the dose deposition coefficients before the optimization and 15 seconds for computing the final dose distribution after the optimization using a GPU-based dose calculation engine (*Gu et al. (2011, 2009)*), this work provides a major step towards a clinical application of on-line adaptive VMAT therapy.

CHAPTER III

VMAT with Constant Gantry Speed and Dose Rate

3.1 Introduction

Volumetric Modulated Arc Therapy (VMAT) with constant gantry speed and dose rate (VMATC) was first introduced in its initial form in *Yu (1995)* as the Intensity Modulated Arc Therapy (IMAT). In this type of treatment, the gantry rotates continuously around the patient while the radiation source is tuned on. The leaves in the Multileaf Collimator (MLC) system move dynamically to modulate the shape of the exposed field (*aperture*), however, both the gantry speed and the source output rate (*dose rate*) must stay constant during the treatment.

VMATC treatment has attracted the attention from many researchers as well as clinicians. Similar to VMAT, VMATC treatment is delivered while the gantry is in continuous motion, whereas the conventional Intensity Modulated Radiation Therapy (IMRT) delivers radiation sequentially from a number of beam angles in a “step-and-shoot” fashion, and requires setup time between consecutive beam angles. Moreover, IMRT delivery from each beam angle takes a long time because of the intensity modulation required. As a result, VMATC has the potential to shorten the overall treatment time compared to IMRT, which leads to less patient discomfort, less intrafraction motion uncertainties, as well as increased equipment utilization for the clinics. Furthermore, VMATC is able to use a large set of angles around the patient, instead of being restricted to a relatively small number of pre-determined

angles. This may potentially lead to improved treatment plan quality over IMRT.

VMATC treatments can be delivered from an IMRT treatment machine equipped with conventional linear accelerator and MLC system. Because IMRT is commonly used as the primary external beam radiation therapy treatment method, this gives a wide range of clinics access to research and experiments with VMATC treatments. However, VMATC lacks the flexibility to dynamically vary the gantry speed and dose rate. This gives the optimization for VMATC a very different nature from that for VMAT. Therefore it is necessary to develop optimization algorithms specifically for VMATC.

Because of the lack of inverse planning techniques, some early efforts in designing treatment plans for VMATC were focused on using a Beam-Eye-View (BEV) approach, by aligning the apertures at different angles to the shape of the projected target (*Yu et al. (2002); Wong et al. (2002)*). The majority of research was done by converting IMRT plans and making them deliverable in treatments consisting of multiple arcs. *Yu (1995)* relied on an IMRT treatment with beams 5° apart as the starting plan. The intensity profile at every beam was then heuristically decomposed into a number of intensity levels and the associated leaf positions, so that the leaf positions with matching intensity level at different angles are compatible on the same arc. *Crooks et al. (2003)* segmented an IMRT fluence map at each angle into several intensity levels, and distributed these segments on the small arc between two consecutive angles. The time it takes for the gantry to traverse this arc is divided and assigned to each segment so that their delivery time, and thus intensity, matched the original intensity. However, the authors assumed that segment transition can happen instantaneously, and that the leaf travel speed was unbounded, which may lead to undeliverable plans. Moreover, the gantry speed necessary to deliver the treatment may be undesirably slow in the solution. In a study on designing leaf-sequencing methods for converting IMRT fluence map to deliverable leaf settings, *Shepard et al. (2007)* employed a simulated annealing heuristic to find the leaf settings that best resemble the IMRT fluence map. Later in a similar study, *Luan et al. (2008)* applied a shortest path algorithm that minimizes the error

between the original IMRT and the sequenced fluence maps. However, due to the differences between IMRT and VMATC treatments, a treatment plan designed for IMRT and perturbed to satisfy the compatibility constraints is not necessarily a good plan for VMATC. Moreover, the process of converting the IMRT plan will inevitably introduce degradations in plan quality, and the commonly used method of minimizing the difference in fluence maps may not lead to the minimal difference in plan quality.

Using heuristic methods that do not depend on a prior IMRT treatment plan, *Cotrutz et al.* (2000) proposed an approach that combined a number of arcs, each with a manually picked rectangular aperture, to produce a treatment. *Earl et al.* (2003) and *Cameron* (2005) both proposed methods that used simulated annealing to perturb the starting set of apertures, which conformed to the shape of the projected target at different angles (i.e., a BEV approach). In the former article the dose rate was part of the decisions perturbed in the simulated annealing method, whereas in the latter two the dose rate was determined through optimization for the given leaf positions.

The results in these studies showed that either multiple arcs or one arc with long delivery time were required to generate VMATC treatment plans that resemble the quality of a conventional IMRT plan. In favor of this observation, *Bortfeld and Webb* (2009) performed a theoretical study based on a phantom case, and claimed that single-arc VMATC plans may unduly compromise the plan quality compared to IMRT plans if the treatment time was to be kept under 2 minutes.

The above-mentioned studies, among others, have led to the advancement in treatment technologies and the introduction of dedicated commercial VMAT systems that enable variable gantry speed and dose rate during the treatment, and research focus has shifted towards plan optimization problem for full VMAT treatments since (*Otto* (2008); *Bedford* (2009); *Men et al.* (2010b); *Peng et al.* (2012)).

Although VMATC lacks two important features, namely dynamically changing gantry speed and dose rate during the treatment, being able to deliver this type of treatment can

still be very appealing to clinics equipped with conventional IMRT machines, as VMATC treatment provides many the benefits of full VMAT without requiring an expensive system upgrade. This can be particularly beneficial to countries that do not have access to treatment machines with full VMAT capabilities. Moreover, fully realizing the potential of VMATC treatment will require investigation into all degrees of freedom it offers. To the best of our knowledge, none of the existing research studies on VMATC have directly taken the optimization of the constant gantry speed into account. However, as gantry speed directly determines how far the leaves can travel between consecutive control points and thus the allowed change in aperture shapes, it impacts the plan quality in addition to treatment time.

In this research we consider the treatment planning problem for VMATC with a single arc, and we treat both (constant) gantry speed and dose rate as decision variables in the optimization process. We compare our VMATC plans with VMAT on actual patient cases, and show that VMATC can generate single-arc treatment plans with quality close to full VMAT. Our approach can readily be extended to treatments with multiple arcs. The results of our research may help clinics make an informed decision when upgrading to dedicated VMAT systems.

3.2 VMATC optimization problem formulation

We assume that during VMATC treatment the gantry rotates continuously along a pre-specified trajectory. Let S denote the constant gantry speed, chosen from a range $[S^L, S^U]$ that is physically feasible for the treatment machine, and let $R \in [R^L, R^U]$ represent the constant dose rate for the VMATC treatment. The MLC apertures can, in principle, be controlled at every angle on the gantry rotation trajectory. However, for tractability reasons the treatment is usually defined by controlling the apertures at all or a subset of K discrete *control points* along the arc. We define the angular distance between control point $k - 1$ and k as δ_k ($k = 1, \dots, K$), and define an additional control point 0 in the beginning of the treatment arc for convenience. In between control points, the treatment machine will

configure itself so that the leaves in the MLC move at constant speeds to reach their positions at the next control point from the previous control point.

By controlling the machine this way, we can in principle calculate the dose to each voxel as a function of S , R , and the leaf positions along the gantry trajectory. However, computing the dose based on the continuous leaf motion is not only inefficient, but also impractical, since it usually needs to be done numerous times in the optimization process. We therefore consider an approximation that computes the overall dose as the sum of dose associated with the control points. Note that this is the same “step-and-shoot” approximation described in Chapter II, which assumes that the aperture stays constant on the arc spanning control points $k - 1$ and k ($k = 1, \dots, K$). We can calculate the dose delivered to voxel $j \in \mathcal{V}$ as follows:

$$z_j = \sum_{k=1}^k D_{kj}(A_k) \cdot \frac{R \delta_k}{S} \quad \forall j \in \mathcal{V}, \quad (3.1)$$

where $D_{kj}(A_k)$ is the dose deposition coefficient associated with aperture A_k and voxel j , and $\frac{R \delta_k}{S}$ equals the total source output (measured by *Monitor Unit* (MU)) delivered from control point k . *Otto* (2008) showed that z_j calculated by (3.1) will be sufficiently close to the actual dose if K is sufficiently large and δ_k is sufficiently small for all $k \in \{1, \dots, K\}$.

Now given that the VMATC planning problem is essentially a VMAT problem with variables S and R restricted to be constant for all angles, we can easily modify formulation (FP) in Chapter II by equating all s_k to S and all r_k to R , and arrive at the formulation for VMATC:

$$\text{(VC) minimize}_{S,R,A,z} \quad F(\mathbf{z})$$

subject to

$$z_j = \sum_{k=1}^K D_{kj}(A_k) \cdot \frac{R \delta_k}{S} \quad j \in \mathcal{V} \quad (3.2)$$

$$S \leq S_{k-1,k}^U(A_{k-1}, A_k) \quad k = 1, \dots, K \quad (3.3)$$

$$S \in [S^L, S^U]$$

$$R \in [R^L, R^U]$$

$$A_k \in \mathcal{A} \qquad k = 1, \dots, K.$$

Recall that $S_{k-1,k}^U(A_{k-1}, A_k)$ is the maximum speed allowed for the MLC apertures to shift from A_{k-1} to A_k between control points $k-1$ and k , and \mathcal{A} is the set of apertures deliverable by the MLC system.

Although this problem appears similar to problem (FP), it cannot be solved using the same algorithm, where after every aperture is added to the master problem, the fluence rates of all apertures need to be chosen so that the objective function value will improve. For the VMATC problem, the fluence rate must always be the constant value associated with S and R . As a result, adding an aperture with a negative price may actually lead to a worsened objective function value. Therefore, we will develop (heuristic) algorithms specifically for (VC) in this chapter.

3.3 Decomposition approaches to (VC)

Solving (VC) directly is difficult; in what follows, we propose and test two heuristic optimization frameworks that rely on the decomposition approaches to (VC) described in this section. Problem (VC) has two types of decisions: (i) selecting values of continuous variables S and R , and (ii) selecting apertures at control points. To describe our heuristic optimization frameworks, it is convenient to consider the following two “decomposition-based” interpretation of (VC).

3.3.1 Decomposition approach 1: gantry speed and dose rate selection as a sub-problem

Suppose that we have selected one aperture A_k for every control point $k \in \{1, \dots, K\}$, then the gantry speed and dose rate can be optimized. However, not all gantry speeds in the range $[S^L, S^U]$ can be used, as S must be compatible with the apertures in consecutive

control points in order to allow the current aperture sequence to be deliverable. Let M be the set of MLC leaf rows, and let $a_{km} = (\ell_{km}, r_{km})$ be the leaf setting in row $m \in M$ at control point k . Finally, let v be the maximum travel speed for leaves in the MLC. It can be easily seen that S must be chosen so that:

$$\frac{|\ell_{mk} - \ell_{m,k-1}|}{v} \leq \frac{\delta_k}{S} \quad \forall m \in M, k \in \{1, \dots, K\} \quad (3.4)$$

$$\frac{|r_{mk} - r_{m,k-1}|}{v} \leq \frac{\delta_k}{S} \quad \forall m \in M, k \in \{1, \dots, K\}. \quad (3.5)$$

These conditions ensure the gantry rotation time between two consecutive control points is sufficient for the leaf transition. We can define the maximum allowed (constant) speed \bar{S}^U from (3.4) and (3.5) as:

$$\bar{S}^U = \min_{m \in M, k \in \{1, \dots, K\}} \left\{ \frac{\delta_k \cdot v}{|\ell_{mk} - \ell_{m,k-1}|}, \frac{\delta_{k-1} \cdot v}{|r_{mk} - r_{m,k-1}|} \right\}.$$

The optimization problem for the dose rate and gantry speed then reads:

$$(GD^{(A)}) \underset{z, R, S}{\text{minimize}} \quad F(\mathbf{z})$$

subject to:

$$z_j = \sum_{k=1}^K D_{kj}(A_k) \cdot \frac{R \delta_k}{S} \quad \forall j \in \mathcal{V}$$

$$R \in [R^L, R^U] \quad (3.6)$$

$$S \in [S^L, \bar{S}^U]. \quad (3.7)$$

As long as $\bar{S}^U, S^L > 0$, which is apparent as the gantry must rotate at a positive speed, this problem can be transformed into an equivalent optimization problem by defining fluence rate $y = \frac{R}{S}$, and combining the feasible ranges for R and S into the feasible range for y :

$$(DY^{(A)}) \underset{z, y}{\text{minimize}} \quad F(\mathbf{z})$$

subject to:

$$z_j = \sum_{k=1}^K D_{kj}(A_k) \delta_k y \quad \forall j \in \mathcal{V}$$

$$y \in [Y_L, Y_U]$$

where

$$Y_L = \frac{R^L}{\bar{S}^U}, Y_U = \frac{R^U}{S^L}.$$

Note that problem (DY^(A)) is a convex problem that can be easily solved.

3.3.2 Decomposition approach 2: aperture selection as a sub-problem

The gantry speed, together with the fixed maximum leaf travel speed, determines how far individual leaves in the MLC can travel between two control points, which in turn determines the sequence of apertures deliverable in the treatment. Once the gantry speed and dose rate combination is chosen, the achievable plan quality only depends on the aperture selection (i.e., leaf setting) at each control point. We formulate the problem for optimizing the apertures as follows:

$$\begin{aligned}
(\text{LS}^{(S,R)}) \quad H(S, R) = \underset{A, z}{\text{minimize}} \quad & F(\mathbf{z}) \\
\text{subject to} \quad & \\
z_j = \sum_{k=1}^K D_{kj}(A_k) \cdot \frac{R \delta_k}{S} \quad & j \in \mathcal{V} \\
S \leq S_{k-1, k}^U(A_{k-1}, A_k) \quad & k = 1, \dots, K \\
A_k \in \mathcal{A} \quad & k = 1, \dots, K.
\end{aligned}$$

We observe that solving (VC) is the same as minimizing $H(S, R)$ over S and R , or equivalently, solving the following problem:

$$(\text{SR}) \underset{S, R}{\text{minimize}} \quad H(S, R)$$

subject to

$$S \in [S^L, S^U] \tag{3.8}$$

$$R \in [R^L, R^U]. \tag{3.9}$$

However, directly solving $(\text{LS}^{(S,R)})$ is difficult, as the dose deposition matrix $D_{kj}(A_k)$ has a complex relationship with A_k that is dependent on the patient’s anatomy. We describe a heuristic method for solving this problem in the next section.

3.4 Methods for (approximately) solving $(\text{LS}^{(S,R)})$

We propose an algorithm for solving $(\text{LS}^{(S,R)})$ that iteratively selects and refines apertures at the set of control points, while maintaining a deliverable treatment plan throughout the solution process. We start with an outline of the algorithm:

1. Initialize with an “empty” treatment – all apertures are assumed “closed”. Let $\mathcal{C} = \emptyset$ be the set of control points that currently have apertures selected
2. Selecting an initial sequence of apertures:
 - (a) determine a sequence of control points, which we will follow to select the apertures;
 - (b) for the next control point in the sequence, say \bar{k} , set $\mathcal{C} \leftarrow \mathcal{C} \cup \{\bar{k}\}$, select aperture $A_{\bar{k}}$;
 - (c) if $|\mathcal{C}| = K$, go to Step 3; otherwise go to Step 2(b).
3. Refining the collection of apertures selected in Step 2:
 - (a) determine a sequence of control points to follow;
 - (b) perform the following refinement step:
 - i. for the next control point in the sequence, say \bar{k} , refine aperture $A_{\bar{k}}$;
 - ii. if we are at the end of the sequence, go to Step 3(c); otherwise go to i.
 - (c) terminate if the objective does not improve; otherwise repeat Step 3(b).

In this algorithm, every time an aperture is selected or refined, it is done in a way that will improve or maintain the quality of the current solution. We discuss this in detail in the

following sections. In our implementation in Section 3.6, we used the following sequence of control point in both Steps 2 and 3: the first two apertures are selected at control points 1 and K , the next aperture is selected halfway between 1 and K to reduce the distance between selected apertures by half, then two more apertures are selected to reduce the distance between selected apertures to $1/4$ or the entire arc, etc., until all the control points have apertures selected.

3.4.1 Determining apertures at all control points given set \mathcal{C}

Recall that when an aperture is selected at every control point and S and R are given, the same “step-and-shoot” approximation we used for VMAT can be used to perform the dose calculation. However, during Step 2 of algorithms that fit the above general description, we will encounter intermediate steps in which only a subset of the control points have an aperture specified. Unlike in VMAT, where we would set dose rate at such control points to zero, varying the dose rate is not an option in VMATC if we want to maintain a deliverable plan at every step of the algorithm, and care needs to be taken when calculating the delivered dose in this situation.

Let $\mathcal{C} \subset \{1, \dots, K\}$ be the subset of control points that have apertures specified in the solution process. Then the apertures at all other angles, including control points $\{1, \dots, K\} \setminus \mathcal{C}$, can be found by linearly interpolating the leaf positions at control points in \mathcal{C} . We can define function \mathcal{F} , which calculates leaf settings a_{mk} at control point $k \in \{1, \dots, K\}$ and row $m \in M$, given those specified in \mathcal{C} :

$$\mathcal{F}_{km}^{\mathcal{C}}(a_{mk'} : k' \in \mathcal{C}) = \left(\ell_{mk^-} + \frac{\ell_{mk^+} - \ell_{mk^-}}{\sum_{i=k^-+1}^{k^+} \delta_i} \cdot \sum_{i=k^-+1}^k \delta_i, r_{mk^-} + \frac{r_{mk^+} - r_{mk^-}}{\sum_{i=k^-+1}^{k^+} \delta_i} \cdot \sum_{i=k^-+1}^k \delta_i \right). \quad (3.10)$$

Here k^- and k^+ are the indices of control points in \mathcal{C} that are immediately before and after control point k , respectively:

$$k^- = \max\{c, c \in \mathcal{C} : c \leq k\} \text{ and } k^+ = \min\{c, c \in \mathcal{C} : c \geq k\}$$

It is easy to see that k^- , k^+ and k coincide if k is itself included in \mathcal{C} .

In the discussion so far we assumed that we used at least two control points 1 and K . Under this assumption, the control points k^- and k^+ are well-defined for any $k \in \{1, \dots, K\}$, allowing us to evaluate \mathcal{F} via a simple interpolation. However, we can extend the definition of function \mathcal{F} to cover other situations: if control point k has no aperture specified before (after) itself, it will use the same aperture as the one specified at the control point after (before). And if no aperture has been specified yet, it is natural to define the treatment plan as empty where all leaves are closed at the edge of the beam. Once the apertures at all control points are determined with function \mathcal{F} , we can readily calculate the dose distribution using (3.1).

3.4.2 Methods for selecting/refining apertures

In Steps 2 and 3 of the aperture selection algorithm, we select apertures sequentially at the set of control points so that the treatment quality can be improved. In this section we describe two types of methods, one based on a local search procedure that can only examine a finite set of options when selecting the leaf setting, the other based on optimization, which can find the leaf setting in one row optimally but takes longer to execute. Moreover, the optimization based method can only be applied in refining the apertures.

3.4.2.1 Local search heuristic

This method can be used in both Steps 2 and 3 of the above algorithm. Suppose that the set of control points that currently have an aperture selected is $\mathcal{C} \subset \{1, \dots, K\}$, and now for more control over the treatment delivery process and/or for improving the solution quality, we want to increase the number of selected apertures by 1, say at $k \in \{1, \dots, K\} \setminus \mathcal{C}$. The leaves at k used to occupy the interpolated positions determined by the selected control points k^- and k^+ in \mathcal{C} . To choose a potentially better leaf setting, first we need to determine which ones are feasible in each row. This is again dictated by the constant gantry speed and

the leaf travel speed:

$$|\ell_{mk^-} - \ell_{mk}| \leq v \cdot \frac{\sum_{i=k^-+1}^k \delta_i}{S} \quad \forall m \in M \quad (3.11)$$

$$|r_{mk^-} - r_{mk}| \leq v \cdot \frac{\sum_{i=k^-+1}^k \delta_i}{S} \quad \forall m \in M \quad (3.12)$$

$$|\ell_{mk^+} - \ell_{mk}| \leq v \cdot \frac{\sum_{i=k+1}^{k^+} \delta_i}{S} \quad \forall m \in M \quad (3.13)$$

$$|r_{mk^+} - r_{mk}| \leq v \cdot \frac{\sum_{i=k+1}^{k^+} \delta_i}{S} \quad \forall m \in M \quad (3.14)$$

$$\ell_{mk} \leq r_{mk} \quad (3.15)$$

Here (3.15) is the physical restriction that leaves in the same row cannot overlap. Note that these restrictions are written assuming k^- and k^+ both exist. If either or none of them exist because control point k does not have a preceding or succeeding control point, we can simply ignore the associated constraints.

Let $\tilde{a}_{mk} = (\tilde{\ell}_{mk}, \tilde{r}_{mk})$ ($m \in M$), be the original interpolated leaf positions at k . Constraints (3.11) — (3.14) provide a feasible range, say $[\underline{\ell}_m, \bar{\ell}_m]$ and $[\underline{r}_m, \bar{r}_m]$, for the left and the right leaves (ℓ_{mk}, r_{mk}) in each row $m \in M$. Suppose that the dimension of each row is $[0, N]$, and we divide it into N beamlets along the leaf travel direction in order to allow computation of the dose deposition coefficients in a pre-processing step. We propose the following method for selecting/refining the aperture at k :

1. set $(\ell_{mk}, r_{mk}) = (\tilde{\ell}_{mk}, \tilde{r}_{mk}) \quad \forall m \in M$
2. for each row $m \in M$ sequentially, perform the following local search:
 - (a) for every combination in the set:

$$\left\{ \left(\hat{\ell}_{mk}, \hat{r}_{mk} \right) : \hat{\ell}_{mk} \leq \hat{r}_{mk}, \hat{\ell}_{mk} \in \{ \underline{\ell}_m, [\underline{\ell}_m], \dots, [\bar{\ell}_m], \bar{\ell}_m, \tilde{\ell}_{mk} \}, \right. \\ \left. \hat{r}_{mk} \in \{ \underline{r}_m, [\underline{r}_m], \dots, [\bar{r}_m], \bar{r}_m, \tilde{r}_{mk} \} \right\}$$

$$\text{set } (\ell_{mk}, r_{mk}) = \left(\hat{\ell}_{mk}, \hat{r}_{mk} \right)$$

- (b) find the interpolated leaf positions in row $m \in M$ for control points $k^-, \dots, k-1$

and $k + 1 \dots, k^+$ using (3.10)

- (c) calculate the corresponding dose distribution using (3.1), and evaluate the objective function

set (ℓ_{mk}, r_{mk}) to be the leaf setting that provides the minimum objective.

The leaf position in any beam row found in this procedure will either be the original interpolated position, a boundary point, or one that occupies one of the beamlet edges. This ensures that the treatment plan quality can be improved or maintained after an aperture is selected/refined.

3.4.2.2 Exact optimization

This is an alternative method for refining the aperture selection in Step 3 of the algorithm above, and is based on identifying the feasible ranges for leaf positions in each row given S and the apertures currently selected at $k - 1$ and $k + 1$, and partitioning those ranges into small, usually beamlet-sized intervals. Within each interval, the dose vector is an affine function of the leaf position, and thus the optimal leaf positions within that interval can be obtained by solving a continuous optimization problem. By enumerating all such intervals, we can obtain an optimal leaf setting (for given S and R , and apertures selected at other control points). However, the solution time of all of these optimization problems will significantly exceed that of the local search heuristic presented above. We will study the tradeoffs between solution quality and running time of the two aperture selection methods. Moreover, we will only utilize this method in Step 3 of the algorithm, when $k^- = k - 1$ and $k^+ = k + 1$, since otherwise changing the leaf positions at k affects all control points between k^- and k^+ , and the number of intervals we need to consider may be too large to execute the procedure.

Assuming all control points $k \in \{1, \dots, K\}$ has one aperture assigned already, we study the optimization of leaf settings at control point $\bar{k} \in \{1, \dots, K\}$ and row $\bar{m} \in M$. Note that dose delivered from all other control points is constant, and that dose delivered from rows

other than \bar{m} at \bar{k} is also constant. For simplicity we define this constant dose as

$$\bar{z}_j = \sum_{k \in \{1, \dots, K\} \setminus \{\bar{k}\}} \sum_{m \in M} D_{kmj}(a_{mk}) \cdot \frac{R \delta_k}{S} + \sum_{m \in M \setminus \bar{m}} D_{\bar{k}mj}(a_{m\bar{k}}) \cdot \frac{R \delta_{\bar{k}}}{S} \quad \forall j \in \mathcal{V}.$$

Here $D_{kmj}(a_{mk})$ is the dose deposition coefficient associated with the leaf setting a_{mk} at control point k , row m , and voxel j . Suppose that the left and right leaf positions in row \bar{m} at \bar{k} are l and r , respectively, the optimization problem then reads:

$$\text{(LO) minimize}_{z, \ell, r} \quad F(\mathbf{z})$$

subject to:

$$z_j = \bar{z}_j + D_{\bar{k}\bar{m}j}(\ell, r) \cdot \frac{R \delta_{\bar{k}}}{S} \quad (3.16)$$

$$\ell \in [\underline{\ell}, \bar{\ell}], \quad r \in [\underline{r}, \bar{r}], \quad 0 \leq \ell \leq r \leq N \quad (3.17)$$

Again we use $[\underline{\ell}, \bar{\ell}]$ and $[\underline{r}, \bar{r}]$ to represent the feasible ranges for the left and the right leaves derived from constraints (3.11) — (3.14).

Recall that we discretize a continuous beam row into N different beamlets. In order to compute the dose deposition coefficient associated with leaf setting (ℓ, r) , we assume that a partially covered beamlet will deliver dose proportional to the fraction of the beamlet area that is exposed. As a result this coefficient can be computed as:

$$D_{\bar{k}\bar{m}j}(\ell, r) = \left(([\ell] - \ell) \cdot D_{\bar{k}\bar{m}[\ell]j} + \sum_{n=[\ell]}^N D_{\bar{k}\bar{m}nj} \right) - \left(([r] - r) \cdot D_{\bar{k}\bar{m}[r]j} + \sum_{n=[r]}^N D_{\bar{k}\bar{m}nj} \right), \quad (3.18)$$

where D_{kmnj} is the beamlet-based dose deposition coefficient associated with beamlet n in row m at control point k .

Both terms in (3.18) are piece-wise linear and neither convex nor concave in general, therefore problem (LO) is not directly solvable. However, when ℓ and r are each restricted to be within a beamlet, problem (LO) is convex. We can then independently solve (LO) for every combination of intervals:

$$\ell \in [\underline{\ell}, [\underline{\ell}]], \quad \text{or} \quad [[\underline{\ell}], [\underline{\ell}] + 1], \dots, \quad \text{or} \quad [[\bar{\ell}], \bar{\ell}], \quad \text{and}$$

$$r \in [\underline{r}, \lceil \underline{r} \rceil], \text{ or } [\lceil \underline{r} \rceil, \lceil \underline{r} \rceil + 1], \dots, \text{ or } [\lceil \bar{r} \rceil, \bar{r}],$$

for which $\ell \leq r$ is satisfied. The solution to the optimization problem that yields the best objective function value is then the optimal solution to (LO).

3.5 Two frameworks for finding (approximate) solutions to (VC)

With the technical preliminaries explained, we are ready to describe the two frameworks for solving problem (VC).

3.5.1 Framework A: approximately solving (SR)

Recall that if we solve problem (SR) to optimality, we will have the optimal solution to (VC). However, the methods for evaluating the objective function of (SR) by solving (LS^(S,R)) described in Section 3.4 are heuristics, and thus solving (LS^(S,R)) only yields an approximate solution. Let $\tilde{H}(S, R)$ be the objective value obtained by applying a heuristic solution method to (LS^(S,R)), we can then approximately solve (SR) by considering the following problem:

$$(\text{SR}^{\text{approx}}) \underset{S, R}{\text{minimize}} \tilde{H}(S, R)$$

subject to (3.8) and (3.9). Note that one single evaluation of function $\tilde{H}(S, R)$ is equivalent to performing the entire aperture selection procedure, and therefore is computationally expensive. As a result, performing a large number of function evaluations is not realistic. Moreover, the fact that we solve problem (LS^(S,R)) with a heuristic method means that $\tilde{H}(S, R)$ does not provide any reliable information about the smoothness, differentiability or convexity of $H(S, R)$. Therefore we have to solve problem (SR^{approx}) with a method that does not rely on derivatives. Optimization algorithms of this type usually fall into one of the following categories:

1. Line-search methods based on estimated gradients. This type of method estimates

the gradient value with approximation techniques, such as finite differences (*Gilmore and Kelley (1995)*), then performs a line search along a direction dependent on the gradient to find a improving solution. However, in cases when the objective function is not smooth or continuous, the approximated descent direction is not necessarily a good direction to search along. Moreover, finding the step size may involve many function evaluations, making it unattractive for expensive functions like ours (*Fowler et al. (2008)*).

2. Search heuristics that mimic processes and phenomenon that happen in nature, such as simulated annealing and genetic algorithm. Simulated annealing (*Kirkpatrick et al. (1983)*; *Van Laarhoven and Aarts (1987)*) models the annealing process in metallurgy. In each step the algorithm probabilistically determines if it will move to a neighboring point of the current point, and the probability that controls this process is dynamically and usually slowly changed to lead the system to points with better objective function values. Genetic algorithm as well as related heuristics (see, for example, *Goldberg (1989)*; *Gen and Cheng (1999)*), mimics the natural evolution process. It maintains a population of solutions, and by selecting, mutating and varying stochastically individuals that are more fit, i.e., have better objective values, produces a new population from the current one. These algorithms have performed very well in finding good (but not necessarily optimal) solutions in practice. However, the optimization process usually requires a large number of function evaluations.
3. Direct search algorithms. This includes a family of pattern search methods (*Audet and Dennis Jr (2004)*; *Conn et al. (2009)*), and simplicial search methods, (most notably the Nelder-Mead algorithm *Nelder and Mead (1965)*), among others. These methods decide where the next function evaluation will take place entirely based on the function values encountered in history. These methods, especially the Nelder-Mead method, have been popular among scientists faced with an expensive function and a limited

computational budget because of their ability to find good solutions within a relatively small number of function evaluations (*Lagarias et al. (1998)*).

4. Trust region based methods. These methods rely on approximating the objective function with analytical and usually quadratic functions (a.k.a. a surrogate model) that are easy to evaluate and optimize. By maintaining a trust region, in which the surrogate model is believed to be a good approximation of the underlying function, and optimizing the surrogate model over the trust region, the algorithm proceeds to new solutions believed to provide sufficient decrease for the underlying function. The surrogate model and the trust region are continuously modified as the algorithm proceeds. These methods are attractive because the quadratic surrogate models used make them very robust when dealing with curvature information. Extensive research has been done to investigate the convergence properties and application of these methods (*Conn et al. (2000, 2009)*).
5. Response surface methods (RSM). RSM is a collection of methodologies useful in applications such as optimization, and selection of conditions to meet requirements (*Montgomery and Myers (2002)*). The fact that RSM can fulfill the users' goals by using a limited number of observations makes it a desirable choice for many. RSM assumes the underlying function of interest is noisy because of an underlying statistical error component, and is usually based on an approximation of the function by a low order polynomial. Optimization with RSM in general involves an iterative process of fitting a model to available data, finding improving directions, moving along the direction to a new solution, and updating the model.

Although many of these methods and/or their variants, with the exception of those in the second category, have been shown to converge locally or globally for continuously differentiable functions (see, for example, *Lagarias et al. (1998)*; *Conn et al. (2009)*), in cases when no information is available about the objective function, none of these methods can

guarantee convergence even to locally optimal solutions. Moreover, given that the evaluation of function $\tilde{H}(S, R)$ is likely to take a nontrivial amount of time, it is highly desirable that we chose a method that can efficiently reduce the objective function in small number of iterations, and one that can work with function shapes that do not necessarily resemble linear or quadratic functions. For these reasons, we choose the Nelder-Mead algorithm to apply in our experiments.

We give a brief introduction of the Nelder-Mead method following the notations in *Conn et al.* (2009). For an optimization problem in n -dimensions, the Nelder-Mead algorithm is based on maintaining and changing a simplex of $n + 1$ vertices $Y = \{y^0, \dots, y^n\}$. The method iteratively replaces the worst vertex, say y^n , from the simplex with a point on the line connecting y^n and the centroid of the best n vertices y^c . The new vertex of each iteration can be one of the following types (illustrated in Figure 3.1 (a)): a direct reflection y^r , an expansion y^e , an outside contraction y^{oc} , or an inside contraction y^{ic} . However, in some cases the algorithm can perform a shrink, which only keeps the best vertex, and shrinks the simplex by moving all other vertices close to the best one (illustrated in Figure 3.1 (b)). In each iteration, the algorithm first evaluates objective value at y^r , then potentially at y^e , y^{oc} , y^{oc} or y^{ic} . The type of outcome is determined by the objective values at these points relative to the objective value of the $n + 1$ vertices in the current simplex. The algorithm can be terminated when the diameter of the simplex reduces below a predetermined threshold, or when the difference between the minimum and maximum objective values at all vertices becomes smaller than a threshold, or alternatively, as the original authors suggested in *Nelder and Mead* (1965), when the variance in the objective function values at all vertices drops below certain threshold, at which point the best vertex in the final simplex is used as the solution. In the experiments we use a modification of the implementation in the NLOpt library (*Johnson*), which includes the Nelder-Mead variant described in *Box* (1965) that incorporates bound constraints.




Figure 3.1: Possible outcomes of one Nelder-Mead iteration, starting from initial simplex $\{y^0, y^1, y^2\}$

3.5.2 Framework B: alternating optimization

An alternative framework for solving (VC) is by using an *alternating optimization* (AO) scheme. The AO method, sometimes also referred to as the Gauss-Seidel method (), is suitable for situations where the problem is difficult to solve directly, and there is a natural partitioning of variables for which the optimization is easy to carry out for each group of variables. The problem is then solved by alternating among each subset of variables. Our second framework, illustrated in Figure 3.2, is as follows:

1. Generate an initial (S, R) pair, solve problem $(LS^{(S,R)})$ to obtain a set of apertures.
2. Given the apertures, optimize the gantry speed and dose rate by solving problem $(GD^{(A)})$ described in Section 3.3.1.
3. If the objective function value does not improve, terminate; otherwise update gantry speed and dose rate and go to Step 4.
4. Refine the apertures given the updated (S, R) pair until no more objective improvements can be made; go to Step 2.

Note that the solution to problem $(GD^{(A)})$ will provide us with an optimal fluence rate y^* . However, there may be multiple (S, R) combinations that correspond to this ratio and

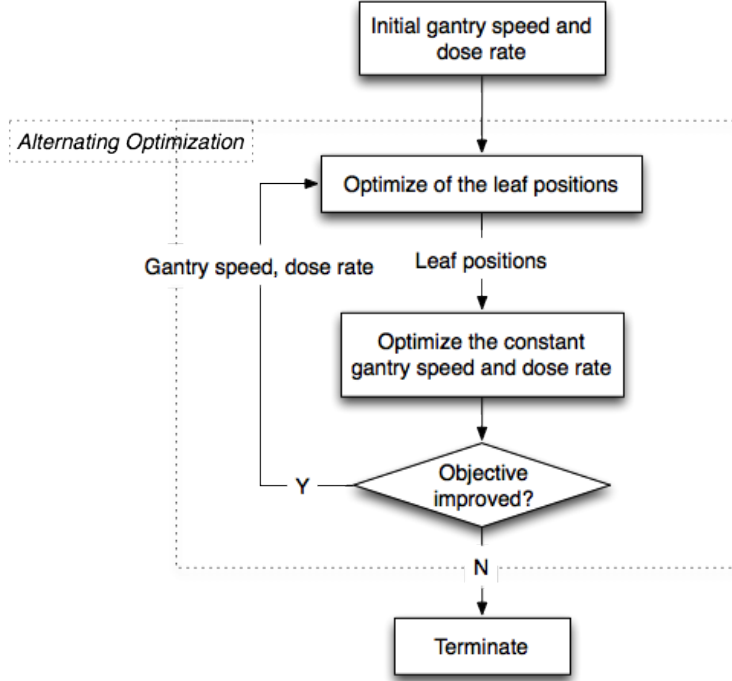


Figure 3.2: Optimization scheme for VMATC treatment plans

satisfy (3.6) and (3.7). Although it is not apparent which (S, R) pair will lead to better final solutions, selecting a slower gantry speed will inevitably result in longer treatment times. Moreover, the aperture selection procedure will take advantage of a slower gantry speed, and allow the MLC leaves to travel longer distances between two control points. This will in turn lead to smaller \bar{S}^U value the next time we solve problem $(GD^{(A)})$. As a result, from an initial (S, R) pair, treatment plans generated later in framework B will have the same or longer treatment times than those seen earlier. This can be mitigated by choosing the resulting (S, R) values in such a way that the treatment time is minimized, in other words, choose values:

$$\begin{aligned}
 S = \bar{S}^U, R = \bar{S}^U y^* & & \text{if } y^* \leq \frac{R^U}{\bar{S}^U} \\
 S = \frac{R^U}{y^*}, R = R^U & & \text{if } y^* > \frac{R^U}{\bar{S}^U}.
 \end{aligned}$$

Under conditions on the differentiability and/or the convexity of the objective function, and the feasible set being a Cartesian product of sets, global convergence results for AO have

been established for both the unconstrained and constrained cases (*Grippo and Sciandrone (2000); Tseng (2001)*). However, in our case the compatible sequence of apertures depends on the choice of gantry speed. And because our aperture selection is done with a heuristic procedure, performing Step 2 is not guaranteed to yield the optimal apertures for an (S, R) combination. Therefore, we do not expect the AO to produce a globally optimal solution for (VC).

Under this framework, the method can potentially inspect more (S, R) pairs since refining the apertures will in general be faster than the initial aperture selection problem. However, no guidance is available in the selection of new (S, R) points when the AO procedure terminates. We augment the AO with a multi-start method that randomly generates initial S and R values in the feasible region, and runs the AO procedure repeatedly from each initial value to search for the best solution.

3.6 Experiments

Our test dataset includes five clinical prostate cancer patients cases (the same five cases used in Chapter II), gathered from University of California at San Diego (UCSD) Moores Cancer Center. There is one target (prostate) with a prescription dose of 79.2 Gy over 44 fractions, and three critical structures (bladder, rectum, femoral heads) besides normal tissue. The dose volume histogram (DVH) criteria for evaluating the treatment plans are summarized in Table 3.1. The patient volume is divided into a grid of voxels of size $4 \times 4 \times 2.5$ mm³. This entire grid is used in the final dose calculation to evaluate the delivered dose under a given treatment plan. However, in the optimization problem we consider a downsampled grid, in which we keep one in every 8 voxels in the critical organs and one in every 64 voxels in the normal tissue. For each patient the treatment includes one arc with 177 total control points, and the treatment beam is discretized into 1×1 cm² beamlets. The resulting problem dimensions for all 5 patients are presented in Table 3.2. We use a set of machine parameters (shown in Table 3.3) that resembles that of a Varian treatment machine.

Structure	Threshold (Gy)	Percentage (%)
PTV	73.7	≥ 99
	79.2	≥ 95
	87.1	≤ 10
Bladder	65	≤ 17
	40	≤ 35
Femoral heads	50	≤ 10
	45	≤ 25
	40	≤ 40
Rectum	75	≤ 15
	70	≤ 25
	65	≤ 35
	40	≤ 45

Table 3.1: Clinical criteria for the prostate cancer cases.

Case	voxels	beamlets	nonzero D_{bmnj}	distance between control points ($^\circ$)		
				$\delta_3, \dots, \delta_{176}$	δ_2, δ_{177}	δ_1
1	18,375	14,337	37,224,301	$\frac{330}{174}$	$\frac{1}{2} \cdot \frac{330}{174}$	0
2	14,392	11,328	26,751,546	$\frac{330}{174}$	$\frac{1}{2} \cdot \frac{330}{174}$	0
3	18,469	21,417	52,850,633	$\frac{330}{174}$	$\frac{1}{2} \cdot \frac{330}{174}$	0
4	9,747	14,337	47,707,351	$\frac{330}{174}$	$\frac{1}{2} \cdot \frac{330}{174}$	0
5	17,342	17,700	61,606,811	$\frac{340}{174}$	$\frac{1}{2} \cdot \frac{340}{174}$	0

Table 3.2: Problem dimensions of the downsampled cases.

Our objective F is a smooth piece-wise quadratic function, which penalizes both over- and under-dosing of each individual voxel in the target, and only overdosing of voxels in the other structures. To facilitate comparison with the full VMAT plans, we use the same weighting factors for different structures as those in Chapter II, and for the same of making a comparison, the resulting treatment plans are scaled so that the dose distribution achieves 95% target coverage. Our algorithm is implemented in C++, with all convex optimization problems solved using CPLEX 12.5. The evaluation of delivered dose and objective function in the aperture selection procedure is implemented with the Parallel Programming and Computing Platform (CUDA) to take advantage of the parallel computing power of the Graphics

Processing Unit (GPU).

Dose rate		gantry speed		physical leaf speed	distance from source	
(MU/sec)		(deg/sec)		(cm/sec)	(cm)	(cm)
R^L	R^U	S^L	S^U	v	SAD	SCD
0	10	0.83	6.0	2.25	100	53.9

Table 3.3: Machine parameters used in VMATC experiments.

3.6.1 Exact vs direct search methods for refining apertures

In Section 3.4 we proposed two different methods for fine-tuning the leaf positions given a complete plan with K control points. The exact method requires us to solve a number of optimization problems and potentially requires much longer solution time compared to the direct search method, however it guarantees the optimal leaf setting in each optimization problem, whereas the direct search method is fast but only explores a finite number of options. First we explore the tradeoff between solution time and quality by employing the exact method in refining the set of selected apertures.

Expecting a long runtime for the exact method, we chose Case 2, the smallest case of the five in the number of beamlets per beam, as the sample case for our experiment here. The size of the problem is further reduced by using $K = 34$ total control points, each 10° apart. We set $S = S^U$ so that the range of feasible leaf positions in each optimization problem is the smallest, and chose $R = 7.34$ MU/sec randomly from the feasible range. Both methods are tested under the same setting. We compare the objective function value in different iterations (where one iteration determines the leaf setting in one row and one control point at a time). Figure 3.3 shows the progression of the objective function vs iteration number for the two methods.

The exact method outperforms the direct search method in each iteration in reducing the objective function value. We stop this test when the apertures at every control point have been refined once, at which point the final objective function value obtained by the exact




Figure 3.3: Objective function value in exact and direct search methods for refining apertures

method is 1.8% better than that obtained by the direct search method. However, solving the optimization problems for all refinement iterations only took 1.9 seconds for the direct search method, whereas the exact method took 30,111.4 seconds (more than 8.3 hours). The significant reduction in solution time clearly outweighs the slight loss of plan quality in using the direct search method, which will become more pronounced if we apply these methods to a case with a larger input dataset. Moreover, if we repeat the refinement procedure with the direct search method for more iterations, the objective quickly reduces to below the value obtained in the exact method. This comparison is shown in Figure 3.4. The direct search method took a total of 21.1 seconds to converge, at which point no more than 0.1% of objective improvements can be made after refining all leaf pairs at all control points. For these reasons we will focus exclusively on the direct search method in the aperture selection in the experiments below.

3.6.2 Overall strategy under frameworks A and B

For the remainder of this section we focus our attention on designing a treatment plan within 30 minutes, a reasonable estimate of the maximum time available for treatment planning in clinics. Within frameworks A and B, our strategy depends on how fast we can solve the aperture selection problem. For framework A in particular, if this time is short




Figure 3.4: Objective function value in exact and direct search methods for refining apertures, with the direct search method run until convergence

enough to allow the Nelder-Mead algorithm to converge quickly, we can pursue a multi-start approach and explore different areas in the feasible region. However, if the aperture selection time is long, we need to then look for a suitable starting simplex for the Nelder-Mead algorithm that in general provides good solutions. On the other hand, in framework B, the alternating optimization procedure is always embedded in a multi-start approach, where whenever the AO converges, another (S, R) pair will be chosen randomly to initialize the AO procedure again, until the time limit is reached.

Table 3.4 summarizes the average runtime of the aperture selection described in Section 3.4, calculated from 30 independent runs each associated with a random (S, R) pair. Note that in framework B, from every starting (S, R) pair the AO procedure consists of performing the aperture selection as well as (S, R) optimization potentially multiple times, the average runtime may take longer than these reported values. We set the termination criteria so that if no more than 0.1% improvement in the objective function can be made by repeating Step 3(b) in the algorithm, we terminate the search. The resulting runtimes range from just over 2 minutes to more than 6 minutes, which will only allow a few to a maximum of just over a dozen (S, R) pairs to be explored in either frameworks.

Case	1	2	3	4	5
Average runtime (s)	252.6	136.5	379.5	260.8	381.8

Table 3.4: Average runtime of the aperture selection problem from 30 random (S, R) pairs

3.6.2.1 Starting (S, R) point under Framework A

Given the observation above, the Nelder-Mead algorithm in framework A is not likely to start and converge multiple times within the time limit, and therefore we need to look for one particular simplex in the feasible region that in general leads to good solutions. The number of vertices used to construct the simplex in the Nelder-Mead algorithm here equals three, resulting in a triangle.

Again, without any information about the shape of function $\tilde{H}(S, R)$, no theoretical results exist to guide the choice. Therefore we need to test different configurations for the triangle in terms of size and the area they cover. We terminate the Nelder-Mead algorithm if the maximum difference in the objective values from all vertices is less than 0.1%. Figure 3.5 (a) shows the initial 15 triangles used in our tests. We take two large, four medium and nine small triangles and spread them out to cover the entire feasible region. After the Nelder-Mead algorithm reaches 30 minutes or terminates early (which in our case does not happen in any test) from each starting triangle for each case, we put all final solutions in a pool and plot the best 80%. From Figure 3.5 (b) we can see that these solutions are all concentrated on the lower right half of the feasible region, which is reasonable since neither low dose rate combined with high gantry speed, or high dose rate combined with low gantry speed are likely to generate good quality plans. We fit these solutions to a line, and expand the line in both directions to encompass the set of solutions. From this region we then randomly generate five additional triangles. We repeat the generation process to find a set of triangles that together cover most of this region, the results are shown in Figure 3.6, and these triangles are tested in the Nelder-Mead algorithm.

At the end of the tests we compare the 20 solutions, started with the original 15 and the

additional 5 triangles, for every individual case in terms of objective, treatment time and total MU values in Figure 3.7, where all values are scaled to the average of the 20 solutions. If we require that the objective function value as well as the treatment time be lower than the average for all cases, only the medium triangle on the right bottom corner in Figure 3.5 (a), and the red triangle in Figure 3.6 can meet the requirement. Both are able to keep the total MU under 105% of the average for all cases. In the subsequent tests, we use the latter triangle for its ability to provide better objective function values in comparison.

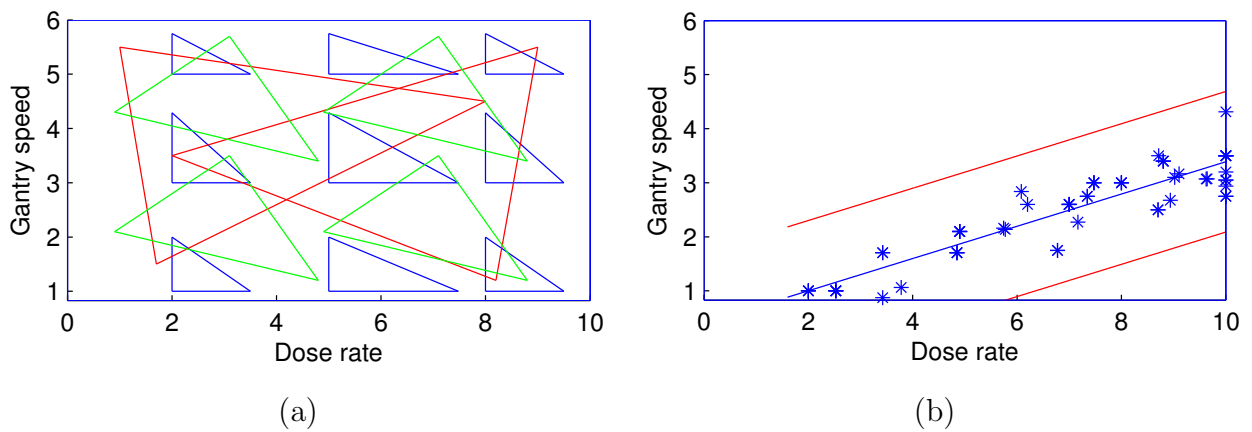


Figure 3.5: (a): 15 initial triangles tested in the Nelder-Mead algorithm; (b): (S, R) pairs (stars) corresponding to the best 80% solutions




Figure 3.6: 5 additional triangles tested in the Nelder-Mead algorithm

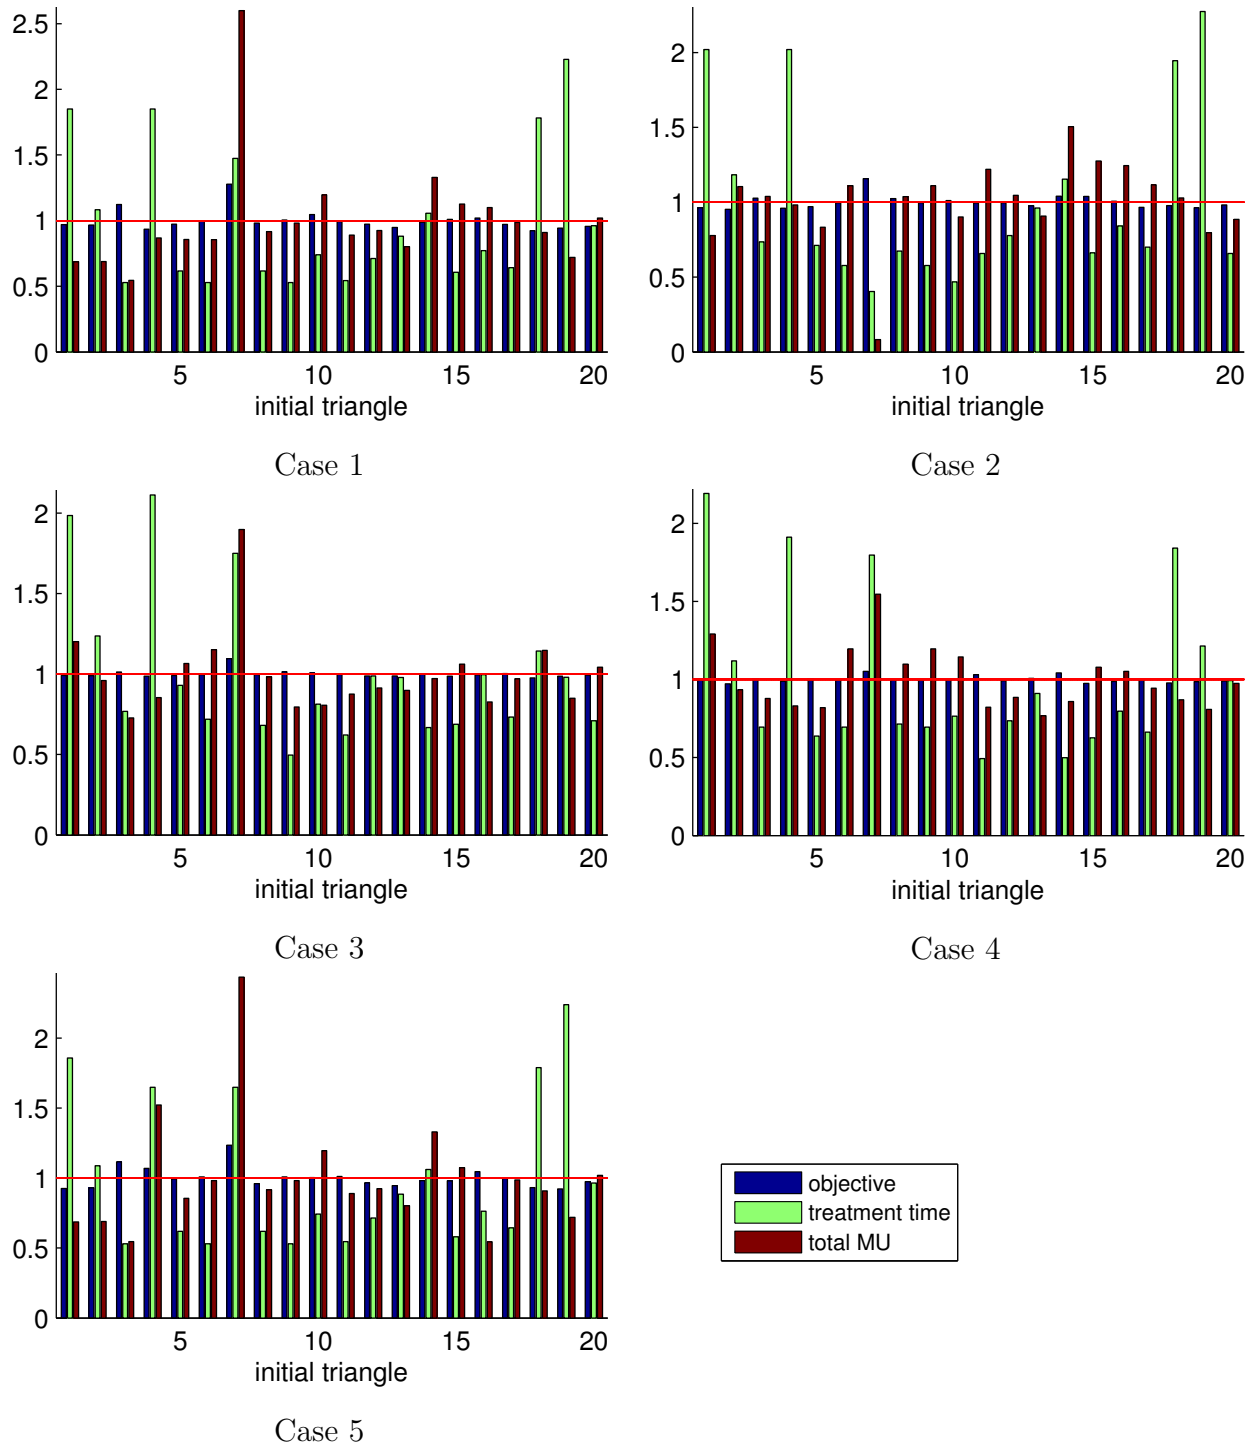


Figure 3.7: Objective, treatment time and total MU comparison for all 20 triangles tested in the Nelder-Mead algorithm. 1 — 9: small; 10 — 13: medium; 14 — 15: large; 16 — 20: additional five triangles

3.6.2.2 Performance under framework B

Next we compare the performance of framework B and framework A. In implementing framework B, we terminate the the alternating optimization when no more than 0.1% improvement can be made in the optimization of (S, R) . Once the AO stops, another initial (S, R) pair is chosen to start the optimization again, until the time limit is reached.

We ran framework B independently for ten different times, each time randomly and uniformly choosing an initial (S, R) pair. We keep the best solution encountered in each run for all cases, and again plot the (S, R) pairs corresponding to the 80% of the best solutions on the feasible region. We can see in Figure 3.8 that the solutions again in general lie within a band close to the bottom right corner.

Based on this information, we adjust the way that the initial (S, R) pairs are sampled: we first fit a line to these solutions, and uniformly generate a point along the line. Then the point is moved perpendicularly to the line according to a normal distribution with mean 0 and variance σ^2 , where σ^2 is set to be equal to the variance of the distance from the set of points to the fitted line. The (S, R) pairs generated this way stay close to the band of good solutions, and still have a chance of exploring different areas in the feasible region. We ran framework B again 10 times under this setting, and record the range of resulting objective, treatment time and total MU values corresponding to these plans. Figure 3.9 shows these quantities for all five patient cases, normalized to the solutions from framework A. Even though the objective values under Framework B are better for case 5, for all other cases framework A clearly outperforms framework B. Framework A also achieves good performance in terms of treatment time and total MU values. Therefore we choose framework A over B in all the remaining experiments below.

3.6.3 Comparison with VMAT plans

Finally, we compare the VMATC plans obtained from framework A to VMAT treatments, for which we already have a method of optimizing the treatment plans (described in Chapter

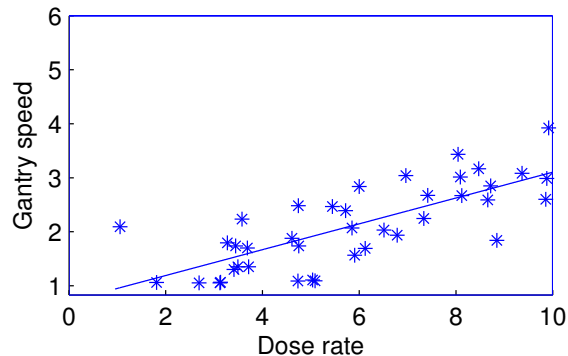


Figure 3.8: (S, R) pairs (star) corresponding to the best 80% solutions from 10 initial runs under Framework B

II). In Table 3.5 we summarize the statistics of the VMATC plans versus those of the corresponding VMAT plans. Out of the five test cases, VMATC actually outperforms the full VMAT plans for 4 cases in terms of objective value and for 2 cases in terms of treatment time. Moreover, the dose distributions for the VMATC plans are able to satisfy the same set of DVH criteria as the corresponding VMAT plans with no significant deterioration in quality, as can be seen from Figure 3.11. The only compromise in VMATC treatments is that the total MU is in general considerably higher under VMATC. Without the ability to dynamically control the gantry speed during the treatment, VMATC cannot quickly open up or close apertures, or selectively turn off the source during the treatment. Therefore it has to use smaller apertures throughout the treatment to avoid giving large amounts of dose from unfavorable angles. Figure 3.10 shows the average size of active apertures (i.e., those that have positive fluence rates), calculated by adding up the area of each open leaf row, for the corresponding VMAT and VMATC treatment plans. The apertures used in VMAT plans are significantly larger than those in VMATC plans on average, ranging between 42% — 90% larger for the five cases tested. Having to use small apertures makes VMATC less effective in delivering the prescribed dose, and leads to higher total MU values in comparison.

Case	modality	treatment time (s)	total MU	runtime (s)	objective value
1	VMAT	149.3	611.5	78.8	22279
	VMATC	107.4	751.6	1800	21166
2	VMAT	149.3	667.7	71.5	41763
	VMATC	110.8	1017.8	1800	42589
3	VMAT	141.1	539.0	59.2	33865
	VMATC	171.4	979.1	1800	30985
4	VMAT	142.6	557.6	54.2	54344
	VMATC	171.4	979.1	1800	52451
5	VMAT	155.5	610.4	67.6	32819
	VMATC	176.6	1008.8	1800	30693

Table 3.5: VMATC and full VMAT treatment plans for cases 1-5

3.7 Conclusions

We developed two solution frameworks for the optimization problem for VMATC treatments. We considered the optimization of the constant gantry speed and dose rate as well as the apertures. We recommend the configuration under framework A that consistently provides high quality solution in terms of objective function value, total MU and total treatment time. A comparison with VMAT on actual patient cases shows that VMATC is capable of producing very high quality plans. However, the lack of the ability to dynamically change the dose rate and gantry speed typically leads to higher total MU values for the VMATC plans, and the treatment planning time is much longer compared to VMAT, for which the optimization usually only takes around 1 minute to complete. Moreover, one set of optimization parameters may not be suitable for multiple cancer sites and prescription doses, and we may need to “retune” the implementation in these cases.

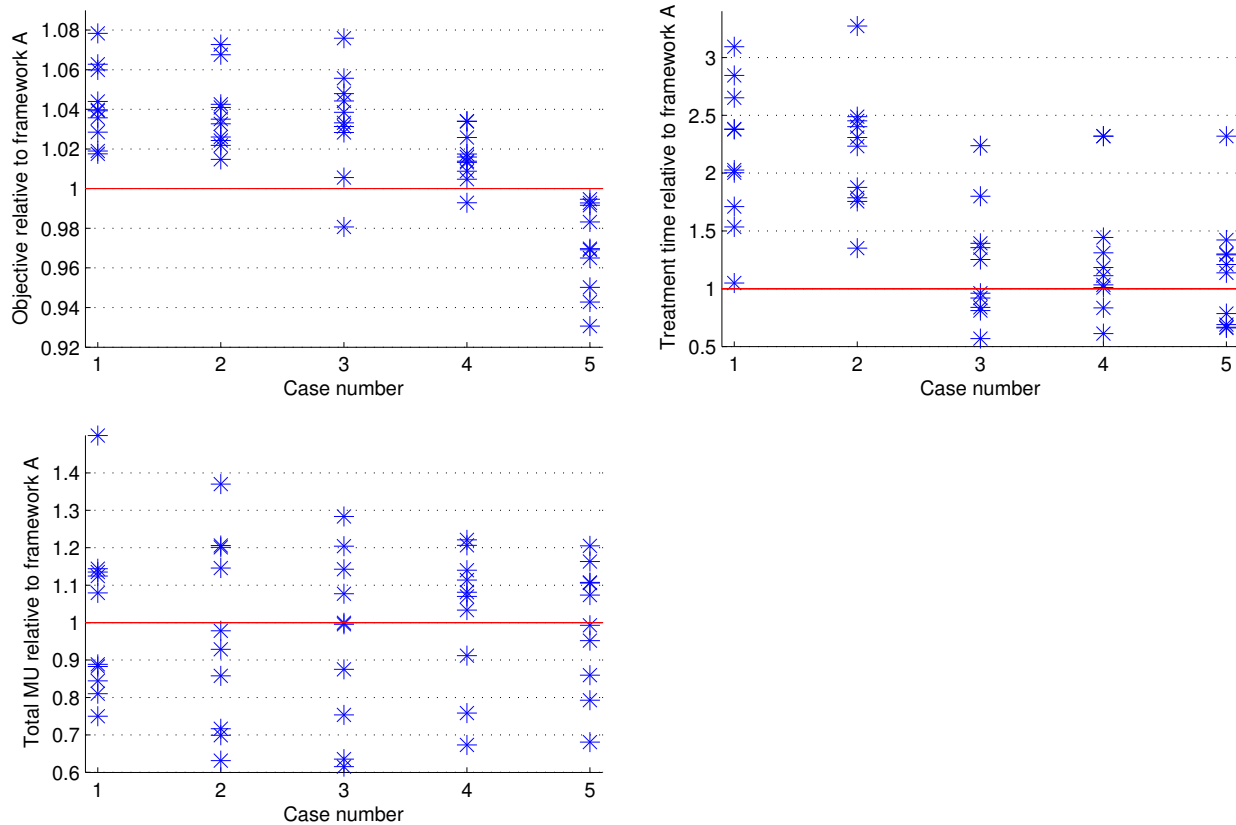


Figure 3.9: Objective, treatment time and total MU of 10 plans under framework B (star) relative to the best plan from framework A




Figure 3.10: Comparison of average aperture size for VMAT (orange) vs VMATC (green) plans

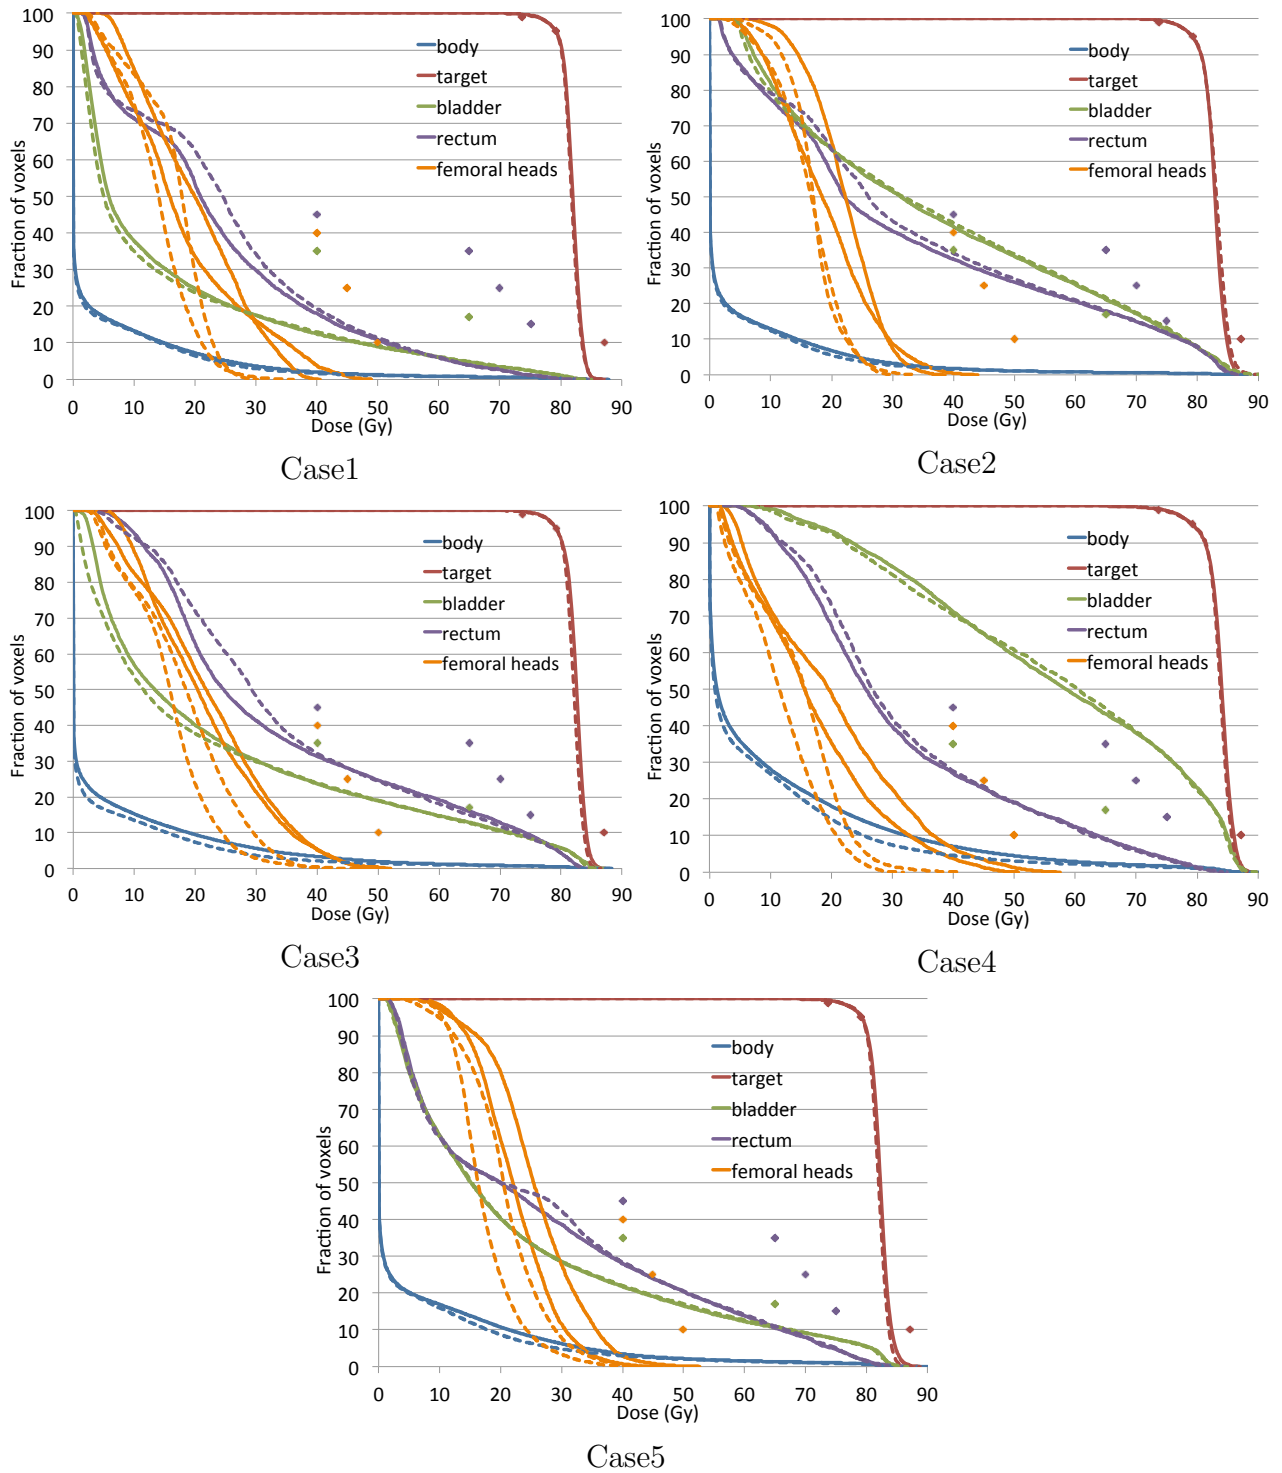


Figure 3.11: DVH of VMATC (dashed) and VMAT plans (solid) for Cases 1-5

CHAPTER IV

Adaptive Radiation Therapy

4.1 Introduction

Many aspects of Intensity Modulated Radiation Therapy (IMRT) treatments are subject to uncertainty, of which the most common is the uncertainty in the positions of the tumor and its surrounding organs with respect to the treatment beams. The treatment is subject to (i). uncertainties caused by the day-to-day variations in and after the setup process and changes in the patient's anatomy; and (ii). uncertainties during the treatment, caused by the shifts in structures due to motions such as breathing and filling up of internal organs (see, e.g., *Langen and Jones (2001)*, *Van Herk et al. (2004)*). The former is usually referred to as inter-fraction uncertainty, and the latter – intra-fraction uncertainty. These uncertainties can lead to deviations of delivered dose from dose designed in the treatment planning process, and compromise the quality of the actual treatment.

The conventional approach to accounting for uncertainties is to first define the Clinical Target Volume (CTV), which includes the tumor and its immediately surrounding tissues that likely contains cancerous cells, then add an internal margin to account for intra-fraction motion which leads to the Internal Target Volume (ITV), and finally add a margin around the ITV to produce the Planning Target Volume (PTV). The PTV is then used as the target in the treatment planning process. The margin is intended to account for both inter- and intra-fraction motions, and its size is determined based on statistics of the patient population (*van*

Herk et al. (2000); Van Herk et al. (2004)). However, the use of margin might not adequately model the dose delivered to different structures under uncertainty for several reasons: First, shifts in the patient and/or organ position do not usually result in the same rigid shifts in the dose distribution. In fact motions may heavily distort the dose distribution and lead to under-dosing of the targets and/or over-dosing of the healthy structures; secondly, the use of a uniform margin may not be appropriate, particularly when the target is closely surrounded by critical structures. Overlap of the margin and other structures will likely cause the normal structures to be over-dosed in delivering the treatment. Moreover, the margin-based model uses only the aggregate dose of the entire treatment and disregards the per fraction dose, which is important for certain targets and critical structures. As a result the delivered dose in certain fractions may differ from the prescribed dose even when the aggregate planned dose is adequate.

Many researchers have therefore explored various alternative approaches. In the field of robust optimization, *Chu et al. (2005)* approximated the dose delivered to every voxel as normally distributed assuming Central Limit Theorem (CLT), and calculated its mean and variance using 7 pre-calculated scenarios and their probabilities. The authors then proposed a robust LP with constraints on the probability of over- or under-dosing every voxel, and solved the problem by transforming it first to a Second Order Cone Program (SOCP). Similarly, *Olafsson and Wright (2006)* assumed normal cumulative dose and calculated its mean and variance from 5 scenarios. An optimization problem was then formulated as an LP with bounds on the delivered dose, and solved through an SOCP reformulation. However, imposing the constraint that the probability of every single voxel meeting certain criteria does not guarantee the collection of voxels will all meet those criteria, especially when the number of voxels is very large. In a study to investigate the benefits of robust optimization for intra-fraction breathing motion, *Chan et al. (2006)* and *Bortfeld et al. (2008)* associated breathing phases with probability density function (pdf) and bounds on the pdf to model uncertainty in motion patterns. The optimization of beam weight profile was done by solving

an LP that took into account the uncertainty set. Relying on CLT, *Sobotta et al.* (2010) assumed that the tail of the dose distribution under uncertainty resembled the tail of a Gaussian distribution, and thus the mean, variance and the probability of the dose lying in certain intervals of interest can be calculated. The authors suggested an optimization problem that maximized the probability of finding the penalty function within the interval, while imposing constraints on violation of that interval. Finally, *Fredriksson et al.* (2011) applied min-max optimization and minimized the worst case penalty function in a number of scenarios. However, the fact that the treatment is usually delivered over many fractions means that the realized cumulative dose is generally far different from the dose in any of the extreme cases, and the worst case optimization is too pessimistic.

Stochastic programming is also a method of choice when protection is sought against uncertainty. *Löf et al.* (1999) solved a model that maximizes the expected value of a function of the probability of tumor control and normal tissue damage. However the authors had to resort to Monte Carlo simulation methods for approximating the objective function when the number of fractions N is above 5 because of computational difficulties. *Li and Xing* (2000) modeled the expected dose to a voxel as the sum of the dose to a different position, weighted by the probability of the voxel occupying that position, and penalized the difference between the prescribed and expected dose for every voxel. The experiments showed that the model achieved improved sparing of critical structures and similar target coverage. *Unkelbach and Oelfke* (2005) employed a similar dose calculation method, but incorporated the variance of dose in target voxels as part of the penalty function to ensure homogeneous target dose. *Baum et al.* (2006) used coverage probability to represent the probability of a point in space being covered by a certain organ under uncertainty, and the objective function consisted of the sum of target and critical organ penalties weighted by their respective coverage probability. These studies all relied on one assumption that may not be appropriate in real-world situations: the dose distribution in space is assumed to be constant regardless of motion and error. In reality any motion will inevitably change

the dose distribution, sometimes substantially. *McShan et al.* (2006) proposed a technique called Multiple Instance Geometry Approximation (MIGA), which calculated the dose distribution for a number of scenarios, and the treatment plan optimization is performed using the expected dose calculated from these scenarios. However, one potential issue with this approach is that the number of scenarios used (seven) may not be large enough to capture the distribution of uncertainty, or computational difficulties may become apparent. In a tumor-tracking based approach, *Nohadani et al.* (2010) took advantage of the 4-D CT data and separated the patient's breathing motion into 6 phases, and proposed a model that optimized the intensity-phase combination and minimized the expected penalty. Finally, *Men et al.* (2011) proposed an optimization problem that included expected penalties for both the per-fraction dose and the cumulative dose. Assuming CLT, the expected penalty for the cumulative dose was replaced by the penalty of the expected dose, and the per-fraction dose was approximated by a large number of scenarios.

Many of the above-mentioned approaches relied on CLT, which was reasonable when the number of fractions in the treatment was large. However, for treatments with only a small number of fractions (a good example is Stereotactic Body Radiation Therapy (SBRT), which involves treatments up to only 5 fractions), this assumption may not be appropriate anymore. This also applies to the re-optimization of the treatment plans, where only the remaining fractions in the treatment are considered.

Technological advances have made it possible to monitor the patient's anatomy change, fraction-to-fraction setup variations, as well as to reconstruct the actual dose distribution using daily images of the patient (*Mohan et al.* (2005); *Langen et al.* (2005); *Oldham et al.* (2005); *Yang et al.* (2007)). To exploit the additional information provided by these technologies, many techniques, which we all refer to as Adaptive Radiation Therapy (ART), have been developed. *Ferris and Voelker* (2004) proposed a DP framework for determining the optimal dose in each fraction of the treatment under random errors. The authors proposed a neuro-dynamic programming (NDP) based heuristic that, at every stage, minimized

the expected approximate terminal penalty assuming the same policy is applied henceforth. The optimization is done by way of simulation, where 5 different shifts were used to model the random errors in any scenario. *Deng and Ferris (2008)* employed the same solution technique, except that a neural network was used in approximating the terminal penalty. However, the strategies used are chosen from a pool of pre-selected decisions, which may not be the best achievable. *de la Zerda et al. (2007)* combined closed-loop control theory with ART, and proposed two adaptive optimization methods that, based on the knowledge of the deviation in delivered dose and the geometry change, tried to correct the deviation in the immediate next fraction and the remaining fractions respectively. However, the optimizations were done without considering any uncertainty factors, even though motion was expected in the treatment. *Saka et al. (2011)* followed a similar approach and optimized the constant plan for all remaining fractions after taking the delivered dose and observed geometry change into account. However the optimization model did not include any protection against uncertainty in the dose to be delivered. Finally, *Sir et al. (2012)* employed the MIGA (*McShan et al. (2006)*) approach in modeling inter-fraction uncertainties, and integrated stochastic programming with open-loop feedback control to optimize the constant plan to the end of the treatment. The nature of the stochastic programming approach means that a larger number of fractions inevitably translates to significantly larger problem size. Also realistic model of uncertainty requires potentially many scenarios, which can easily lead to intractable optimization problems.

The goal of this research is to build upon the previous work done by the fore-mentioned researchers, and achieve four things that, to our knowledge, have not all been considered by any single research simultaneously:

- Adequately model the patient’s inter-fraction motion uncertainties;
- Take both per-fraction and cumulative dose evaluation criteria into account in the optimization model;

- Solve the stochastic programming problem with high, verifiable quality for real-world cases with long treatment courses;
- Extend the optimization approach to cover adaptive radiation therapy, where central limit theorem may not be an appropriate assumption in modeling the total dose to be delivered.

The remainder of this chapter is organized as follows: first we describe the general assumptions and the optimization model in Section 4.2. We explain in detail our proposed solution approach, and extend the approach to ART problems in Sections 4.4 and 4.4.4. In Section 4.5 we compare our model with the model proposed in *Men et al.* (2011), which is simpler to implement in practice, under different settings. Finally we present numerical experiments with real-world patient cases and results in Section 4.6.

4.2 The adaptive IMRT optimization model

Let N denote the total number of fractions in the treatment, and n the number of fractions already delivered. Let \mathbf{s} be a random scenario, defined by a rigid shift of the patient along the three axes, and let $\mathbf{p}^{(n)} = (\mathbf{s}^1, \dots, \mathbf{s}^n)$ be a random path consisting of n consecutive random scenarios $\mathbf{s}^1, \dots, \mathbf{s}^n$. Note that throughout this chapter we use boldface letters to represent random parameters, and regular font for parameters that are deterministic. In addition, let $D(\cdot)$ denote the dose deposition matrix. If the argument is a scenario, it corresponds to the dose deposition matrix under that scenario. On the other hand, if the argument is a path, $D(\cdot)$ then equals the sum of all the dose deposition matrices under the series of scenarios in that path. Moreover, let $\bar{z}^{(n)}$ denote the dose already delivered in the first n fractions, and let $z^{(N-n)}(\cdot)$ denote the dose that is yet to be delivered in the next $N - n$ fractions. Moreover, to distinguish between dose delivered in one and multiple fractions, let $z^k(\cdot)$ denote the dose to be delivered in fraction k , $k = 1, \dots, N$. Finally, let L and C be the sets of evaluation criteria that apply to fractional and cumulative dose, respectively, and let $G_\ell(\cdot)$, $\ell \in L$ and

$G_c(\cdot)$, $c \in C$ be the the (convex) penalty term that applies to fractional and cumulative dose, respectively.

The problem for adaptive IMRT under uncertainty can then be formulated as:

$$(P) \ g^* = \min_x \quad E \left[\sum_{\ell \in L} G_\ell(z^{n+1}(x, \mathbf{s})) + \sum_{c \in C} G_c \left(\frac{1}{N} \cdot (\bar{z}^{(n)} + z^{(N-n)}(x, \mathbf{p}^{(N-n)})) \right) \right]$$

subject to:

$$z^{n+1}(x, \mathbf{s}) = D(\mathbf{s})^\top x \quad (4.1)$$

$$z^{(N-n)}(x, \mathbf{p}^{(N-n)}) = D(\mathbf{p}^{(N-n)})^\top x \quad (4.2)$$

$$x \in X. \quad (4.3)$$

Here X is a compact feasible region for variable x . We minimize the sum of two expected penalty terms: one that applies to the dose to be delivered in the immediate next fraction, and another that applies to the average dose in the entire treatment course. Note that since \mathbf{s} and $\mathbf{p}^{(N-n)}$ are random quantities, the dose to be delivered, namely $z^{n+1}(x, \mathbf{s})$ and $z^{(N-n)}(x, \mathbf{p}^{(N-n)})$, are also random. In this formulation, if $n = 0$ and thus $\bar{z}^{(n)} = 0$, the entire treatment is yet to be delivered. On the other hand if $n > 0$, this formulation becomes an adaptive model that takes delivered dose $\bar{z}^{(n)}$ into account and re-optimizes the treatment plan for each fraction in the remainder of the treatment.

Note that functions $G_\ell(z^{n+1}(x, \mathbf{s}))$ and $G_c \left(\frac{1}{N} \cdot (\bar{z}^{(n)} + z^{(N-n)}(x, \mathbf{p}^{(N-n)})) \right)$ implicitly depend on matrices $D(\mathbf{s})$ and $D(\mathbf{p}^{(N-n)})$. Since it depends on the patient's internal geometry, the relationship between the patient setup scenario and the dose deposition matrices cannot be easily characterized mathematically. Therefore instead of solving problem (P) directly, we apply a technique known as the Sample Average Approximation (SAA) (*King and Wets (1991); Kleywegt et al. (2002); Verweij et al. (2003); Linderoth et al. (2006)*). We first randomly draw samples from the distribution of \mathbf{p} to replace the random parameters in the penalty function. The true objective function value is then approximated with the average of these samples. By doing this we essentially transformed problem (P) into a deterministic

problem, where the samples are known parameters:

$$(SAA) \ g_m = \min_x \frac{1}{m} \sum_{i=1}^m \left[\sum_{\ell \in L} G_\ell(z^{n+1}(x, s_i)) + \sum_{c \in C} G_c \left(\frac{1}{N} \cdot (\bar{z}^{(n)} + z^{(N-n)}(x, p_i^{(N-n)})) \right) \right]$$

subject to:

$$z^{n+1}(x, s_i) = D(s_i)^\top x, \quad \forall i = 1, \dots, m$$

$$z^{(N-n)}(x, p_i^{(N-n)}) = D^{(N-n)}(p_i^{(N-n)})^\top x, \quad \forall i = 1, \dots, m$$

$$x \in X.$$

Here s_i is the first in the series of scenarios in $p_i^{(N-n)}$ for all i .

Note that if the number of samples $m = \infty$, problem (SAA) will become equivalent to problem (P). However, in reality it is not possible to take an infinite number of samples, and we must decide on the value of m before problem (SAA) can be formulated or solved. Therefore we need to ensure that m is large enough so that our approximation of (P) with (SAA) is adequate, i.e., the resulting objective function g_m is close enough to g^* .

4.3 Setting m a-priori

One possible option for selecting m is to choose a sufficiently large value a-priori, so that the resulting objective function value at the optimal solution in the (SAA) has a high probability of being a nearly-optimal solution to (P).

Assuming that X is bounded and the expected penalty:

$$g(x) \doteq E \left[\sum_{\ell \in L} G_\ell(z^{n+1}(x, \mathbf{s})) + \sum_{c \in C} G_c \left(\frac{1}{N} \cdot (\bar{z}^{(n)} + z^{(N-n)}(x, \mathbf{p}^{(N-n)})) \right) \right]$$

is finite for all $x \in X$, *Ruszczynski and Shapiro* (2003) derived the number of samples m needed for the solution to (SAA) to be have a high probability of being a nearly-solution to (P). Let \hat{S}_m^δ and S^ϵ represent the set of approximate solutions to (SAA) and (P) respectively, defined by:

$$\hat{S}_m^\delta = \left\{ x \in X : \frac{1}{m} \sum_{i=1}^m \left[\sum_{\ell \in L} G_\ell(z^{n+1}(x, s_i)) + \sum_{c \in C} G_c \left(\frac{1}{N} \cdot (\bar{z}^{(n)} + z^{(N-n)}(x, p_i^{(N-n)})) \right) \right] \right\}$$

$$\leq g_m + \delta\},$$

and

$$S^\epsilon = \{x \in X : E \left[\sum_{\ell \in L} G_\ell(z^{n+1}(x, \mathbf{s})) + \sum_{c \in C} G_c \left(\frac{1}{N} \cdot (\bar{z}^{(n)} + z^{(N-n)}(x, \mathbf{p}^{(N-n)})) \right) \right] \leq g^* + \epsilon\}.$$

The number of samples m required for $P(\hat{S}_m^\delta \subset S^\epsilon) \geq 1 - \alpha$, for small $\epsilon > \delta \geq 0$, is then:

$$m \geq \frac{12\sigma_{\max}^2}{(\epsilon - \delta)^2} \left(\eta \log \frac{2DL}{\epsilon - \delta} - \log \alpha \right), \quad (4.4)$$

where $D := \sup_{x, y \in X} \|x - y\|$ is the diameter of the feasible region X , η is the dimensionality of the feasible region, and L is the Lipschitz constant of the penalty function, defined such that:

$$|g(x) - g(y)| \leq L\|x - y\|$$

for all $x, y \in X$. Moreover, define $u(x)$ as a mapping from the set X/S^ϵ to S^0 . σ_{\max}^2 is then defined as:

$$\begin{aligned} \sigma_{\max}^2 \doteq & \max_{x \in X/S^\epsilon} \text{Var} \left[\sum_{\ell \in L} G_\ell(z^{n+1}(u(x), \mathbf{s})) + \sum_{c \in C} G_c \left(\frac{1}{N} \cdot (\bar{z}^{(n)} + z^{(N-n)}(u(x), \mathbf{p}^{(N-n)})) \right) \right] \\ & - \text{Var} \left[\sum_{\ell \in L} G_\ell(z^{n+1}(x, \mathbf{s})) + \sum_{c \in C} G_c \left(\frac{1}{N} \cdot (\bar{z}^{(n)} + z^{(N-n)}(x, \mathbf{p}^{(N-n)})) \right) \right]. \end{aligned}$$

To estimate m with (4.4), we need to know the values of L and σ_{\max}^2 . The main difficulty in finding these values is that the distribution of the patient's position does not directly translate to the distribution of the delivered dose. Because the calculation of the dose deposition matrix involves applying a model of the physical interaction of the radiation particle with the patient's body, it cannot be effectively accounted for in calculating σ_{\max}^2 . Moreover, as the authors mentioned in *Ruszczynski and Shapiro (2003)*, the bound resulting from (4.4) is typically too conservative to be used practically.

For these reasons, rather than searching for value m a-priori, we use an alternative method that sets m dynamically.

4.4 Monte-Carlo bounding techniques

Before discussing how we can choose the value of m in problem (SAA), we need to examine in detail how we can find the range in which the optimal objective value of (P) resides. This is done by finding both an upper and lower bound for g^* . We first state the analysis for a general problem:

$$(\text{SP}) \quad h^* = \min_{y \in Y} E[f(y, \boldsymbol{\xi})], \quad \text{with } y^* = \arg \min_{y \in Y} E[f(y, \boldsymbol{\xi})],$$

where y is the vector of decision variables and Y is a compact feasible region. It is reasonable to assume that $E[f(y, \boldsymbol{\xi})] < \infty$ for every $y \in Y$, or we can shrink the region Y to only include points that satisfy this requirement. Moreover, $\boldsymbol{\xi}$ is the vector of random parameters, and assume that $f(y, \xi)$ is a convex function for any given ξ .

Assuming that function $E[f(y, \boldsymbol{\xi})]$ is too complex or large to be computed or optimized directly, we take advantage of the SAA approach. Taking m i.i.d. samples ξ^i from the distribution of $\boldsymbol{\xi}$, we construct the following approximating problem:

$$(\text{SP}_m) \quad h_m = \min_{y \in Y} \frac{1}{m} \sum_{i=1}^m f(y, \xi^i), \quad \text{with } y_m = \arg \min_{y \in Y} \frac{1}{m} \sum_{i=1}^m f(y, \xi^i).$$

The Monte-Carlo bounding method focuses on verifying the quality of the solution to $(\text{SP})_m$ by establishing a confidence interval on h^* .

4.4.1 Lower bound

The lower bound on h^* can be found from the following theorem (*Madansky (1960); Mak et al. (1999)*):

$$E[h_m] = E \left[\min_{y \in Y} \frac{1}{m} \sum_{i=1}^m f(y, \xi^i) \right] \leq h^*.$$

Here we paraphrase the proof as follows: first note that

$$h^* = \min_{y \in Y} E[f(y, \boldsymbol{\xi})] = \min_{y \in Y} E \left[\frac{1}{m} \sum_{i=1}^m f(y, \boldsymbol{\xi}^i) \right].$$

Let $y_m(\xi^1, \dots, \xi^m)$ be the minimizer of $\frac{1}{m} \sum_{i=1}^m f(y, \xi^i)$, it immediately follows that for every set of m random samples ξ^i ($i = 1, \dots, m$), we have:

$$\frac{1}{m} \sum_{i=1}^m f(y^*, \xi^i) \geq \frac{1}{m} \sum_{i=1}^m f(y_m(\xi^1, \dots, \xi^m), \xi^i).$$

As a result:

$$E \left[\frac{1}{m} \sum_{i=1}^m f(y^*, \xi^i) \right] = h^* \geq E \left[\frac{1}{m} \sum_{i=1}^m f(y_m(\xi^1, \dots, \xi^m), \xi^i) \right] = E[h_m].$$

This leads to the lower bound on h^* .

For the same reasons that we cannot solve problem (SP), in practice the value $E(h_m)$ is found by approximation using a batch-mean method:

$$E[h_m] \approx L(M) \doteq \frac{1}{M} \sum_{i=1}^M h_m^i.$$

Here we solve a total of M (SP _{m}) problems, each of which is formulated with independent sample groups $\xi^{i1}, \xi^{i2}, \dots, \xi^{iM}$, $i = 1, \dots, m$, and h_m^i is the optimal objective function value of the i -th such problem, obtained at y_m^i . Then the average of these objective values can be used as an estimator of the lower bound $E[h_m]$.

4.4.2 Upper bound

Because y^* is the optimal solution to problem (SP), we can readily see that:

$$E[f(\hat{y}, \xi)] \geq E[f(y^*, \xi)] = h^* \text{ for any } \hat{y} \in Y.$$

Therefore we can find an upper bound on h^* by finding $E[f(\hat{y}, \xi)]$ for some \hat{y} . Of course we want the resulting upper bound $E[f(\hat{y}, \xi)]$ to be as close to g^* as possible. This is done by selecting \hat{y} from the set of points that are already promising in coming close to y^* . Recall that the lower bound on h^* is found by solving M SAA problems. The solutions y_m^i are our best attempt at finding the approximate optimal solution to problem (SP), and we use these solutions here to calculate the upper bound on h^* .

Again, in practice the expected value $E[f(\hat{y}, \xi)]$ can be estimated by taking a large

number of samples, and using the sample mean estimator:

$$E[f(\hat{y}, \boldsymbol{\xi})] \approx U(K) \doteq \frac{1}{K} \sum_{i=1}^K f(\hat{y}, \xi^i)$$

where K is a large number.

4.4.3 Calculating the confidence intervals

In the above analysis, both the upper and lower bounds are calculated using samples. This can help us find a confidence interval on the value h^* . As long as we use sufficiently large number of samples M and K , it is appropriate to assume that Central Limit Theorem applies. We have:

$$\begin{aligned} \sqrt{M}(L(M) - E[h_m]) &\sim N(0, \sigma_m^2), \text{ and} \\ \sqrt{K}(U(K) - E[f(\hat{y}, \boldsymbol{\xi})]) &\sim N(0, \bar{\sigma}^2(\hat{y})) \end{aligned}$$

where $\sigma_m^2 = \text{var}(h_m)$, and $\sigma^2(\hat{y}) = \text{var}(f(\hat{y}, \boldsymbol{\xi}))$. These variance values can be approximated by the sample variances:

$$\begin{aligned} \sigma_m^2 &\approx \bar{\sigma}_m^2 \doteq \frac{1}{M-1} \sum_{j=1}^M (h_m^j - L(M))^2, \\ \text{and } \sigma^2(\hat{y}) &\approx \bar{\sigma}^2(\hat{y}) \doteq \frac{1}{K-1} \sum_{k=1}^K (f(\hat{y}, \xi^k) - U(K))^2. \end{aligned}$$

Now let α be the confidence level we want to achieve, and let T_n be a random variable that has a t-distribution with degree of freedom $n - 1$, and let $t_{n-1, \alpha}$ satisfy:

$$P\{T_n \leq t_{n-1, \alpha}\} = 1 - \alpha.$$

We can then define:

$$\epsilon_u = \frac{t_{K-1, \frac{\alpha}{2}}}{\sqrt{K}} \bar{\sigma}(\hat{z}), \text{ and } \epsilon_\ell = \frac{t_{M-1, \frac{\alpha}{2}}}{\sqrt{M}} \bar{\sigma}_m.$$

Assuming that the samples in estimating the upper and lower bounds are independent of each other, we have:

$$P\{L(M) - \epsilon_\ell \leq E[h_m] \leq h^* \leq E[f(\hat{y}, \boldsymbol{\xi})] \leq U(K) + \epsilon_u\}$$

$$\begin{aligned}
&= 1 - P\{L(M) - \epsilon_\ell \geq E[h_m], \text{ or } E[f(\hat{y}, \xi)] \geq U(K) + \epsilon_u\} \\
&\geq 1 - (P\{L(M) - \epsilon_\ell \geq E[h_m]\} + P\{E[f(\hat{y}, \xi)] \geq U(K) + \epsilon_u\}) \\
&= 1 - (1 - P\{L(M) - \epsilon_\ell \leq E[h_m]\}) - (1 - P\{E[f(\hat{y}, \xi)] \leq U(K) + \epsilon_u\}) \\
&\approx 1 - \frac{\alpha}{2} - \frac{\alpha}{2} \\
&= 1 - \alpha.
\end{aligned}$$

Therefore we can use

$$[L(M) - \epsilon_\ell, U(K) + \epsilon_u] \quad (4.5)$$

as the confidence interval for h^* with confidence level α .

4.4.4 Applying the bounding technique to problem (P)

Now we can apply the Monte-Carlo bounding technique to problem (P), and use (SAA) to find the approximate solution to (P). First, define the penalty corresponding to a single path $p_i^{(N-n)}$ as:

$$\mathcal{G}(x, p_i^{(N-n)}) = \sum_{\ell \in L} G_\ell(z^{n+1}(x, s_i)) + \sum_{c \in C} G_c\left(\frac{1}{N} \cdot (\bar{z}^{(n)} + z^{(N-n)}(x, p_i^{(N-n)}))\right) \quad i = 1, \dots, m.$$

Then the objective function of (SAA) can be simplified as:

$$g_m = \min_x \sum_{i=1}^m \frac{1}{m} \mathcal{G}(x, p_i^{(N-n)}).$$

We can see that the results in Section 4.4 can be applied directly here:

$$E[g_m] \leq g^* \leq E[\mathcal{G}(\hat{x}, \mathbf{p}^{(N-n)})], \text{ for any } \hat{x} \in X. \quad (4.6)$$

This gives rise to the following procedure for determining the sample size and solving problem (P) with (SAA).

Dynamic sampling procedure

1. Determine initial sample size m , batch size M and sample size K . These can be picked arbitrarily as the starting point. Set the confidence level α , and determine the stopping

criteria for this procedure based on the relative size of the confidence interval.

2. Randomly sample M batches of m sample paths $p_{1,j}^{(N-n)}, p_{2,j}^{(N-n)}, \dots, p_{m,j}^{(N-n)}, j = 1, \dots, M$.

Formulate an (SAA) problem with every batch of sample paths, solve a total of M (SAA) problems and obtain:

$$g_m^j = \min_{x \in X} \frac{1}{m} \sum_{i=1}^m \mathcal{G}(x, p_{i,j}^{(N-n)}) \quad \text{for } j = 1, \dots, M$$

$$x_m^j = \arg \min_{x \in X} \frac{1}{m} \sum_{i=1}^m \mathcal{G}(x, p_{i,j}^{(N-n)}) \quad \text{for } j = 1, \dots, M.$$

By averaging the resulting objective values from all of the (SAA) problems, obtain the lower bound on g^* :

$$L(M) = \frac{1}{M} \sum_{j=1}^M g_m^j.$$

3. Obtain K independent samples $p_1^{(N-n)}, \dots, p_K^{(N-n)}$, and calculate for all candidate solutions x_m^j ($j = 1, \dots, M$) found in Step 2:

$$u_j = \frac{1}{K} \sum_{k=1}^K \mathcal{G}(x_m^j, p_k^{(N-n)}), \quad \forall j = 1, \dots, M.$$

Take the best, i.e., the smallest of these, to be our upper bound:

$$U(K) = \min_{j=1, \dots, M} u_j.$$

Moreover, define the corresponding solution as:

$$\hat{x} \doteq x_m^{\hat{j}} \text{ such that } u_{\hat{j}} = U(K), \quad \hat{j} \in \{1, \dots, M\}.$$

4. Calculate:

$$\bar{\sigma}_m^2 = \frac{1}{M-1} \cdot \sum_{j=1}^M (g_m^j - L(M))^2,$$

$$\text{and } \bar{\sigma}^2(\hat{x}) = \frac{1}{K-1} \cdot \sum_{k=1}^K (\mathcal{G}(\hat{x}, p_k^{(N-n)}) - U(K))^2,$$

and calculate:

$$\epsilon_\ell = \frac{t_{M-1, \frac{\alpha}{2}}}{\sqrt{M}} \bar{\sigma}_m, \quad \epsilon_u = \frac{t_{K-1, \frac{\alpha}{2}}}{\sqrt{K}} \bar{\sigma}(\hat{x}).$$

From (4.5) we know that the $1 - \alpha$ confidence interval on g^* is:

$$[L(M) - \epsilon_\ell, U(K) + \epsilon_u].$$

5. If the confidence interval obtained in Step 4 satisfies the stopping criteria, terminate; otherwise increase M or m or K , and go to Step 2.

When the algorithm terminates, we can use the solution \hat{x} as the approximate solution to problem (P). We show in Appendix C that when the numbers m , M , and $K \rightarrow \infty$, g_m as well as both $L(M), U(K)$ will converge to g^* . Therefore this procedure will eventually converge for any confidence level α and any arbitrarily tight stopping criteria. Note that, as we show in the numerical experiments below, the actual number of samples m , M and K do not have to be very large in order for us to find a satisfactory solution.

4.5 Comparison with a model in the literature

Men et al. (2011) proposed a similar model in dealing with inter-fraction uncertainties in the plan optimization for the entire treatment. When the number of fractions in the treatment is large enough to justify the law of large numbers (LLN) assumption, the authors replaced the path-based expected penalty term with penalty of the expected dose:

$$\begin{aligned} & \min_x \sum_{\ell \in L} E [G_\ell(z^1(x, \mathbf{s}))] + \sum_{c \in C} E \left[G_c \left(\frac{1}{N} \cdot z^{(N)}(x, \mathbf{p}^{(N)}) \right) \right] \\ \approx \text{(CH)} & \min_x \sum_{\ell \in L} E [G_\ell(z^1(x, \mathbf{s}))] + \sum_{c \in C} G_c \left(E \left[\frac{1}{N} \cdot z^{(N)}(x, \mathbf{p}^{(N)}) \right] \right), \end{aligned}$$

while imposing the same constraints (4.1) (4.2) and (4.3). The assumption behind the approximation is the convergence of the average dose to the expected dose. In solving this model, the two terms in model (CH) were further approximated using a large number m' scenarios:

$$\min_x \frac{1}{m'} \sum_{i=1}^{m'} \sum_{\ell \in L} [G_\ell(z^1(x, s_i))] + \sum_{c \in C} G_c \left(\frac{1}{m'} \sum_{i=1}^{m'} z^1(x, s_i) \right),$$

i.e., the expected penalty associated with fractional dose is approximated with the average of penalty in m' scenarios, and the expected dose is approximated with the average dose of m' scenarios.

This model is promising and easy to use in that it only requires solving one single optimization problem, and the solution can be used as the treatment plan. However, it could not provide a measure of the quality the approximating solution and objective compared to the original problem. Moreover, the LLN assumption relies on a large number of fractions in the path \mathbf{p} , which means it might not work as well in designing adaptive treatments where the path may only contain a few fractions. We want to compare our SAA framework with approach (CH), and explore two things in particular:

1. is model (CH) good enough in approximating the true objective g^* and obtaining a good solution?
2. is the above still true for $n > 0$, i.e., in the re-optimization framework?

To answer question 1, note that since G_c is a convex function, applying Jensen's inequality shows that:

$$\sum_{c \in C} G_c \left(E \left[\frac{1}{N} \cdot z^{(N)}(x, \mathbf{p}^{(N)}) \right] \right) \leq \sum_{c \in C} E \left[G_c \left(\frac{1}{N} \cdot z^{(N)}(x, \mathbf{p}^{(N)}) \right) \right].$$

In other words, the objective in (CH) is an under-estimator of the true objective at any solution. However, we are most interested in the quality of the objective at its optimal solution \bar{x} . We can evaluate it by performing the paired t-test. The paired t-test compares the mean of two random variables, which in our case are functions $\mathcal{G}(\hat{x}, \mathbf{p}^{(N)})$ and $\mathcal{G}(\bar{x}, \mathbf{p}^{(N)})$. Here \hat{x} is the terminal solution provided by the dynamic sampling procedure. The null hypothesis for the test is that the mean difference between the paired observations is zero:

$$E [\mathcal{G}(\hat{x}, \mathbf{p}^{(N)})] = E [\mathcal{G}(\bar{x}, \mathbf{p}^{(N)})].$$

To carry out the test, first we draw a large number, e.g., K' , of sample paths from the

distribution of $\mathbf{p}^{(N)}$, and calculate:

$$d_k = \mathcal{G}(\hat{x}, \mathbf{p}_k^{(N)}) - \mathcal{G}(\bar{x}, \mathbf{p}_k^{(N)}) \quad \forall k = 1, \dots, K'$$

$$\bar{d} = \frac{1}{K'} \sum_{k=1}^{K'} d_k, \text{ and}$$

$$s_d = \sqrt{\frac{1}{K' - 1} \sum_{k=1}^{K'} (d_k - \bar{d})^2}.$$

Then we can find the t-statistic

$$T = \frac{\bar{d}}{s_d / \sqrt{K'}}.$$

Under the null hypothesis, T follows the t-distribution with degree of freedom $K' - 1$. Given confidence level α , we can look up the t-distribution table and find the corresponding p -value for T . If the p -value is higher than our confidence level α , we then accept the null hypothesis, and say that the two solutions \hat{x} and \bar{x} produce the same objective function value; otherwise we reject the null hypothesis. The t-test can also provide the $1 - \alpha$ confidence interval for the difference:

$$\left[\bar{d} - t_{K'-1, \frac{\alpha}{2}} \frac{s_d}{\sqrt{K'}}, \bar{d} + t_{K'-1, \frac{\alpha}{2}} \frac{s_d}{\sqrt{K'}} \right].$$

The authors did not apply this model in the adaptive setting, because the approximation:

$$\sum_{c \in C} E \left[G_c \left(\frac{1}{N} \cdot z^{(N)}(x, \mathbf{p}^{(N)}) \right) \right] \approx \sum_{c \in C} G_c \left(E \left[\frac{1}{N} \cdot z^{(N)}(x, \mathbf{p}^{(N)}) \right] \right).$$

depends on the (LLN) assumption: the average of a large number N independent and identically distributed (i.i.d.) random variables is close to the expected value. When the number of fractions to be delivered becomes smaller, this assumption may not apply anymore.

However, we can still extend model (CH) to the adaptive optimization setting:

$$\min_x \sum_{\ell \in L} E [G_\ell(z^\ell(x, \mathbf{s}))] + \sum_{c \in C} E \left[G_c \left(\frac{\bar{z}^{(n)} + z^{(N-n)}(x, \mathbf{p}^{(N-n)})}{N} \right) \right]$$

$$\approx \min_x \sum_{\ell \in L} E [G_\ell(z^\ell(x, \mathbf{s}))] + \sum_{c \in C} G_c \left(\frac{z^{(n)}}{N} + \frac{(N-n)}{N} E [z^{(N)}(x, \mathbf{p}^{(N)})] \right)$$

We test the performance of this model in solving the adaptive optimization problems,

and compare its solutions to the (SAA) solutions in Section 4.6.3.

4.6 Experiments

We compare our algorithms on four prostate patient cases, exported from an imaging and treatment planning system at Department of Radiation Oncology at the University of Florida. All patients are treated with 9 equispaced ^{60}Co beams at angles 0° , 40° , 80° , 120° , 160° , 200° , 240° , 280° and 320° respectively. The size of the beam is $40 \times 40 \text{ cm}^2$, and the beamlets are of size $1 \times 1 \text{ cm}^2$. In discretizing the patients' geometry, we used a grid consisting of voxels of size $4 \times 4 \times 4 \text{ mm}^3$. In the optimization model we used a downsampled grid of size $8 \times 8 \times 8 \text{ mm}^3$ in unspecified tissues. However, in evaluating the final dose we always use the full resolution grid. All of our test cases have one target (prostate) with a prescription dose of 73.8 Gy to be delivered over 41 daily fractions. The critical structures include rectum, bladder and femoral heads in addition to normal tissue.

The patient setup location is assumed to be jointly normally distributed with mean of 0 and standard deviation of 3mm along each of the three axes (see *Ekberg et al. (1998)*; *Tinger et al. (1998)*; *Van Herk et al. (2004)*). Assuming that large errors will be detected and corrected prior to the treatment, we restrict the error along each axis to be within $\pm 6\text{mm}$. Note, however, that our model and solution procedure does not require the normal distribution assumption, and they work in principle with any distribution appropriate in modeling the specific motion type. For each case, we randomly sample 500 setup scenarios by perturbing the patient's position according to this distribution, and use this collection of scenarios as an approximation of the continuous space of random setup locations. All of the sample scenarios and paths used in the optimization problems are then drawn from this pool of 500 samples. The sizes of relevant parameters are listed in Figure 4.1. The number of non-zero elements in the dose deposition matrix is not a constant any more, as shifts in the patient position relative to the beams will affect the number of voxels irradiated by the same beamlet, and thus affecting the total number of non-zero entries in the dose deposition

matrix.

Case	# voxels (total / down-sampled)	# beamlets	# non-zero elements in $D(\mathbf{s})$
1	226,324 / 38,731	2,318	2,908,444 – 3,103,615
2	205,911 / 35,988	2,336	3,320,876 – 3,432,446
3	131,216 / 22,330	2,002	2,074,564 – 2,124,810
4	162,628 / 29,828	2,413	3,232,972 – 3,273,917

Table 4.1: Problem dimensions of our test cases.

The penalty function used in all of our experiments are one-sided quadratic functions that are based on voxel dose, having the same structure as the ones used in *Men et al.* (2011). The penalty functions are as follows:

$$G_\ell(z) = w_t^\ell \cdot \frac{1}{|v_t|} \sum_{j \in v_t} \max\{T_t - z_j, 0\}^2 \quad \forall t \in T, \ell \in L$$

$$G_c(z) = w_t^c \cdot \frac{1}{|v_t|} \sum_{j \in v_t} \max\{z_j - T_t, 0\}^2 \quad \forall t \in T \cup O, c \in C,$$

where T and O are the set of targets and critical structures, respectively, v_t is the collection of voxels in structure t , and T_t is the threshold value for the dose delivered to a voxel in structure t . Moreover, w_t^ℓ and w_t^c are the weights assigned to structure t in function G_ℓ and G_c ($\ell \in L, c \in C$), respectively. Our experiments are performed on a Mac Pro with 2×2.8 GHz Quad-core Xeon CPUs and 14 GB of memory. All of the optimization problems are solved with an in-house interior point method based solver (see *Wright* (1987); *Aleman et al.* (2010) for a description of the interior-point method), implemented in Matlab version 7.14.

The DVH constraints and criteria for evaluating the treatment plans are established based on RTOG protocols (*Radiation Therapy Oncology Group* (2004)). Those defined for the critical structures are summarized in Table 4.2. In establishing the constraints for the target, for the margin-based model we followed the traditional approach and required that 95% of the PTV voxels receive the prescription dose, $\geq 99\%$ receive 93% of the prescription dose, and $\leq 10\%$ receive more than 110% of the prescription dose. For the stochastic model,

since we removed the margin and only took the ITV into account, we want to cover the entire ITV with prescription dose under all circumstances; in addition, we require that none of the ITV voxels receive more than 110% of the prescription dose.

Structure	Rectum				Bladder		Femoral heads		
Threshold dose (Gy)	75	70	65	40	65	40	50	45	40
Upper limit on volume (%)	15	25	35	45	17	35	10	25	40

Table 4.2: DVH criteria for critical structures.

One of the major challenges in comparing different plan optimization models is that the process of adjusting the parameters to meet the DVH constraints is highly subjective and model dependent. The margin built around the ITV in the conventional models leads to a larger target volume, and usually also introduces overlaps with the surrounding critical structures. These differences in the structure sizes, combined with different dose constraints for the target, necessitate tuning the parameters for individual models. This is done to best meet the DVH requirements while taking into account the trade-off between structures and clinical considerations.

After the tuning is done, the same models are applied to all patients to eliminate further variations that may be introduced by additional tuning for individual patient cases. In order to facilitate the comparison across different cases, the treatment plans obtained from the margin-based model are scaled so that 95% of the total PTV volume receive the prescription dose under the nominal scenario.

After the plans from the margin-based model and the stochastic model have been generated, we apply them to the same set of 5 simulated treatment paths, each consisting of 41 fractions. For the following comparisons, we only inspect the target dose distributions based on ITV dose, as the ITV is the region we are interested in covering with prescription dose.

4.6.1 Comparing the stochastic model with the conventional model

We use our first experiment to show how directly taking randomness into account affects the treatment quality in a random setting. We first set $n = 0$ and $\bar{z}^{(n)} = 0$ and find an approximate solution to problem (P) using the procedure described in Section 4.4.4. This is equivalent to designing an initial plan at the beginning of the treatment, without performing any adaptive re-planning. In Table 4.3 we summarize the sizes of different samples used in the solution process, and the runtime for one (SAA) problem as well as the bound obtained on the objective g^* for both cases. For the same case, the runtime for the (SAA) problems are not identical in different runs, however they are very similar to each other because of the similar problem sizes, and therefore we report only one representative statistic. Observe that the number of samples m that we need in the (SAA) problem is rather small, showing that (SAA) is very efficient in approximating problem (P) and in finding high quality solutions. However, doing this requires solving the (SAA) problem M times, which may lead to long treatment planning time.

Case	sample size m	batch size M	sample size K	bound on g^*	(SAA) runtime
1	30	20	100	2.1%	235 s
2	30	25	100	0.9%	307 s
3	30	30	100	5.4%	181 s
4	30	25	100	5.5%	319 s

Table 4.3: Sample sizes and bounds in the initial (SAA) problem.

The solutions are then applied to the 5 sample treatments, and the resulting DVH is compared with DVH obtained from applying the margin-based solutions. We combine the DVH curves for each patient of 5 sample treatments, and the result is a “cloud” of curves, presented in Figure 4.1 for our test cases.

It can be observed that the treatments benefit from taking advantage of the stochastic model. The stochastic plans are in general providing similar or better target coverage without introducing additional hotspots, and can reduce the dose delivered to especially the rectum

and the bladder which in general receive high doses. Although the femoral heads and skin receive higher doses at the very low dose ranges, these structures under the stochastic plans can still easily meet the same set of criteria compared to the conventional plans.

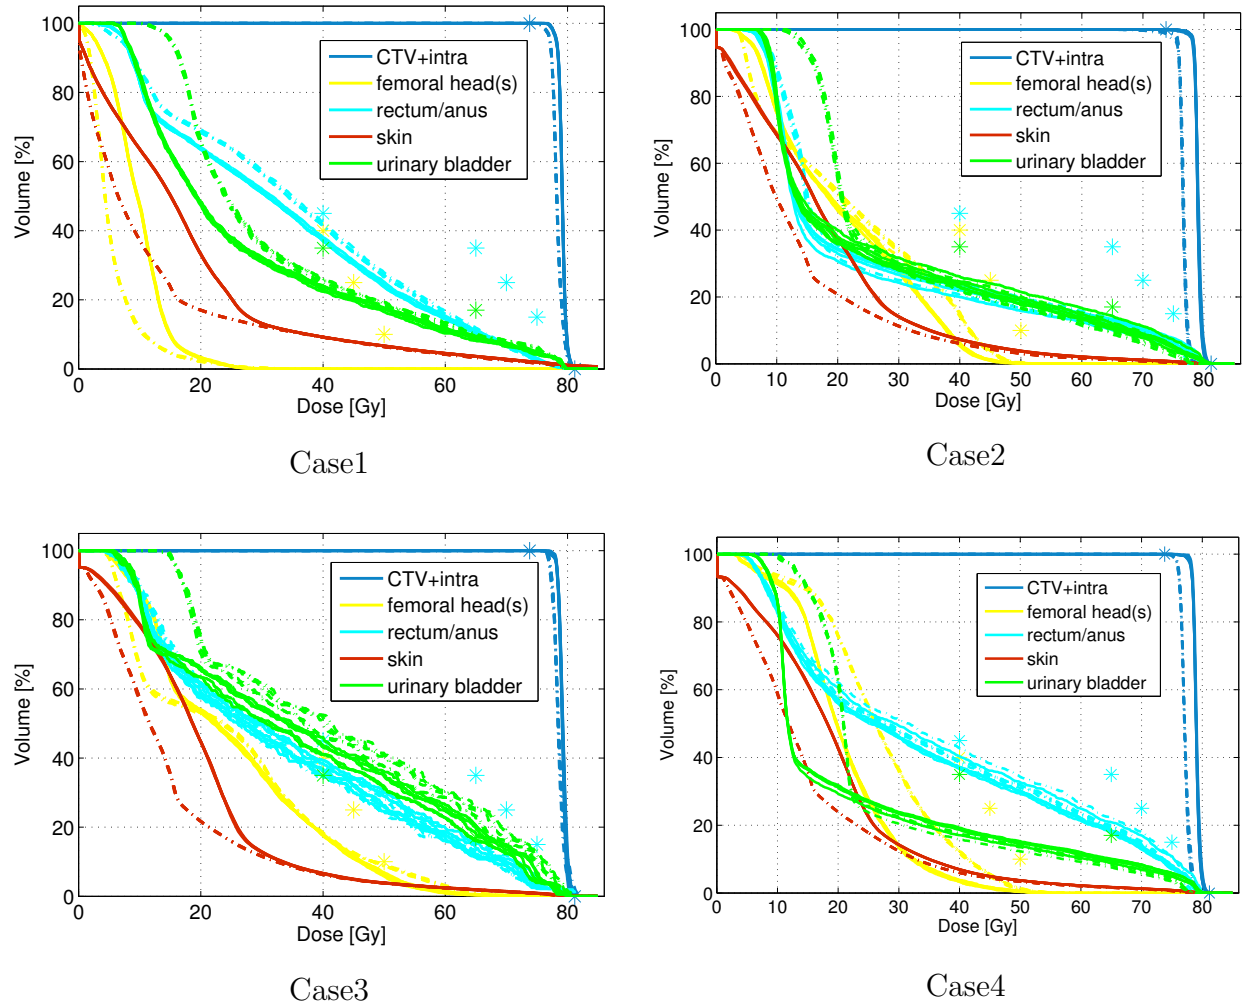


Figure 4.1: DVH clouds for the stochastic plans (solid) and the conventional plans (dashed) for 5 sample treatments

Another criterion important to the clinicians is the probability of delivering the prescribed dose to the ITV in each fraction. As mentioned before, ensuring the target receives the prescription dose not only over the entire treatment, but also in each fraction is critical in a successful treatment. For both the conventional plan and the stochastic plan, we use 100 randomly sampled scenarios to evaluate the probability of a particular fraction of tar-

get voxels receiving the prescription dose. The results are shown in Figure 4.2. In general the stochastic model is able to achieve higher target coverage compared to the conventional model. For Case 3, more than 24% of the bladder and more than 16% of the rectum overlap with the PTV. By ignoring the normal tissue in the overlapped regions, the conventional model compromises the rectum and bladder to provide better target target coverage. However, if we adjust the stochastic model to put more emphasis on the per-fraction effects, even for Case 3 a higher coverage probability can be achieved (Figure 4.3).

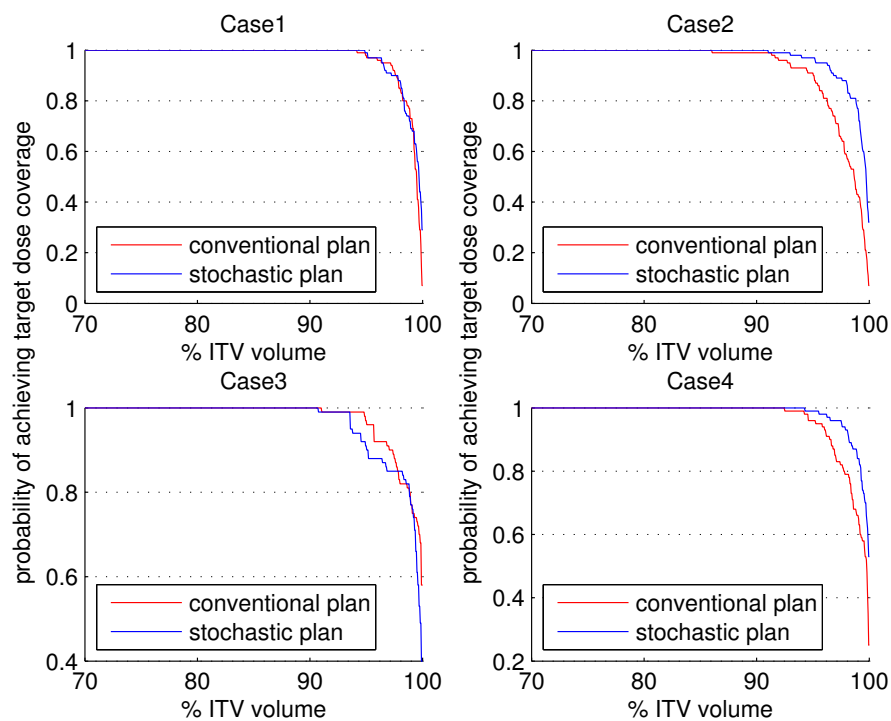


Figure 4.2: Probability of covering a % volume of the ITV in one fraction for conventional vs. stochastic plans

4.6.2 Re-optimization and adaptive radiation therapy

Being able to collect and exploit additional information about the delivered dose provides us the opportunity to explore the benefits of adaptive radiation therapy. Due to the randomness in the treatments the delivered dose distribution may differ from the planned one, and adaptive re-optimization makes it possible to react to and correct the deviations in the

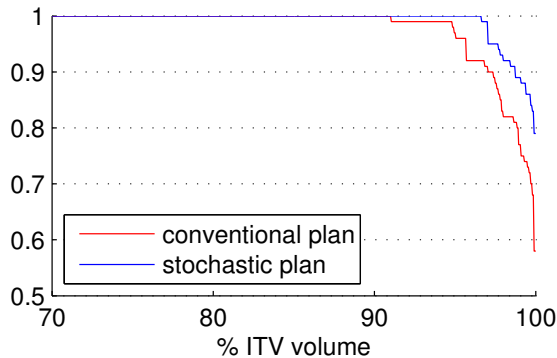


Figure 4.3: Coverage probability for conventional vs. adjusted stochastic plan, Case 3

delivered dose and achieve potential additional benefits compared to performing stochastic optimization alone. As mentioned in Section 4.2, our dynamic sampling framework also works for designing adaptive treatments when updated information about delivered dose becomes available during the treatment course.

We next investigate the effects of applying re-optimization during the simulated treatment. Recall that if $n > 0$, which means n fractions have been carried out, solving model (P) will be re-optimizing the treatment plan for the rest of the $N - n$ fractions, taking into account the delivered dose $\bar{z}^{(n)}$. Here we assign to $\bar{z}^{(n)}$ the dose delivered to the patient in the first n fractions of the simulated treatments, and re-optimize the plan given this information. We then apply the updated treatment plan to the remaining fractions in the treatments. The results are compared with the stochastic optimization plans on the same simulated treatment courses.

In order to compare different re-optimization schedules we tested two schemes: one where re-optimization is done once at the midpoint of the treatment, and another where we perform weekly re-optimization for every treatment course. We will focus on for Case 1, for which the trade-off between the target and the critical structures is more evident and the differences in the results can be more readily seen. Table 4.4 shows the sample numbers m , M and K used and the associated bounds in these experiments. It is interesting to notice that the runtime depends heavily on the number of samples m but not much on the number of fractions left

in the treatment course.

Scheme	fraction	sample path	m	M	K	bound on g^*	(SAA) runtime (s)
Reopt once	20	1	30	20	100	2.4%	297
		2	25	25		2.1%	240
		3	30	20		2.0%	293
		4	30	20		1.8%	290
		5	30	25		2.4%	280
Reopt weekly	7	1	30	25	100	5.1%	272
		2	25	20		5.2%	229
		3	25	20		3.8%	
		4	25	20		5.0%	
		5	25	25		5.3%	
	14	1	30	25	100	5.3%	274
		2	25	20		4.5%	226
		3	25	20		4.7%	
		4	25	20		5.8%	
		5	25	25		4.6%	
	21	1	30	25	100	4.2%	278
		2	25	20		1.8%	229
		3	25	20		2.5%	
		4	25	20		5.9%	
		5	25	25		4.8%	
	28	1	30	25	100	2.1%	320
		2	25	20		3.0%	261
		3	25	20		1.9%	
		4	25	20		1.1%	
		5	25	25		4.9%	
	35	1	30	25	100	1.1%	380
		2	25	20		1.2%	274
		3	25	20		1.3%	
		4	25	20		0.9%	
		5	25	25		1.9%	

Table 4.4: Sample sizes and bounds on the objective function in the adaptive reoptimization problems for Case 1.

Figure 4.4 shows the DVH cloud for both cases. Notice that the rectum and the bladder both receive higher dose under the adaptive optimization schemes. However, from Figure 4.5 we can clearly see that the later the reoptimization takes place, the more the target coverage probability for the reoptimized plans improves. Because the per-fraction and cumulative dose

criteria are conflicting, this results in more dose being delivered to the critical structures close to the ITV.

We offer a qualitative explanation of why performing reoptimization gives more dose to the target region. Recall that the objective function in the adaptive optimization model takes the following form:

$$\frac{1}{m} \sum_{i=1}^m \left[\sum_{\ell \in L} G_{\ell}(z^{n+1}(x, s_i)) + \sum_{c \in C} G_c \left(\frac{1}{N} \cdot (\bar{z}^{(n)} + z^{(N-n)}(x, p_i^{(N-n)})) \right) \right].$$

For the cumulative term, the intensity vector x only affects future fractions $n + 1$ to N . Therefore as more fractions are delivered, increasing x has less effect on the cumulative dose and thus on the penalty term associated with it. This can be verified by examining the (SAA) solutions for reoptimization at fractions 8, 15, 22, 29 and 36, respectively, in one sample treatment for Case 1. For each adaptive optimization problem, we increase the optimal intensity by 10% and calculate the average increase in the cumulative dose for each structure over m different samples. The results in Table 4.5 confirm that x has much smaller effect on the cumulative dose as n grows. On the other hand, x affects the per-fraction penalty term the same way regardless of n , as the penalty depends on dose in fraction $n + 1$ only. As a result, re-optimization towards the end of the treatment course will increase the dose in the target region and thus the dose to the surrounding structures.

Fraction number	Increase in cumulative dose (%)				
	normal tissue	femoral heads	rectum	bladder	ITV
8	8.3	8.3	8.3	8.2	8.3
15	6.6	6.6	6.6	6.7	6.6
22	4.8	4.8	5.0	4.9	4.9
29	3.1	3.0	3.3	3.5	3.2
36	1.4	1.3	1.6	1.6	1.5

Table 4.5: Effect on cumulative dose of increasing the intensity by 10% at different stages of the treatment course.

It can be concluded that under the current framework, adaptive optimization puts more

emphasis on target coverage than critical structure sparing, especially when the number of remaining fractions is small. This is an important factor to consider when applying adaptive optimization with this model.

4.6.3 Comparison with model (CH)

In this section, we compare our model with the one proposed in (*Men et al. (2011)*). As discussed in Section 4.5, model (CH) is easier to implement, and does not require solving multiple optimization problems. However, it does not provide any measure of the solution quality. On the same set of cases, we compare the solution quality of the plans obtained from our model and model (CH).

4.6.3.1 Initial optimization problem

We follow the authors' method in implementing model (CH), and use $m' = 100$ samples in estimating the expected penalty and the expected dose in model (CH), and apply the same set of parameters in defining the penalty terms. Table 4.6 includes results of the objective value comparison for our solution and the solution from (CH) for the initial optimization problem. For all cases except Case 3, the p -values are much smaller than the critical value 0.05, and it is statistically significant enough to show that the true objective value at \hat{x} is lower than that at \bar{x} . For Case 3, the p -value is also very close to 0.05. However, the difference in objective values are only around a few percent, which can be expected as (CH) is designed to be a good approximation to (P) when N is large. The solution time for model (CH) is much shorter compared to (SAA). In Figures 4.6 and 4.7 we compare the simulated dose and coverage probability to our plans. While the cumulative dose to the critical structures, most notably rectum and bladder, is higher for all cases under model (CH), the target coverage probabilities are in general better for the corresponding solutions.

Case	estimated obj. $U(K)$	p -value	95% c.i. for $\mathcal{G}(\hat{x}, \mathbf{p}^{(N)}) - \mathcal{G}(\bar{x}, \mathbf{p}^{(N)})$	(CH) runtime
1	3709.1	0.003	[-114.2, -24.1]	48
2	3132.6	0.094	[-189.0, 15.2]	48
3	4663.3	<0.1%	[-277.9, -107.0]	29
4	3596.1	<0.1%	[-210.0, -82.7]	45

Table 4.6: Comparison of (SAA) solution \hat{x} with solution from model (CH) for the initial stochastic optimization problem

4.6.3.2 Adaptive optimization problems

Finally, we perform the same tests to compare the expected objective function values for the adaptive optimization problems. In order to fairly compare the outcomes from different optimization models, the input dose used in the respective problem, in other words, $\bar{z}^{(n)}$, must be consistent. We therefore apply the initial and adaptive solutions from our model until the fraction the comparison takes place. Then model (CH) is solved and its solutions compared to our model. The results in Tables 4.7 and 4.8 show that for all but a few cases, our solutions are similar to or better than the solutions from model (CH). It is interesting to see that model (CH) performed very well in all cases, using much shorter optimization time than (SAA).

Fraction	sample	estimated objective $U(K)$	p -value	95% c.i	(CH) runtime (s)
20	1	3613.0	0.013	[-69.9, -8.3]	62
	2	3788.1	0.001	[26.0, 98.7]	58
	3	3811.8	0.008	[23.6, 151.7]	58
	4	3662.0	<0.1%	[-97.3, -39.1]	54
	5	3809.0	<0.1%	[59.8, 103.1]	60

Table 4.7: Comparison of (SAA) solutions with those from model (CH) for reoptimization problems under one-time reoptimization for Case 1

Intuitively, the LLN assumption that the average of i.i.d. random samples equals the expected value becomes less accurate the fewer the number of samples. However, we do not observe any deterioration in solution quality in the adaptive solutions from applying model (CH). This is because as n increases, a larger portion of the overall dose is delivered. Even

Fraction	sample	estimated objective $U(K)$	p -value	95% c.i	(CH) runtime (s)
8	1	3760.1	0.063	[-25.5, 0.70]	46
	2	3661.1	<0.1%	[-119.8, -87.0]	52
	3	3705.1	<0.1%	[-99.9, -75.4]	49
	4	3785.9	<0.1%	[-110.4, -101.9]	49
	5	3798.1	0.98	[-85.1, 87.7]	48
15	1	3741.8	0.018	[-29.8, -2.9]	49
	2	3735.9	<0.1%	[-85.0, -27.7]	54
	3	3691.7	<0.1%	[-88.3, -62.6]	51
	4	3708.1	<0.1%	[-86.1, -54.7]	50
	5	3805.1	0.22	[-101.1, 23.9]	53
22	1	3782.0	0.25	[-37.5, 9.9]	52
	2	3703.1	0.001	[-80.1, -23.3]	62
	3	3697.8	<0.1%	[-111.2, -42.3]	63
	4	3817.7	0.001	[-83.2, -21.1]	54
	5	3782.5	0.86	[-42.1, 50.0]	57
29	1	3786.7	0.57	[-70.2, 38.7]	67
	2	3716.3	<0.1%	[-82.3, -27.0]	49
	3	3757.6	0.612	[-18.1, 30.6]	61
	4	3825.9	0.08	[-65.0, 3.9]	70
	5	3754.2	0.468	[-53.2, 24.6]	65
36	1	3763.0	0.17	[-83.6, 14.7]	111
	2	3736.4	0.547	[-66.2, 35.3]	125
	3	3688.9	<0.1%	[-164.7, -73.7]	108
	4	3731.5	0.44	[-65.8, 28.9]	96
	5	3842.3	<0.1%	[91.4, 134.7]	109

Table 4.8: Comparison of (SAA) solutions with those from model (CH) for weekly reoptimization problems for Case 1

though the difference between the actual and expected dose over the remaining fractions becomes bigger, this difference is “diluted” by the delivered dose. We test this empirically with one sample path: at fractions 8, 15, 22, 29 and 36 respectively, calculate the expected dose over the remaining $N - n$ fractions based on the optimal solution obtained from model (CH). Then we randomly sample 10 paths of $N - n$ fractions, and calculate the actual dose over each path with the same solution from (CH). The absolute difference between the expected and actual doses for each path is calculated, and averaged over the 10 paths. Table 4.9 lists the relative difference for all five structures. The relative difference increases

with n for all structures, indicating that the LLN assumption is less and less applicable when n grows. However, comparing the expected and actual cumulative dose over the entire treatment in Table 4.10 reveals that the impact on the overall dose becomes much smaller.

Fraction number $n + 1$	Difference in expected and actual dose (%)				
	normal tissue	femoral heads	rectum	bladder	ITV
8	0.05	0.7	1.2	1.9	0.04
15	0.07	1.0	1.2	1.7	0.04
22	0.05	1.3	1.1	2.6	0.07
29	0.11	2.3	1.6	2.8	0.09
36	0.09	3.6	1.5	2.7	0.1

Table 4.9: Relative difference between actual and expected dose over the undelivered fractions at different stages of the treatment course

Fraction number $n + 1$	Difference in expected and actual dose (%)				
	normal tissue	femoral heads	rectum	bladder	ITV
8	0.04	0.6	1.0	1.6	0.03
15	0.04	0.6	0.8	1.1	0.02
22	0.02	0.6	0.6	1.3	0.03
29	0.03	0.7	0.5	0.9	0.03
36	0.01	0.5	0.2	0.4	0.02

Table 4.10: Relative difference between actual and expected dose over the entire treatment at different stages of the treatment course

Finally, we apply model (CH) in performing the adaptive optimization, and obtain the DVH clouds for the overall dose in the sample treatments, as well as the target coverage probabilities in Figures 4.8 and 4.9. Similar to the observation in the initial optimization, model (CH) tends to emphasize the per-fraction target coverage, and results in more dose to the critical organs in general.

4.7 Conclusions

We investigated a stochastic programming based model that minimizes the expected penalty based on both the per-fraction and cumulative doses. By solving a sample average approximation model we can greatly reduce the complexity required in solving the proposed model. Also, a statistical bound on the true objective function can be established by applying the dynamic sampling procedure. Experiments with real patient data show that our approach achieves improvements over the conventional, margin-based model in target coverage as well as overall dose distribution. Our model can easily be extended to adaptive treatment planning. However, in the adaptive optimization problems the solution lays more stress on the target coverage and therefore results in higher dose to the surrounding structures. This effect can be undesirable and needs to be factored in when using this model. Moreover, our model requires solving multiple optimization problems, which is time consuming without using any parallel computation. A comparable but simpler model proposed in the literature is shown to provide moderately deteriorated but still satisfactory results.

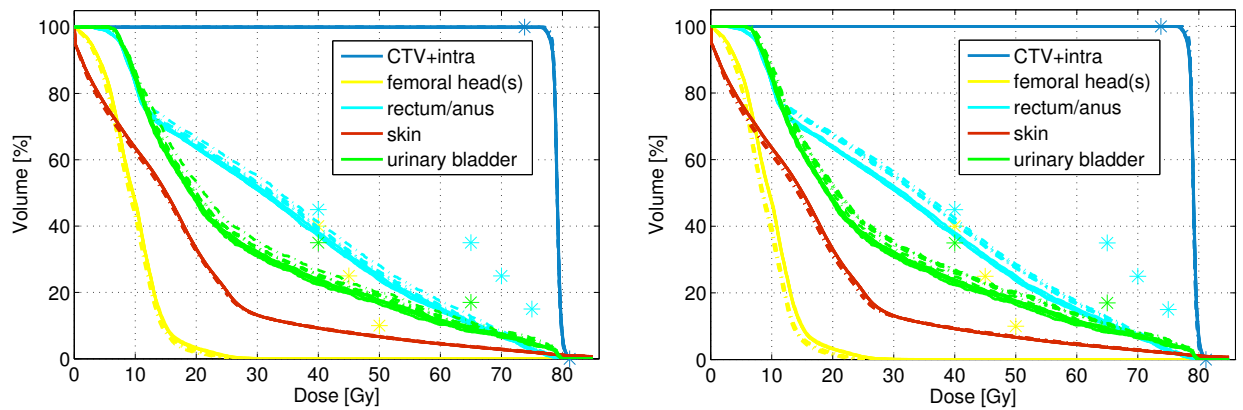


Figure 4.4: DVH clouds for the initial stochastic plans (solid) and adaptive plans (dashed) for Case 1 in five sample treatments. Left: reoptimizing once, right: weekly reoptimization




Figure 4.5: Coverage probability for initial stochastic and adaptive plans for Cases 1. (a): reoptimizing once at fraction 20; (b-1) — (b-5): weekly reoptimization at fractions 8, 15, 22, 29, and 36, respectively

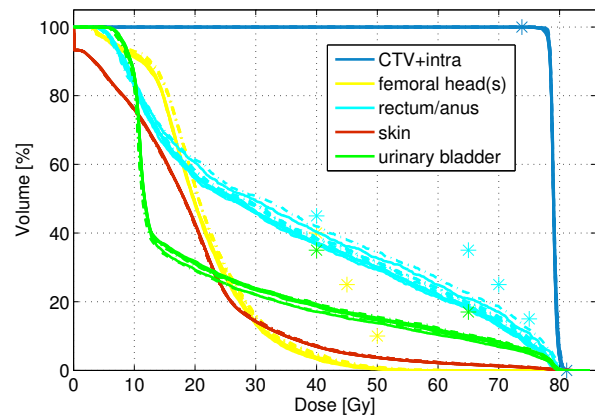
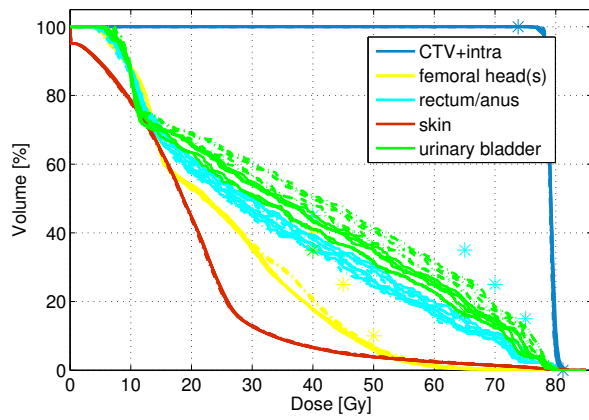




Figure 4.6: DVH clouds for the stochastic plans (solid) and the plans from model (CH) (dashed) without adaptive reoptimization for 5 sample treatments

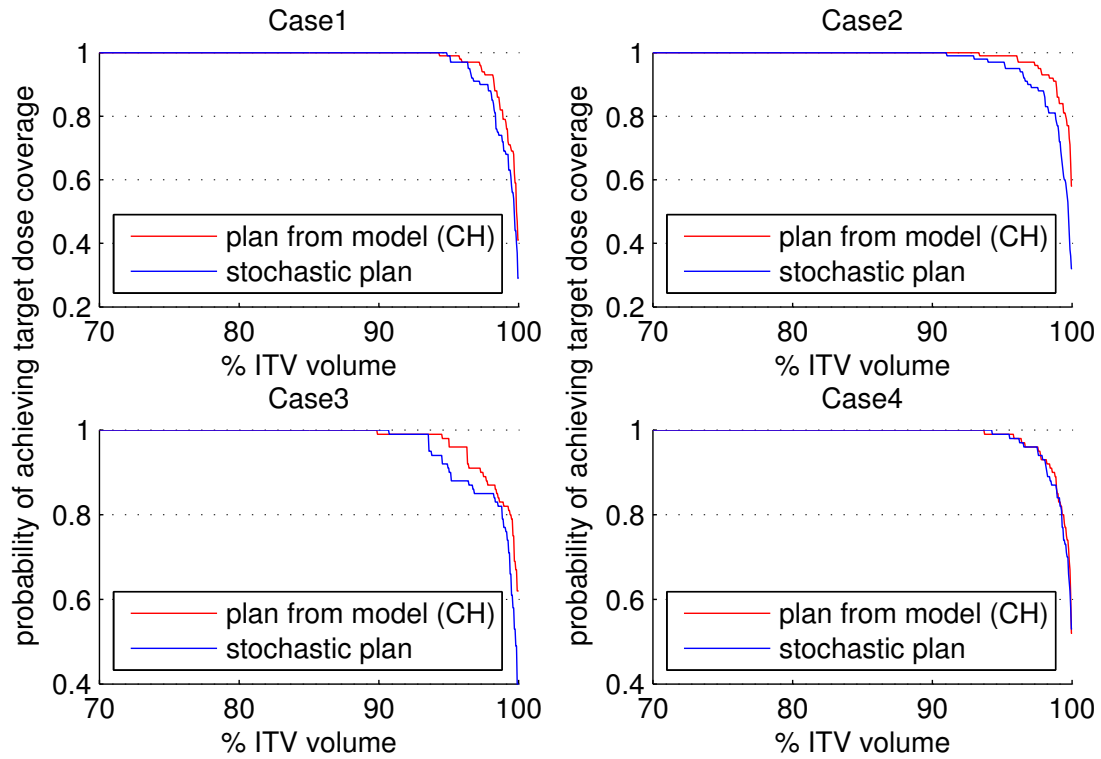


Figure 4.7: Probability of covering a certain volume of the ITV in one fraction for stochastic plans vs. plans obtained by solving model (CH)

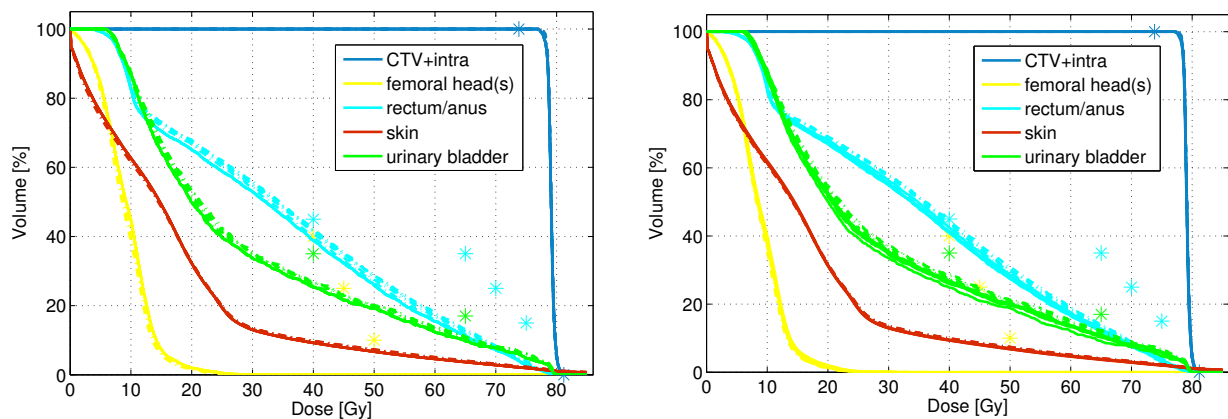


Figure 4.8: DVH clouds for the adaptive (SAA) plans (solid) and the adaptive plans from model (CH) (dashed) of 5 sample treatments for Case 1. Left: reoptimization once; right: weekly reoptimization




Figure 4.9: Probability of covering a certain volume of the ITV in one fraction for adaptive stochastic plans vs. adaptive plans from model (CH) for Case 1. (a): reoptimizing once at fraction 20; (b-1) — (b-5): weekly reoptimization at fractions 8, 15, 22, 29, and 36, respectively

CHAPTER V

Conclusions and Future Research

Radiation therapy is a complex yet exciting field that interacts with medicine, physics, operations research as well as many other subjects. While it has been an active field for many years, new treatment methods and technologies emerge, and bring with them problems that require unconventional optimization models and algorithms. In particular, the treatment plan optimization problem for VMAT is still a difficult one due to the large problem size and the restrictions due to the continuous gantry movement. While most optimization methods resort to heuristics, our research can serve as a step towards explicitly incorporating all machine constraints, and eventually an exact solution algorithm for VMAT.

On the other hand, the tradeoff between the large capital expenditure required in updating to a dedicated VMAT system and the potential benefits is an important factor to consider when clinics are transitioning to VMAT treatment systems. By developing a treatment planning method for VMATC, we offer an alternative treatment that can be implemented on the existing IMRT systems, and provides similar plan quality to VMAT treatments.

Accounting for uncertain events that happen between and during radiation therapy treatments is important to ensure the outcome of the actual treatment. Even though advances in imaging and patient monitoring technologies have made it possible to detect small changes in the patient setup process as well as their internal geometries, radiation therapy treatment can still benefit from directly incorporating uncertainty in the optimization.

Finally, many of the algorithms we developed for solving radiation therapy treatment planning problems can be applied or modified to solve other problems that involve large-scale, nonlinear optimization systems.

Despite our best effort, we can only address a limited number of issues in this thesis. There are still many interesting problems to be studied.

5.1 VMAT treatment plan optimization problem

For the VMAT treatment plan optimization problem, the greedy nature of the algorithm dictates that apertures added earlier in the column generation process may not be good choices later. A possible improvement is to study how ineffective apertures, in addition to those with zero intensities, can be detected during the solution process.

In theory, the optimal VMAT plans will always outperform VMATC plans due to the additional degrees of freedom. However, compared to VMATC, many of the VMAT plans actually performed worse in terms of objective value, as seen in Table 3.5. We can potentially find better solutions by introducing a refinement step similar to that used in VMATC optimization, which improves the aperture shapes and/or fluence rates during or after the column generation process, although this means that the optimization procedure will require more time to complete. Moreover, using the VMATC solution as the starting point for VMAT optimization can also be a promising direction to pursue.

Reducing the treatment time and total MU values are not considered explicitly in the column generation procedure. However, these are important factors when evaluating the plan quality. Although the post-processing step minimizes the treatment time, the final gantry speed and thus treatment time depend on the apertures and fluence rates that are already determined. We want to, in the future, incorporate constraints that impose upper bounds on the treatment time and total MU so as to explore the tradeoff between the objective value and these other quality measures.

Finally, we would like to be able to include transmission effects in the pricing problems,

so that we can more accurately evaluate the prices of potential apertures. However, this means that the pricing problem cannot be solved with the current method anymore, and we will need to look for other appropriate methods.

5.2 Treatment planning for VMATC

In the optimization of VMATC plans, our method for finding the apertures is very local in nature. In each iteration the selection is only based on changes in the leaf setting in one row. We would like to study algorithms that considers multiple leaf rows at a time so that we can search a larger area of the feasible region.

Moreover, in the current leaf setting optimization process, we have only studied one deterministic strategy for determining the sequence of control points and leaf rows. There are potentially alternative deterministic strategies, as well as strategies that involve randomization in this process. We would like to test these strategies, compare them to the current approach, and determine the best option to be used in the optimization process.

Furthermore, the optimal way of performing the refinement step, which theoretically provides better solutions than the heuristic approach used currently, is too slow to be practically feasible. Exploring parallel solvers to speed up the solution of the optimization problems can benefit the solution quality.

The prostate cancer cases we tested are relatively less complicated compared to cancer sites such as head and neck, which sometimes involves multiple targets with different target prescription dose levels. It may not be true anymore that one starting simplex works for different cases. In that case, finding a good treatment plan in small amount of tests may become a bigger challenge. It will be interesting to see how the current strategies perform for a different set of cases.

Finally, the experiments in Chapter III showed similar qualities for VMATC plans compared to the corresponding full VMAT plans. These results are highly algorithm dependent, and cannot be generalized to arrive at the same conclusions about these treatment modali-

ties, since we cannot solve either treatment planning problem to optimality. However, they do reveal that the current VMATC plans are of high quality. It will be interesting to see the same comparison if we restrict VMATC even further, and require that they use the same amount of treatment time and/or total MU compared to the VMAT plans.

5.3 IMRT plan optimization under uncertainty

The dynamic sampling procedure used in solving the stochastic optimization model is slow, as it needs the solution from multiple optimization problems to generate a bound on the true objective. We would like to explore other implementation options, such as parallel computing, to make it more practical in the clinical setting.

The current uncertainty scenarios only include rigid shifts in the patient setup position. It may be beneficial to also take geometry changes into account in the uncertainty models. This, however, will require a thorough understanding of the potential variations in the sizes and locations of the patients' internal structures. Studying the variations in the patient population with the same cancer type that underwent the same type of treatment can help us extract information in such modeling attempts. Moreover, we would like to study the number of samples used in approximating the random setup error. The number should be sufficiently large to be an appropriate approximation, and ideally be determined through a less arbitrarily process.

Finally, our experiments showed that under the adaptive optimization framework, the emphasis of the solutions shifts away from the cumulative dose penalty, and towards the per-fraction dose penalty as more fractions are delivered. It is important to be able to quantify this tradeoff, and investigate modifications to our model to take advantage of it. We want to also study alternative models that can benefit the treatment plan under the adaptive planning, without making compromises in the process.

APPENDICES

APPENDIX A

Relationship between Moving Speed of Radiation Field Edge and Moving Speed of MLC Leaf

The isocenter plane in radiation therapy is defined as a plane perpendicular to the ray in the center of a beam that crosses the center of the target. The distance from the radiation source to the isocenter plane, called SAD, is usually much larger than the size of the target, therefore the radiation rays emitted from the source are usually considered to be parallel to each other. However, the fact the source is a point means that the rays are not parallel, or perpendicular to the treatment machine's isocenter plane. The primary consequence is that the positions of the leaf ends are not the same as the position of the radiation field edges. Moreover, the speed at which the leaves move are not equal to that of the radiation field edge.

As shown in Figure A.1, when the leaf is at position A which is at a distance w_1 from the center of the beam, the edge of the radiation field on the isocenter plane is at W_1 . Let SCD be the distance from the source to the center of the leaf, we have:

$$\frac{w_1}{W_1} = \frac{SCD}{SAD}.$$

Now if the leaf moves to a point B w_2 away from A, a similar analysis shows that the

displacement of the radiation field edge W_2 satisfies:

$$\frac{w_2}{W_2} = \frac{\text{SCD}}{\text{SAD}} \Rightarrow W_2 = x_2 \cdot \frac{\text{SAD}}{\text{SCD}}.$$

In other words, the radiation field edge moves faster than the speed of the leaf movement by a factor of $\frac{\text{SAD}}{\text{SCD}}$.

This means that any aperture formed by the MLC is larger when projected to the isocenter plane. Since we are interested in the radiation exposure at the isocenter plane where the patient is located, we can use the projected leaf speed, i.e., the radiation field edge speed, instead of the actual leaf speed in the optimization problem. Let \hat{v} be the maximum physical MLC leaf speed, the maximum projected leaf speed can then be calculated as:

$$v = \hat{v} \cdot \frac{\text{SAD}}{\text{SCD}}.$$

Doing this will result in all “leaf positions” being radiation field edge positions, and “apertures” being exposed radiation field shapes. These can be easily converted back to the corresponding MLC leaf positions at the end of the optimization process.

Note that here we made an approximation that the radiation field edge is defined by the ray that passes the tip of the leaf, whereas the ray that is tangent to the leaf is usually slightly closer to the beam center. The difference in radiation field edge position is shown in Figure A.2 as the distance between D and E. Because $\text{SAD} \gg W$, we can assume that angle $\angle\text{CSD} = \angle\text{CSE} = \theta$. Let R be the radius of the leaf end, we have:

$$x = w + R - R \cos \theta$$

$$z = \text{SCD} + R \sin \theta.$$

Therefore:

$$\begin{aligned} \frac{x}{X} &= \frac{z}{\text{SAD}} \\ &= \frac{\text{SCD} + R \sin \theta}{\text{SAD}} \end{aligned}$$

as a result:

$$\begin{aligned}
 X &= \frac{SAD}{SCD + R \sin \theta} \\
 &= \frac{w \cdot SAD + R(1 - \cos \theta) \cdot SAD}{SCD + R \sin \theta} \\
 &= \frac{W \cdot SCD + R(1 - \cos \theta) \cdot SAD}{SCD + R \sin \theta}
 \end{aligned}$$

The difference between X and W equals:

$$\begin{aligned}
 W - X &= W - \frac{W \cdot SCD + R(1 - \cos \theta) \cdot SAD}{SCD + R \sin \theta} \\
 &= \frac{WR \sin \theta - R(1 - \cos \theta) \cdot SAD}{SCD + R \sin \theta}.
 \end{aligned}$$

Because $SAD \gg W$, $\cos \theta$ is close to 1 and $\sin \theta$ is close to 0. As result, $W - X$ is also close to 0.

Therefore we define the radiation field edge with the ray that passes the tip of the leaf. This allows us to use a simple representation of the projected leaf speed in the optimization problem, without unduly compromising the quality of the solution.

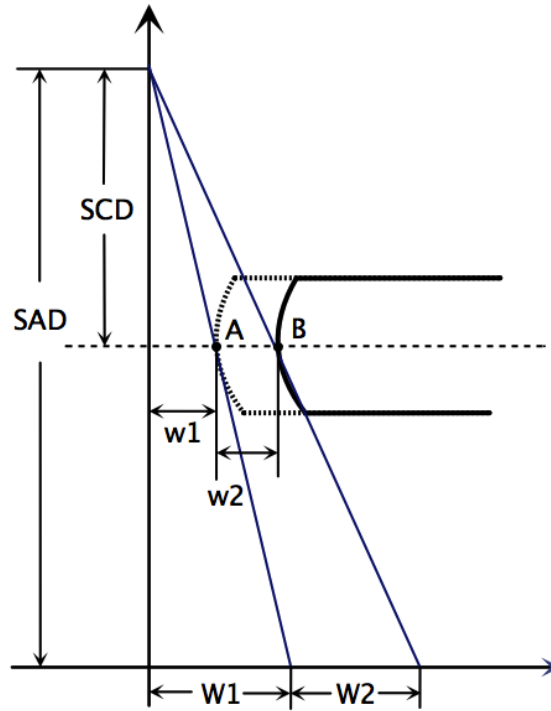


Figure A.1: A schematic diagram of the source, MLC leaf, and isocenter plane




Figure A.2: A ray passing the MLC tip and a second ray tangent to the leaf end

APPENDIX B

Derivation of Transmission Coefficient for a Beamlet Outside the Radiation Field

Part of this section follows the analysis in *Boyer and Li* (1997). Consider two rays ℓ_1 and ℓ_2 emitted from the source, with ℓ_1 tangent to the leaf end, and ℓ_2 intersecting the leaf for a length of ℓ . Shown in Figure B.1, ℓ_1 and ℓ_2 are considered approximately parallel, and their distance

$$d = R - \sqrt{R^2 - \left(\frac{\ell}{2}\right)^2},$$

and

$$d' = \frac{d}{\cos \theta}.$$

Following a similar analysis as Appendix A, we can calculate the length of the projection of d' on the isocenter plane

$$\begin{aligned} \delta &= d' \cdot \frac{\text{SAD}}{\text{SCD} + R \sin \theta} \\ &= \frac{R - \sqrt{R^2 - \left(\frac{\ell}{2}\right)^2}}{\cos \theta} \frac{\text{SAD}}{\text{SCD} + R \sin \theta} \end{aligned} \tag{B.1}$$

The Beer-Lambert law describes that the radiation intensity I after transmitted through a material of thickness y depends on the initial intensity I_0 and the material-specific coefficient




Figure B.1: Schematic of a ray tangent to the leaf end and a ray intersecting the leaf

λ :

$$\frac{I}{I_0} = e^{-\lambda y}.$$

For each material, we can define a Half Value Thickness (HVT) as the thickness required to reduce the intensity to half of the original magnitude. Then:

$$\begin{aligned} \frac{1}{2} &= e^{-\lambda \text{HVT}} \\ \Rightarrow \lambda &= \frac{\ln 2}{\text{HVT}}. \end{aligned}$$

For ray ℓ_2 the attenuation factor then equals:

$$\begin{aligned} \alpha &= e^{-\lambda \ell} = e^{-\frac{\ln 2}{\text{HVT}} \cdot \ell} = \left(\frac{1}{2}\right)^{\ell/\text{HVT}} < 1 \\ \Rightarrow \ell &= -\frac{\ln \alpha}{\ln 2} \cdot \text{HVT}. \end{aligned} \quad (\text{B.2})$$

Plugging (B.1) into (B.2), we can eliminate ℓ and find an relationship between α and δ :

$$\begin{aligned} \delta &= \frac{R - \sqrt{R^2 - \left(\frac{-\ln \alpha}{\ln 2} \cdot \text{HVT}\right)^2}}{\cos \theta} \frac{\text{SAD}}{\text{SCD} + R \sin \theta} \\ \Leftrightarrow R - \delta \cdot \frac{\cos \theta (\text{SCD} + R \sin \theta)}{\text{SAD}} &= \sqrt{R^2 - \left(\frac{-\ln \alpha}{\ln 2} \cdot \text{HVT}\right)^2} \end{aligned}$$

$$\Leftrightarrow \delta^2 \cdot \left[\frac{\cos \theta(\text{SCD} + R \sin \theta)}{\text{SAD}} \right]^2 - 2R\delta \cdot \frac{\cos \theta(\text{SCD} + R \sin \theta)}{\text{SAD}} = -\frac{1}{4} \left(\frac{\ln \alpha}{\ln 2} \right)^2 \text{HVT}.$$

Let:

$$\begin{aligned} \gamma_1 &= \frac{\cos \theta(\text{SCD} + R \sin \theta)}{\text{SAD}} = \frac{\frac{\text{SAD}}{\sqrt{\text{SAD}^2 + W^2}} (\text{SCD} + R \frac{W}{\sqrt{\text{SAD}^2 + W^2}})}{\text{SAD}} \\ &= \text{SCD} \cdot \sqrt{\text{SAD}^2 + W^2} + RW \\ \gamma_2 &= \frac{\text{HVT}}{2 \ln 2}, \end{aligned}$$

we have:

$$\delta^2 \gamma_1^2 - 2R\delta \gamma_1 = -\gamma_2^2 \ln \alpha^2.$$

Because $\alpha < 1$, $\ln \alpha < 0$, we get:

$$\begin{aligned} \alpha &= \exp \left[-\frac{1}{\gamma_2} (2R\delta \gamma_1 - \gamma_1^2 \delta^2)^{\frac{1}{2}} \right] \\ &\approx \exp \left[-\frac{1}{\gamma_2} (2R\delta \gamma_1)^{\frac{1}{2}} \right] && \text{(because } \delta \text{ is small)} \\ &\doteq \exp \left[-\mu \sqrt{\delta} \right] && \text{(define } \mu = \frac{\sqrt{2R\gamma_1}}{\gamma_2} \text{)}. \end{aligned}$$

Now we can calculate the transmission coefficient for a beam that lies between τ_1 and τ_2 from the edge of the radiation field:

$$\begin{aligned} \alpha(\tau_1, \tau_2) &= \frac{1}{\tau_2 - \tau_1} \int_{\tau_1}^{\tau_2} \alpha(\delta) d\delta \\ &= \frac{1}{\tau_2 - \tau_1} \int_{\tau_1}^{\tau_2} e^{-\mu \sqrt{\delta}} d\delta \\ &= \frac{2}{\tau_2 - \tau_1} \int_{\tau_1}^{\tau_2} e^{-\mu \sqrt{\delta}} \sqrt{\delta} d\sqrt{\delta} \\ &= \frac{2}{\tau_2 - \tau_1} \int_{\sqrt{\tau_1}}^{\sqrt{\tau_2}} e^{-\mu t} t dt \\ &= -\frac{2}{\mu(\tau_2 - \tau_1)} \left[e^{-\mu \sqrt{\tau_2}} \left(\sqrt{\tau_2} + \frac{1}{\mu} \right) - e^{-\mu \sqrt{\tau_1}} \left(\sqrt{\tau_1} + \frac{1}{\mu} \right) \right]. \end{aligned}$$

APPENDIX C

Convergence of the SAA Approach

In this section we show that when m , M , and $K \rightarrow \infty$, the objective function value of problem (SAA), namely g_m , as well as the bounds $L(M)$ and $U(K)$ converge to the optimal objective value g^* .

First note that as M and $K \rightarrow \infty$, the estimates $L(M) \rightarrow E[g_m]$ and $U(K) \rightarrow E[\mathcal{G}(\hat{x}, \mathbf{p})]$. If the solution to the (SAA) problem, namely \hat{x} , converges to the optimal solution x^* of (P), it then follows that g_m converges to g^* , and both $E[g_m]$ and $E[\mathcal{G}(\hat{x}, \mathbf{p})]$ converge to g^* , and as a result $L(M)$ and $U(K)$ also converge to g^* . We will show below that, as $m \rightarrow \infty$, the solution of the (SAA) problem converges to the solution of (P).

We start by examining the general problem (SP) and (SP_{*m*}). *Robinson* (1996) showed that with probability 1, a minimizer of (SP_{*m*}) will, for large m , exist and be as close as desired to the set of minimizers of (SP) if the following three conditions are satisfied:

1. Y is compact
2. $Ef(y, \boldsymbol{\xi})$ is continuous
3. $\frac{1}{m} \sum_{i=1}^m f(y, \boldsymbol{\xi}^i)$ converges to $Ef(y, \boldsymbol{\xi})$ uniformly with probability 1.

Instead of proving Condition 2 directly, we can take advantage of the Uniform Convergence Theorem, which states that:

Uniform Convergence Theorem. *If f_n is a series of continuous functions that uniformly converges to a function f . Then f is continuous.*

If we can prove that Condition 3 is satisfied, we can alternatively show that the following condition holds in order for Condition 2 to hold:

Condition. $\frac{1}{m} \sum_{i=1}^m f(y, \xi^i)$ is continuous for all $m > 0$.

We leave the first two conditions aside for now, and first tackle the third condition. This proof can be done using the following theorem:

Theorem. (Jennrich (1969)) *If*

- Y is compact
- $f(y, \boldsymbol{\xi})$ is continuous at each $y \in Y$ for all $\boldsymbol{\xi}$, and a measurable function of $\boldsymbol{\xi}$ at each y
- there exists a dominating function $d(\cdot)$ such that $E[d(\boldsymbol{\xi})] < \infty$, and

$$\|f(y, \boldsymbol{\xi})\| \leq d(\boldsymbol{\xi}) \quad \forall y \in Y, \boldsymbol{\xi} \in \Xi$$

Then

$$\sup_{y \in Y} \left\| \frac{1}{n} \sum_{i=1}^n f(y, \boldsymbol{\xi}_i) - E[f(y, \boldsymbol{\xi})] \right\| \rightarrow 0 \text{ almost surely,}$$

which implies uniform convergence.

The first two conditions in the Theorem above are easy to see for our problem, since the feasible region X is compact, and given the realized path $p^{(N-n)}$, the objective function is the sum of quadratic terms, and is therefore continuous. For the last condition, note that assuming the setup error of the patient is restricted to be within a certain region, i.e., large displacement will be detected and corrected, our objective function is bounded above by a finite number, say \mathcal{M} , and below by 0. Therefore if we define $d(\cdot) = \mathcal{M}$, the third condition

will also be satisfied. As a result, we know that as $m \rightarrow \infty$, the solution to problem (SAA) will be arbitrarily close to the set of solutions to (P).

BIBLIOGRAPHY

BIBLIOGRAPHY

- Ahnesjö, A., M. Saxner, and A. Trepp (1992), A pencil beam model for photon dose calculation, *Medical physics*, 19, 263.
- Aleman, D., D. Glaser, H. Romeijn, and J. Dempsey (2010), Interior point algorithms: guaranteed optimality for fluence map optimization in imrt, *Physics in medicine and biology*, 55(18), 5467.
- Audet, C., and J. Dennis Jr (2004), A pattern search filter method for nonlinear programming without derivatives, *SIAM Journal on Optimization*, 14(4), 980–1010.
- Baum, C., M. Alber, M. Birkner, and F. Nüsslin (2006), Robust treatment planning for intensity modulated radiotherapy of prostate cancer based on coverage probabilities, *Radiotherapy and oncology*, 78(1), 27–35.
- Bazaraa, M. S., H. D. Sherali, and C. M. Shetty (2006), *Nonlinear programming: theory and algorithms*, Wiley-interscience.
- Bedford, J. (2009), Treatment planning for volumetric modulated arc therapy, *Medical Physics*, 36(11), 5128–5138.
- Bentel, G. C. (1998), *Patient Positioning and Immobilization in Radiation Oncology*, McGraw-Hill Professional Publishing.
- Bertelsen, A., C. R. Hansen, J. Johansen, C. Brink, et al. (2010), Single arc volumetric modulated arc therapy of head and neck cancer., *Radiotherapy and oncology: journal of the European Society for Therapeutic Radiology and Oncology*, 95(2), 142.
- Bortfeld, T., and S. Webb (2009), Single-arc IMRT?, *Physics in Medicine and Biology*, 54, N9–N20.
- Bortfeld, T., T. Chan, A. Trofimov, and J. Tsitsiklis (2008), Robust management of motion uncertainty in intensity-modulated radiation therapy, *Operations Research*, 56(6), 1461–1473.
- Box, M. (1965), A new method of constrained optimization and a comparison with other methods, *The Computer Journal*, 8(1), 42–52.
- Boyer, A., P. Biggs, J. Galvin, E. Klein, T. LoSasso, D. Low, K. Mah, and C. Yu (2001), Basic applications of multileaf collimators. report of the aapm radiation therapy committee task group no. 50, *Med Phys*.

- Boyer, A. L., and S. Li (1997), Geometric analysis of light-field position of a multileaf collimator with curved ends, *Medical physics*, *24*, 757.
- Bzdusek, K., H. Friberger, K. Eriksson, B. Hårdemark, D. Robinson, and M. Kaus (2009), Development and evaluation of an efficient approach to volumetric arc therapy planning, *Medical Physics*, *36*(6), 2328–2339.
- Cameron, C. (2005), Sweeping-window arc therapy: an implementation of rotational IMRT with automatic beam-weight calculation, *Physics in Medicine and Biology*, *50*, 4317–4336.
- Cao, D., T. W. Holmes, M. Afghan, D. M. Shepard, et al. (2007), Comparison of plan quality provided by intensity-modulated arc therapy and helical tomotherapy., *International journal of radiation oncology, biology, physics*, *69*(1), 240.
- Cao, D., M. Afghan, J. Ye, F. Chen, and D. Shepard (2009), A generalized inverse planning tool for volumetric-modulated arc therapy, *Physics in Medicine and Biology*, *54*, 6725–6738.
- Chan, T., T. Bortfeld, and J. Tsitsiklis (2006), A robust approach to imrt optimization, *Physics in medicine and biology*, *51*(10), 2567.
- Choi, B., and J. O. Deasy (2002), The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning, *Physics in medicine and biology*, *47*(20), 3579.
- Chu, M., Y. Zinchenko, S. Henderson, and M. Sharpe (2005), Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty, *Physics in Medicine and Biology*, *50*(23), 5463.
- Conn, A. A. R., K. Scheinberg, and L. N. Vicente (2009), *Introduction to derivative-free optimization*, vol. 8, Siam.
- Conn, A. R., N. I. Gould, and P. L. Toint (2000), Trust-region methods.
- Cotrutz, C., C. Kappas, and S. Webb (2000), Intensity modulated arc therapy (imat) with centrally blocked rotational fields, *Physics in medicine and biology*, *45*(8), 2185.
- Cozzi, L., et al. (2008), A treatment planning study comparing volumetric arc modulation with rapidarc and fixed field imrt for cervix uteri radiotherapy, *Radiotherapy and Oncology*, *89*(2), 180–191.
- Craft, D., D. McQuaid, J. Wala, W. Chen, and T. Bortfeld (2011), Multicriteria VMAT optimization, *Technical report*, Massachusetts General Hospital, Boston, Massachusetts.
- Crooks, S., L. McAven, D. Robinson, and L. Xing (2002), Minimizing delivery time and monitor units in static imrt by leaf-sequencing, *Physics in medicine and biology*, *47*(17), 3105.
- Crooks, S., X. Wu, C. Takita, M. Watzich, and L. Xing (2003), Aperture modulated arc therapy, *Physics in Medicine and Biology*, *48*, 1333–1344.

- 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.
- Deng, G., and M. Ferris (2008), Neuro-dynamic programming for fractionated radiotherapy planning, *Optimization in Medicine*, pp. 47–70.
- Earl, M., D. Shepard, S. Naqvi, X. Li, and C. Yu (2003), Inverse planning for intensity-modulated arc therapy using direct aperture optimization, *Physics in Medicine and Biology*, 48, 1075–1089.
- Ekberg, L., O. Holmberg, L. Wittgren, G. Bjelkengren, and T. Landberg (1998), What margins should be added to the clinical target volume in radiotherapy treatment planning for lung cancer?, *Radiotherapy and oncology*, 48(1), 71–77.
- Ferris, M., and M. Voelker (2004), Fractionation in radiation treatment planning, *Mathematical programming*, 101(2), 387–413.
- Fogliata, A., G. Nicolini, A. Clivio, E. Vanetti, L. Cozzi, et al. (2011), Dosimetric evaluation of acuros xb advanced dose calculation algorithm in heterogeneous media, *Radiat Oncol*, 6, 82.
- Fowler, K. R., et al. (2008), Comparison of derivative-free optimization methods for groundwater supply and hydraulic capture community problems, *Advances in Water Resources*, 31(5), 743–757.
- Fredriksson, A., A. Forsgren, and B. Hårdemark (2011), Minimax optimization for handling range and setup uncertainties in proton therapy, *Medical Physics*, 38, 1672.
- Galvin, J. M., et al. (1992), Evaluation of multileaf collimator design for a photon beam, *International Journal of Radiation Oncology* Biology* Physics*, 23(4), 789–801.
- Gen, M., and R. Cheng (1999), *Genetic algorithms and engineering optimization*, vol. 7, Wiley-interscience.
- Gilmore, P., and C. T. Kelley (1995), An implicit filtering algorithm for optimization of functions with many local minima, *SIAM Journal on Optimization*, 5(2), 269–285.
- Goldberg, D. E. (1989), Genetic algorithms in search, optimization, and machine learning.
- Gözbasi, H. (2010), Optimization approaches for planning external beam radiotherapy, Ph.D. thesis, Georgia Institute of Technology.
- Grippo, L., and M. Sciandrone (2000), On the convergence of the block nonlinear gauss–seidel method under convex constraints, *Operations Research Letters*, 26(3), 127–136.
- Gu, X., D. Choi, C. Men, H. Pan, A. Majumdar, and S. Jiang (2009), GPU-based ultra-fast dose calculation using a finite size pencil beam model, *Physics in Medicine and Biology*, 54, 6287–6297.

- Gu, X., U. Jelen, J. Li, X. Jia, and S. Jiang (2011), A GPU-based finite-size pencil beam algorithm with 3D-density correction for radiotherapy dose calculation, *Physics in Medicine and Biology*, *56*, 3337–3350.
- Halperin, E. C., C. A. Perez, and L. W. Brady (2008), *Perez and Brady's principles and practice of radiation oncology*, Lippincott Williams & Wilkins.
- Intensity Modulated Radiation Therapy Collaborative Working Group (2001), Intensity-modulated radiotherapy: current status and issues of interest, *International Journal of Radiation Oncology Biology Physics*, *51*(4), 880–914.
- Jeleń, U., M. Söhn, and M. Alber (2005), A finite size pencil beam for imrt dose optimization, *Physics in medicine and biology*, *50*(8), 1747.
- Jennrich, R. I. (1969), Asymptotic properties of non-linear least squares estimators, *The Annals of Mathematical Statistics*, pp. 633–643.
- Jia, X., X. Gu, J. Sempau, D. Choi, A. Majumdar, and S. B. Jiang (2010), Development of a gpu-based monte carlo dose calculation code for coupled electron–photon transport, *Physics in medicine and biology*, *55*(11), 3077.
- Jia, X., C. Men, Y. Lou, and S. B. Jiang (2011), Beam orientation optimization for intensity modulated radiation therapy using adaptive l2, 1-minimization, *Physics in medicine and biology*, *56*(19), 6205.
- Johnson, S. G. (), The nlopt nonlinear-optimization package, <http://ab-initio.mit.edu/nlopt>, accessed: 02/2013.
- King, A., and R. Wets (1991), Epi-consistency of convex stochastic programs, *Stochastics and Stochastic Reports*, *34*(1-2), 83–92.
- Kirkpatrick, S., D. G. Jr., and M. P. Vecchi (1983), Optimization by simulated annealing, *science*, *220*(4598), 671–680.
- Kleywegt, A., A. Shapiro, and T. Homem-de Mello (2002), The sample average approximation method for stochastic discrete optimization, *SIAM Journal on Optimization*, *12*(2), 479–502.
- Lagarias, J. C., J. A. Reeds, M. H. Wright, and P. E. Wright (1998), Convergence properties of the nelder–mead simplex method in low dimensions, *SIAM Journal on Optimization*, *9*(1), 112–147.
- Langen, K., and D. Jones (2001), Organ motion and its management, *International Journal of Radiation Oncology* Biology* Physics*, *50*(1), 265–278.
- Langen, K., S. Meeks, D. Poole, T. Wagner, T. Willoughby, P. Kupelian, K. Ruchala, J. Haimerl, and G. Olivera (2005), The use of megavoltage ct (mvct) images for dose recomputations, *Physics in medicine and biology*, *50*(18), 4259.

- Li, J., and L. Xing (2000), Inverse planning incorporating organ motion, in *Engineering in Medicine and Biology Society, 2000. Proceedings of the 22nd Annual International Conference of the IEEE*, vol. 4, pp. 3269–3272, IEEE.
- Linderoth, J., A. Shapiro, and S. Wright (2006), The empirical behavior of sampling methods for stochastic programming, *Annals of Operations Research*, 142(1), 215–241.
- Löf, J., B. Lind, and A. Brahme (1999), Optimal radiation beam profiles considering the stochastic process of patient positioning in fractionated radiation therapy, *Inverse Problems*, 11(6), 1189.
- Luan, S., C. Wang, D. Cao, D. Chen, D. Shepard, and C. Yu (2008), Leaf-sequencing for intensity-modulated arc therapy using graph algorithms, *Medical Physics*, 35(1), 61–69.
- Lyman, J. T. (1985), Complication probability as assessed from dose-volume histograms, *Radiation Research*, 104(2s), 13–19.
- Ma, L., et al. (2001), Optimized intensity-modulated arc therapy for prostate cancer treatment, *International journal of cancer*, 96(6), 379–384.
- Mackie, T., J. Scrimger, and J. Battista (1985), A convolution method of calculating dose for 15-mv x rays, *Medical physics*, 12, 188.
- Madansky, A. (1960), Inequalities for stochastic linear programming problems, *Management science*, 6(2), 197–204.
- Mak, W., D. Morton, and R. Wood (1999), Monte carlo bounding techniques for determining solution quality in stochastic programs, *Operations Research Letters*, 24, 47–56.
- Matuszak, M. M., D. Yan, I. Grills, A. Martinez, et al. (2010), Clinical applications of volumetric modulated arc therapy., *International journal of radiation oncology, biology, physics*, 77(2), 608.
- McShan, D., M. Kessler, K. Vineberg, and B. Fraass (2006), Inverse plan optimization accounting for random geometric uncertainties with a multiple instance geometry approximation (miga), *Medical physics*, 33, 1510.
- Men, C., X. Gu, D. Choi, A. Majumdar, Z. Zheng, K. Mueller, and S. Jiang (2009), GPU-based ultra-fast IMRT plan optimization, *Physics in Medicine and Biology*, 54, 6565–6573.
- Men, C., X. Jia, and S. Jiang (2010a), GPU-based ultra-fast direct aperture optimization for online adaptive radiation therapy, *Physics in Medicine and Biology*, 55, 4309–4319.
- Men, C., H. Romeijn, X. Jia, and S. Jiang (2010b), Ultra-fast treatment plan optimization for volumetric modulated arc therapy (VMAT), *Medical Physics*, 37(11), 5787–5791.
- Men, C., H. Edwin Romeijn, A. Saito, and J. Dempsey (2011), An efficient approach to incorporating interfraction motion in imrt treatment planning, *Computers & Operations Research*.

- Mohan, R., X. Zhang, H. Wang, Y. Kang, X. Wang, H. Liu, K. Ang, D. Kuban, and L. Dong (2005), Use of deformed intensity distributions for on-line modification of image-guided imrt to account for interfractional anatomic changes, *International Journal of Radiation Oncology* Biology* Physics*, 61(4), 1258–1266.
- Montgomery, D. C., and R. H. Myers (2002), Response surface methodology: process and product optimization using designed experiments, *Response Surface Methodology: Process and Product Optimization Using Designed Experiments*.
- Nelder, J. A., and R. Mead (1965), A simplex method for function minimization, *The computer journal*, 7(4), 308–313.
- Niemierko, A. (1997), Reporting and analyzing dose distributions: a concept of equivalent uniform dose, *Medical physics*, 24, 103.
- Niemierko, A. (1999), A generalized concept of equivalent uniform dose (eud), *Med Phys*, 26(6), 1100.
- Nohadani, O., J. Seco, and T. Bortfeld (2010), Motion management with phase-adapted 4d-optimization, *Physics in medicine and biology*, 55(17), 5189.
- Olafsson, A., and S. Wright (2006), Efficient schemes for robust imrt treatment planning, *Physics in medicine and biology*, 51(21), 5621.
- Oldham, M., et al. (2005), Cone-beam-ct guided radiation therapy: a model for on-line application, *Radiotherapy and oncology*, 75(3), 271–E1.
- Otto, K. (2008), Volumetric modulated arc therapy: IMRT in a single gantry arc, *Medical Physics*, 35(1), 310–317.
- Otto, K. (2009), Letter to the editor on ‘Single-Arc IMRT?’, *Physics in Medicine and Biology*, 54, L37–L41.
- Palma, D., E. Vollans, K. James, S. Nakano, V. Moiseenko, R. Shaffer, M. McKenzie, J. Morris, and K. Otto (2008), Volumetric modulated arc therapy for delivery of prostate radiotherapy: comparison with intensity-modulated radiotherapy and three dimensional conformal radiotherapy, *International Journal of Radiation Oncology Biology Physics*, 72(4), 996–1001.
- Papanikolaou, N., T. R. Mackie, C. Meger-Wells, M. Gehring, and P. Reckwerdt (1993), Investigation of the convolution method for polyenergetic spectra, *Medical physics*, 20, 1327.
- Partridge, M., M. Ebert, and B. Hesse (2002), Imrt verification by three-dimensional dose reconstruction from portal beam measurements, *Medical physics*, 29, 1847.
- Peng, F., X. Jia, X. Gu, M. A. Epelman, H. E. Romeijn, and S. B. Jiang (2012), A new column-generation-based algorithm for vmat treatment plan optimization, *Physics in Medicine and Biology*, 57(14), 4569.

- 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., J. G. Li, A. L. Boyer, S. L. Hancock, Q.-T. Le, S. S. Donaldson, and L. Xing (2001), Role of beam orientation optimization in intensity-modulated radiation therapy, *International Journal of Radiation Oncology* Biology* Physics*, 50(2), 551–560.
- Radiation Therapy Oncology Group (2004), H-0126: A phase iii randomized study of high dose 3d-crt/imrt versus standard dose 3d-crt/imrt in patients treated for localized prostate cancer, <http://www.rtog.org/members/protocols/0126/p0126.pdf>.
- Robinson, S. (1996), Analysis of sample-path optimization, *Mathematics of Operations Research*, pp. 513–528.
- Rogers, D., and A. Bielajew (1990), Monte carlo techniques of electron and photon transport for radiation dosimetry, *The dosimetry of ionizing radiation*, 3, 427–539.
- 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.
- Romeijn, H. E., and J. F. Dempsey (2008), Intensity modulated radiation therapy treatment plan optimization, *Top*, 16(2), 215–243.
- Romeijn, H. E., R. K. Ahuja, J. F. Dempsey, A. Kumar, and J. G. 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.
- Romeijn, H. E., J. F. Dempsey, and J. G. Li (2004), A unifying framework for multi-criteria fluence map optimization models, *Physics in medicine and biology*, 49(10), 1991.
- Ruszczynski, A., and A. Shapiro (2003), Stochastic programming, handbook in or & ms, vol. 10.
- Saka, B., R. Rardin, M. Langer, and D. Dink (2011), Adaptive intensity modulated radiation therapy planning optimization with changing tumor geometry and fraction size limits, *IIE Transactions on Healthcare Systems Engineering*, 1(4), 247–263.
- 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.
- Shepard, D. M., M. C. Ferris, G. H. Olivera, and T. R. Mackie (1999), Optimizing the delivery of radiation therapy to cancer patients, *Siam Review*, 41(4), 721–744.
- Shepard, D. M., D. Cao, M. Afghan, and M. Earl (2007), An arc-sequencing algorithm for intensity modulated arc therapy, *Medical Physics*, 34(2), 464–470.
- Sir, M., M. Epelman, and S. Pollock (2012), Stochastic programming for off-line adaptive radiotherapy, *Annals of Operations Research*, 196(1), 767.

- Sobotta, B., M. Söhn, and M. Alber (2010), Robust optimization based upon statistical theory, *Medical physics*, *37*, 4019.
- Stein, J., R. Mohan, X.-H. Wang, T. Bortfeld, Q. Wu, K. Preiser, C. C. Ling, and W. Schlegel (1997), Number and orientations of beams in intensity-modulated radiation treatments, *Medical Physics*, *24*, 149.
- Taşkın, Z. C., J. C. Smith, H. E. Romeijn, and J. F. Dempsey (2010), Optimal multileaf collimator leaf sequencing in imrt treatment planning, *Operations research*, *58*(3), 674–690.
- Thieke, C., T. Bortfeld, A. Niemierko, and S. Nill (2003), From physical dose constraints to equivalent uniform dose constraints in inverse radiotherapy planning, *Medical physics*, *30*, 2332.
- Tinger, A., J. M. Michalski, A. Cheng, D. A. Low, W. BOSCH, J. PURDY, C. PEREZ, et al. (1998), A critical evaluation of the planning target volume for 3-d conformal radiotherapy of prostate cancer, *International journal of radiation oncology, biology, physics*, *42*(1), 213–221.
- Tseng, P. (2001), Convergence of a block coordinate descent method for nondifferentiable minimization, *Journal of optimization theory and applications*, *109*(3), 475–494.
- Unkelbach, J., and U. Oelfke (2005), Incorporating organ movements in imrt treatment planning for prostate cancer: minimizing uncertainties in the inverse planning process, *Medical physics*, *32*, 2471.
- Van Esch, A., et al. (2006), Testing of the analytical anisotropic algorithm for photon dose calculation, *Medical physics*, *33*, 4130.
- van Herk, M., P. Remeijer, C. Rasch, J. V. Lebesque, et al. (2000), The probability of correct target dosage: dose-population histograms for deriving treatment margins in radiotherapy., *International journal of radiation oncology, biology, physics*, *47*(4), 1121.
- Van Herk, M., et al. (2004), Errors and margins in radiotherapy, in *Seminars in radiation oncology*, vol. 14, p. 52.
- Van Laarhoven, P. J., and E. H. Aarts (1987), *Simulated annealing*, Springer.
- Verbakel, W., J. Cuijpers, D. Hoffmans, M. Bieker, B. Slotman, and S. Senan (2009), Volumetric intensity-modulated arc therapy vs. conventional IMRT in head-and-neck cancer: a comparative planning and dosimetric study, *International Journal of Radiation Oncology Biology Physics*, *74*(1), 252–259.
- Verellen, D., N. Linthout, G. Soete, S. Van Acker, P. De Roover, and G. Storme (2002), Considerations on treatment efficiency of different conformal radiation therapy techniques for prostate cancer, *Radiotherapy and oncology*, *63*(1), 27–36.

- Verweij, B., S. Ahmed, A. Kleywegt, G. Nemhauser, and A. Shapiro (2003), The sample average approximation method applied to stochastic routing problems: a computational study, *Computational Optimization and Applications*, 24(2), 289–333.
- Wang, C., S. Luan, G. Tang, D. Chen, M. Earl, and C. Yu (2008), Arc-modulated radiation therapy (AMRT): a single-arc form of intensity-modulated arc therapy, *Physics in Medicine and Biology*, 53, 6291–6303.
- Wang, L., C.-S. Chui, and M. Lovelock (1998), A patient-specific monte carlo dose-calculation method for photon beams, *Medical physics*, 25, 867.
- Wolbarst, A. B. (1984), Optimization of radiation therapy ii: The critical-voxel model, *International Journal of Radiation Oncology* Biology* Physics*, 10(5), 741–745.
- Wong, E., J. Z. Chen, and J. Greenland (2002), Intensity-modulated arc therapy simplified, *International Journal of Radiation Oncology* Biology* Physics*, 53(1), 222–235.
- Wright, S. (1987), *Primal-dual interior-point methods*, vol. 54, Society for Industrial Mathematics.
- Wu, Q., D. Djajaputra, H. H. Liu, L. Dong, R. Mohan, and Y. Wu (2005), Dose sculpting with generalized equivalent uniform dose, *Medical physics*, 32, 1387.
- Yang, Y., E. Schreibmann, T. Li, C. Wang, and L. Xing (2007), Evaluation of on-board kv cone beam ct (cbct)-based dose calculation, *Physics in medicine and biology*, 52(3), 685.
- Yu, C. (1995), Intensity-modulated arc therapy with dynamic multileaf collimation: an alternative to tomotherapy, *Physics in Medicine and Biology*, 40(1435–1449).
- Yu, C. X., et al. (2002), Clinical implementation of intensity-modulated arc therapy., *International journal of radiation oncology, biology, physics*, 53(2), 453.
- Zaider, M., and G. Minerbo (1999), Tumour control probability: a formulation applicable to any temporal protocol of dose delivery, *Physics in medicine and biology*, 45(2), 279.