

ADAPTIVE AND ROBUST RADIATION THERAPY OPTIMIZATION FOR
LUNG CANCER

by

Velibor V. Mišić

A thesis submitted in conformity with the requirements
for the degree of Master of Applied Science
Graduate Department of Mechanical and Industrial Engineering
University of Toronto

Copyright © 2012 by Velibor V. Mišić

Abstract

Adaptive and robust radiation therapy optimization for lung cancer

Velibor V. Mišić

Master of Applied Science

Graduate Department of Mechanical and Industrial Engineering

University of Toronto

2012

A previous approach to robust intensity-modulated radiation therapy (IMRT) treatment planning for moving tumours in the lung involves solving a single planning problem before treatment and using the resulting solution in all of the subsequent treatment sessions. In this thesis, we develop two adaptive robust IMRT optimization approaches for lung cancer, which involve using information gathered in prior treatment sessions to guide the reoptimization of the treatment for the next session. The first method is based on updating an estimate of the uncertain effect, while the second is based on additionally updating the dose requirements to account for prior errors in dose. We present computational results using real patient data for both methods and an asymptotic analysis for the first method. Through these results, we show that both methods lead to improvements in the final dose distribution over the traditional robust approach, but differ greatly in their daily dose performance.

Acknowledgements

First of all, I would like to thank my advisor and friend, Timothy Chan. I cannot begin to express my gratitude for his thoughtfulness, patience, encouragement and support over the last two years. One of the things that I have learned about being a researcher is that you get a front-row seat to the full spectrum of human emotion: the dizzying highs that come with an exciting new theoretical result or computational result, but also the dark lows when you just cannot push that proof any further or your computational experiments do not work out as you had hoped they would. Whenever I walked into Tim's office with fears and doubts, though, I always came out with a new way of looking at my problems and a renewed determination to solve them. From discussions on which line of research to pursue to effective writing to academia and life in general, I have learned an immeasurable amount from Tim. If I do become a professor one day, I hope that I can be the kind of advisor to my students that Tim has been to me.

I would like to thank Dionne Aleman, my undergraduate advisor, and Roy Kwon, for being on my Master's defense committee and for providing many helpful suggestions on this work; I would also like to thank Dionne for introducing me to the radiation therapy optimization research field. I would like to thank Thomas Bortfeld of Massachusetts General Hospital for providing the patient data used in the computational studies in Chapters 4 and 5. I would like to thank Troy Long, Marina Epelman and Edwin Romeijn of the University of Michigan for helpful discussions and for providing the encouragement that led to the dose-reactive work that appears as Chapter 5.

I would also like to thank the MIE graduate office staff for their help in various matters. I would especially like to thank Donna Liu, who was there during my first year and who always had an answer for every question that I came in with, and Brenda Fung, who helped me with my various questions in my second year.

I have had the pleasure of knowing many wonderful graduate students in the MIE department. Within the Applied Optimization Lab in RS206, I would like to thank my labmates Derya Demirtas, Taewoo Lee, Sarina Turner, Houra Mahmoudzadeh, Auyon Siddiq, Heyse Li, Brendan Eagen and Yifang Liu. Outside of AOL, I would like to thank my UTOrg crew Michael Kim, Jonathan Li, Kimia Ghobadi and Jenya Doudareva, as well as Jason Lee, Konstantin Shestopaloff, Hamid Ghaffari, Jane Ip and Peter Zhang.

My friends from Toronto have been my biggest fans and deserve great thanks for keeping me sane. The following is by no means an exhaustive list: thank you Tiffany Woo, Matthew Gordon, Jarrtan Naphtal, Emma Tsui, Geoffrey Vishloff, Lily Qiu, Mike Lennox, Archis Deshpande, Oren Kraus, Anne Zhang, Flavia Franco, Jordan DiCarlo and Torin Gillen. I would especially like to thank my former roommate Eric Bradshaw for constantly reminding me why we are here and for always believing in me. I would also like to thank Chen Chen, my former classmate, for the many enjoyable conversations on research and academia and for his constant encouragement.

I would like to thank Elena for being a wonderfully supportive and understanding roommate, for introducing me to such delights as *solyanka* and for making our evening viewings of episodes of *The Wire*, *The Walking Dead* and (most recently) *Breaking Bad* unforgettable. Lastly, I would like to thank my family, Bata, Voja and Lela for their boundless love, support and advice. Their skill at dynamic decision making under uncertainty far exceeds anything that you will see in this thesis, and without them, all of this would have been in vain.

The research conducted for this thesis was supported by a Canada Graduate Scholarship Master's level (CGS-M) award from the Natural Sciences and Engineering Research Council (NSERC) of Canada, an Ontario Graduate Scholarship Master's level (OGS-M) award from the Ontario Ministry of Training, Colleges and Universities, and by the Collaborative Health Research Projects (CHRP) grant #398106-2011 from NSERC and the Canadian Institutes for Health Research (CIHR).

Contents

1	Introduction	1
1.1	Motivation	1
1.2	Contributions	4
1.3	Organization	5
2	Literature review	6
2.1	Uncertainty in radiation therapy treatments	6
2.2	Adaptive radiation therapy	8
2.3	Dose-reactive methods	9
2.4	Contribution to the literature	10
3	Static robust IMRT optimization	14
3.1	Model	15
3.2	Solving the robust problem	17
4	Adaptive and robust IMRT optimization	19
4.1	Method	20
4.1.1	Update algorithms	21
4.1.2	Prescient solutions	22
4.2	Asymptotic optimality of the adaptive robust method	24
4.2.1	An alternative robust problem	26
4.2.2	Convergence of optimal solution sets and dose distributions under convex-convergent update algorithms	27
4.2.3	Three theoretical insights	30
4.3	Computational study	32
4.3.1	Background	32

4.3.2	Results	34
4.3.3	Discussion	39
4.4	Conclusions	43
5	Dose-reactive methods	45
5.1	Method	47
5.2	Results	51
5.2.1	Background	51
5.2.2	Final cumulative dose results	54
5.2.3	Daily dose results	62
5.2.4	Discussion	71
5.3	Theoretical insight into daily dose performance	74
5.4	Conclusions	78
6	Conclusions and future work	80
A	Proofs for Chapter 4	84
A.1	Assumptions	84
A.2	An alternative robust problem	89
A.3	Convergence of optimal solutions sets and dose distributions under convex-convergent update algorithms	95
B	Pathological PMF sequence example for Chapter 4	108
C	Proofs for Chapter 5	113
D	Second PMF sequence results for Chapter 5	115
	Bibliography	128

List of Figures

4.1	Graphical representation of the steps of the adaptive robust method. . .	20
4.2	Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the first PMF sequence (full).	36
4.3	Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the first PMF sequence (zoomed-in). .	37
4.4	Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the second PMF sequence (full). . . .	37
4.5	Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the second PMF sequence (zoomed-in). .	38
4.6	Dose volume histogram comparing the treatments obtained by exponential smoothing with $\alpha = 0.5$ started from the robust uncertainty set (implementation (ES(0.5),R)) and the static method with the robust uncertainty set (implementation (S,R)) with the first PMF sequence.	38
4.7	Plot of mean and maximum underdose by fraction for the three exponential smoothing with $\alpha = 0.9$ treatments ((ES(0.9),N), (ES(0.9),R), (ES(0.9),M)) under the first PMF sequence.	42
5.1	Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0$ applied to the first PMF sequence.	57
5.2	Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.1$ applied to the first PMF sequence.	58
5.3	Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.5$ applied to the first PMF sequence.	59

5.4	Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.9$ applied to the first PMF sequence.	60
5.5	Mean and max. underdose and overdose by fraction for implementations using the reactive, reactive ⁻ or reactive ⁺ exponential smoothing methods with $\alpha = 0$ or the static (non-adaptive, non-reactive) method, under the first PMF sequence.	63
5.6	Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive [±] , reactive ⁻ or reactive ⁺ exponential smoothing methods with $\alpha = 0.1$, under the first PMF sequence.	64
5.7	Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive [±] , reactive ⁻ or reactive ⁺ exponential smoothing methods with $\alpha = 0.5$, under the first PMF sequence.	65
5.8	Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive [±] , reactive ⁻ or reactive ⁺ exponential smoothing methods with $\alpha = 0.9$, under the first PMF sequence.	66
5.9	Mean and max. underdose and overdose by fraction for the non-reactive, reactive [±] , reactive ⁻ and reactive ⁺ daily prescient methods under the first PMF sequence.	67
5.10	Daily delivered dose and daily target minimum and maximum doses by fraction for the (ES(0.1),R), (R [±] ES(0.1),R), (R ⁻ ES(0.1),R) and (R ⁺ ES(0.1),R) implementations under the first PMF sequence.	69
B.1	Plot of observed PMF in fraction 29 versus lower/upper bound vectors for fraction 29 for three different exponential smoothing treatments.	109
B.2	Plot of observed PMF in fraction 30 versus lower/upper bound vectors for fraction 30 for three different exponential smoothing treatments.	109
B.3	Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the pathological PMF sequence (full).	111
B.4	Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the pathological PMF sequence (zoomed-in).	112

D.1	Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0$ applied to the second PMF sequence.	116
D.2	Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.1$ applied to the second PMF sequence.	117
D.3	Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.5$ applied to the second PMF sequence.	118
D.4	Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.9$ applied to the second PMF sequence.	119
D.5	Mean and max. underdose and overdose by fraction for implementations using the reactive, reactive ⁻ or reactive ⁺ exponential smoothing methods with $\alpha = 0$ or the static (non-adaptive, non-reactive) method, under the second PMF sequence.	120
D.6	Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive [±] , reactive ⁻ or reactive ⁺ exponential smoothing methods with $\alpha = 0.1$, under the second PMF sequence.	121
D.7	Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive [±] , reactive ⁻ or reactive ⁺ exponential smoothing methods with $\alpha = 0.5$, under the second PMF sequence.	122
D.8	Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive [±] , reactive ⁻ or reactive ⁺ exponential smoothing methods with $\alpha = 0.9$, under the second PMF sequence.	123
D.9	Mean and max. underdose and overdose by fraction for the non-reactive, reactive [±] , reactive ⁻ and reactive ⁺ daily prescient methods under the second PMF sequence.	124
D.10	Daily delivered dose and daily target minimum and maximum doses by fraction for the (ES(0.1),R), (R [±] ES(0.1),R), (R ⁻ ES(0.1),R) and (R ⁺ ES(0.1),R) implementations under the second PMF sequence.	125

List of Tables

4.1	Dose statistics for the first PMF sequence.	34
4.2	Dose statistics for the second PMF sequence.	35
5.1	Reactive [±] , reactive ⁻ and reactive ⁺ underdose example.	51
5.2	Reactive [±] , reactive ⁻ and reactive ⁺ overdose example.	51
5.3	Dose statistics for the first PMF sequence under the non-reactive and reactive methods.	55
5.4	Dose statistics for the first PMF sequence under the reactive ⁻ and reactive ⁺ methods.	56
B.1	Dose statistics for the pathological PMF sequence.	111
D.1	Dose statistics for the second PMF sequence under the non-reactive and reactive [±] methods.	126
D.2	Dose statistics for the second PMF sequence under the reactive ⁻ and reactive ⁺ methods.	127

And she spoke unto me saying,
'Fear not the movement of the heavens above or the earth below
For change is what we are, my child.

Righteous are those
Who look up and sway with the wind,
Who look down and dance with the shifting of the soil,
Who swim with the movement of the tides,
Who seek the truth around them
And discover that we are
And have always been in paradise:
The reflections of heaven on earth. Amen!'

– Puscifer, *Sour Grapes*

Chapter 1

Introduction

1.1 Motivation

Lung cancer is the leading cause of death due to cancer in North America, killing an estimated 180,000 people in 2010 (American Cancer Society [2010], Canadian Cancer Society's Steering Committee [2010]) and accounting for over 25% of all cancer deaths. Lung cancer is often treated using radiation therapy (Toschi et al. [2007]). Of the different types of radiation therapy, one of the most commonly used in practice for treating cancer in general is intensity-modulated radiation therapy (IMRT) (Mell et al. [2005]). In an IMRT treatment, the patient is irradiated from multiple beams, each of which is decomposed into a large number of small beamlets. The beamlet intensities can be controlled through the use of a multileaf collimator (MLC) that moves metal leaves in and out of the beam field in order to block certain parts of the beam (Spirou and Chui [1994]). By appropriately setting the beamlet intensities, the volume that is irradiated can be made to closely conform to the shape of the target. The basic problem in planning an IMRT treatment is to determine how the beamlet intensities or weights should be set so that the target receives an adequate dose while the healthy tissue receives a minimal dose. This is known as the beamlet weight optimization problem or the fluence map optimization problem. Since the inception of IMRT, much research has focused on modelling and solving this problem as a mathematical program (see Romeijn and Dempsey [2008] for a comprehensive overview).

In practice, the beamlet weight optimization problem is complicated by the presence of uncertainties, such as those arising from errors in beam positioning and patient placement, internal organ motion during treatment, and changes in organ position

between treatment sessions. All of these factors affect the relative position of the tumour with respect to the beams, which in turn affects how much dose is deposited in the tumour and the healthy tissue. For tumours in the lung, the most significant uncertainty arises from breathing motion. During treatment, the patient is constantly breathing, and the tumour moves with the expansion and contraction of the patient's lungs. Furthermore, the patient's breathing pattern during treatment is not known exactly beforehand and can vary from day to day. If a treatment is planned with a specific breathing pattern in mind but a different pattern is realized during treatment, the tumour may end up being underdosed and the quality of the treatment may thus be greatly compromised (Lujan et al. [2003], Sheng et al. [2006]). At the same time, if the treatment is designed to deliver the prescription dose to the tumour under any breathing pattern, regardless of how unlikely some of those patterns may be, then an unnecessarily high amount of damage will be done to the healthy tissue. There exist approaches that aim to mitigate the effect of uncertainty by only activating the beam when the patient is in a certain phase of the breathing cycle (Ohara et al. [1989]), but these approaches have the disadvantage of prolonging the treatment time, leading to reduced throughput in the treatment center.

One methodology that allows for tumour coverage under uncertainty to be balanced with healthy tissue sparing, without extending the treatment time, is robust optimization. There are many robust optimization approaches to IMRT treatment planning. In this thesis, we build on the method developed in Bortfeld et al. [2008] and Chan et al. [2006], which is specifically designed to manage the instantaneous breathing motion uncertainty that may be realized during treatment sessions. Although robust optimization has been shown to be a valuable methodology for managing uncertainty due to breathing motion, it can be further improved by leveraging the fact that treatments are typically fractionated. Fractionation refers to the practice of dividing up the prescription dose and delivering small amounts in daily treatment sessions over multiple weeks, in order to exploit the more capable repair mechanism that healthy tissue has over cancerous tissue (Thames and Hendry [1987]).

To date, the typical approach to fractionated treatment planning and delivery has been to determine what the beamlet intensities should be to deliver a certain prescription dose, scale them by $1/n$ where n is the number of fractions, and use the resulting intensities in every treatment fraction. This, however, does not take advantage of the dynamics that fractionation injects into IMRT treatment planning and the underlying

uncertainty. For example, it may be the case that the pre-treatment estimates of the patient's breathing pattern are erroneous, and the breathing patterns that are planned for differ from those realized over the treatment. In this case, the wrong beamlet intensities would be used for every fraction, without the possibility of correcting them. On the other hand, the beamlet intensities may be optimized in such a way that the prescription dose is delivered under a very large range of breathing patterns, when in actuality the patient's breathing pattern does not vary significantly from fraction to fraction. In this case, there would be no way to correct the conservatism of the treatment, and the patient's healthy tissue will incur more damage than necessary. Another example comes from the possibility that the patient's breathing pattern may change over the course of treatment as the patient becomes more comfortable in the treatment room or as the evolution of the underlying disease affects the patient's respiratory ability. As a result, latter treatment sessions may deliver insufficient dose to the tumour and more dose than anticipated to healthy tissue. These issues motivate the incorporation of dynamic re-optimization into robust optimization models.

To deal with fractionation in more general IMRT treatment planning contexts, much research has been conducted in the medical physics community in what is known as off-line adaptive radiation therapy (ART). In off-line ART, after the current fraction is delivered, new information obtained from imaging devices is fed back into the planning process to design the beamlet intensities of the next fraction (Yan et al. [1997a]). While off-line ART research has been conducted concurrently with robust optimization for IMRT, these two methods have largely been considered independently of each other.

In light of this separation between robust optimization and adaptive radiation therapy, this thesis broadly seeks to answer two questions:

1. How can we combine the robust optimization methodology of Chan et al. [2006] and Bortfeld et al. [2008] with off-line ART?
2. What can be gained from combining these two distinct methodologies?

The clinical value of a method combining robust optimization and adaptive radiation therapy is that it has the potential to improve upon the non-adaptive robust optimization method in both tumour coverage and healthy tissue sparing. This is significant because patient survival rates have been shown to improve when the tumour dose is increased (see Perez et al. [1986]). Unfortunately, dose escalation to the tumour

is limited by the radiation tolerance of healthy tissue, particularly the lungs (Cox et al. [1990]). If it is possible to increase tumour dose and reduce healthy tissue dose from the levels achieved by non-adaptive robust optimization, further dose escalation will become viable, and improvements in patient survival rate may be realized.

1.2 Contributions

The contributions of this thesis are as follows. In the first part of this thesis, we present a method that combines off-line adaptive radiation therapy and the robust optimization methodology of Bortfeld et al. [2008] and Chan et al. [2006] to treat lung cancer in the presence of breathing motion uncertainty. This method is based on adaptively updating an uncertainty set that models the breathing motion uncertainty and then solving a robust optimization problem with the new uncertainty set prior to each fraction. As a result, this method is able to combat both the instantaneous uncertainty realized in each fraction as well as changes in this uncertainty that occur from day to day. We show computationally, using real lung cancer patient data, that this method improves on the non-adaptive robust optimization method in both tumour coverage and healthy tissue sparing. We also derive theoretical results, based on the convergence of optimal solution sets of optimization problems, that explain the performance we observe in our computational study.

In the second part of this thesis, we study adaptive and robust methods for lung cancer IMRT that also incorporate *dose-reaction*. In the dose-reactive adaptive robust method, the uncertainty set is still updated with each fraction, but in addition, the target dose for the remaining fractions is also updated based on the dose delivered in prior fractions. The rationale behind such a method is that by incorporating the prior delivered dose in each fraction, the method will be able to correct relative underdosages and overdosages experienced in earlier fractions, and ensure a more satisfactory final dose to the tumour. We show that although methods that employ dose-reaction can lead to improvements in the final dose distribution over non-reactive methods, they behave very differently from non-reactive methods from the perspective of daily dose – in particular, they exhibit growing underdose and/or overdose on a daily level, which may be detrimental to the treatment from a clinical point of view. Using the same computational setup, we show that in contrast to dose-reactive

methods, non-reactive methods do not exhibit this type of behavior. Finally, using the theoretical framework developed in the first part of the thesis, we provide some theoretical intuition to explain why dose-reactive and non-reactive methods differ in daily dose performance.

1.3 Organization

This thesis is organized as follows. In Chapter 2, we review the literature on uncertainty in IMRT treatments, adaptive radiation therapy and dose-reactive methods. In Chapter 3, we review the robust IMRT optimization methodology of Chan et al. [2006] and Bortfeld et al. [2008] on which this thesis builds. In Chapter 4, we develop an adaptive and robust IMRT optimization method for lung cancer. In Chapter 5, we develop dose-reactive adaptive robust IMRT optimization methods. Finally, we provide concluding remarks and directions for future work in Chapter 6.

Chapter 2

Literature review

In this chapter, we discuss some of the prior research that is related to this thesis. In Section 2.1, we survey the literature on uncertainty in radiation therapy treatments. In Section 2.2, we survey the literature on adaptive radiation therapy. In Section 2.3, we survey the literature on dose-reactive methods, which are the topic of the second part of this thesis (Chapter 5). Finally, we discuss the contribution of this thesis to the landscape of prior research in Section 2.4.

2.1 Uncertainty in radiation therapy treatments

In the last decade, there has been much interest in the medical physics community to consider the effect of uncertainty on IMRT treatment plans. Uncertain effects such as intrafraction organ motion (organ motion that occurs during a treatment session), interfraction organ motion (organ position changes that occur between treatment sessions), patient positioning error and beam placement error are of critical importance to IMRT treatment planning as they have the potential to severely compromise the quality of treatments.

If the nature of an uncertain effect is well understood, then it can be properly accounted for during treatment using probability distributions. For instance, if the uncertain effect is interfraction target motion or setup uncertainty, then the spatial probability distribution of the target position can be included in the inverse planning process (Unkelbach and Oelfke [2004], Li and Xing [2000]). For intrafraction organ motion, one can similarly define a probability density function for the displacement of each voxel (Bortfeld et al. [2002]) or a probability mass function for different motion

“states” (Bortfeld et al. [2004], Trofimov et al. [2005]), which can then be used to design the treatment. It is also possible to plan for an uncertain effect even if the underlying probability distribution is not known precisely (Unkelbach and Oelfke [2004, 2005]).

A very common method for dealing with an uncertain effect in that situation is to incorporate a margin around the target, and many studies have considered the optimal design of margins (see for example van Herk et al. [2000], van Herk [2004]). Some margin solutions, however, can be overly conservative, in the sense that they may protect against natures of the uncertain effect (i.e., probability distributions) that are unlikely to be realized. For example, when dealing with interfraction breathing motion, one type of margin solution may be obtained by prescribing the target dose to the union of all tumour-shaped volumes over every possible displacement from inhale to exhale. This solution would then protect against the case where the tumour spends the entire treatment at one displacement; although this may occur for some patients and some treatments, it is highly improbable, and planning for this possibility greatly increases the damage done to healthy tissue.

One methodology that allows for target coverage under uncertainty to be balanced with damage to healthy tissue is robust optimization. In the Operations Research (OR) community, the methodology of robust optimization has been extensively developed (see Ben-Tal et al. [2009] and Bertsimas et al. [2011]) and brought to bear on numerous practical problems outside of the medical physics domain such as structural design (Ben-Tal and Nemirovski [1997]), financial portfolio optimization (Ben-Tal and Nemirovski [1999], Bertsimas and Sim [2004]) and supply chain management (Ben-Tal et al. [2005], Bertsimas and Thiele [2006]). In the robust optimization approach to IMRT, the treatment planner defines an uncertainty set, which is a subset of all possible realizations of the uncertain effect (Bortfeld et al. [2008]). The planner then solves a mathematical program that yields a treatment that is “robust” to the uncertain effect – that is, the treatment meets clinical requirements for all realizations of the uncertain effect in the uncertainty set.

A number of studies have shown the power of this approach. In Chu et al. [2005], tumour position uncertainty due to interfraction motion and setup error was modelled using an ellipsoidal uncertainty set and the authors showed that the resulting robust solution achieves the same level of tumour coverage as a clinical margin solution, in which the volume that receives the prescription dose is expanded to account for changes in position, but with less healthy tissue dose. In Ólafsson and Wright

[2006], the uncertainty due to dose calculation error and interfraction organ motion was also modelled using an ellipsoidal uncertainty set, and the authors showed for a nasopharyngeal case that a robust solution achieves better tumour coverage than the nominal solution (one which assumes a dose matrix known with certainty) and leads to better organ sparing than the margin solution that was clinically prescribed for the case. In Chan et al. [2006] and Bortfeld et al. [2008], uncertainty due to intrafraction breathing motion was modelled using a polyhedral uncertainty set, and it was shown that the robust solution provides improved tumour coverage over the nominal solution, while providing tumour coverage comparable to the margin solution at reduced healthy tissue dose. Robust optimization has also been used in intensity modulated proton therapy (IMPT), a more advanced external beam radiation therapy modality (Unkelbach et al. [2007], Fredriksson et al. [2011]).

2.2 Adaptive radiation therapy

In all of the studies listed above in Section 2.1, the planning problem is solved once, and the resulting treatment is delivered in every fraction, without any further replanning. At the same time, there has also been research done into what is known as adaptive radiation therapy (ART) and in particular, off-line ART. With regard to off-line ART, a variety of procedures have been proposed for how to treat a patient based on information obtained in the most recent fraction. Many studies have proposed using daily computed tomography (CT) images as a form of feedback. For example, Mohan et al. [2005], Lu et al. [2006], Wu et al. [2006] and Wu et al. [2008] all studied schemes where daily CT images are used to periodically assess the treatment plan for errors in dose delivery due to organ deformation and intrafraction motion, and to aid either the re-adjustment or complete re-optimization of the treatment for subsequent fractions. Other imaging methods have also been proposed for off-line ART. For instance, de la Zerda et al. [2007] proposed algorithms for generating new treatment plans from fraction to fraction in response to changes in the patient geometry detected from cone beam CT (CBCT) images, as well as to the cumulative delivered dose. Another example is Yan et al. [1998], who tested an ART method where the daily setup error is measured using portal imaging and is used to correct the treatment if the error becomes sufficiently large. It is worth noting that some studies have considered

treatments where the beamlet intensities are altered in real-time *during* a treatment session, in response to changes in the patient geometry (see for example Mestrovic et al. [2007]). Treatments of this type are known as online ART and are not studied in this thesis.

Some ART studies have also incorporated probabilistic models of uncertainty into their proposed planning methods. Yan et al. [1997a] studied an ART method where the daily treatment target and beam placement variation are measured and used to modify the treatment dose and field margin in each fraction under the assumption that the error is normally distributed. Löf et al. [1998] studied an ART scheme where the daily setup position is measured in every fraction, but assumed that the probabilistic dynamics of the internal organ motion are known beforehand. Rehbinder et al. [2004] applied linear quadratic control to ART planning and showed that using this form of adaptation removes the need for a margin, allowing for safer dose escalation. The problem of ART can also be seen as a problem of sequential decision making under uncertainty and as such, several studies have considered dynamic programming techniques. For example, Sir et al. [2010] showed that an adaptive open-loop feedback control (OLFC) policy achieves better performance than a non-adaptive OLFC policy. Ferris and Voelker [2004] and Deng and Ferris [2008] applied neuro-dynamic programming (NDP) to ART and similarly showed that their NDP policy offers an improvement over a constant, non-adaptive policy.

2.3 Dose-reactive methods

Within the set of methods that can be considered off-line ART methods, a subset of these can be classified as dose-reactive methods. Within the operations research community, several ART methods that can be considered dose-reactive have been developed. As stated above, Ferris and Voelker [2004] and Deng and Ferris [2008] apply neuro-dynamic programming (NDP) to ART, and consider NDP policies constructed from simpler heuristic policies. One such policy is a “reactive” policy that at each stage adjusts the dose to account for how much has been delivered up to that point; our choice of the name “dose-reactive” stems from these studies. The work of Sir et al. [2010] mentioned above considers certainty-equivalent control and open-loop feedback control for performing ART, where prior to each fraction the beamlet

intensities are re-optimized and take into account the dose delivered in the previous fractions.

Within the medical physics community, several papers have proposed incorporating information about the dose delivered in prior fractions into the planning process for subsequent fractions. Birkner et al. [2003] study an ART method where daily CT and portal images are collected, and after a pre-specified number of fractions, the treatment is updated based on the dose that was delivered prior to that fraction; they show on prostate cancer cases that such a method improves tumour coverage and reduces dose to healthy tissues. Wu et al. [2006] study an ART method for prostate cancer where the cumulative dose distribution is calculated on the basis of daily CT images, and errors in the dose are compensated after every week or at the end of treatment; they show that dose compensation allows for better target coverage and reduced dose to healthy tissues such as the bladder and the rectal wall. de la Zerda et al. [2007] consider several algorithms for ART that correct for errors in dose that occur due to set-up error or changes in the patient geometry. Webb [2008] proposes a framework for dealing with variable intrafraction motion where, in each fraction i of an n -fraction treatment, one delivers $1/(n - i + 1)$ of the residual dose distribution – the difference of the dose delivered so far and the target dose distribution. In addition to this, Webb [2008] proposes a framework where in each fraction i , one delivers $1/(n - i + 1)$ of the residual fluence map (beam intensity profile), which is the difference of the full n -fraction fluence map and the $i - 1$ delivered fluence maps, adjusted for the motion. Webb [2008] notes that although this algorithm is better than performing no adaptation of any kind, it is still problematic because if some parts of the tumour are consistently underdosed, then the algorithm will compensate by delivering larger and larger overall doses in future fractions.

2.4 Contribution to the literature

As stated in Chapter 1, this thesis can be thought of as consisting of two parts: the first part, in Chapter 4, which is concerned with developing a method that combines robust optimization with off-line ART; and the second part, in Chapter 5, which is concerned with adaptive and robust methods that are also dose-reactive and capable of correcting for errors in dose realized in earlier fractions.

The contribution of the first part of this thesis to the uncertainty in radiation therapy and ART literatures can be understood by considering the weaknesses of existing off-line ART methods. In particular, the limitation of the majority of algorithms for off-line ART is in how they have approached uncertainty. Many studies have not explicitly included uncertain effects into their treatment planning process. Of the studies that have, most have assumed that realizations of the uncertain effect from fraction to fraction are independent and identically distributed – although this may be appropriate for setup error and patient positioning, it may not be suitable for other types of uncertainty. The most significant assumption present in studies that have incorporated probabilistic models of uncertainty is that the underlying probability distribution of the uncertain effect is known precisely before the start of the treatment. In actuality, the distribution which underlies an uncertain effect for a particular patient in a particular treatment is never known precisely, and as already discussed, the quality of a treatment designed with a specific probability distribution in mind can deteriorate significantly if a different distribution is realized.

The fundamental difference between the approach developed in Chapter 4 and the majority of prior ART work lies in the assumptions that are made about the uncertainty. In our approach, we do not assume that the nature of the uncertainty is known precisely *a priori* nor that it needs to stay constant from fraction to fraction. There has been some prior research into using previous measurements of uncertainty to update an estimate of the uncertain effect, using such methods as Kalman filtering (Yan et al. [1997b], Keller et al. [2003]) and Bayesian updating (Lam et al. [2005], Sir [2007]). Very few studies have considered these methods with treatment re-planning (Yan et al. [1997b], Sir [2007] are exceptions) and these studies assume that the uncertain effect follows a particular distribution. No studies have combined such estimation methods with the robust optimization approach of Chan et al. [2006] and Bortfeld et al. [2008].

We note here that there has been research in the OR community in extending robust optimization methods to multi-stage decision making settings. For example, Ben-Tal et al. [2004] studied the affinely adjustable robust counterpart (AARC) framework, in which decisions in later stages are constrained to be affine functions of previous realizations of the uncertainty; this framework was applied to a multi-period supply chain planning problem in Ben-Tal et al. [2005]. Bertsimas and Thiele [2006] applied the robust optimization methodology of Bertsimas and Sim [2004] to a multi-stage inventory problem. Bertsimas and Caramanis [2007] considered adaptable robust op-

timization through sampling, which allows for decisions in later stages to be piecewise polynomial functions of prior realizations of the uncertainty. Bertsimas and Caramanis [2010] considered a finite adaptation approach, in which decisions in later stages are constrained to be piecewise constant functions of previous realizations of the uncertainty. The method that we present in Chapter 4 differs from these prior methods in that the uncertainty set is adaptively updated in each stage using the most recent measurement of the uncertainty, and the corresponding single-stage robust optimization problem is solved to obtain the decision to implement in that stage.

The second part of this thesis is concerned with dose-reactive methods, and its contribution to the existing literature is multi-faceted. The dose-reactive methods we present in Chapter 5 build on the method described in Chapter 4 and thus (1) are specialized for lung cancer and breathing motion uncertainty, (2) do not make any distributional assumptions about the uncertainty in each fraction, and (3) do not assume that the nature of the uncertainty stays the same throughout the treatment. A notable difference between our work and the methods considered in Ferris and Voelker [2004] and Deng and Ferris [2008] is the decision variable in each stage: in the latter two studies, the decision for each fraction is the dose distribution to be delivered, while in our method, the decision for each fraction is the collection of beamlet intensities to use in that fraction (from which the dose distribution is determined via a linear transformation). By focusing on beamlet intensities, we allow only for dose distributions that are physically deliverable. More importantly, as we will show in Section 5.2, the use of beamlet intensities uncovers an important insight into the performance of dose-reactive methods that stems directly from the ability, or rather the inability, to deliver highly homogeneous tumour dose distributions. With respect to Webb [2008], a major difference between the method presented therein and the methods that we study here is that our methods also update an estimate of the uncertainty by updating an uncertainty set, whereas the method in Webb [2008] always assumes the tumour to be in a fixed reference position. As we will show in our computational study, the uncertainty set adaptation aspect of our method is able, to some extent, to mitigate the explosive growth in daily underdose and overdose that is alluded to in Webb [2008].

Perhaps the most notable difference between our work and extant research in this area is that, while prior dose-reactive studies focused on the performance of their methods in terms of cumulative dose at the end of the treatment horizon, we also provide detailed results on the daily dose performance of dose-reactive and non-reactive

methods. It is important to consider how much dose is delivered on a per fraction basis because interfractional variations in dose, as well as intrafractional variations (i.e., heterogeneity in the dose distribution that is achieved in each day), can have a significant impact on the final treatment outcome (Ebert [2000]). Our computational results show that dose-reactive methods lead to higher levels of daily underdose and overdose, which increases the chance of treatment failure, even when the final cumulative dose is acceptable. To the best of our knowledge, no previous work has provided the same kind of results on the daily dose performance of an ART method, nor has highlighted the characteristic behavior of dose-reactive methods in this respect.

Chapter 3

Static robust IMRT optimization

In this chapter, we review the robust optimization approach of Chan et al. [2006] and Bortfeld et al. [2008]. In their approach, the planner decides on an uncertainty set before treatment and solves the robust optimization problem corresponding to that uncertainty set. The optimal solution of the robust problem is a vector that specifies the intensity of each beamlet in the ensemble of beams to be used for treatment. These optimized beamlet intensities are robust to the uncertain effect whenever it takes on values from the uncertainty set: whenever the uncertain effect realizes a value inside the uncertainty set while the patient is being irradiated according to the optimized intensities, the dose to every part of the tumour will be within the prescribed bounds. We will refer to this approach as the *static* robust optimization approach, as the beamlet intensity vector is fixed and does not change over the course of the treatment (in each fraction, $1/n$ of the intensity vector is delivered, where n is the number of fractions). For lung cancer, the primary uncertainty comes from breathing motion. More specifically, the uncertain effect is modelled using the patient's breathing motion probability mass function (PMF). The breathing motion PMF for a patient in a particular fraction specifies the proportion of time the patient spends in each of a finite number of breathing motion states during that fraction. Each breathing motion state corresponds to a particular displacement or snapshot of the patient geometry during the patient's respiratory cycle. For example, a possible set of motion states may be in-hale, mid-ventilation and exhale. The uncertainty set is then a set of breathing motion PMFs that we anticipate could be realized during treatment. We note here that in the method of Bortfeld et al. [2008], the set of breathing motion states is determined determined from pre-treatment breathing-correlated 4D computed tomography (4D-

CT) images. This assumes that the amplitude of the tumour motion observed in these pre-treatment images represents the full extent of the tumour motion, and that the amplitude of the motion that will be observed during the ensuing treatment sessions will not exceed these limits. This potential issue can be circumvented by having the patient take deep breaths during the pre-treatment 4D-CT imaging session, allowing the treatment planner to more accurately construct the set of breathing motion states.

3.1 Model

We now define the static robust IMRT optimization problem. Let \mathcal{B} be the set of beamlets and let w_b denote the intensity of beamlet $b \in \mathcal{B}$; the w_b variables are the decision variables of the robust optimization problem. The patient geometry (the healthy organs and tumour tissue) is divided into three dimensional voxels, or volume elements. Let \mathcal{V} denote the set of all voxels and \mathcal{T} denote the set of tumour voxels. Let X denote the finite set of breathing motion states. Each breathing motion state is associated with a particular displacement of the patient geometry (the voxels), so to each beamlet $b \in \mathcal{B}$, breathing motion state $x \in X$ and voxel $v \in \mathcal{V}$, we associate a dose deposition coefficient $\Delta_{v,x,b}$. This coefficient specifies the amount of dose deposited in voxel v when the patient is in breathing motion state x and the beamlet b is at unit intensity. Let $p(x)$ be the probability of the patient being in motion state x at any given time during a treatment session. The dose that is delivered to voxel v is then given by

$$\sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b,$$

which is the sum of the doses to voxel v that would be delivered under each breathing motion state weighted by the corresponding proportions of time spent in those states. Due to the use of beamlet intensities, this model does not account for the interplay effect between the motion of the MLC leaves and the tumour motion (see, for example, Yu et al. [1998] and Bortfeld et al. [2002]); the impact of such an effect, however, is small in practice (Bortfeld et al. [2004]). For each tumour voxel $v \in \mathcal{T}$, let θ_v be the prescribed minimum dose and let $\gamma\theta_v$ be the prescribed maximum dose, where $\gamma \geq 1$.

Let \mathcal{P} be the set of all PMFs on the finite set X , defined as

$$\mathcal{P} = \left\{ \mathbf{p} \in \mathbb{R}^{|X|} \mid \text{For each } x \in X, p(x) \geq 0; \sum_{x \in X} p(x) = 1 \right\}.$$

The set \mathcal{P} defines the $(|X| - 1)$ -dimensional unit simplex. Let $P \subseteq \mathcal{P}$ be the uncertainty set, which is a bounded polyhedron defined by a lower bound vector ℓ and an upper bound vector \mathbf{u} :

$$P = \{ \mathbf{p} \in \mathcal{P} \mid \text{For each } x \in X, \ell(x) \leq p(x) \leq u(x) \}.$$

With these definitions, the static robust problem can be written as

$$\begin{aligned} & \text{minimize} && \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{p}(x) w_b \\ & \text{subject to} && \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b \geq \theta_v, \quad \forall v \in \mathcal{T}, \mathbf{p} \in P, \\ & && \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b \leq \gamma \theta_v, \quad \forall v \in \mathcal{T}, \mathbf{p} \in P, \\ & && w_b \geq 0, \quad \forall b \in \mathcal{B}, \end{aligned} \tag{3.1}$$

where $\bar{\mathbf{p}}$ is some nominal PMF chosen by the treatment planner before the start of treatment as approximately representative of the patient's overall breathing pattern. It is assumed that problem (3.1) is feasible for the choice of $P = \mathcal{P}$, in which case problem (3.1) will be feasible for any uncertainty set $P \subseteq \mathcal{P}$; in Appendix A we provide specific technical assumptions that ensure problem (3.1) is feasible and that are necessary for the theoretical analysis of the adaptive robust method of Chapter 4. Though not currently written as such, problem (3.1) can be converted to a finite linear program and solved using a standard optimization solver; we provide the corresponding formulation in Section 3.2. The optimal solution of (3.1) is a beamlet intensity vector \mathbf{w}^* that meets minimum and maximum tumour dose requirements no matter which breathing motion PMF $\mathbf{p} \in P$ is realized, while ensuring the lowest possible total dose to the patient under the nominal PMF. In the case that $\mathbf{p} \notin P$, it was shown empirically in Bortfeld et al. [2008] that \mathbf{w}^* remains feasible with high probability. Due to the linearity of this model and the use of polyhedral uncertainty sets, this model is capable of accomodating robust versions of many other types of clinically relevant

constraints, such as partial volume constraints based on conditional-value-at-risk (see Romeijn et al. [2006]) or constraints involving α EUD (Thieke et al. [2002]), which is a linear approximation to the popular equivalent uniform dose (EUD) metric. For simplicity, we consider only lower and upper dose bounds on the tumour.

There are two special cases of (3.1) that will be referred to later. The first is the case when P consists of a single PMF. The resulting problem is called the *nominal problem*; when $P = \{\mathbf{p}\}$, we refer to problem (3.1) as the *nominal problem with respect to \mathbf{p}* . The solution of the nominal problem is the least conservative treatment we can deliver, as we are assuming that the PMF that will be realized during treatment will be the exact PMF that we are guarding against. The second is the case when P consists of all possible PMFs on X , that is, $P = \mathcal{P}$; this problem is called the *margin problem*. The solution of the margin problem is the most conservative treatment that we can deliver, because we assume that any PMF is possible. These two types of formulations represent two extremes of a continuum of robustness (Chan et al. [2006]), with the nominal problem being the least robust and the margin problem being the most robust. The robust problem with a general uncertainty set (which is neither a singleton nor \mathcal{P}) is somewhere in between. With increasing levels of robustness, characterized by larger uncertainty sets, the level of tumour coverage increases, but at the cost of increased dose to healthy tissue.

3.2 Solving the robust problem

Problem (3.1) generally has an uncountable number of constraints because the set P is usually a polyhedron with more than one element. It can be transformed into an equivalent linear program of finite size using duality (Bortfeld et al. [2008]):

$$\begin{aligned}
& \text{minimize} && \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{p}(x) w_b \\
& \text{subject to} && q_v + \sum_{x \in X} \ell(x) r_{v,x} - \sum_{x \in X} u(x) s_{v,x} \geq \theta_v, \quad \forall v \in \mathcal{T}, \\
& && q_v + r_{v,x} - s_{v,x} = \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b, \quad \forall v \in \mathcal{T}, x \in X, \\
& && q'_v - \sum_{x \in X} \ell(x) r'_{v,x} + \sum_{x \in X} u(x) s'_{v,x} \leq \gamma \theta_v, \quad \forall v \in \mathcal{T},
\end{aligned}$$

$$q_v - r_{v,x} + s_{v,x} = \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b, \quad \forall v \in \mathcal{T}, x \in X,$$

$$q_v, q'_v \text{ free}, \quad \forall v \in \mathcal{T},$$

$$r_{v,x}, r'_{v,x}, s_{v,x}, s'_{v,x} \geq 0, \quad \forall v \in \mathcal{T}, x \in X,$$

$$w_b \geq 0, \quad \forall b \in \mathcal{B}.$$

Chapter 4

Adaptive and robust IMRT optimization

In the previous chapter, we discussed the robust optimization approach of Chan et al. [2006] and Bortfeld et al. [2008], which is a static approach, i.e., the beamlet intensities are fixed throughout the treatment. In this chapter, we develop a robust optimization method that is also adaptive, and allows the beamlet intensities to change by accounting for changes in the uncertainty. Our contributions are as follows:

1. We develop a method that is the first to combine robust optimization with adaptive radiation therapy. This method has modest computational requirements and can be easily generalized to other types of cancer and other types of uncertainty. Most importantly, this method can be readily implemented with resources available in most clinical environments.
2. We prove an asymptotic optimality result: if the patient breathing pattern converges over the treatment course, then the dose distribution produced by our adaptive robust method converges to an “ideal” set of dose distributions that exhibit no tumour underdose or overdose and low healthy tissue dose. This result suggests that our method should 1) outperform the non-adaptive robust method, 2) perform almost as well as optimal “prescient” solutions, which correctly anticipate the patient’s breathing over the entire treatment, and 3) be relatively insensitive to the choice of initial uncertainty set.
3. Using clinical patient data, we computationally demonstrate the performance of

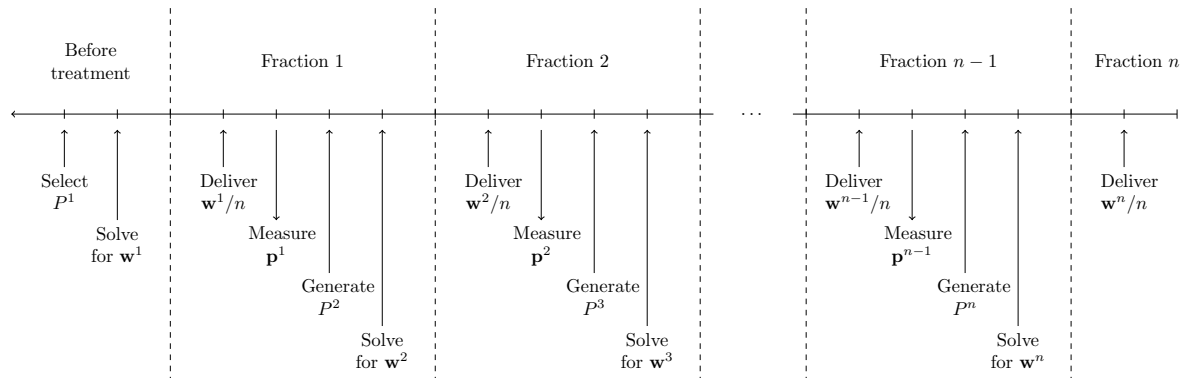


Figure 4.1: Graphical representation of the steps of the adaptive robust method.

this method and verify the behavior predicted by our theoretical development. These computational results are presented in Section 4.3.

This chapter is organized as follows. We describe our method in Section 4.1. We present our theoretical analysis in Section 4.2. We present and discuss the results of our computational study in Section 4.3. Finally, in Section 4.4, we provide some concluding remarks and directions for future work in this area.

1

4.1 Method

In the adaptive robust approach that we propose, we allow the beamlet intensities to be re-optimized from fraction to fraction in the following way. On a given day, the patient is irradiated with a beamlet intensity vector obtained by solving (3.1) with the current uncertainty set as input. During or immediately after the fraction is delivered, the uncertain effect is measured. This measurement, together with the current uncertainty set, is used to generate a new uncertainty set. The robust problem is then solved with this new uncertainty set, leading to a new beamlet intensity vector to be used on the next day. This process is repeated for each fraction until the end of the treatment. This procedure, adapted for the form of breathing motion uncertainty that we study, is presented as Algorithm 1. Figure 4.1 shows the steps of the adaptive robust method graphically.

Although in theory one could measure the patient's PMF before the delivery of a fraction and plan the fraction using the resulting PMF, this approach would constrain the amount of time that would be available for treatment planning and quality assurance, as the patient would have to wait between measurement and fraction delivery.

By measuring the patient's PMF during or after the delivery of the fraction, this temporal constraint can be avoided by having treatment planning and quality assurance done offline between treatment fractions, which would be more desirable from an operational point of view. Furthermore, we will show in Section 4.2 that these two methods have the same asymptotic performance.

Algorithm 1 Adaptive robust optimization method

Require: Total number of fractions n , initial uncertainty set P^1

- 1: Initialize $i = 1$
 - 2: Solve problem (3.1) with $P = P^1$ to obtain a beamlet intensity vector \mathbf{w}^1
 - 3: Deliver \mathbf{w}^1/n to the patient
 - 4: **while** $i \leq n - 1$ **do**
 - 5: Measure the patient's breathing motion and construct the PMF \mathbf{p}^i
 - 6: Generate the new uncertainty set P^{i+1} from P^i, \mathbf{p}^i
 - 7: Solve problem (3.1) with $P = P^{i+1}$ to obtain a new beamlet intensity vector \mathbf{w}^{i+1}
 - 8: Deliver \mathbf{w}^{i+1}/n to the patient
 - 9: Set $i = i + 1$
 - 10: **end while**
-

4.1.1 Update algorithms

One major step in the adaptive robust approach specified as Algorithm 1 is the generation of the new uncertainty set from the old uncertainty set and the most recent measurement of the patient's breathing motion PMF. We refer to any procedure that can be used to perform this step as an *uncertainty set update algorithm*. In our computational study we consider two such update algorithms, exponential smoothing and running average, though our theoretical analysis generalizes to a larger class of update algorithms.

Recall that we define an uncertainty set P in terms of a lower bound vector ℓ and an upper bound vector \mathbf{u} . Therefore, the uncertainty set for fraction i , denoted P^i , is defined by lower and upper bound vectors ℓ^i and \mathbf{u}^i , respectively. The update algorithms we present will determine lower and upper bound vectors $\ell^{i+1}, \mathbf{u}^{i+1}$ from the most recent vectors ℓ^i, \mathbf{u}^i and the most recent PMF \mathbf{p}^i .

The *exponential smoothing* update algorithm generates a new uncertainty set by taking a convex combination of the most recent lower and upper bound vectors with the

most recent PMF. Specifically, given ℓ^i , \mathbf{u}^i and \mathbf{p}^i for day i , ℓ^{i+1} and \mathbf{u}^{i+1} are defined as

$$\ell^{i+1}(x) = (1 - \alpha)\ell^i(x) + \alpha p^i(x), \quad (4.1)$$

$$u^{i+1}(x) = (1 - \alpha)u^i(x) + \alpha p^i(x), \quad (4.2)$$

for each $x \in X$, where $\alpha \in [0, 1]$.

The *running average* update algorithm generates a new uncertainty set by averaging the PMFs that have been observed so far, together with the very first lower and upper bound vectors. Specifically, given ℓ^1 , \mathbf{u}^1 and $\mathbf{p}^1, \dots, \mathbf{p}^i$ on day i , ℓ^{i+1} and \mathbf{u}^{i+1} are defined as

$$\ell^{i+1}(x) = \frac{1}{i+1} \left(\ell^1(x) + \sum_{j=1}^i p^j(x) \right),$$

$$u^{i+1}(x) = \frac{1}{i+1} \left(u^1(x) + \sum_{j=1}^i p^j(x) \right).$$

for each $x \in X$. These expressions can be rewritten as functions of only $\ell^i(x)$, $u^i(x)$ and $p^i(x)$:

$$\ell^{i+1}(x) = \frac{i}{i+1}\ell^i(x) + \frac{1}{i+1}p^i(x), \quad (4.3)$$

$$u^{i+1}(x) = \frac{i}{i+1}u^i(x) + \frac{1}{i+1}p^i(x). \quad (4.4)$$

4.1.2 Prescient solutions

Because the update algorithms are myopic and are not guaranteed to correctly anticipate the patient's breathing motion PMF in every fraction, it is important to be able to measure the loss of optimality with respect to a "prescient" solution – one that would be generated if we knew $\mathbf{p}^1, \dots, \mathbf{p}^n$ ahead of time.

The best possible treatment – one that minimizes the total dose delivered to the patient after n fractions under the nominal PMF $\bar{\mathbf{p}}$, subject to the constraint that the dose to every tumour voxel by the end of the treatment is between the prescribed

bounds – is obtained by solving the following linear program:

$$\begin{aligned}
& \text{minimize} && \sum_{i=1}^n \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{p}(x) \frac{w_b^i}{n} \\
& \text{subject to} && \sum_{i=1}^n \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n} \geq \theta_v, \quad \forall v \in \mathcal{T}, \\
& && \sum_{i=1}^n \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n} \leq \gamma \theta_v, \quad \forall v \in \mathcal{T}, \\
& && w_b^i \geq 0, \quad \forall b \in \mathcal{B}, i \in \{1, \dots, n\}.
\end{aligned} \tag{4.5}$$

We refer to the solution of problem (4.5) as the *overall prescient* solution. While the overall prescient solution is in theory the best possible treatment for this problem, it is possible that this solution delivers little to no dose on certain days. This possibility exists because there is no constraint that forces the delivery of a minimum amount of dose on every day and indeed, we can construct simple examples where all of the prescription dose is delivered to a single voxel in one fraction and zero dose is delivered in all of the other fractions; such solutions defeat the purpose of fractionation. Although the overall prescient solution is the best possible solution, the fact that the dose delivered on a given day may be very low reduces its value as a benchmark for our adaptive and static robust treatments.

For these reasons, we consider two other kinds of prescient solutions. The first type of prescient solution that we define is the *daily prescient* solution. It is obtained by solving the following linear program:

$$\begin{aligned}
& \text{minimize} && \sum_{i=1}^n \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{p}(x) \frac{w_b^i}{n} \\
& \text{subject to} && \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n} \geq \frac{\theta_v}{n}, \quad \forall v \in \mathcal{T}, i \in \{1, \dots, n\}, \\
& && \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n} \leq \frac{\gamma \theta_v}{n}, \quad \forall v \in \mathcal{T}, i \in \{1, \dots, n\}, \\
& && w_b^i \geq 0, \quad \forall b \in \mathcal{B}, i \in \{1, \dots, n\}.
\end{aligned} \tag{4.6}$$

It is straightforward to see that the solution to this problem is obtained by setting each w^i to the solution of problem (3.1) with $P = P^i = \{\mathbf{p}^i\}$. Using this treatment, the cumulative dose to every tumour voxel v is within the prescribed bounds θ_v and $\gamma \theta_v$.

Furthermore, because each \mathbf{w}^i is obtained from a nominal problem, the overall dose delivered to the patient should be low compared to other treatments that also deliver the minimum required dose to every tumour voxel. Unlike the overall prescient treatment, this treatment is guaranteed to deliver a dose of at least θ_v/n to every tumour voxel v in every fraction.

The second type of prescient solution that we define is the *average prescient* solution. It is obtained by solving the following linear program:

$$\begin{aligned}
& \text{minimize} && \sum_{i=1}^n \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{p}(x) \frac{w_b^i}{n} \\
& \text{subject to} && \sum_{i=1}^n \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n} \geq \theta_v, \quad \forall v \in \mathcal{T}, \\
& && \sum_{i=1}^n \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n} \leq \gamma \theta_v, \quad \forall v \in \mathcal{T}, \\
& && \mathbf{w}^i = \mathbf{w}^j, \quad \forall i, j \in \{1, \dots, n\}, i \neq j, \\
& && w_b^i \geq 0, \quad \forall b \in \mathcal{B}, i \in \{1, \dots, n\}.
\end{aligned} \tag{4.7}$$

This formulation is simply problem (4.5) with the additional constraint that the beamlet intensity vector be the same in every fraction. The ‘‘average’’ in the name comes from the fact (easily verifiable) that the common intensity vector \mathbf{w} is the solution of the nominal problem with $P = \{\mathbf{p}_{\text{avg}}\}$, where $\mathbf{p}_{\text{avg}} = 1/n \cdot \sum_{i=1}^n \mathbf{p}^i$. Like the daily prescient solution, the final dose to every tumour voxel is within the prescribed bounds for that voxel. Furthermore, because the intensity vector of every fraction is obtained from a nominal problem, the overall dose delivered to the patient should be low compared to other treatments that also deliver the minimum required dose to every tumour voxel. While each tumour voxel $v \in \mathcal{T}$ is not guaranteed to receive θ_v/n in each fraction, the fact that $\mathbf{w}^i = \mathbf{w}^j$ for any two fractions i and j ensures that the dose delivered by this solution in any fraction cannot be too low.

4.2 Asymptotic optimality of the adaptive robust method

In this section, we develop our core result: if the sequence of breathing motion PMFs $(\mathbf{p}^i)_{i=1}^{\infty}$ converges to \mathbf{p}^* and the uncertainty set update algorithm belongs to a special class of update algorithms, then as the number of fractions n tends to infinity, the dose

distribution eventually enters the epsilon neighborhood of a set \mathbf{D} of “ideal” dose distributions. The set \mathbf{D} is the set of dose distributions obtained when any optimal solution to the nominal problem with respect to \mathbf{p}^* is delivered while the patient is actually breathing according to \mathbf{p}^* .

The arc of reasoning we will take to prove this result is as follows. We begin by defining an alternative robust problem, where instead of enforcing the minimum and maximum tumour dose constraints for a (possibly infinite) set of PMFs P , we enforce those constraints at finitely many PMFs $\mathbf{q}^1, \dots, \mathbf{q}^M$. This problem is equivalent to the original robust problem under a simple condition. Next, we show that if the defining PMFs of the alternative robust problem, $\mathbf{q}^1, \dots, \mathbf{q}^M$, all converge to a single PMF \mathbf{p}^* in the usual mathematical sense, then the feasible regions of the alternative robust problems along the way converge to the feasible region of the alternative robust problem with all M defining PMFs set to \mathbf{p}^* . We then build on this result to show that the corresponding sets of optimal solutions also converge.

At this point, we define the notion of a convex-convergent update algorithm – any update algorithm where the update is a convex combination of the most recent lower/upper bound vector with the most recent PMF, and where the lower/upper bound vector sequences inherit the convergence of the underlying PMF sequence. We show that if the sequence of breathing motion PMFs $(\mathbf{p}^i)_{i=1}^{\infty}$ converges to a PMF \mathbf{p}^* and the sequences of lower and upper bound vectors $(\ell^i)_{i=1}^{\infty}$ and $(\mathbf{u}^i)_{i=1}^{\infty}$ are obtained by any convex-convergent update algorithm, then the optimal solution sets of the original robust problems enter the epsilon neighborhood of the set of optimal solutions to the nominal problem with respect to \mathbf{p}^* . Finally, we prove our core result of convergence of the dose distribution. We also establish the same convergence result for the dose distributions obtained from the daily and average prescient methods.

We begin with the definition and properties of the alternative robust problem. All of the proofs, auxiliary results and assumptions are stated in Appendix A. The norm $\|\cdot\|$ is the 1-norm on the appropriate finite dimensional real vector space, although technically any p -norm can be used. The set \mathbb{Z}_+ is the set of positive integers.

4.2.1 An alternative robust problem

Given M PMFs $\mathbf{q}^1, \dots, \mathbf{q}^M$, the alternative robust problem is defined as

$$\begin{aligned}
& \text{minimize} && \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{p}(x) w_b \\
& \text{subject to} && \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} q^j(x) w_b \geq \theta_v, \quad \forall v \in \mathcal{T}, j \in \{1, \dots, M\}, \\
& && \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} q^j(x) w_b \leq \gamma \theta_v, \quad \forall v \in \mathcal{T}, j \in \{1, \dots, M\}, \\
& && w_b \geq 0, \quad \forall b \in \mathcal{B}.
\end{aligned} \tag{4.8}$$

We will refer to the M PMFs $\mathbf{q}^1, \dots, \mathbf{q}^M$ as the *defining PMFs* of the corresponding problem (4.8). The convex hull of these PMFs is denoted as $\text{conv} \{\mathbf{q}^1, \dots, \mathbf{q}^M\}$. Under a very simple condition, the alternative robust problem is equivalent to the original robust problem.

Lemma 1 *If $P = \text{conv} \{\mathbf{q}^1, \dots, \mathbf{q}^M\}$, then problems (3.1) and (4.8) are equivalent.*

Now, let $H(\mathbf{q}^1, \dots, \mathbf{q}^M)$ denote the feasible region of the alternative robust problem (4.8) with the M defining PMFs $\mathbf{q}^1, \dots, \mathbf{q}^M$.

Proposition 1 *Suppose that the i th instance of the alternative robust problem (4.8) is defined by the M defining PMFs $\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}$, for all $i \in \mathbb{Z}_+$. If $\mathbf{q}^{j,i} \rightarrow \mathbf{p}^*$ as $i \rightarrow \infty$ for all $j \in \{1, \dots, M\}$, then*

$$\lim_{i \rightarrow \infty} H(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}) = H(\mathbf{p}^*, \dots, \mathbf{p}^*),$$

where the notion of set convergence is that of Dantzig et al. [1967].

In other words, if the defining PMFs of (4.8) all converge to some PMF \mathbf{p}^* , then the feasible region of (4.8) with those defining PMFs converges to the feasible region of (4.8) with all defining PMFs set to \mathbf{p}^* . Note that when all M defining PMFs are set to \mathbf{p}^* , problem (4.8) is equivalent to the nominal problem with respect to \mathbf{p}^* . The proof of Proposition 1 uses Theorem II.2.2 of Dantzig et al. [1967] and the fact that the affine functions that define the $H(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})$ sets converge to the affine function that defines $H(\mathbf{p}^*, \dots, \mathbf{p}^*)$.

Next, we show that the sets of optimal solutions converge. Define the function $Z : \mathbb{R}^{|\mathcal{B}|} \rightarrow \mathbb{R}$ as

$$Z(\mathbf{w}) = \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{p}(x) w_b;$$

Z is the objective function of both problems (3.1) and (4.8). Suppose that V is a subset of \mathbb{R}^n and $\phi : \mathbb{R}^n \rightarrow \mathbb{R}$ is a function. Following Dantzig et al. [1967], we define the set $M(\phi | V)$ as

$$M(\phi | V) = \{x \in V \mid \phi(x) = \inf\{\phi(y) \mid y \in V\}\}.$$

When the infimum can be attained, $M(\phi | V)$ is the set of optimal solutions to the constrained minimization problem $\min_{x \in V} \phi(x)$. We define the *epsilon neighborhood* $U(V, \epsilon)$ of a subset $V \subseteq \mathbb{R}^n$ as

$$U(V, \epsilon) = \bigcup_{x \in V} B(x, \epsilon),$$

where $B(x, \epsilon)$ is the open ball of radius ϵ about x ($B(x, \epsilon) = \{x' \in \mathbb{R}^n \mid \|x - x'\| < \epsilon\}$).

The next result states that, if the defining PMFs of problem (4.8) all converge to a PMF \mathbf{p}^* , then the corresponding optimal solution sets eventually enter the epsilon neighborhood of the optimal solution set of problem (4.8) with all of the defining PMFs set to \mathbf{p}^* (which, by our note above, is the optimal solution set of the nominal problem with respect to \mathbf{p}^*).

Proposition 2 *Suppose that for every $j \in \{1, \dots, M\}$, $\mathbf{q}^{j,i} \rightarrow \mathbf{p}^*$ as $i \rightarrow \infty$. Then for every $\epsilon > 0$, there exists an $N \in \mathbb{Z}_+$ such that for $i > N$, $M(Z \mid H(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})) \subseteq U(M(Z \mid H(\mathbf{p}^*, \dots, \mathbf{p}^*)), \epsilon)$.*

The proof of this proposition uses Theorem I.3.3 of Dantzig et al. [1967].

4.2.2 Convergence of optimal solution sets and dose distributions under convex-convergent update algorithms

Having developed convergence results for the alternative robust problem (4.8), we now connect these results to the robust problem (3.1). We say that an uncertainty set update algorithm is a *convex-convergent update algorithm* if it satisfies two conditions:

1. For every $i \in \mathbb{Z}_+$, ℓ^{i+1} and \mathbf{u}^{i+1} can be written as convex combinations of \mathbf{p}^i with

ℓ^i and \mathbf{u}^i respectively; that is, there exists an $\alpha_i \in [0, 1]$ such that

$$\begin{aligned}\ell^{i+1} &= (1 - \alpha_i)\ell^i + \alpha_i\mathbf{p}^i, \\ \mathbf{u}^{i+1} &= (1 - \alpha_i)\mathbf{u}^i + \alpha_i\mathbf{p}^i.\end{aligned}$$

2. If $\mathbf{p}^i \rightarrow \mathbf{p}^*$ as $i \rightarrow \infty$, then $\ell^i \rightarrow \mathbf{p}^*$ and $\mathbf{u}^i \rightarrow \mathbf{p}^*$ as $i \rightarrow \infty$.

In Proposition 6 in Appendix A, we verify that both the exponential smoothing and running average update algorithms are convex-convergent.

Let $\mathbf{w}^*(\ell, \mathbf{u})$ denote the set of optimal solutions for the robust problem (3.1) with its uncertainty set P defined by lower and upper bound vectors ℓ and \mathbf{u} . Let $\mathbf{w}^*(\mathbf{p}^*)$ denote the set of optimal solutions for the nominal problem with respect to \mathbf{p}^* . The next result states that $\mathbf{w}^*(\ell^i, \mathbf{u}^i)$ eventually enters the epsilon neighborhood of $\mathbf{w}^*(\mathbf{p}^*)$.

Theorem 1 *Let $(\mathbf{p}^i)_{i=1}^\infty$ be a sequence of PMFs that converges to \mathbf{p}^* . Let $(\ell^i)_{i=1}^\infty$ and $(\mathbf{u}^i)_{i=1}^\infty$ be lower and upper bound sequences generated from $(\mathbf{p}^i)_{i=1}^\infty$ by any convex-convergent update algorithm. Then for every $\epsilon > 0$, there exists an $N \in \mathbb{Z}_+$ such that for all $i > N$, $\mathbf{w}^*(\ell^i, \mathbf{u}^i) \subseteq U(\mathbf{w}^*(\mathbf{p}^*), \epsilon)$.*

The key insight required in the proof is that the robust problem in fraction i with P^i specified by ℓ^i and \mathbf{u}^i is equivalent to the alternative robust problem with defining PMFs $\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}$, where $\mathbf{q}^{j,i}$ is obtained by repeatedly applying the update algorithm on $\mathbf{q}^{j,1}$ with the observed PMFs $\mathbf{p}^1, \dots, \mathbf{p}^{i-1}$. This equivalence holds because of Lemma 1 and the fact that modifying an uncertainty set P by applying the update algorithm to ℓ and \mathbf{u} is equivalent to applying the update algorithm to the M PMFs $\mathbf{q}^1, \dots, \mathbf{q}^M$ whose convex hull is P (we formally state and prove this fact in Appendix A as Lemma 2).

We now arrive at our core result. Define the set \mathbf{D} , the set of dose distributions when \mathbf{p}^* is realized while $\mathbf{w} \in \mathbf{w}^*(\mathbf{p}^*)$ is being delivered, as

$$\mathbf{D} = \{\mathbf{d} \in \mathbb{R}^{|\mathcal{V}|} \mid \mathbf{d} = \Delta\mathbf{p}^*\mathbf{w} \text{ for some } \mathbf{w} \in \mathbf{w}^*(\mathbf{p}^*)\}$$

where the product $\Delta\mathbf{p}\mathbf{w}$ is defined as

$$\Delta\mathbf{p}\mathbf{w} = \left[\sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b \right]_{v \in \mathcal{V}}.$$

For an n fraction treatment during which the intensity vectors were $\mathbf{w}^1, \dots, \mathbf{w}^n$ and the realized PMFs were $\mathbf{p}^1, \dots, \mathbf{p}^n$, the final dose distribution is given by $1/n \cdot \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i$.

Theorem 2 (*Adaptive robust dose convergence.*) *Let $(\mathbf{p}^i)_{i=1}^\infty$ be a sequence of PMFs that converges to \mathbf{p}^* . Let $(\ell^i)_{i=1}^\infty$ and $(\mathbf{u}^i)_{i=1}^\infty$ be lower and upper bound sequences generated from $(\mathbf{p}^i)_{i=1}^\infty$ by a convex-convergent update algorithm. For each $i \in \mathbb{Z}_+$, let $\mathbf{w}^i \in \mathbf{w}^*(\ell^i, \mathbf{u}^i)$. Then for every $\epsilon > 0$, there exists an $N \in \mathbb{Z}_+$ such that for all $n > N$,*

$$\frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i \in U(\mathbf{D}, \epsilon).$$

The theorem states that the dose distribution obtained via any convex-convergent update algorithm approaches a set of optimal dose distributions \mathbf{D} . The intuition behind this result is that as $i \rightarrow \infty$, \mathbf{w}^i is increasingly similar to some solution in $\mathbf{w}^*(\mathbf{p}^*)$, and \mathbf{p}^i is increasingly similar to \mathbf{p}^* . Therefore, as n tends to infinity, each term in the tail of the sum $\sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i$ becomes increasingly similar to some $\mathbf{d} \in \mathbf{D}$.

This set \mathbf{D} of optimal dose distributions is the same set that is approached by both the daily prescient and average prescient dose distributions.

Theorem 3 *Let $(\mathbf{p}^i)_{i=1}^\infty$ be a sequence of PMFs that converges to \mathbf{p}^* .*

(a) (*Daily prescient dose convergence.*) *For each $i \in \mathbb{Z}_+$, let $\mathbf{w}^i \in \mathbf{w}^*(\mathbf{p}^i, \mathbf{p}^i)$. Then for every $\epsilon > 0$, there exists an $N \in \mathbb{Z}_+$ such that for all $n > N$,*

$$\frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i \in U(\mathbf{D}, \epsilon).$$

(b) (*Average prescient dose convergence.*) *For each $n \in \mathbb{Z}_+$, let $\mathbf{w}^n \in \mathbf{w}^*(1/n \cdot \sum_{i=1}^n \mathbf{p}^i, 1/n \cdot \sum_{i=1}^n \mathbf{p}^i)$. Then for every $\epsilon > 0$, there exists an $N \in \mathbb{Z}_+$ such that for all $n > N$,*

$$\frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^n \in U(\mathbf{D}, \epsilon).$$

When the nominal problem with respect to \mathbf{p}^* has a unique optimal solution, the difference between the adaptive robust and either prescient dose distribution tends to zero.

Corollary 1 *Let $(\mathbf{p}^i)_{i=1}^\infty$ be a sequence of PMFs that converges to \mathbf{p}^* . For each $i \in \mathbb{Z}_+$, let $\mathbf{w}_{\text{AR}}^i \in \mathbf{w}^*(\ell^i, \mathbf{u}^i)$ (the adaptive robust intensity vector for fraction i), where $(\ell^i)_{i=1}^\infty$ and $(\mathbf{u}^i)_{i=1}^\infty$ are obtained by a convex-convergent update algorithm, and let $\mathbf{w}_{\text{DP}}^i \in \mathbf{w}^*(\mathbf{p}^i, \mathbf{p}^i)$ (the daily prescient intensity vector for fraction i). For each $n \in \mathbb{Z}_+$, let $\mathbf{w}_{\text{AP}}^n \in \mathbf{w}^*(1/n \cdot \sum_{i=1}^n \mathbf{p}^i, 1/n \cdot \sum_{i=1}^n \mathbf{p}^i)$ (the average prescient intensity vector for fraction i). If the nominal problem with respect to \mathbf{p}^* has a unique optimal solution (i.e., $\mathbf{w}^*(\mathbf{p}^*)$ is a singleton), then*

$$\lim_{n \rightarrow \infty} \left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}_{\text{AR}}^i - \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}_{\text{DP}}^i \right\| = \lim_{n \rightarrow \infty} \left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}_{\text{AR}}^i - \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}_{\text{AP}}^n \right\| = 0.$$

4.2.3 Three theoretical insights

These theoretical results provide us with three insights about the limiting behavior of the adaptive robust method. The first insight is that adaptive robust methods dominate static robust treatments. To see this, suppose that \mathbf{w} is a solution of the robust problem with uncertainty set P . The static robust dose distribution after n fractions is given by

$$\frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w} = \Delta \sum_{i=1}^n \frac{\mathbf{p}^i}{n} \mathbf{w}.$$

If $(\mathbf{p}^i)_{i=1}^\infty$ converges to \mathbf{p}^* , then this dose distribution converges to $\Delta \mathbf{p}^* \mathbf{w}$. Whether or not this dose distribution is satisfactory depends on whether or not \mathbf{p}^* is in P . If $\mathbf{p}^* \in P$, then every tumour voxel will receive sufficient dose in the limiting dose distribution. However, if $\mathbf{p}^* \notin P$, then there is no guarantee that every tumour voxel will receive sufficient dose. In contrast to the static method, the adaptive robust method always leads to sufficient tumour dose in the limit. The reason for this is that every dose distribution \mathbf{d} in \mathbf{D} has the property that all tumour voxels receive the minimum prescription dose. This property is a consequence of the fact that $\mathbf{d} = \Delta \mathbf{p}^* \mathbf{w}$ for some intensity vector \mathbf{w} corresponding to the singleton uncertainty set $\{\mathbf{p}^*\}$; by the definition of problem (3.1), it is straightforward to see that the dose to tumour voxel v satisfies

$$d_v = \sum_{x \in X} \sum_{b \in B} \Delta_{v,x,b} p^*(x) w_b \geq \theta_v.$$

Furthermore, observe that every beamlet intensity vector in $\mathbf{w}^*(\mathbf{p}^*)$ is an optimal

solution of the nominal problem with respect to \mathbf{p}^* . As we discussed earlier, solutions to nominal problems are ones that generally result in the lowest level of normal tissue damage. Therefore, not only do the dose distributions in \mathbf{D} meet the tumour dose requirements, but they are also likely to lead to low levels of healthy tissue dose. In contrast, a static robust treatment will generally not possess both of these qualities simultaneously. If the treatment planner selects P to be a singleton (i.e., w is the solution of the nominal problem with a specific PMF), then the treatment is likely to result in low healthy tissue dose, but the treatment planner is taking a major risk by assuming what \mathbf{p}^* will be. If the treatment planner selects a larger P , then there is a better chance that \mathbf{p}^* will be contained in P , but at the cost of increased healthy tissue dose.

The second insight is that the adaptive robust and prescient treatments, in the limit, converge to the same set of ideal dose distributions. When the nominal problem with respect to \mathbf{p}^* has a unique optimal solution, the difference between the adaptive robust dose distribution and either prescient dose distribution must tend to zero by Corollary 1. In the case that there are multiple optimal solutions to the nominal problem with respect to \mathbf{p}^* , it seems reasonable to expect that the dose distributions that are achieved when these solutions are delivered and the patient breathes according to \mathbf{p}^* should not differ significantly. Therefore, the difference between either prescient dose distribution and the adaptive robust dose distribution should also become small in the limit.

The third insight is that in the limit, the adaptive robust method's performance is independent of the choice of the initial uncertainty set. In the proof of Theorem 2, the choice of ℓ^1 and \mathbf{u}^1 is not important: the dose distribution always approaches the set \mathbf{D} , no matter what ℓ^1 and \mathbf{u}^1 are. Again, if \mathbf{D} is a singleton, then any two dose distributions that are obtained using the same adaptive method but different ℓ^1 and \mathbf{u}^1 vectors will converge to the single distribution in \mathbf{D} . Since the two dose distributions converge to the same limiting dose distribution, the difference between those dose distributions must tend to zero. In practice, we will still see some dependence on the initial uncertainty set due to the finite number of fractions, but with a large number of fractions, the effect of the initial uncertainty set will become diminished.

While it is unlikely for a real patient PMF sequence to converge in the mathematical sense, as long as the patient's PMF sequence stabilizes, we would expect to observe the performance suggested by the convergence results. The real patient PMF sequences used in Section 4.3 are relatively stable. In Appendix B, we empirically show that even

with PMFs that do not stabilize, good performance can still be achieved with the right update algorithm.

4.3 Computational study

4.3.1 Background

In order to compare our results with previous static robust optimization results, the patient geometry and beam geometry used for this computational study were exactly the same as in Bortfeld et al. [2008]. The patient geometry consisted of 110,275 voxels, of which 5,495 were tumour voxels, and each voxel was of size $2.93 \text{ mm} \times 2.5 \text{ mm} \times 2.93 \text{ mm}$. The robust problem solved in each fraction of the dynamic method had 122,515 variables and 65,940 constraints, and took approximately 7 minutes to solve on average. The parameter γ was set to 1.1, while the minimum dose θ_v was set to 72Gy for every tumour voxel v . The set of motion states X consisted of five motion states. The nominal PMF used in the objective function was the same nominal PMF used previously. The adaptive robust and the static robust methods were tested using two different PMF sequences derived from clinical lung cancer patient data obtained at Massachusetts General Hospital. These PMF sequences were calculated using the method of Lujan et al. [1999] from breathing data obtained using Varian's real-time position management (RPM) system (Varian Medical Systems, Inc., Palo Alto, CA). In each imaging session, the RPM system was used to monitor the displacement over time of an external marker on the patient's abdomen. The full range of this displacement was divided into subintervals corresponding to the motion states obtained from the pre-treatment 4D-CT images. The proportion of the total time that the marker spent in each of the subintervals was then calculated, leading to a breathing motion PMF for that session. While the RPM system measures the external motion of a marker, which is not necessarily the same as the internal tumour motion, studies have shown that they are generally highly correlated (see, for example, Tsunashima et al. [2004] and Gierga et al. [2005]). We emphasize here that our adaptive robust method is not predicated on the use of the RPM system; any method is acceptable so long as it can be used to measure the patient's breathing motion and ultimately estimate the patient's breathing motion PMF. An imaging modality that can be used is 4D cone-beam CT (4D-CBCT), first introduced in Sonke et al. [2005]. Using daily

4D-CBCT imaging, the treatment planner can directly measure the internal motion of the tumour and obtain accurate PMFs after each fraction.

The two PMF sequences consisted of the first 30 PMFs from the two experiments, respectively, used in Bortfeld et al. [2008]. This number of PMFs was chosen to simulate a realistic six-week, Monday-to-Friday treatment course. For the adaptive robust method, we tested the exponential smoothing update algorithm with α equal to 0.1, 0.5, 0.9 and 1, and the running average update algorithm. Three different initial uncertainty sets were used for the dynamic and static robust methods. The first type of uncertainty set was a nominal uncertainty set, with ℓ and \mathbf{u} both equal to the nominal PMF. The second uncertainty set was the margin uncertainty set ($P = \mathcal{P}$). The third uncertainty set was a robust (intermediately-sized) uncertainty set, which was neither a singleton nor the whole space of PMFs \mathcal{P} ; for each PMF sequence we used the corresponding robust uncertainty set from Bortfeld et al. [2008]. In addition to the dynamic and static robust methods, both prescient solutions were calculated for each of the two PMF sequences.

The experiments were performed on a Dell workstation with a quad-core Intel Xeon 2.67GHz processor and 6GB of memory. All of the steps of the adaptive robust method – with the exception of the step where the robust problem is solved to obtain \mathbf{w}^{i+1} – were implemented in MATLAB (The MathWorks, Inc., Natick, Massachusetts). The robust problem in each fraction was solved using CPLEX 12.1 (IBM Corp., Armonk, NY) via AMPL (AMPL Optimization LLC, Albuquerque, NM).

After the dynamic and static robust methods were executed with each of the PMF sequences described above, the resulting treatments were evaluated by calculating the final dose distribution. The final dose distribution is a vector $\mathbf{d} \in \mathbb{R}^{|\mathcal{V}|}$ of doses delivered to the voxels in the patient geometry, defined component-wise as

$$d_v = \sum_{i=1}^n \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n}.$$

From the final dose distribution, we calculated three key statistics: the minimum tumour dose, $d_{\mathcal{T}}^{\min} = \min_{v \in \mathcal{T}} d_v$; the mean left lung dose, $d_{\mathcal{L}}^{\text{mean}} = \sum_{v \in \mathcal{L}} d_v / |\mathcal{L}|$, where \mathcal{L} is the set of voxels belonging to the left lung; and the mean normal tissue dose, $d_{\mathcal{N}}^{\text{mean}} = \sum_{v \in \mathcal{N}} d_v / |\mathcal{N}|$, where \mathcal{N} is the set of normal tissue (non-tumour) voxels.

Implementation	Min. tumour dose		Mean lung dose		Mean n. tissue dose	
	Gy	%	Gy	%	Gy	%
(S,N)	67.25	93.40	17.39	85.26	9.04	88.99
(ES(0.1),N)	70.79	98.32	17.50	85.85	9.05	89.11
(ES(0.5),N)	71.74	99.63	17.57	86.16	9.06	89.13
(ES(0.9),N)	71.83	99.77	17.57	86.17	9.06	89.13
(ES(1),N)	71.84	99.77	17.57	86.16	9.06	89.13
(RA,N)	71.47	99.26	17.52	85.91	9.06	89.12
(S,R)	71.37	99.13	18.19	89.20	9.44	92.93
(ES(0.1),R)	71.80	99.72	17.83	87.43	9.19	90.46
(ES(0.5),R)	71.99	99.98	17.65	86.55	9.09	89.50
(ES(0.9),R)	71.98	99.97	17.60	86.33	9.07	89.28
(ES(1),R)	71.97	99.95	17.59	86.29	9.07	89.26
(RA,R)	71.88	99.84	17.66	86.61	9.12	89.71
(S,M)	72.04	100.06	20.39	100.00	10.16	100.00
(ES(0.1),M)	72.05	100.07	18.31	89.78	9.44	92.91
(ES(0.5),M)	72.03	100.04	17.80	87.27	9.16	90.10
(ES(0.9),M)	72.02	100.02	17.71	86.87	9.11	89.67
(ES(1),M)	71.99	99.98	17.67	86.65	9.09	89.49
(RA,M)	72.03	100.04	17.94	87.97	9.23	90.86
(DLYP)	72.00	100.00	17.57	86.18	9.06	89.13
(AVGP)	72.00	100.00	17.54	86.02	9.06	89.13

Table 4.1: Dose statistics for the first PMF sequence.

4.3.2 Results

Table 4.1 displays dose statistics for the first PMF sequence. Under “Implementation”, the first term indicates the type of method: static (S), exponential smoothing with parameter α (ES(α)), running average (RA), daily prescient (DLYP) and average prescient (AVGP). For the non-prescient methods, the second term indicates which of the uncertainty sets described in Section 4.3.1 was used as the initial uncertainty set: nominal (N), robust (R) and margin (M). The percentage under “Min. tumour dose” is the minimum tumour dose as a percentage of the prescription dose (72Gy). The percentage under “Mean lung dose” is the mean left lung dose as a percentage of the mean left lung dose delivered in the static margin treatment, indicated as implementation (S,M). The percentage under “Mean n. tissue dose” is defined analogously. Table 4.2 displays the same statistics for the second PMF sequence.

In Figures 4.2 and 4.3, we plot the minimum tumour dose as a percentage of 72Gy versus the mean left lung dose as a percentage of the static margin solution for each method applied to the first PMF sequence. Figure 4.2 shows all of the treatments,

Implementation	Min. tumour dose		Mean lung dose		Mean n. tissue dose	
	Gy	%	Gy	%	Gy	%
(S,N)	64.37	89.40	17.18	85.11	8.99	88.92
(ES(0.1),N)	70.46	97.87	17.51	86.78	9.00	88.99
(ES(0.5),N)	71.61	99.46	17.55	86.96	8.98	88.87
(ES(0.9),N)	71.69	99.56	17.54	86.92	8.98	88.85
(ES(1),N)	71.68	99.56	17.54	86.90	8.98	88.84
(RA,N)	71.49	99.30	17.55	86.97	8.99	88.90
(S,R)	71.99	99.98	17.99	89.12	9.39	92.90
(ES(0.1),R)	71.98	99.98	17.68	87.59	9.12	90.20
(ES(0.5),R)	72.00	100.00	17.59	87.15	9.01	89.16
(ES(0.9),R)	71.96	99.94	17.58	87.08	9.00	89.03
(ES(1),R)	71.94	99.92	17.57	87.04	8.99	88.97
(RA,R)	71.96	99.95	17.59	87.16	9.04	89.40
(S,M)	72.05	100.07	20.18	100.00	10.11	100.00
(ES(0.1),M)	72.05	100.07	18.07	89.54	9.36	92.62
(ES(0.5),M)	72.01	100.01	17.69	87.67	9.06	89.65
(ES(0.9),M)	71.98	99.97	17.65	87.43	9.02	89.27
(ES(1),M)	71.96	99.95	17.64	87.40	9.02	89.21
(RA,M)	72.02	100.03	17.75	87.94	9.15	90.48
(DLYP)	72.00	100.00	17.55	86.97	8.98	88.84
(AVGP)	72.00	100.00	17.57	87.06	8.98	88.85

Table 4.2: Dose statistics for the second PMF sequence.

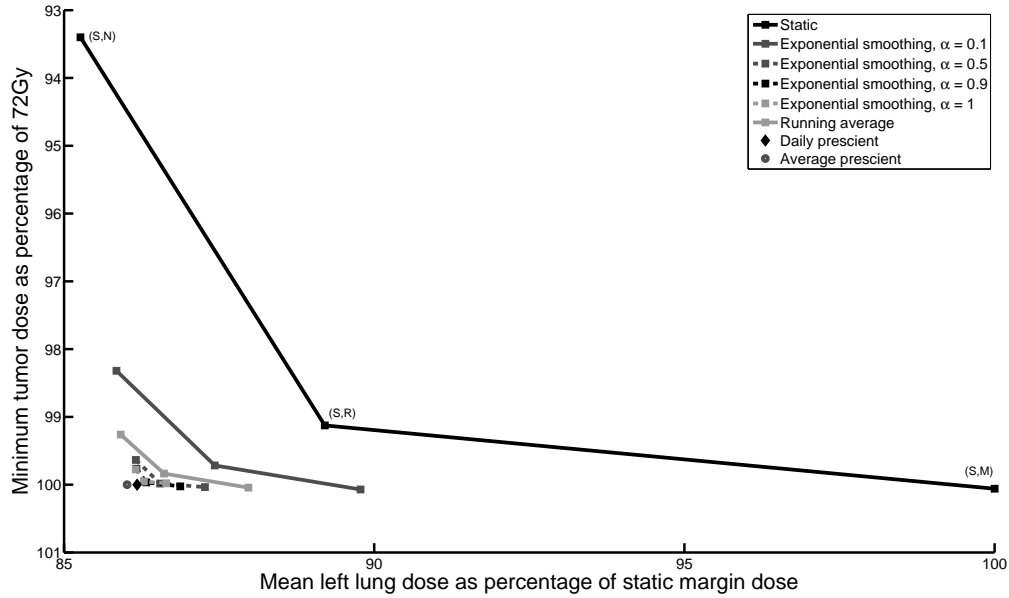


Figure 4.2: Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the first PMF sequence (full). The points in the bottom left part of the plot are the adaptive robust and prescient treatments.

while Figure 4.3 zooms in on the dynamic and prescient treatments. On these plots, points obtained by the same method (static, exponential smoothing with parameter α or running average) from different initial uncertainty sets are joined together by straight lines to illustrate the approximate Pareto frontiers generated by the corresponding method. Figures 4.4 and 4.5 display analogous plots for the second PMF sequence.

In Figure 4.6, we also provide dose volume histograms (DVHs) for two treatments obtained from the first patient sequence – the exponential smoothing algorithm with smoothing factor 0.5 and robust uncertainty set (implementation (ES(0.5),R)) and the static method with the same uncertainty set (implementation (S,R)). A DVH is a plot for a particular structure which indicates, for a given level of dose, what volume of the structure receives at least that dose. Ideally, the healthy organ DVHs would be step functions from 100% to 0% at 0Gy (i.e., no dose is delivered to the healthy organs), and the tumour DVH would be a step function from 100% to 0% at 72Gy (i.e., the whole tumour volume receives exactly 72Gy).

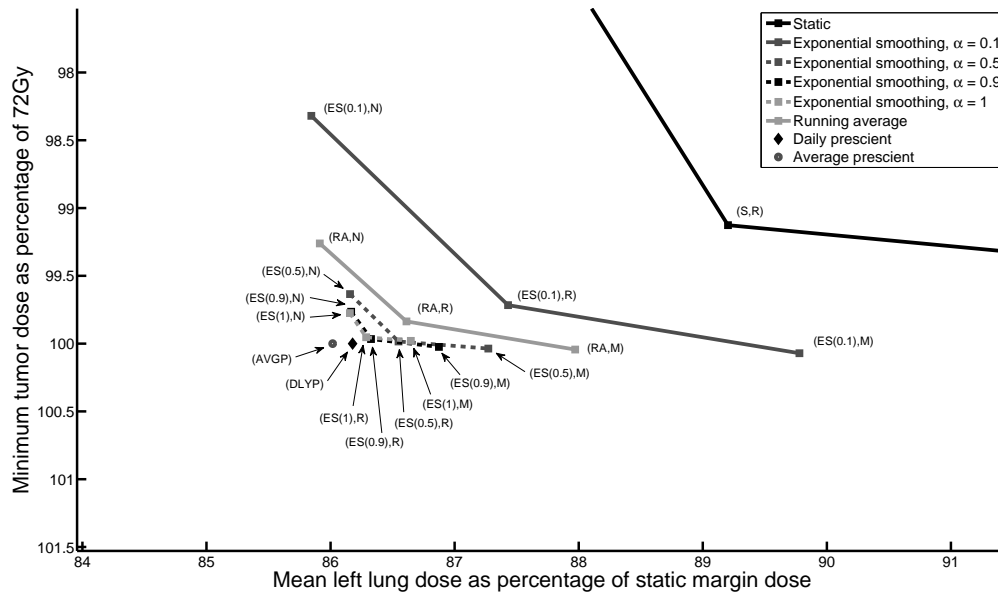


Figure 4.3: Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the first PMF sequence (zoomed-in).

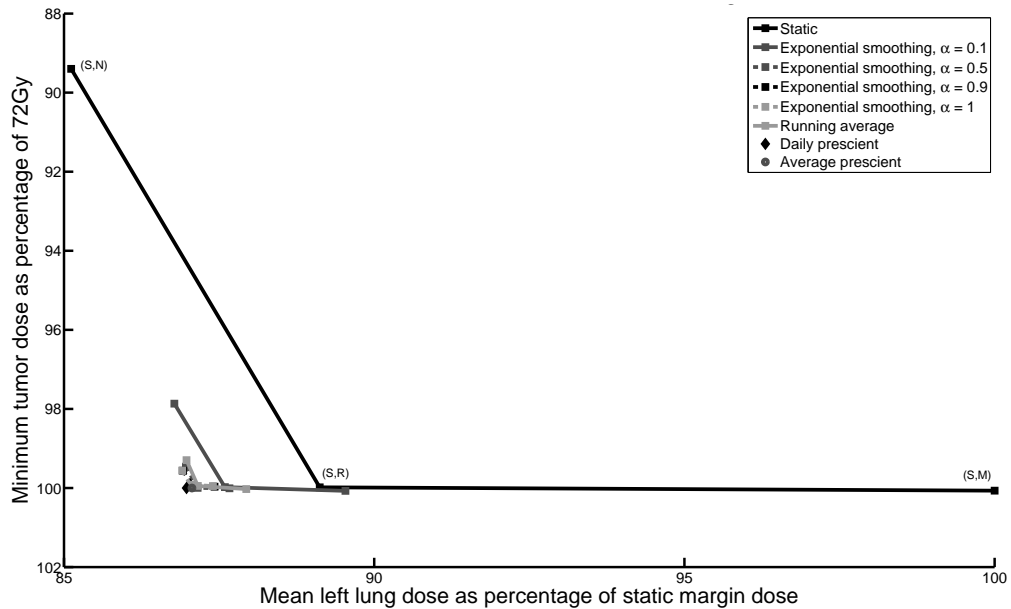


Figure 4.4: Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the second PMF sequence (full). The points in the bottom left part of the plot are the adaptive robust and prescient treatments.

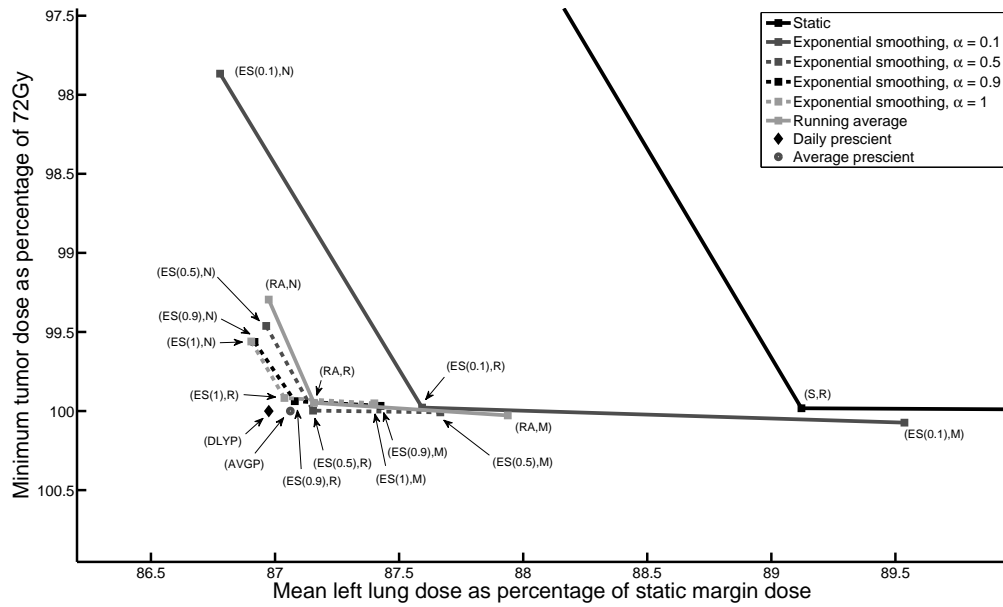


Figure 4.5: Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the second PMF sequence (zoomed-in).




Figure 4.6: Dose volume histogram comparing the treatments obtained by exponential smoothing with $\alpha = 0.5$ started from the robust uncertainty set (implementation (ES(0.5),R)) and the static method with the robust uncertainty set (implementation (S,R)) with the first PMF sequence.

4.3.3 Discussion

The computational results support the three insights discussed in Section 4.2.3. First, the adaptive robust method generally dominates the static robust method, generating solutions that can simultaneously improve on tumour coverage and healthy tissue dose. As an example of this, compare the dose statistics of the static and the exponential smoothing with $\alpha = 0.5$ methods for the first PMF sequence in Table 4.1. For the margin (M) uncertainty set, moving from the static (implementation (S,M)) to the exponential smoothing with $\alpha = 0.5$ (implementation (ES(0.5),M)) treatment results in a reduction in the mean left lung dose of 12.73% of the static margin dose, with virtually no change in tumour coverage. For the robust (R) uncertainty set, moving from the static (implementation (S,R)) to the exponential smoothing with $\alpha = 0.5$ (implementation (ES(0.5),R)) treatment results in an increase in minimum tumour dose of 0.85% of the prescribed dose, and a decrease in the mean left lung dose of 2.65% of the static margin dose – both objectives are improved. For the nominal (N) uncertainty set, moving from the static (implementation (S,N)) to the exponential smoothing with $\alpha = 0.5$ (implementation (ES(0.5),N)) treatment results in a significant increase in the minimum tumour dose of 6.23% of the prescription dose with only a marginal increase in mean left lung dose of 0.91% of the static margin dose. Clinically, the marginal increase in lung dose (0.18Gy) would be acceptable given the increase in minimum tumour dose (4.49Gy). Similar results are seen for the second PMF sequence.

The dominance of the adaptive robust method over the static robust method is evident in Figures 4.2 and 4.3. On these plots, we can see that the frontiers of all of the adaptive methods are generally below and to the left of the static frontier. The DVHs in Figure 4.6 further support this insight. From Figure 4.6, we can see that the left lung, esophagus and heart curves for the adaptive treatment are mostly to the left of the corresponding curves for the static treatment, indicating an overall reduction of dose in these structures. The maximum spinal cord dose is slightly higher in the adaptive treatment, but is still clinically acceptable. The homogeneity of the tumour dose is slightly better in the static treatment, as can be seen in the sharper decline of the static treatment tumour dose-volume curve. However, the adaptive treatment has less tumour underdose and still satisfies the maximum tumour dose constraint. Again, Figure 4.6 supports the observation that the adaptive robust method generally leads to simultaneous improvement in tumour coverage and healthy tissue sparing.

These simultaneous improvements in tumour dose and lung dose together allow for clinically significant dose escalation. For example, if the prescription dose for the static robust treatment (implementation (S,R)) for the first PMF sequence is scaled up so that the mean left lung dose matches that of the static margin treatment (implementation (S,M)), the resulting minimum tumour dose is 80.00Gy. Doing the same for the exponential smoothing treatment with $\alpha = 0.5$ and the robust uncertainty set (implementation (ES(0.5),R)) for the first PMF sequence, the minimum tumour dose is 83.17Gy. Using the estimated relationship between the 5-year local tumour control rate and tumour dose from Kong et al. [2005] for non-small cell lung cancer (NSCLC), this difference of 3.17Gy translates to a 4% increase in local tumour control. In addition to the mean lung dose, dose can also be escalated according to the commonly used lung V20 metric, which is defined as the volume of the lung tissue that receives a dose of 20Gy or higher. If the prescription dose is escalated so that the (ES(0.5),R) and (S,R) treatments for the first PMF sequence have the same left lung V20 as the (S,M) treatment, then the difference in minimum tumour dose grows to 6.40Gy; using the relationship from Kong et al. [2005], this translates approximately to an 8% increase in 5-year local tumour control. Given that the 5-year control rates for tumour doses in the 63-69Gy and 74-84Gy ranges are on the order of 12% and 35% respectively (Kong et al. [2005]), the gains in local control rate that we have described are significant. Depending on the stage of lung cancer in question, the benefits of dose escalation become even greater: in Rengan et al. [2004], which studied the effect of dose escalation on local failure rates for patients with stage III NSCLC, a 10Gy increase in dose was found to decrease local failure rates by 36%.

The second observation that can be made is that even though the adaptive robust method is myopic and does not know what PMFs will be realized in future fractions, its performance is comparable to solutions that can fully anticipate future PMFs. This is evident in the fact that the treatments obtained from the adaptive robust method are very close in tumour coverage and healthy tissue dose to the two types of prescient solutions. For example, for the first PMF sequence, when the robust (R) uncertainty set is chosen and exponential smoothing with $\alpha = 0.9$ is applied (implementation (ES(0.9),R) in Table 4.1), the minimum tumour dose as a percentage is 99.97% and the mean left lung dose as a percentage is 86.33%. For both the daily and average prescient treatments, the minimum tumour dose as a percentage is 100%; for the daily prescient treatment, the mean left lung dose as a percentage is 86.18%, whereas for the average

prescient treatment, the mean left lung dose as a percentage is 86.02%. Even at slower rates of adaptation, the adaptive robust treatments perform similarly to the prescient solutions. For instance, using exponential smoothing with $\alpha = 0.1$ and starting with the robust (R) uncertainty set (implementation (ES(0.1),R)), the minimum tumour dose as a percentage of 72Gy is within 0.5% of the prescient solutions in minimum tumour dose and within 1.5% of the prescient solutions in mean left lung dose.

The third observation that can be made is that the final dose distribution from any of the adaptive robust methods is far less sensitive to the choice of initial uncertainty set compared to the static method. For instance, for the first PMF sequence, consider the treatments that use exponential smoothing with $\alpha = 0.9$ (implementations (ES(0.9),N), (ES(0.9),R) and (ES(0.9),M)). For the nominal (N), robust (R) and margin (M) uncertainty sets, the minimum tumour doses achieved using exponential smoothing with $\alpha = 0.9$ are 71.83, 71.98 and 72.02Gy respectively, while the mean left lung doses are 17.57, 17.60 and 17.71Gy respectively. Visually, this insight is evident in the fact that the adaptive method frontiers are much smaller than the static frontier on both Figures 4.2 and 4.4. This insight is significant because it indicates that the success of this method is not contingent on having accurate knowledge of what the PMFs will be before the treatment begins. From a clinical standpoint, this eliminates an important barrier to clinical implementation, as the treatment planner may not have enough detailed knowledge of a particular patient's breathing patterns before the start of treatment to create a single uncertainty set suitable for the entire treatment course.

Given the third observation, it may seem that adding adaptation diminishes the value of robust optimization, as the adaptive treatments that begin with a nominal uncertainty set have similar performance to the adaptive treatments that begin with the other two larger uncertainty sets. Furthermore, the treatments from exponential smoothing with $\alpha = 1$ ((ES(1),N), (ES(1),R) and (ES(1),M)), for which the uncertainty set from fraction 2 on is always a singleton, perform almost as well as the prescient solutions with respect to the final dose distribution. However, even though the final dosimetry of the treatment after all of the fractions have been delivered may appear acceptable, we must still evaluate the dose delivered in each fraction. In Figure 4.7, we plot the mean and maximum underdose by fraction as a percentage of the prescription dose per fraction ($72/30 = 2.4\text{Gy}$) for the three treatments obtained with exponential smoothing with $\alpha = 0.9$ (implementations (ES(0.9),N), (ES(0.9),R) and (ES(0.9),M)) under the first PMF sequence. From this plot, we can see that from fraction 8 on, the three




Figure 4.7: Plot of mean and maximum underdose by fraction for the three exponential smoothing with $\alpha = 0.9$ treatments ((ES(0.9),N), (ES(0.9),R), (ES(0.9),M)) under the first PMF sequence.

treatments are identical in mean and maximum underdose. However, before fraction 8, there is a difference in the mean underdose across the treatments – in general, the (ES(0.9),M) treatment has the least mean underdose, followed by the (ES(0.9),R) treatment, followed by the (ES(0.9),N) treatment. The same behavior is observed for the maximum underdose in each fraction. High underdose is undesirable in any fraction. It is especially undesirable in early fractions because the effect of delivering insufficient dose to parts of the tumour is compounded in the remaining fractions, potentially leading to tumour growth that is not accounted for by the end of treatment. Due to such biological considerations, there is a clinical desire to ensure that a sufficient and consistent dose is delivered each day. By starting from a sufficiently large uncertainty set, the robust optimization component of our method ensures that this clinical objective is met in the early fractions of the treatment. Thus, robust optimization is vital to ensuring the clinical acceptability of the final treatment.

In addition to biological considerations, there may be other situations in which robustness is necessary for good performance from an adaptive method. The PMF sequences that we have used here are well-behaved and relatively stable. However, sequences that exhibit more variability may cause poor performance if there is too

much adaptation and too little robustness (i.e., α is set close to 1 and the initial uncertainty set is small); we explore such a sequence in Appendix B. Similarly, sequences that exhibit transient behavior and drift outside of the initial uncertainty set towards a new value could also necessitate the use of a large uncertainty set early on while the lower and upper bound vectors are still “homing in” on the PMFs. Robustness is also crucial when the number of treatment fractions is low, as in stereotactic body radiation therapy. The lower number of fractions translates to an increased dose per fraction, and so the underdose or overdose that is realized when the PMF during a given fraction falls outside of the uncertainty set of that fraction is much greater. Furthermore, the low number of fractions leaves limited room for those errors in dose to be corrected in the future. In this situation, robustness and adaptation together clearly lead to more clinical gains than adaptation or robustness alone.

Although the insights that we have described here hold for our two real patient PMF sequences, they may not necessarily hold for all PMF sequences and all update algorithms. In Appendix B, we give an example of a pathological sequence under which the exponential smoothing treatments are significantly different from the pre-scient treatments, but the running average treatments still exhibit comparable performance.

4.4 Conclusions

In this chapter, we developed an adaptive robust optimization method for fractionated IMRT treatment planning that is able to combat breathing motion uncertainty. We studied the effectiveness of our method on two PMF sequences obtained from clinical lung cancer patients and showed that our adaptive robust optimization method improves upon the static method in both tumour coverage and healthy tissue sparing simultaneously. We also studied the effectiveness of our method from a theoretical standpoint and proved that the adaptive robust dose distribution approaches a limiting set of optimal dose distributions. Even though this adaptive approach is myopic, it performs well when the PMF sequence converges, and for some uncertainty set update algorithms, performs well even when the PMF sequence does not converge. Overall, the clinical value of this method is that it allows for the tumour dose to be safely escalated without leading to additional healthy tissue toxicity, which may ulti-

mately improve the rate of patient survival.

From an implementation standpoint, this method is feasible in most clinical environments. Determining the beamlet intensity vector for each fraction (i.e., solving the robust optimization problem) takes no more than 10 minutes, which is certainly clinically acceptable. A variety of imaging technologies that provide the capability for measuring the patient's breathing motion and determining the patient's breathing motion PMF already exist. As we have shown empirically and theoretically, our method does not require perfectly accurate information to be available before the start of treatment. The limiting factor for implementing this type of method in a clinical environment is the additional time that physicians and medical physicists will need to spend each day to perform quality assurance and to evaluate the dose that will be delivered on that day. A relevant follow-up to this study, therefore, is a study to quantify the costs of implementing this method on different adaptation schedules and to compare them to the benefit to the patient and to the entire healthcare system.

Although the computational results from our two PMF sequences are very encouraging, testing with more real PMF sequences is necessary to fully determine the benefit of this method. It would also be interesting to further stress-test the method with artificial sequences, such as sequences where the PMF alternates among finitely many PMFs (like the one considered in Appendix B), sequences that are randomly drawn from some distribution on the probability simplex (e.g., the Dirichlet distribution) and sequences where the PMFs are chosen by an "adversary" (e.g., the PMF that is realized in a fraction is the worst possible one for the uncertainty set of that fraction). In the theoretical direction, future work should consider what results can be obtained when the convergence assumption is relaxed and the PMFs are instead drawn from a distribution on the probability simplex.

Chapter 5

Dose-reactive methods

In Chapter 4, we developed an adaptive robust method that operates by solving a sequence of robust optimization problems, one for each fraction, with the uncertainty set updated each time with the most recent measurement of the patient's breathing motion PMF.

While the approach in Chapter 4 is adaptive with respect to new observations of the uncertainty, it does not adjust the target dose requirements based on knowledge of the dose delivered up until that point. That is, the beamlet intensities of each fraction are always designed to deliver $1/n$ of the prescription dose to the whole tumour, where n is the number of fractions, regardless of whether some parts of the tumour received significantly more or significantly less than $1/n$ of the prescription dose in previous fractions. Such underdose or overdose may occur in the adaptive robust approach because the uncertainty set in a given fraction may not necessarily contain the breathing pattern that is realized in that fraction, and thus the beamlet intensities used in that fraction may not maintain the tumour dose within specified upper and lower bounds. Although the previous adaptive robust method performs quite well in comparison to the static robust method, it is an open question whether or not accounting for the dose delivered in prior fractions will improve this performance. In principle, such a method, which we will call a *dose-reactive adaptive robust* method, or *dose-reactive* method for short, should be able to correct for relative underdosages or overdosages in earlier fractions and lead to a better final dose distribution. However, it has been hypothesized that the performance of such a dose-reactive method, from the perspective of the daily delivered dose, may be quite different to and potentially worse than the performance of non-reactive adaptive robust methods. To distinguish it from the

dose-reactive methods, we will refer to the adaptive robust method of Chapter 4 as the *non-reactive adaptive robust* method, or *non-reactive* method for short.

In this chapter, we consider dose-reactive methods for lung cancer IMRT that are able to react to the dose delivered in prior fractions. We consider three methods: the *reactive*[±] method, which adjusts the target dose distribution in response to both underdose and overdose in previous fractions; the *reactive*⁻ method, which adjusts the target dose distribution in response to underdose only; and the *reactive*⁺ method, which adjusts the target dose distribution in response to overdose only.

Our contributions in this chapter as follows:

1. We formalize and develop three dose-reactive methods that adjust dose requirements over the course of the treatment in response to the dose that has been delivered in prior fractions. These methods are computationally tractable, as they only involve solving a sequence of linear programming problems. We describe these methods in Section 5.1.
2. We show, through computational studies using real patient data, that while these dose-reactive methods can lead to an improvement over the non-reactive adaptive robust method with respect to the final dose distribution, they do not perform as well with respect to daily dose. More precisely, dose-reactive methods generally lead to higher levels of daily underdose and/or overdose than the non-reactive method. We also show that the dose-reactive methods can result in growth in the daily underdose and/or overdose and growth in the heterogeneity of the daily tumour dose distribution over the course of a multi-week treatment horizon. This is the first study to highlight the characteristic behavior of dose-reactive methods in this respect. We analyze these results in Section 5.2.
3. We provide a theoretical analysis to explain why non-reactive and dose-reactive methods differ in their daily dose performance. We prove that the non-reactive method guarantees diminishing daily underdose and overdose in an asymptotic sense. We use the same theoretical framework to argue why dose-reactive methods will generally have growing daily underdose and/or overdose.

This chapter is organized as follows. In Section 5.1, we develop the *reactive*[±], *reactive*⁻ and *reactive*⁺ methods. In Section 5.2, we provide the results of our computational study with these methods. In Section 5.3 we present our theoretical analysis.

Finally, in Section 5.4 we conclude and provide some directions for future work on the dose-reactive methods studied here.

5.1 Method

The dose-reactive methods build on Algorithm 1 by allowing the dose requirements to change from fraction to fraction. To define the dose-reactive methods, we introduce the modified robust optimization problem, which is defined as

$$\begin{aligned}
& \text{minimize} && \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{p}(x) w_b \\
& \text{subject to} && \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b \geq \underline{\delta}_v, \quad \forall v \in \mathcal{T}, \mathbf{p} \in P, \\
& && \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b \leq \bar{\delta}_v, \quad \forall v \in \mathcal{T}, \mathbf{p} \in P, \\
& && w_b \geq 0, \quad \forall b \in \mathcal{B}.
\end{aligned} \tag{5.1}$$

Here, $\underline{\delta}_v$ is the target minimum dose for voxel v and $\bar{\delta}_v$ is the target maximum dose for voxel v . This model is identical to the problem (3.1), except for $\underline{\delta}_v$ and $\bar{\delta}_v$ replacing θ_v and $\gamma\theta_v$, respectively. We will use $\underline{\delta}$ and $\bar{\delta}$ to represent the vectors of $\underline{\delta}_v$ and $\bar{\delta}_v$ values, respectively.

In this method, the target minimum dose $\underline{\delta}_v$ and target maximum dose $\bar{\delta}_v$ represent the minimum and maximum “dose-to-go” for a given tumour voxel v . After each fraction i , the uncertainty set P^i is updated using the most recent PMF \mathbf{p}^i , as in the adaptive robust method. In addition to P^i , the target minimum and maximum doses $\underline{\delta}^i$ and $\bar{\delta}^i$ are adjusted according to how much dose was delivered to the tumour in that fraction. With these new problem data – P^{i+1} , $\underline{\delta}^{i+1}$ and $\bar{\delta}^{i+1}$ – the modified robust problem (5.1) is solved to obtain a vector of beamlet intensities \mathbf{w}^{i+1} . Since the dose-to-go vectors $\underline{\delta}^{i+1}$ and $\bar{\delta}^{i+1}$ are for the remaining $n - i$ fractions, the intensity vector $\mathbf{w}^{i+1}/(n - i)$ is delivered in fraction $i + 1$. This process repeats until the end of the treatment. This procedure is presented in algorithmic form as Algorithm 2. As stated at the start of this chapter, we will refer to a method of this type as a *dose-reactive adaptive robust* method, or *dose-reactive* method for short.

We consider three methods for updating $\underline{\delta}$ and $\bar{\delta}$ (Step 7 in Algorithm 2) after each fraction. In the *reactive*[±] method, after observing the PMF \mathbf{p}^i in fraction i , we set $\underline{\delta}_v^{i+1}$

Algorithm 2 Dose-reactive adaptive robust method

Require: Total number of fractions n , initial uncertainty set P^1 , initial target minimum dose $\underline{\delta}_v^1 = \theta_v$ for all $v \in \mathcal{T}$, initial target maximum dose $\bar{\delta}_v^1 = \gamma\theta_v$ for all $v \in \mathcal{T}$.

- 1: Initialize $i = 1$
- 2: Solve problem (5.1) with $P = P^1$, $\underline{\delta} = \underline{\delta}^1$ and $\bar{\delta} = \bar{\delta}^1$ to obtain a beamlet intensity vector \mathbf{w}^1
- 3: Deliver \mathbf{w}^1/n to the patient
- 4: **while** $i \leq n - 1$ **do**
- 5: Measure the patient's breathing motion and construct the PMF \mathbf{p}^i
- 6: Generate the new uncertainty set P^{i+1} from P^i , \mathbf{p}^i
- 7: Generate the new target minimum and maximum dose vectors $\underline{\delta}^{i+1}$ and $\bar{\delta}^{i+1}$ from $\underline{\delta}^i$, $\bar{\delta}^i$ and \mathbf{w}^i
- 8: Solve problem (5.1) with $P = P^{i+1}$, $\underline{\delta} = \underline{\delta}^{i+1}$ and $\bar{\delta} = \bar{\delta}^{i+1}$ to obtain a new beamlet intensity vector \mathbf{w}^{i+1}
- 9: Deliver $\mathbf{w}^{i+1}/(n - i)$ to the patient
- 10: Set $i = i + 1$
- 11: **end while**

for fraction $i + 1$ as

$$\underline{\delta}_v^{i+1} = \max \left\{ 0, \underline{\delta}_v^i - \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n - i + 1} \right\}$$

and $\bar{\delta}_v^{i+1}$ for fraction $i + 1$ as

$$\bar{\delta}_v^{i+1} = \max \left\{ 0, \bar{\delta}_v^i - \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n - i + 1} \right\},$$

for each voxel $v \in \mathcal{T}$. This update reduces the minimum and maximum dose-to-go for a given tumour voxel v by the actual dose delivered to that voxel in fraction i ; the more dose a voxel v has accumulated by fraction i , the smaller $\underline{\delta}_v^i$ and $\bar{\delta}_v^i$ will be, and the less dose will be delivered to voxel v over the remaining $n - i$ fractions. If a given voxel v has accumulated its prescribed minimum dose of θ_v , the use of $\max\{0, \cdot\}$ ensures that the minimum dose-to-go is set to zero, so that subsequent beamlet intensity vectors are not forced to deliver any dose to that voxel. If a voxel v has accumulated a dose equal to or higher than its prescribed maximum dose $\gamma\theta_v$, both $\underline{\delta}_v$ and $\bar{\delta}_v$ are set to zero, forcing subsequent beamlet intensity vectors to deliver no dose to voxel v . Note that the dose that is delivered in a fraction i , in addition to depending on the dose

requirements $\underline{\delta}^i$ and $\bar{\delta}^i$ of fraction i , also depends on whether or not the PMF \mathbf{p}^i falls in or outside the uncertainty set P^i ; if $\mathbf{p}^i \notin P^i$, it is possible that the dose delivered to a tumour voxel v falls outside of the interval $[\underline{\delta}_v^i/(n-i+1), \bar{\delta}_v^i/(n-i+1)]$.

In the *reactive*⁻ method, after observing the PMF \mathbf{p}^i in fraction i , we set $\underline{\delta}_v^{i+1}$ for fraction $i+1$ as

$$\underline{\delta}_v^{i+1} = \max \left\{ \left(\frac{n-i}{n} \right) \theta_v, \underline{\delta}_v^i - \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n-i+1} \right\} \quad (5.2)$$

and $\bar{\delta}_v^{i+1}$ for fraction $i+1$ as

$$\bar{\delta}_v^{i+1} = \max \left\{ \left(\frac{n-i}{n} \right) \gamma \theta_v, \bar{\delta}_v^i - \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n-i+1} \right\}, \quad (5.3)$$

for each voxel $v \in \mathcal{T}$. The *reactive*⁻ method reacts to underdose: if a voxel v is underdosed, the second term in the max expression in equation (5.2) will dominate and the requirement $\underline{\delta}_v^{i+1}$ will be larger than $(n-i)/n \cdot \theta_v$, allowing for more than θ_v/n units of dose to be delivered to voxel v in the upcoming fraction. In the case that a tumour voxel v is overdosed, however, the *reactive*⁻ update ensures that the minimum and maximum dose-to-go $\underline{\delta}_v^i$ and $\bar{\delta}_v^i$ for a given fraction i and a given voxel $v \in \mathcal{T}$ never drop below $(n-i+1)/n \cdot \theta_v$ and $(n-i+1)/n \cdot \gamma \theta_v$, respectively. As a result, the *reactive*⁻ method cannot exploit situations where a voxel v has received more than $i/n \cdot \theta_v$ dose by the end of fraction i in order to deliver less than $(n-i)/n \cdot \theta_v$ over the remaining $n-i$ fractions. By its definition, the *reactive*⁻ method always creates beamlet intensities that are designed to deliver at least θ_v/n to each voxel $v \in \mathcal{T}$ in each fraction whenever the realized PMF falls inside the uncertainty set.

In the *reactive*⁺ method, we set $\underline{\delta}_v^{i+1}$ in each fraction i as

$$\underline{\delta}_v^{i+1} = \max \left\{ 0, \min \left\{ \left(\frac{n-i}{n} \right) \theta_v, \underline{\delta}_v^i - \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n-i+1} \right\} \right\} \quad (5.4)$$

and $\bar{\delta}_v^{i+1}$ in each fraction i as

$$\bar{\delta}_v^{i+1} = \max \left\{ 0, \min \left\{ \left(\frac{n-i}{n} \right) \gamma \theta_v, \bar{\delta}_v^i - \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n-i+1} \right\} \right\}, \quad (5.5)$$

for each voxel $v \in \mathcal{T}$. The reactive⁺ method is the opposite of the reactive⁻ method, as it only responds to overdose. If a voxel v is overdosed, the second term in the inner min expression in equation (5.4) will dominate, and so $\underline{\delta}_v^{i+1}$ will be smaller than $(n-i)/n \cdot \theta_v$, allowing for less than θ_v/n units of dose to be delivered to voxel v in the upcoming fraction. If a voxel v is underdosed, however, the reactive⁺ update ensures that the minimum and maximum dose-to-go $\underline{\delta}_v^i$ and $\bar{\delta}_v^i$ for a given fraction i and a given voxel $v \in \mathcal{T}$ never go above $(n-i+1)/n \cdot \theta_v$ and $(n-i+1)/n \cdot \gamma\theta_v$, respectively. In this way, the reactive⁺ update ignores situations where too little dose is delivered in a fraction.

The following example illustrates the differences between the three different reactive methods.

Example Suppose that for a given tumour voxel v , the prescribed minimum dose is $\theta_v = 10$ and the prescribed maximum dose is $\gamma\theta_v = 11$ ($\gamma = 1.1$). Suppose also that the number of fractions in the treatment is 10, in which case the daily dose in each fraction should ideally be between $\theta_v/n = 1$ and $\gamma\theta_v/n = 1.1$. Suppose that the dose delivered to voxel v in fraction 1 (mathematically, $1/(n-1+1) \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^1(x) w_b^1$) is 0.9 – less than the daily prescribed minimum dose of 1. In this situation we would have the values shown in Table 5.1. From this table, we can see that the target minimum and maximum dose for the reactive[±] and reactive⁻ methods match ($\underline{\delta}_v^2 = 9.1$, $\bar{\delta}_v^2 = 10.1$), because both methods account for underdose, but the target minimum and maximum doses for the reactive⁺ method are lower ($\underline{\delta}_v^2 = 9.0$, $\bar{\delta}_v^2 = 9.9$), because the reactive⁺ method ignores underdose. In the next fraction, the reactive[±] and reactive⁻ methods would strive to deliver more dose to voxel v than the reactive⁺ method because their target minimum dose requirements are higher; by doing so, they attempt to correct the underdose to voxel v that occurred in fraction 1.

Now let us take the same problem setup, and suppose that in fraction 1, we deliver 1.2, which is higher than the daily prescribed maximum dose of 1.1. This would lead us to the values shown in Table 5.2. Now, the target minimum and maximum dose for the reactive[±] and reactive⁺ methods match ($\underline{\delta}_v^2 = 8.8$, $\bar{\delta}_v^2 = 9.8$), because both of those methods respond to overdose, but the target minimum and maximum dose for the reactive⁻ method are higher ($\underline{\delta}_v^2 = 9.0$, $\bar{\delta}_v^2 = 9.9$), because the reactive⁻ method does not account for overdose. In the next fraction, the reactive[±] and reactive⁺ methods would strive to deliver less dose to voxel v than the reactive⁻ method because their

Quantity	Expression	Reactive [±]	Reactive ⁻	Reactive ⁺
Reduced target min. dose	$\underline{\delta}_v^1 - \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^1(x) \frac{w_b^1}{n-1+1}$	9.1	9.1	9.1
Reduced target max. dose	$\bar{\delta}_v^1 - \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^1(x) \frac{w_b^1}{n-1+1}$	10.1	10.1	10.1
Underdose threshold	$\frac{n-1}{n} \theta_v$	9.0	9.0	9.0
Overdose threshold	$\frac{n-1}{n} \gamma \theta_v$	9.9	9.9	9.9
New target min. dose	$\underline{\delta}_v^2$	9.1	9.1	9.0
New target max. dose	$\bar{\delta}_v^2$	10.1	10.1	9.9

Table 5.1: Reactive[±], reactive⁻ and reactive⁺ underdose example.

Quantity	Expression	Reactive [±]	Reactive ⁻	Reactive ⁺
Reduced target min. dose	$\underline{\delta}_v^1 - \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^1(x) \frac{w_b^1}{n-1+1}$	8.8	8.8	8.8
Reduced target max. dose	$\bar{\delta}_v^1 - \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^1(x) \frac{w_b^1}{n-1+1}$	9.8	9.8	9.8
Underdose threshold	$\frac{n-1}{n} \theta_v$	9.0	9.0	9.0
Overdose threshold	$\frac{n-1}{n} \gamma \theta_v$	9.9	9.9	9.9
New target min. dose	$\underline{\delta}_v^2$	8.8	9.0	8.8
New target max. dose	$\bar{\delta}_v^2$	9.8	9.9	9.8

Table 5.2: Reactive[±], reactive⁻ and reactive⁺ overdose example.

target minimum dose requirements are lower, thus allowing them to take advantage of the overdose experienced in fraction 1.

5.2 Results

5.2.1 Background

We tested the dose-reactive approaches described above using the same computational setup as in Section 4.3. Two PMF sequences consisting of 30 PMFs, obtained from real lung cancer patients, were used to test the dose-reactive and non-reactive methods. In this section, we only focus on results for the first PMF sequence, as the results obtained using the second PMF sequence were qualitatively very similar; results for the second PMF sequence are given in Appendix D. The prescribed minimum tumour dose θ_v for each tumour voxel v was set to 72Gy, while the prescribed maximum tumour dose $\gamma \theta_v$ was set to 79.2Gy (corresponding to $\gamma = 1.1$). The reactive[±],

reactive⁻ and reactive⁺ target dose updates for $\underline{\delta}$ and $\bar{\delta}$ were tested in conjunction with the exponential smoothing update for the uncertainty set. The initial uncertainty sets used for the resulting reactive[±], reactive⁻ and reactive⁺ methods were the same as those used in Section 4.3. Specifically, we considered a nominal uncertainty set (N) consisting of only the nominal PMF ($P = \{\bar{p}\}$); a margin uncertainty set (M) consisting of all PMFs ($P = \mathcal{P}$); and a robust uncertainty set (R) that is in between the nominal and margin uncertainty set in size. In addition to these implementations, we also evaluated the performance of several different daily prescient treatments. Each of the daily prescient treatments that we consider incorporates one of the different dose-reactive target dose updates, thus resulting in non-reactive daily prescient, reactive[±] daily prescient, reactive⁻ daily prescient and reactive⁺ daily prescient treatments.

In Section 5.2.2, we will discuss tables and figures that describe the final cumulative dose distribution for each of the implementations that we test. For the non-reactive method, given a sequence of beamlet intensity vectors $\mathbf{w}^1, \dots, \mathbf{w}^n$, the final dose to voxel $v \in \mathcal{V}$ is calculated as

$$d_v = \sum_{i=1}^n \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n}.$$

For the dose-reactive methods, the final dose is calculated as

$$d_v = \sum_{i=1}^n \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n - i + 1}.$$

Given the final dose d_v to each voxel v in the patient geometry, the minimum tumour dose, $d_{\mathcal{T}}^{\min}$ is calculated as

$$d_{\mathcal{T}}^{\min} = \min_{v \in \mathcal{T}} d_v, \quad (5.6)$$

while the maximum tumour dose, $d_{\mathcal{T}}^{\max}$, is calculated as

$$d_{\mathcal{T}}^{\max} = \max_{v \in \mathcal{T}} d_v. \quad (5.7)$$

The mean left lung dose, $d_{\mathcal{L}}^{\text{mean}}$, is calculated as

$$d_{\mathcal{L}}^{\text{mean}} = \frac{1}{|\mathcal{L}|} \sum_{v \in \mathcal{L}} d_v, \quad (5.8)$$

where \mathcal{L} is the subset of voxels in \mathcal{V} belonging to the left lung. The mean normal tissue dose, $d_{\mathcal{N}}^{\text{mean}}$, is calculated as

$$d_{\mathcal{N}}^{\text{mean}} = \frac{1}{|\mathcal{N}|} \sum_{v \in \mathcal{N}} d_v, \quad (5.9)$$

where \mathcal{N} is the set of all normal tissue (non-tumour) voxels (i.e., $\mathcal{N} = \mathcal{V} \setminus \mathcal{T}$).

In Section 5.2.3, we will discuss figures that describe the daily dose distributions for each of the implementations that we test. For the non-reactive method and a given fraction i , the dose delivered to voxel $v \in \mathcal{V}$ in fraction i is

$$d_{i,v} = \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n}.$$

For the dose-reactive methods, the dose delivered to voxel v in fraction i is

$$d_{i,v} = \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p^i(x) \frac{w_b^i}{n - i + 1}.$$

To calculate underdose and overdose, we first define two sets of tumour voxels. We define \mathcal{T}_i^- , the set of tumour voxels underdosed in fraction i , as

$$\mathcal{T}_i^- = \{v \in \mathcal{T} \mid d_{i,v} < \theta_v/n\},$$

and in a similar fashion, we define \mathcal{T}_i^+ , the set of tumour voxels overdosed in fraction i as

$$\mathcal{T}_i^+ = \{v \in \mathcal{T} \mid d_{i,v} > \gamma \theta_v/n\}.$$

With these definitions, the mean tumour underdose in fraction i is defined as

$$g_i^{-,\text{mean}} = \begin{cases} \frac{1}{|\mathcal{T}_i^-|} \sum_{v \in \mathcal{T}_i^-} (\theta_v/n - d_{i,v}), & \text{if } \mathcal{T}_i^- \neq \emptyset, \\ 0, & \text{if } \mathcal{T}_i^- = \emptyset, \end{cases} \quad (5.10)$$

while the maximum tumour underdose in fraction i is defined as

$$g_i^{-,\text{max}} = \begin{cases} \max_{v \in \mathcal{T}_i^-} (\theta_v/n - d_{i,v}), & \text{if } \mathcal{T}_i^- \neq \emptyset, \\ 0, & \text{if } \mathcal{T}_i^- = \emptyset. \end{cases} \quad (5.11)$$

The definitions for mean tumour overdose and maximum tumour overdose are analogous to those for mean and maximum tumour underdose; the mean tumour overdose

in fraction i is

$$g_i^{+, \text{mean}} = \begin{cases} \frac{1}{|\mathcal{T}_i^+|} \sum_{v \in \mathcal{T}_i^+} (d_{i,v} - \gamma \theta_v / n), & \text{if } \mathcal{T}_i^+ \neq \emptyset, \\ 0, & \text{if } \mathcal{T}_i^+ = \emptyset, \end{cases} \quad (5.12)$$

and the maximum tumour overdose in fraction i is defined as

$$g_i^{+, \text{max}} = \begin{cases} \max_{v \in \mathcal{T}_i^+} (d_{i,v} - \gamma \theta_v / n), & \text{if } \mathcal{T}_i^+ \neq \emptyset, \\ 0, & \text{if } \mathcal{T}_i^+ = \emptyset. \end{cases} \quad (5.13)$$

5.2.2 Final cumulative dose results

Table 5.3 presents dose statistics for the final dose distribution obtained from the non-reactive and reactive[±] implementations, while Table 5.4 presents analogous dose statistics for the reactive⁻ and reactive⁺ implementations. Under “Implementation”, the first term in the parentheses indicates the method used: static (S); exponential smoothing with smoothing constant α and no dose reaction (ES(α)); exponential smoothing with smoothing constant α and the reactive[±] target dose update (R[±]ES(α)); exponential smoothing with smoothing constant α and the reactive⁻ target dose update (R⁻ES(α)); exponential smoothing with smoothing constant α and the reactive⁺ target dose update (R⁺ES(α)); non-reactive daily prescient (DLYP); daily prescient with the reactive target dose update (R[±]DLYP); daily prescient with the reactive⁻ target dose update (R⁻DLYP); and daily prescient with the reactive⁺ target dose update (R⁺DLYP). The second term in the parentheses, where provided, indicates the initial uncertainty set: nominal (N), robust (R) and margin (M). Under “Min. tumour dose”, the minimum tumour dose (equation (5.6)) is provided for each implementation both as an absolute value in Gy and as a percentage of the prescribed minimum dose of 72Gy. Under “Max. tumour dose”, the maximum tumour dose (equation (5.7)) is provided for each implementation as an absolute value in Gy and as a percentage of the prescribed maximum dose of 79.2Gy. Under “Mean lung dose” and “Mean n. tissue dose” the mean left lung dose (equation (5.8)) and the mean normal tissue dose (equation (5.9)) are given, respectively, as absolute values in Gy and as percentages of the corresponding metric for the static margin implementation (implementation (S,M)).

Figures 5.1, 5.2, 5.3 and 5.4 plot the minimum and maximum tumour doses versus mean left lung doses for the implementations that use exponential smoothing with α

Implementation	Min. tumour dose		Max. tumour dose		Mean lung dose		Mean n. tissue dose	
	Gy	%	Gy	%	Gy	%	Gy	%
(S,N)	67.25	93.40	82.35	103.98	17.39	85.26	9.04	88.99
(S,R)	71.37	99.13	79.28	100.10	18.19	89.20	9.44	92.93
(S,M)	72.04	100.06	78.40	98.99	20.39	100.00	10.16	100.00
(ES(0.1),N)	70.79	98.32	80.14	101.18	17.50	85.85	9.05	89.11
(ES(0.1),R)	71.80	99.72	79.24	100.05	17.83	87.43	9.19	90.46
(ES(0.1),M)	72.05	100.07	78.84	99.55	18.31	89.78	9.44	92.91
(ES(0.5),N)	71.74	99.63	79.37	100.21	17.57	86.16	9.06	89.13
(ES(0.5),R)	71.99	99.98	79.16	99.94	17.65	86.55	9.09	89.50
(ES(0.5),M)	72.03	100.04	79.06	99.83	17.80	87.27	9.16	90.10
(ES(0.9),N)	71.83	99.77	79.30	100.13	17.57	86.17	9.06	89.13
(ES(0.9),R)	71.98	99.97	79.19	99.99	17.60	86.33	9.07	89.28
(ES(0.9),M)	72.02	100.02	79.14	99.92	17.71	86.87	9.11	89.67
(DLYP)	72.00	100.00	79.20	100.00	17.57	86.18	9.06	89.13
(R [±] ES(0),N)	71.86	99.81	79.34	100.18	17.57	86.19	9.10	89.53
(R [±] ES(0),R)	71.94	99.91	79.22	100.03	17.76	87.08	9.19	90.46
(R [±] ES(0),M)	72.00	100.00	79.20	99.99	18.57	91.06	9.50	93.54
(R [±] ES(0.1),N)	72.00	100.00	79.20	100.00	17.55	86.08	9.06	89.17
(R [±] ES(0.1),R)	72.00	100.00	79.20	100.00	17.72	86.89	9.08	89.41
(R [±] ES(0.1),M)	72.00	100.00	79.20	100.00	17.92	87.89	9.18	90.32
(R [±] ES(0.5),N)	71.99	99.99	79.21	100.01	17.58	86.21	9.06	89.14
(R [±] ES(0.5),R)	71.99	99.99	79.21	100.01	17.62	86.42	9.06	89.21
(R [±] ES(0.5),M)	71.99	99.99	79.21	100.01	17.77	87.13	9.11	89.62
(R [±] ES(0.9),N)	71.99	99.99	79.21	100.01	17.58	86.22	9.06	89.13
(R [±] ES(0.9),R)	71.99	99.99	79.21	100.01	17.59	86.27	9.06	89.17
(R [±] ES(0.9),M)	71.99	99.99	79.21	100.01	17.66	86.61	9.08	89.33
(R [±] DLYP)	72.00	100.00	79.20	100.00	17.59	86.25	9.06	89.12


Table 5.3: Dose statistics for the first PMF sequence under the non-reactive and reactive methods.

set to 0, 0.1, 0.5 and 0.9 respectively. In each figure, the top plot displays the minimum tumour dose versus the mean left lung dose, while the bottom plot displays the maximum tumour dose versus the mean left lung dose. The minimum and maximum tumour doses are shown as percentages of the prescribed minimum and maximum doses of 72Gy and 79.2Gy, respectively, while the mean left lung dose is expressed as a percentage of the static margin (implementation (S,M)) mean left lung dose.


From these results, we can see that incorporating dose reaction into the adaptive robust framework leads to treatments that are clinically viable with respect to the final dose distribution. With regard to the tumour dose, the three types of dose-reactive

Implementation	Min. tumour dose		Max. tumour dose		Mean lung dose		Mean n. tissue dose	
	Gy	%	Gy	%	Gy	%	Gy	%
(R ⁻ ES(0),N)	71.89	99.85	81.60	103.03	17.52	85.93	9.16	90.14
(R ⁻ ES(0),R)	71.95	99.92	79.23	100.04	18.07	88.60	9.36	92.08
(R ⁻ ES(0),M)	72.11	100.16	79.19	99.99	19.37	94.98	9.90	97.47
(R ⁻ ES(0.1),N)	72.00	100.00	80.03	101.05	17.57	86.18	9.09	89.44
(R ⁻ ES(0.1),R)	72.00	100.00	79.25	100.06	17.78	87.19	9.16	90.12
(R ⁻ ES(0.1),M)	72.07	100.10	79.20	100.00	18.14	88.97	9.37	92.24
(R ⁻ ES(0.5),N)	72.00	100.00	79.39	100.24	17.57	86.18	9.06	89.20
(R ⁻ ES(0.5),R)	72.01	100.02	79.21	100.01	17.70	86.79	9.11	89.63
(R ⁻ ES(0.5),M)	72.07	100.10	79.21	100.01	17.85	87.56	9.18	90.38
(R ⁻ ES(0.9),N)	72.00	99.99	79.32	100.15	17.57	86.17	9.06	89.18
(R ⁻ ES(0.9),R)	72.00	100.00	79.21	100.01	17.59	86.26	9.07	89.26
(R ⁻ ES(0.9),M)	72.01	100.02	79.21	100.01	17.68	86.70	9.10	89.55
(R ⁻ DLYP)	72.00	100.00	79.20	100.00	17.57	86.17	9.06	89.13
(R ⁺ ES(0),N)	67.74	94.09	79.31	100.14	17.37	85.17	8.99	88.48
(R ⁺ ES(0),R)	71.36	99.12	79.17	99.97	17.88	87.68	9.28	91.30
(R ⁺ ES(0),M)	72.00	100.00	78.19	98.73	19.07	93.52	9.69	95.34
(R ⁺ ES(0.1),N)	70.84	98.38	79.20	100.00	17.54	86.02	9.03	88.86
(R ⁺ ES(0.1),R)	71.78	99.69	79.18	99.97	17.77	87.15	9.12	89.78
(R ⁺ ES(0.1),M)	72.00	100.00	78.69	99.35	18.00	88.29	9.25	91.04
(R ⁺ ES(0.5),N)	71.72	99.61	79.20	100.00	17.57	86.17	9.05	89.07
(R ⁺ ES(0.5),R)	71.96	99.94	79.11	99.88	17.70	86.79	9.10	89.53
(R ⁺ ES(0.5),M)	71.99	99.99	78.99	99.74	17.80	87.28	9.13	89.84
(R ⁺ ES(0.9),N)	71.81	99.74	79.20	100.01	17.57	86.19	9.05	89.09
(R ⁺ ES(0.9),R)	71.96	99.95	79.19	99.98	17.60	86.30	9.07	89.22
(R ⁺ ES(0.9),M)	72.00	99.99	79.13	99.91	17.67	86.67	9.09	89.43
(R ⁺ DLYP)	72.00	100.00	79.20	100.00	17.57	86.19	9.06	89.13

Table 5.4: Dose statistics for the first PMF sequence under the reactive⁻ and reactive⁺ methods.




(a) Min. tumour dose vs. mean left lung dose.




(b) Max. tumour dose vs. mean left lung dose.

Figure 5.1: Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0$ applied to the first PMF sequence.

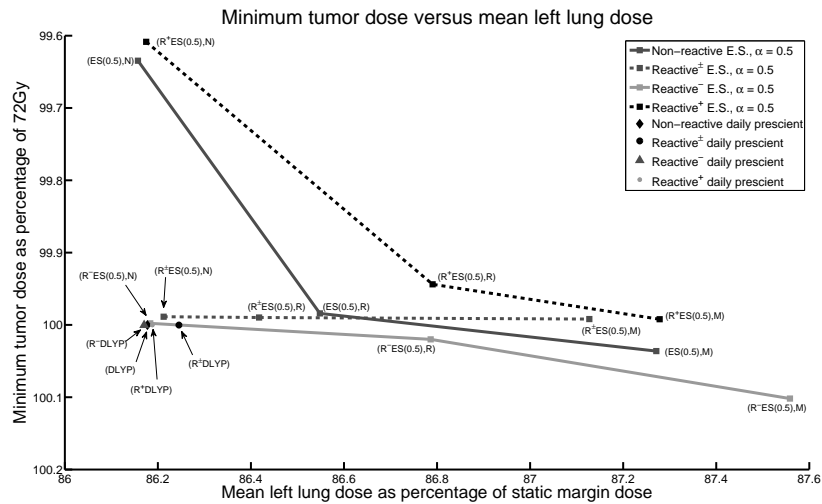


(a) Min. tumour dose vs. mean left lung dose.




(b) Max. tumour dose vs. mean left lung dose.

Figure 5.2: Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.1$ applied to the first PMF sequence.




(a) Min. tumour dose vs. mean left lung dose.




(b) Max. tumour dose vs. mean left lung dose.

Figure 5.3: Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.5$ applied to the first PMF sequence.



(a) Min. tumour dose vs. mean left lung dose.



(b) Max. tumour dose vs. mean left lung dose.

Figure 5.4: Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.9$ applied to the first PMF sequence.

methods – reactive[±], reactive⁻ and reactive⁺ – in general achieve minimum and maximum tumour doses that are very close to 100% of the prescribed minimum and maximum tumour doses of 72Gy and 79.2Gy, respectively. For example, for reactive⁻, every implementation listed in Table 5.4 achieves a minimum tumour dose that is above 99% of 72Gy, and only two achieve maximum tumour doses that are higher than 100.5% of 79.2Gy. When α is not set to zero (i.e., the uncertainty set is updated on a daily basis), Tables 5.3 and 5.4 show that all of the reactive[±], reactive⁻ and reactive⁺ implementations have mean left lung doses that are less than 90% and mean normal tissue doses that are less than 91% of their respective static margin treatment values. We can also see from Figures 5.1, 5.2, 5.3 and 5.4 that the reactive[±], reactive⁻ and reactive⁺ treatments are in general close to the daily prescient (DLYP), reactive daily prescient (R[±]DLYP), reactive⁻ daily prescient (R⁻DLYP) and reactive⁺ daily prescient (R⁺DLYP) points, indicating that treatments obtained using dose reaction are good in an absolute sense.

With respect to the minimum tumour dose specifically, the upper plots of Figures 5.1, 5.2, 5.3 and 5.4 show that the relative ordering of the dose-reactive methods, from best to worst, is reactive⁻, reactive[±], reactive⁺. This is because the reactive⁻ method ignores overdose, and always aims to deliver at least θ_v/n to each voxel v in each fraction, regardless of whether certain voxels have been previously overdosed. The reactive[±] method reacts to underdose but does not ignore overdose, so it will take advantage of situations where a voxel has previously been overdosed to deliver less dose to that voxel in future fractions, which results in a decrease in the final cumulative dose to the tumour. The reactive⁺ method ignores underdose, so naturally it results in the lowest minimum tumour dose of the three dose-reactive methods. With respect to the maximum tumour dose, the lower plots of Figures 5.1, 5.2, 5.3 and 5.4 show that the relative ordering of the dose-reactive methods, from best to worst, is reactive⁺, reactive[±], reactive⁻. This is the reverse of the previous ordering, as expected.


These results also show us that the performance of the dose-reactive methods relative to the non-reactive method seems to vary. In some cases, the dose-reactive methods lead to improvement in both tumour coverage and healthy tissue dose. For example, from the top plot of Figure 5.2, we can see that the reactive[±] point for $\alpha = 0.1$ with the robust uncertainty set (R[±]ES(0.1),R) is below and to the left of the corresponding non-reactive point (ES(0.1),R). In other cases, the dose-reactive methods lead to a worsening of both tumour coverage and healthy tissue dose. For example, for

$\alpha = 0.5$, we can see in the top plot of Figure 5.3 that the reactive⁺ points are all above and to the right of the corresponding non-reactive points, indicating a worsening of both minimum tumour dose and mean left lung dose. In yet other cases, the performance change brought about by dose-reaction is somewhat mixed. For example, from the bottom (maximum tumour dose) plot of Figure 5.4, we can see that the reactive[±] method point (R[±]ES(0.9,M)) is above and to the left of the non-reactive method point (ES(0.9),M); in other words, the mean left lung dose was improved, but the maximum tumour dose increased slightly.


5.2.3 Daily dose results

Figures 5.5, 5.6, 5.7 and 5.8 display the mean and maximum underdose by fraction (equations (5.10) and (5.11), respectively) and the mean and maximum overdose by fraction (equations (5.12) and (5.13), respectively) for the non-reactive, reactive[±], reactive⁻ and reactive⁺ implementations with different exponential smoothing constants: $\alpha = 0$ (Figure 5.5), $\alpha = 0.1$ (Figure 5.6), $\alpha = 0.5$ (Figure 5.7) and $\alpha = 0.9$ (Figure 5.8). Figure 5.9 displays analogous results for the non-reactive, reactive[±], reactive⁻ and reactive⁺ daily prescient implementations. In each of these figures, the rows of plots correspond to different target dose updates (from top to bottom, non-reactive, reactive[±], reactive⁻ and reactive⁺). In each row of each figure, the left-hand plot displays the mean and maximum underdose performance of the relevant implementations, while the right-hand plot displays the mean and maximum overdose of the relevant implementations. The mean and maximum underdose values are given as percentages of the daily prescribed minimum dose of 2.4Gy (= 72Gy/30 fractions), while the mean and maximum overdose values are given as percentages of the daily prescribed maximum dose of 2.64Gy (= 79.2Gy/30 fractions).


Figure 5.10 displays the 5th, 50th and 95th percentiles of the delivered dose for all dose-reactive and non-reactive implementations that use the exponential smoothing uncertainty set update for one example α value (0.1) and one example uncertainty set (R). The left-hand plots of the figure display where these percentiles are with respect to the daily prescribed minimum dose of 2.4Gy and the daily prescribed maximum dose of 2.64Gy; the two dashed horizontal lines correspond to these two dose values. The right-hand plots of the figure display the delivered dose percentiles with the percentiles of the normalized target minimum dose and target maximum dose distri-




(a) Mean and max. underdose, (S,·) implementations.




(b) Mean and max. overdose, (S,·) implementations.




(c) Mean and max. underdose, (R^+ ES(0),·) implementations.




(d) Mean and max. overdose, (R^+ ES(0),·) implementations.




(e) Mean and max. underdose, (R^- ES(0),·) implementations.



(f) Mean and max. overdose, (R^- ES(0),·) implementations.




(g) Mean and max. underdose, (R^+ ES(0),·) implementations.




(h) Mean and max. overdose, (R^+ ES(0),·) implementations.


Figure 5.5: Mean and max. underdose and overdose by fraction for implementations using the reactive, reactive⁻ or reactive⁺ exponential smoothing methods with $\alpha = 0$ or the static (non-adaptive, non-reactive) method, under the first PMF sequence.




(a) Mean and max. underdose, (ES(0.1), \cdot) implementations.




(b) Mean and max. overdose, (ES(0.1), \cdot) implementations.




(c) Mean and max. underdose, (R^{\pm} ES(0.1), \cdot) implementations.




(d) Mean and max. overdose, (R^{\pm} ES(0.1), \cdot) implementations.




(e) Mean and max. underdose, (R^{-} ES(0.1), \cdot) implementations.



(f) Mean and max. overdose, (R^{-} ES(0.1), \cdot) implementations.




(g) Mean and max. underdose, (R^{+} ES(0.1), \cdot) implementations.




(h) Mean and max. overdose, (R^{+} ES(0.1), \cdot) implementations.


Figure 5.6: Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive $^{\pm}$, reactive $^{-}$ or reactive $^{+}$ exponential smoothing methods with $\alpha = 0.1$, under the first PMF sequence.




(a) Mean and max. underdose, (ES(0.5),-) implementations.




(b) Mean and max. overdose, (ES(0.5),-) implementations.




(c) Mean and max. underdose, (R[±]ES(0.5),-) implementations.




(d) Mean and max. overdose, (R[±]ES(0.5),-) implementations.




(e) Mean and max. underdose, (R⁻ES(0.5),-) implementations.



(f) Mean and max. overdose, (R⁻ES(0.5),-) implementations.




(g) Mean and max. underdose, (R⁺ES(0.5),-) implementations.




(h) Mean and max. overdose, (R⁺ES(0.5),-) implementations.


Figure 5.7: Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive[±], reactive⁻ or reactive⁺ exponential smoothing methods with $\alpha = 0.5$, under the first PMF sequence.




(a) Mean and max. underdose, (ES(0.9), \cdot) implementations.




(b) Mean and max. overdose, (ES(0.9), \cdot) implementations.




(c) Mean and max. underdose, (R^{\pm} ES(0.9), \cdot) implementations.




(d) Mean and max. overdose, (R^{\pm} ES(0.9), \cdot) implementations.




(e) Mean and max. underdose, (R^{-} ES(0.9), \cdot) implementations.



(f) Mean and max. overdose, (R^{-} ES(0.9), \cdot) implementations.




(g) Mean and max. underdose, (R^{+} ES(0.9), \cdot) implementations.




(h) Mean and max. overdose, (R^{+} ES(0.9), \cdot) implementations.


Figure 5.8: Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive $^{\pm}$, reactive $^{-}$ or reactive $^{+}$ exponential smoothing methods with $\alpha = 0.9$, under the first PMF sequence.




(a) Mean and max. underdose, (DLYP) implementation.




(b) Mean and max. overdose, (DLYP) implementation.




(c) Mean and max. underdose, (R^+ DLYP) implementation.




(d) Mean and max. overdose, (R^+ DLYP) implementation.




(e) Mean and max. underdose, (R^- DLYP) implementation.



(f) Mean and max. overdose, (R^- DLYP) implementation.



(g) Mean and max. underdose, (R^+ DLYP) implementation.




(h) Mean and max. overdose, (R^+ DLYP) implementation.

Figure 5.9: Mean and max. underdose and overdose by fraction for the non-reactive, reactive[±], reactive⁻ and reactive⁺ daily prescient methods under the first PMF sequence.


butions. For each fraction i , the target minimum and maximum dose distributions for each dose-reactive method are divided by $1/(n-i+1)$ so as to be normalized and comparable to the actual delivered dose distribution of that fraction. For the non-reactive method, the target minimum and maximum dose distributions in each fraction are given by $\underline{\delta}_v^i = \theta_v$ and $\bar{\delta}_v^i = \gamma\theta_v$ for all $v \in \mathcal{T}$; to be comparable with the delivered dose distribution, they are divided by $1/n$. In both the left-hand and right-hand plots of the figure, the lower error bar, midpoint and upper error bar on each line correspond to the 5th, 50th and 95th percentiles respectively for the appropriate distribution.

From the perspective of daily dose performance, the results shown in Figures 5.5 through 5.8 indicate that the reactive[±] method does not perform as well as the non-reactive method. In particular, we can see that the reactive[±] method's daily underdose and overdose grow over the course of the treatment and exceed the levels of underdose and overdose exhibited by the non-reactive method. For example, from Figures 5.6(c) and 5.6(d), starting from the margin uncertainty set with $\alpha = 0.1$, the daily maximum underdose and overdose of the reactive[±] method are 10% of 2.4Gy and 2.64Gy, respectively, in the later fractions of the treatment, whereas the non-reactive method with the same uncertainty set and the same α results in daily underdose and overdose that are both essentially 0% at the end of treatment. This growth in underdose and overdose is also seen for the other reactive[±] implementations in Figures 5.5(c), 5.5(d), 5.7(c), 5.7(d), 5.8(c) and 5.8(d). This is in stark contrast to the non-reactive method, for which the daily underdose and overdose decrease when α is set to a positive value. When $\alpha = 0.1$, for example, the mean and maximum underdose and overdose for all of the non-reactive implementations in Figures 5.6(a) and 5.6(b) drop below 2% and stay below 2% after the eighth treatment fraction. The most striking instance of the growth in daily underdose and overdose of the reactive[±] method is with the margin uncertainty set and $\alpha = 0$, shown in Figure 5.5(c). In this particular case, the mean and maximum underdoses in fraction 30 are over 5% and 40%, respectively, while the mean and maximum overdoses in fraction 30 are over 4% and 18%, respectively.


The performance of the reactive⁻ methods with respect to daily dose is slightly better than that of the reactive[±] method, but still retains some of its flaws. For the reactive⁻ method, the daily underdose decreases as the treatment progresses, even when the uncertainty set is not modified; see, for example, Figures 5.5(e) and 5.6(e). With high values of α , the non-reactive and reactive⁻ underdose curves look almost




(a) Daily dose, $(ES(0.1),R)$ implementation.




(b) Daily dose and target minimum and maximum doses, $(ES(0.1),R)$ implementation.




(c) Daily dose, $(R^{\pm}ES(0.1),R)$ implementation.




(d) Daily dose and target minimum and maximum doses, $(R^{\pm}ES(0.1),R)$ implementation.




(e) Daily dose, $(R^{-}ES(0.1),R)$ implementation.



(f) Daily dose and target minimum and maximum doses, $(R^{-}ES(0.1),R)$ implementation.



(g) Daily dose, $(R^{+}ES(0.1),R)$ implementation.



(h) Daily dose and target minimum and maximum doses, $(R^{+}ES(0.1),R)$ implementation.

Figure 5.10: Daily delivered dose and daily target minimum and maximum doses by fraction for the $(ES(0.1),R)$, $(R^{\pm}ES(0.1),R)$, $(R^{-}ES(0.1),R)$ and $(R^{+}ES(0.1),R)$ implementations under the first PMF sequence. The lower error bar, median and upper error bar for each fraction correspond to the 5th, 50th and 95th percentiles of the appropriate distribution, respectively. On the left-hand plots, the dashed horizontal lines correspond to 2.4Gy and 2.64Gy.

identical: for $\alpha = 0.5$, for instance, compare Figures 5.7(a) and 5.7(e). This is to be expected, as the target minimum dose for each tumour voxel v under the reactive⁻ method in fraction $i+1$ is always above $(n-i)/n \cdot \theta_v$. However, the daily overdose still grows as the treatment progresses and is higher than that found in the non-reactive methods. For $\alpha = 0.1$, for instance, the mean and maximum overdose in fraction 30 for all three reactive⁻ treatments in Figure 5.6(f) appear to be more than double their corresponding values for the corresponding non-reactive treatments in Figure 5.6(a).

In a similar vein, the reactive⁺ method behaves like the reactive[±] method with respect to daily underdose. The daily overdose decreases, since the method explicitly accounts for overdose, but the daily underdose grows as the treatment progresses, and is generally higher than in the non-reactive case. For instance, in Figure 5.6(g), for the reactive⁺ method with $\alpha = 0.1$ and the margin uncertainty set, the maximum underdose in fraction 30 is over 10%. In contrast, in Figure 5.6(a), the maximum underdose in fraction 30 for the non-reactive method with $\alpha = 0.1$ and the margin uncertainty set is essentially zero. For other values of α , the same kind of difference in underdose can be observed.

For the daily prescient treatments, Figures 5.9(a) and 5.9(b) indicate that the non-reactive daily prescient treatment results in no daily underdose and no daily overdose, as we would expect. For the reactive[±] daily prescient treatment, the daily underdose and overdose both grow as the treatment progresses (Figures 5.9(c) and 5.9(d)), while the reactive⁻ and reactive⁺ daily prescient treatments exhibit only growing daily overdose (Figure 5.9(f)) and only growing daily underdose (Figure 5.9(g)), respectively. The reactive⁻ daily prescient treatment does not exhibit any daily underdose because the target minimum dose for each tumour voxel v is always $(n-i+1)/n \cdot \theta_v$ in each fraction i , and due to the correct anticipation of the PMF in each fraction, the delivered beamlet intensities in each fraction i meet this requirement. Similarly, the reactive⁺ daily prescient treatment does not exhibit any daily overdose. We will comment more on the cause of these results and their broader significance in Section 5.2.4.

The plots of delivered and target minimum and maximum dose percentiles in Figure 5.10 lead to several important observations about dose-reactive methods. In particular, from the left-hand plots, we can see that the distance between the 5th and 95th percentiles of the daily delivered dose distribution grows more for the dose-reactive methods than for the non-reactive method. This indicates that the heterogeneity of the daily delivered tumour dose distribution grows more under the dose-reactive meth-

ods than it does under the non-reactive method. From the right-hand plots, we can also see the same percentiles for the target minimum and maximum dose distributions. This group of plots shows us that, in addition to the actual daily tumour dose distribution, the daily target minimum and/or maximum dose distributions (i.e., the dose requirements) become very heterogeneous by the end of treatment under the dose-reactive methods. In contrast, the dose requirements under the non-reactive method are always constant.

5.2.4 Discussion

The results of Section 5.2 allow us to characterize the effect of dose reaction in adaptive robust radiation therapy treatments in several respects. From a final cumulative dose perspective, the dose-reactive methods perform quite well. With respect to the non-reactive method, they are comparable in final treatment quality – in some cases, the dose-reactive methods dominate the non-reactive method, while in others they do not. From a daily dose perspective, however, the dose-reactive methods do not perform as well as the non-reactive method. As the results of Section 5.2.3 show, treatments obtained via dose-reactive methods exhibit growing daily underdose and/or overdose, whereas for the non-reactive method the daily underdose and overdose generally stay roughly constant (for $\alpha = 0$) or decrease as the treatment progresses (for $\alpha > 0$).

Related to the daily dose performance, our computational experiments uncover an important interaction between uncertainty set adaptation and dose reaction. In particular, the dose-reactive implementations with non-zero α , which do adaptively update the uncertainty set, generally result in slower growth in daily underdose and overdose than the dose-reactive implementations with α set to zero, which do not modify the uncertainty set. This result is new to the dose-reactive literature. In Webb [2008], for example, the proposed framework possesses a dose-reactive element, but does not update any estimate of the uncertainty, and the dose distribution planned in each fraction is always planned on the tumour being in the same fixed reference position. Our results indicate that by using dose reaction with uncertainty set adaptation – or more generally, planning the daily dose distribution with some updated estimate of the uncertainty – one can control the growth in daily underdose and overdose to some extent.

There are two important questions that need to be asked in light of the difference

between the non-reactive and dose-reactive methods in daily dose performance. The first question is: what gives rise to the daily dose performance of the dose-reactive implementations? Intuitively, as one repeatedly reacts to overdose, the future target minimum and maximum doses will be adjusted further down, and so in future fractions the tumour will on average receive a dose that is increasingly lower than 2.4Gy. Similarly, as one repeatedly reacts to underdose, the future target minimum and maximum doses will be adjusted further up, resulting in the tumour receiving a dose that is increasingly greater than 2.64Gy.

However, there is a deeper insight that helps us understand this question. Consider Figure 5.9, which displays the daily underdose and overdose performance of each of the daily prescient implementations. This figure provides us with an interesting insight: even when we correctly anticipate the breathing motion PMF in each fraction, performing an update to the target dose distribution can lead to growing daily underdose (reactive⁺ daily prescient, (R⁺DLYP)), growing daily overdose (reactive⁻ daily prescient, (R⁻DLYP)) or both (reactive[±] daily prescient, (R[±]DLYP)). The reason this behavior occurs in the case of these daily prescient algorithms is because for a fixed breathing motion PMF, it is in general not possible to deliver a completely uniform dose to the entire tumour. In addition to this, beamlet intensity vectors that achieve highly homogeneous tumour doses can be very different from beamlet intensity vectors that aim to minimize the total dose delivered to the patient. Therefore, in a given fraction there will be heterogeneity in the delivered dose across the tumour, resulting in heterogeneity in the target minimum and/or maximum doses for the next fraction. The heterogeneity in the target minimum and/or maximum doses for the next fraction then leads to further heterogeneity in the dose delivered in the next fraction. As this process continues, it inevitably results in the daily dose for some voxels deviating below 2.4Gy and for some voxels deviating above 2.64Gy in the later fractions of the treatment.

This process is further exacerbated when the prescience is removed and the uncertainty set is updated according to exponential smoothing. For these methods (R[±]ES, R⁻ES and R⁻ES), the daily PMF may fall outside the uncertainty set, and thus the delivered dose need not satisfy the target minimum and maximum dose requirements. As a result, more underdose and overdose may be realized in each fraction, and the target minimum and maximum dose requirements may grow more heterogeneous as the treatment progresses. A concrete demonstration of the growth in tumour dose

heterogeneity and target minimum and maximum dose requirement heterogeneity is shown in Figure 5.10. We can see that the heterogeneity of the dose distribution grows much more for the reactive[±], reactive⁻ and reactive⁺ methods than for the non-reactive method. On the right-hand side, we see that for the non-reactive method, these requirements stay constant and there is no change in their spread. However, for the dose-reactive implementations, the target minimum and maximum dose distributions change significantly. For instance, for the reactive[±] method, we can see that the median normalized values of the target minimum and maximum dose distributions deviate significantly from their respective starting values of 2.4Gy and 2.6Gy. Furthermore, the spread of the target minimum and maximum dose distributions also increases, as shown in the increased separation between the lower and upper error bars (corresponding to the 5th and 95th percentiles, respectively) as the treatment progresses.

Relative to the reactive[±] method, the reactive⁻ method exhibits similar growth in the target maximum dose but much less change in the heterogeneity of its target minimum dose distribution. This is because the target minimum dose for the reactive⁻ method is always constrained to be at least $(n - i + 1)/n \cdot \theta_v$, and is not allowed to go below this value. At the same time, because the uncertainty set is being updated and the daily uncertainty set is becoming a more accurate representation of the patient's daily PMF, we expect the underdose that we encounter relative to the target minimum dose in each fraction to decrease as the treatment goes on. As a result, we do not expect to see the explosive growth in the target minimum dose heterogeneity for the reactive⁻ method as we do for the reactive[±] method. For the reactive⁺ method, we see that the heterogeneity of the target minimum dose distribution grows as it does for the reactive[±] method, but the target maximum dose distribution does not change significantly; this behavior can be explained in the same way as for the reactive⁻ method.

The second question that needs to be asked with regard to daily dose performance is: why is the daily dose performance important? As noted in Section 2.4, the outcome of a given treatment depends not only on the cumulative tumour dose that was delivered by the end of the treatment, but also on *how* that dose was delivered to the tumour. For example, if parts of the tumour are underdosed early in the treatment, then fewer tumour cells will be killed, and the additional tumour cells that survive will proliferate. This may lead to tumour growth that is not accounted for by the treatment and ultimately, it may result in the failure of the treatment. Indeed, it has

been shown in the medical physics literature that an increase in the intrafraction dose variability (i.e., increases in the heterogeneity of the daily dose distribution) leads to a decrease in tumour control probability and an increase in the probability of treatment failure (Ebert [2000]). We have shown that dose-reactive methods generally lead to increased heterogeneity in the daily dose distribution, as well as increased daily underdose and/or overdose, over the non-reactive method. Therefore, a treatment resulting from a dose-reactive method may be more likely to fail than a treatment resulting from the non-reactive method, even if both treatments achieve similar final dose distributions.

5.3 Theoretical insight into daily dose performance

In addition to the computational results that we have presented, we can gain some more insight into the daily dose performance by building on the theory developed in Section 4.2. There, the non-reactive adaptive robust method was analyzed by assuming that the sequence of breathing motion PMFs, as an infinite sequence, converges to a limiting PMF, and evaluating the dose distribution of an n -fraction treatment as n tends to infinity. The key result in the asymptotic analysis was that for a wide class of uncertainty set update algorithms, the final dose distribution approaches a set of optimal dose distributions that exhibit no tumour underdose or overdose and low healthy tissue dose. The same theoretical framework allows us to characterize the daily performance of the non-reactive method, which we do next. Specifically, we will show that under non-reactive methods, the daily underdose and overdose disappear as the treatment progresses. We also use this framework to explain why the same kind of guarantee is unlikely to extend to dose-reactive methods.

We define the *epsilon neighborhood* of a subset $V \subseteq \mathbb{R}^n$ as

$$U(V, \epsilon) = \bigcup_{y \in V} B(y, \epsilon),$$

where $B(y, \epsilon) = \{y' \in \mathbb{R}^n \mid \|y' - y\| < \epsilon\}$ is the open ball of radius ϵ about y . We take $\|\cdot\|$ to be the 1-norm on the appropriate finite-dimensional real vector space, and \mathbb{Z}_+ to be the set of positive integers. We let $\mathbf{w}^*(\ell^i, \mathbf{u}^i)$ denote the set of optimal solutions to the robust problem (3.1) with uncertainty set P^i defined by the lower and upper bound

vectors ℓ^i and \mathbf{u}^i . We also let the product $\Delta\mathbf{p}\mathbf{w}$ denote the vector defined as

$$\Delta\mathbf{p}\mathbf{w} = \left[\sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b \right]_{v \in \mathcal{V}},$$

i.e., it is the dose distribution that results from delivering the beamlet intensity vector \mathbf{w} in a treatment session while the patient breathes according to the breathing motion PMF \mathbf{p} . Let $\mathbf{w}^*(\mathbf{p}^*)$ denote the optimal solution to the robust problem when $P = \{\mathbf{p}^*\}$ (i.e., the uncertainty set is a singleton), and define the set \mathbf{D} as

$$\mathbf{D} = \{ \mathbf{d} \in \mathbb{R}^{|\mathcal{V}|} \mid \mathbf{d} = \Delta\mathbf{p}^*\mathbf{w} \text{ for some } \mathbf{w} \in \mathbf{w}^*(\mathbf{p}^*) \}.$$

The set of dose distributions \mathbf{D} is a set of optimal dose distributions. In particular, for every dose distribution $\mathbf{d} \in \mathbf{D}$ we have that the prescribed minimum and maximum tumour dose requirements are met; mathematically, $\theta_v \leq d_v \leq \gamma\theta_v$ for every tumour voxel $v \in \mathcal{T}$. Furthermore, the healthy tissue dose associated with the distributions in \mathbf{D} is low because the uncertainty set associated with \mathbf{w} is a singleton ($P = \{\mathbf{p}^*\}$). As discussed in Bortfeld et al. [2008], a beamlet intensity vector protecting against a small uncertainty set will generally result in less healthy tissue dose than one that protects against a larger uncertainty set.

With these definitions, the following result describing the daily performance of the non-reactive adaptive robust method is a consequence of the theory presented in Section 4.2.

Theorem 4 *Let $(\mathbf{p}^i)_{i=1}^\infty$ be a sequence of PMFs that converges to \mathbf{p}^* . Let $(\ell^i)_{i=1}^\infty$ and $(\mathbf{u}^i)_{i=1}^\infty$ be lower and upper bound sequences generated from $(\mathbf{p}^i)_{i=1}^\infty$ by any convex-convergent update algorithm. For each $i \in \mathbb{Z}_+$, let $\mathbf{w}^i \in \mathbf{w}^*(\ell^i, \mathbf{u}^i)$. Then for every $\epsilon > 0$, there exists an $N \in \mathbb{Z}_+$ such that for all $i > N$,*

$$\Delta\mathbf{p}^i\mathbf{w}^i \in U(\mathbf{D}, \epsilon).$$

Proof: See Appendix C. ■

This result states that in the limit, the dose distribution $\Delta\mathbf{p}^i\mathbf{w}^i$ approaches the set of optimal dose distributions \mathbf{D} . As stated above, the dose distributions in \mathbf{D} meet the full tumour dose requirements; since the dose distribution $\Delta\mathbf{p}^i\mathbf{w}^i$ approaches \mathbf{D} , in the

limit it also meets the full tumour dose requirements. Note that in non-reactive treatments, the patient is treated with $(1/n) \cdot \mathbf{w}^i$ and not \mathbf{w}^i ; therefore, the actual delivered dose distribution in fraction i is $(1/n) \cdot \Delta \mathbf{p}^i \mathbf{w}^i$. From the above result, it is clear that the actual daily delivered dose distribution $(1/n) \cdot \Delta \mathbf{p}^i \mathbf{w}^i$ tends to $(1/n) \cdot \mathbf{d}$ for some $\mathbf{d} \in \mathbf{D}$, and the dose distribution $(1/n) \cdot \mathbf{d}$ is such that each tumour voxel v receives exactly between θ_v/n and $\gamma\theta_v/n$. The insight from this result is that in the limit, the non-reactive adaptive robust method will *not* exhibit any daily underdose or overdose in any tumour voxel v with respect to the daily prescribed minimum and maximum doses θ_v/n and $\gamma\theta_v/n$.

For the dose-reactive methods, Theorem 4 does not apply. The proof of Theorem 4 relies on expressing the robust problem (3.1) in an equivalent alternative form in which the minimum and maximum tumour dose constraints are enforced only at the vertices of the uncertainty set P . With this alternative robust problem in hand, one needs to show that the affine functions defining the feasible region of the alternative problem converge to an affine function defining the feasible region of the robust problem with the uncertainty set $P = \{\mathbf{p}^*\}$. To be more specific, in the non-reactive case, one defines for each fraction i an affine function $f^i : \mathbb{R}^{|\mathcal{B}|} \rightarrow \mathbb{R}^{2M|\mathcal{T}|+|\mathcal{B}|+1}$ as

$$f^i(\mathbf{w}) = \mathbf{A}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})\mathbf{w} + \mathbf{c}.$$

In this definition, the PMFs $\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}$ are M PMFs obtained by iterating the chosen uncertainty set update algorithm $i - 1$ times on M PMFs $\mathbf{q}^{1,1}, \dots, \mathbf{q}^{M,1}$ whose convex hull is the initial uncertainty set P^1 (defined by lower and upper bound vectors ℓ^1 and \mathbf{u}^1). The matrix $\mathbf{A}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})$ is a $(2M|\mathcal{T}| + |\mathcal{B}| + 1)$ -by- $|\mathcal{B}|$ matrix that depends on $\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}$ and \mathbf{c} is a $2M|\mathcal{T}| + |\mathcal{B}| + 1$ -element column vector. The vector \mathbf{c} is defined as

$$\mathbf{c} = \begin{bmatrix} -\gamma \mathbf{z} \\ \mathbf{z} \\ \mathbf{0} \\ 0 \end{bmatrix},$$

where $\mathbf{0}$ is a $|\mathcal{B}|$ -component column vector of zeroes, 0 is the scalar zero and \mathbf{z} is a column vector defined (component-wise) as

$$z_{(v,j)} = \theta_v,$$

for each $v \in \mathcal{T}$ and $j \in \{1, \dots, M\}$. The vector \mathbf{c} is used to represent the dose requirements in the right-hand side of the constraints of the robust problem (3.1). In the non-reactive case, we know from Proposition 5 in Appendix A that as $i \rightarrow \infty$, the affine function $f^i(\mathbf{w}) = \mathbf{A}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})\mathbf{w} + \mathbf{c}$ converges to the affine function $f(\mathbf{w}) = \mathbf{A}(\mathbf{p}^*, \dots, \mathbf{p}^*)\mathbf{w} + \mathbf{c}$, which defines the feasible region of the robust problem with its uncertainty set equal to $\{\mathbf{p}^*\}$. This convergence is a direct consequence of the fact that the matrix $\mathbf{A}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})$ converges to the limiting matrix $\mathbf{A}(\mathbf{p}^*, \dots, \mathbf{p}^*)$, and that the offset \mathbf{c} is the same for each affine function and therefore converges.

In the dose-reactive case, the affine function that defines the feasible region of the i th problem will not only have a matrix $\mathbf{A}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})$ that is dependent on the fraction, but also an offset that depends on the fraction, \mathbf{c}^i , because the target minimum and maximum dose requirements are changing in each fraction. In the right-hand side plots of Figure 5.10 we showed how the target minimum and maximum dose requirements, when scaled in each fraction i by a factor of $1/(n-i+1)$, change as the treatment progresses for a given choice of exponential smoothing factor α and initial uncertainty set. From this figure, we can see that the normalized target minimum and maximum dose requirements for all of the voxels do not converge to θ_v/n and $\gamma\theta_v/n$ by the end of the treatment; the median target minimum and maximum dose requirements diverge from θ_v/n and $\gamma\theta_v/n$, and the spread of these dose requirements, as indicated by the distance between the lower and upper error bars, also increases. Since these normalized target minimum and maximum dose requirements do not converge, the affine functions defining the feasible regions of the modified robust problem (5.1) will not converge to the affine function defining the feasible region of the robust problem with the uncertainty set $\{\mathbf{p}^*\}$. As a result, it seems that in the limit, the dose-reactive methods will exhibit underdose or overdose in certain tumour voxels with respect to the daily prescribed minimum and maximum doses θ_v/n and $\gamma\theta_v/n$, respectively.

Furthermore, it seems unlikely that the daily dose distribution under any of the dose-reactive methods will converge under a different set of assumptions. Our computational results show that even when we consider the daily prescient treatments, which perfectly anticipate the daily breathing motion PMF in every fraction, the dose-reactive methods still lead to growing daily underdose and/or overdose (cf. Figure 5.9). These results give us strong reason to believe that in general, the daily dose distribution under any of the dose-reactive methods will not converge, and will not exhibit daily underdose and overdose that diminishes in the limit.

5.4 Conclusions

In this chapter, we considered dose-reactive methods in an adaptive and robust framework for lung cancer IMRT treatment planning. We formally described three different dose-reactive methods: the reactive[±] method, which accounts for both underdose and overdose; the reactive⁻ method, which accounts only for underdose; and the reactive⁺ method, which accounts only for overdose. We compared the performance of these dose-reactive methods with the performance of the non-reactive method computationally using real patient data. Through these computational tests, we found that although dose-reaction can be beneficial in terms of improving both overall end-of-treatment tumour coverage and healthy tissue dose from the levels attained by the non-reactive method, dose-reactive methods result in growing daily underdose and overdose. Moreover, they result in growing daily tumour dose heterogeneity. In contrast, the non-reactive method performs well from a daily dose perspective, with daily underdose and overdose decreasing as the treatment progresses. We also proved theoretically that the daily dose distribution from the non-reactive method will converge to an ideal daily dose distribution that exhibits no tumour underdose or overdose. The same theoretical framework, together with our computational results, suggests that the dose-reactive methods will not exhibit the same convergence.

The characteristic behavior of the dose-reactive methods at the daily level arises from the fact that there is always some heterogeneity in the delivered dose in a given fraction, which causes the dose requirements for subsequent fractions to become more heterogeneous. This effect is further compounded by the fact that there will be some error in the delivered dose relative to the target dose, due to the PMF in each fraction not always falling inside the uncertainty set for that fraction. As shown in prior research from the medical physics community, such effects can have a significant effect on the clinical treatment outcome. Our results thus highlight the need for treatment planners to be cognizant of the type of daily performance we have described when considering dose-reaction as a mechanism for combatting uncertainty in an ART setting. Our results also affirm the current clinical emphasis on tumour dose homogeneity because as we have shown, performance in this objective is easily jeopardized if there is too much focus on both minimizing healthy tissue dose and precisely meeting a final dose target for the tumour.

A relevant direction for future work would be to examine the difference between

the dose-reactive and non-reactive methods from a biological/treatment outcome perspective, rather than from a physical dose perspective. This could be done by using an appropriate tumour control probability (TCP) model, for example. On the theoretical side, a relevant challenge that remains is to characterize the limiting behavior of the final dose distribution arising from a dose-reactive method under the assumption that the PMFs converge. Another challenge is to theoretically quantify how much the daily underdose and/or overdose from dose-reactive methods grow over the treatment course and how this growth depends on the problem data, such as the specific dose deposition coefficients, the anatomical geometry and the stability of the patient's breathing motion PMF sequence.

Chapter 6

Conclusions and future work

In this thesis, we have developed a novel methodology for lung cancer IMRT treatments that combines offline adaptive radiation therapy and robust optimization. In the first part of this thesis, we developed this method and showed, both theoretically and computationally, that it leads to improved performance over the existing static robust optimization approach on both tumour coverage and healthy tissue dose. By improving on both tumour coverage and healthy tissue dose, this method can enable treatment planners to safely escalate the dose, leading to improvements in treatment outcome while maintaining the same level of toxicity to the lungs.

In the second part of this thesis, we considered an adaptive and robust method that also incorporates dose-reaction, and we showed computationally that although such a modification can lead to improvements in the overall dose distribution, it can have a drastic effect on the daily dose distribution, resulting in increasing daily under-dose and/or overdose. We also provided theoretical justification for why treatments that modify the uncertainty set but do not modify dose requirements from fraction to fraction should outperform treatments that also modify the dose requirements from fraction. For biological reasons, the variability in the daily dose distribution may detrimentally affect the treatment outcome, which diminishes the value of dose-reaction as a strategy for managing uncertainty in lung cancer IMRT treatments.

With regard to future work, we have already highlighted some potential directions in Sections 4.4 and 5.4. At a higher level, the treatment of lung cancer using IMRT requires a better understanding of the breathing motion uncertainty that is present. For example, the robust optimization approach that is employed in this paper is based on the use of breathing motion PMFs that describe the proportion of time that the

patient spends in each of finitely many phases of the breathing cycle. These relative frequencies are calculated from a continuous signal that represents the displacement of tumour as a function of time. By aggregating this continuous signal into a single PMF, one loses sight of much information that may be useful. For example, within any given breathing cycle, the time the tumour spends in each phase does not exhibit significant variability; it may be possible to use such information to guide the design of uncertainty sets. Similarly, there has been little research into the statistical properties of patient PMF sequences. With enough data, it may be possible to identify which kinds of assumptions are reasonable and incorporate these into treatment planning models.

Another major direction for lung cancer IMRT research is the use of functional imaging data. In the models that we have considered in this thesis, the treatment plan is designed based on computed tomography (CT) images, which only allow for the geometric delineation of different anatomical structures. However, there have been recent developments in functional imaging modalities, such as positron emission tomography (PET), magnetic resonance imaging (MRI) and single photon emission computed tomography (SPECT). These imaging modalities are distinct from CT in that they provide information about biological properties in different structures. For example, PET and MRI can be used to measure how hypoxic (lacking in oxygen) different parts of the tumour are, and SPECT can be used to measure the perfusion capacity of different areas of lung tissue. These biological properties are of interest because they all affect radiation therapy in some way. For instance, parts of the tumour that are hypoxic are known to be more resistant to radiation, which warrants an escalation of the dose to those parts. Similarly, by knowing the perfusion capacity of different areas of lung tissue, the treatment can be designed to deliver less dose to those functional areas. Functional imaging leads to further questions:

- Given functional images pre-treatment, how should that data be incorporated into an optimization model?
- Given repeated functional images during treatment, how should the treatment be adapted to changes in these functional images? Biological properties can vary over the treatment course; for example, the spatial distribution of hypoxia within a tumour may change due to temporary occlusions of blood vessels and due to reoxygenation induced by radiation damage.

- Given uncertainty in functional imaging data, how should the treatment be designed? For instance, PET can be used to measure the metabolic activity within a tumour, which can then be used to target the most metabolically active areas of the tumour. However, healthy tissue at the boundary of the tumour can exhibit an inflammatory response due to irradiation; when this occurs, the metabolic activity of the healthy tissue also increases. In such cases, there may be uncertainty as to whether a region of high metabolic activity at the tumour boundary is due to new tumour growth, or damaged healthy tissue.

All of these issues remain important problems in incorporating functional imaging information into adaptive radiation therapy treatments for lung cancer.

Lastly, it is worth taking a more comprehensive view of lung cancer. In this thesis, we have focused on how radiation therapy treatments for lung cancer can be improved. However, in optimizing the design of radiation therapy treatments, we must not lose sight of the fact that radiation therapy has certain fundamental limits as a technology and may not be the best solution approach. Perhaps the best illustration of this thought is in Table 4.1 of Chapter 4, where the adaptive robust optimization method is compared to the prescient solutions for the first PMF sequence. We have already presented one way of reading the results in this table: the adaptive robust method attains performance comparable to that of the prescient solutions. Yet there is another, more sobering way of reading these results: even with perfect foresight, as simulated by the prescient solutions, the best we can hope for is a mean left lung dose of 17.54Gy. From an engineering perspective, we should indeed ask how close we can come to such barriers – but we must also ask whether those barriers can be broken.

Already there are technologies that promise this. In this thesis, we have examined the use of conventional photon-based radiation therapy. An alternative to this is proton therapy, where instead of photons, protons are used to damage the DNA of cancer cells. The key difference between protons and photons lies in their energy deposition profile (Schulz-Ertner and Tsujii [2007]). As photons pass through tissue, the energy deposited increases and reaches its peak early, and subsequently falls off slowly. In contrast, as protons pass through tissue, the energy deposited increases slowly and reaches its peak later, and then drops off much more sharply. The clinical consequence of this physical difference is that protons can be used to treat tumours with less damage to healthy tissue than photons. One challenge with using proton

therapy is that due to artifacts in CT images and uncertainties regarding tissue densities, there can be some uncertainty in when the point of maximal energy deposition (known as the Bragg peak) occurs, which has prompted research into appropriate optimization models (Unkelbach et al. [2007], Fredriksson et al. [2011]). As these types of issues are resolved, we can expect proton therapy to become a better option for the treatment of lung cancer than conventional photon therapy.

But even with these possibilities, we need to remember that these are technologies for treatment; they do nothing to prevent new cases of cancer, which numbered approximately 1.6 million in the United States (American Cancer Society [2011]) and 170,000 in Canada (Canadian Cancer Society's Steering Committee [2011]). At the same time, the American Cancer Society estimates that of the 570,000 deaths in the United States that are due to cancer in 2011, approximately 170,000 of those (just under one third) are expected to be attributable to tobacco use and that one third of those 570,000 deaths will also be related to "[obesity], physical inactivity and poor nutrition" (American Cancer Society [2011]). Similarly, Anand et al. [2008] estimates that only 5-10% of all cancer cases originate from genetic defects, while the remaining 90-95% are due to environment and lifestyle. In light of these statistics, one may therefore ask how we can effect appropriate changes in lifestyle and environmental factors to reduce the incidence of cancer. Effecting such changes will by no means be easy, but the benefits to the healthcare system and society at large would be quite significant. The advancement of cancer prevention, therefore, represents an important complementary research effort to the advancement of treatment technology.

Appendix A

Proofs for Chapter 4

In this section, we provide proofs of the results in Section 4.2. The first subsection describes the assumptions that were made in the analysis; all subsequent subsections follow the organization of Section 4.2.

A.1 Assumptions

Assumption 1 *For every voxel $v \in \mathcal{T}$ and motion state $x \in X$, there exists a beamlet $b \in \mathcal{B}$ such that $\Delta_{v,x,b} > 0$.*

This assumption is reasonable because it simply asks for the number of beamlets to not be too small. If this assumption does not hold, then the margin problem is infeasible.

Let γ^* be defined as the optimal objective value of the following linear program:

$$\begin{aligned} & \text{minimize } \gamma \\ & \text{subject to } \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b \geq \theta_v, \quad \forall v \in \mathcal{T}, x \in X, \\ & \quad \quad \quad \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b \leq \gamma \theta_v, \quad \forall v \in \mathcal{T}, x \in X, \\ & \quad \quad \quad \gamma \geq 1, \\ & \quad \quad \quad w_b \geq 0, \quad \forall b \in \mathcal{B}. \end{aligned} \tag{A.1}$$

It is the lowest value of the parameter γ of problem (3.1) under which the margin problem remains feasible. Problem (A.1) is bounded because $\gamma \geq 1$, and feasible because of Assumption 1. Our next assumption is on how the maximum dose parameter γ is

selected:

Assumption 2 *The parameter γ of problem (3.1) satisfies $\gamma > \gamma^*$.*

If γ is strictly less than γ^* , then the margin problem is infeasible. If γ is exactly equal to γ^* , then every feasible solution of the margin is such that there is some $v \in \mathcal{T}$ and some $x \in X$ such that $\sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b = \gamma \theta_v$. By setting γ to be slightly higher than γ^* , we allow for a slightly higher amount of dose in the tumour. The importance of the maximum dose constraint in the formulation of the robust problem is to control the homogeneity of the tumour dose (see Bortfeld et al. [2008]); the homogeneity we lose by increasing γ by a small amount above γ^* is negligible. This assumption is satisfied for the data used in Section 4.3.

An immediate consequence of Assumptions 1 and 2 is the following proposition.

Proposition 3 *There exists a beamlet intensity vector \mathbf{w} such that for every $b \in \mathcal{B}$, $w_b > 0$, and for every $v \in \mathcal{T}$ and every $x \in X$, $\theta_v < \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b < \gamma \theta_v$.*

Proof: Since γ^* is the minimum of problem (A.1), there must be an intensity vector \mathbf{w} that is feasible for γ^* , which means that for every $v \in \mathcal{T}$ and every $x \in X$, $\theta_v \leq \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b \leq \gamma^* \theta_v$ and for every $b \in \mathcal{B}$, $w_b \geq 0$. Since $\gamma > \gamma^*$ by Assumption 2, this means that for every $v \in \mathcal{T}$ and $x \in X$, $\theta_v \leq \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b < \gamma \theta_v$. We now need to carefully modify \mathbf{w} to satisfy the conditions of the proposition. In Step 1, we modify \mathbf{w} to obtain a \mathbf{w}' for which $\theta_v < \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b$ for every $v \in \mathcal{T}$ and $x \in X$. In Step 2, we take the \mathbf{w}' of Step 1, and modify it so that all of the beamlet intensities are positive.

Step 1. Define the set L as

$$L = \left\{ (v, x) \in \mathcal{T} \times X \mid \theta_v = \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b \right\}.$$

L is the set of voxel and motion state pairs for which the constraint $\theta_v \leq \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b$ is active. If L is empty, we go to Step 2, substituting \mathbf{w} for \mathbf{w}' . Otherwise, for each $(v, x) \in L$, Assumption 1 furnishes a beamlet b such that $\Delta_{v,x,b} > 0$. Let $\tilde{\mathcal{B}}$ be the set of all such beamlets; that is,

$$\tilde{\mathcal{B}} = \{b \in \mathcal{B} \mid \Delta_{v,x,b} > 0 \text{ for some } (v, x) \in L\}.$$

Let $\mathbf{e}_{\tilde{\mathcal{B}}}$ be a vector which is 1 at every beamlet $b \in \tilde{\mathcal{B}}$ and 0 at every other beamlet. Let the constant k be defined as

$$k = \min \left\{ \frac{\gamma\theta_v - \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b}{\sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b}} \mid (v,x) \in \mathcal{T} \times X \right\}.$$

The constant k is positive and finite because for each $(v,x) \in L$, $\sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b}$ is non-zero and $\gamma\theta_v - \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b$ is positive.

Now, define a new intensity vector $\mathbf{w}' = \mathbf{w} + (k/2)\mathbf{e}_{\tilde{\mathcal{B}}}$. For every $(v,x) \in L$, we then have

$$\begin{aligned} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b &= \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b + \sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b} \frac{k}{2} \\ &> \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b \\ &= \theta_v, \end{aligned}$$

where the strict inequality is due to $\sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b} > 0$ and $k > 0$. For $(v,x) \notin L$,

$$\begin{aligned} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b &= \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b + \sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b} \frac{k}{2} \\ &\geq \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b \\ &> \theta_v, \end{aligned}$$

because $\sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b} k$ is non-negative. Next we check that \mathbf{w}' still strictly satisfies the upper bound constraint. For every $(v,x) \in \mathcal{T} \times X$ such that $\sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b} > 0$, we have

$$\begin{aligned} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b &= \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b + \sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b} \frac{k}{2} \\ &< \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b + \sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b} k \\ &\leq \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b + \sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b} \cdot \frac{\gamma\theta_v - \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b}{\sum_{b \in \tilde{\mathcal{B}}} \Delta_{v,x,b}} \\ &= \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b + \gamma\theta_v - \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b \\ &= \gamma\theta_v, \end{aligned}$$

where the first inequality follows because k is positive and the second inequality follows from the definition of k . For every $(v, x) \in \mathcal{T} \times X$ such that $\sum_{b \in \bar{\mathcal{B}}} \Delta_{v,x,b} = 0$, we simply have

$$\begin{aligned} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b &= \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b + \sum_{b \in \bar{\mathcal{B}}} \Delta_{v,x,b} \frac{k}{2} \\ &= \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b \\ &< \gamma \theta_v. \end{aligned}$$

Therefore, \mathbf{w}' satisfies

$$\theta_v < \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b < \gamma \theta_v$$

for every tumour voxel v and motion state x .

Step 2. Suppose we have a \mathbf{w}' that satisfies

$$\theta_v < \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b < \gamma \theta_v$$

for every tumour voxel v and motion state x . Let \mathcal{B}_0 be the set of beamlet indices b such that $w'_b = 0$. If \mathcal{B}_0 is empty, we are done; otherwise, let $\mathbf{e}_{\mathcal{B}_0}$ be the vector which is 1 at every beamlet $b \in \mathcal{B}_0$ and 0 at every other beamlet.

If $\sum_{b \in \mathcal{B}_0} \Delta_{v,x,b} = 0$ for every $(v, x) \in \mathcal{T} \times X$, then select any $y > 0$ and define $\bar{\mathbf{w}} = \mathbf{w}' + y \mathbf{e}_{\mathcal{B}_0}$. In this case, we are done; it is easy to check that

$$\theta_v < \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{w}_b < \gamma \theta_v$$

for every $v \in \mathcal{T}$ and $x \in X$, and also that $\bar{w}_b > 0$ for every $b \in \mathcal{B}$.

If $\sum_{b \in \mathcal{B}_0} \Delta_{v,x,b} > 0$ for some $(v, x) \in \mathcal{T} \times X$, then define y as

$$y = \min \left\{ \frac{\gamma \theta_v - \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b}{\sum_{b \in \mathcal{B}_0} \Delta_{v,x,b}} \mid (v, x) \in \mathcal{T} \times X \right\}.$$

The constant y is positive and finite because $\gamma \theta_v - \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b > 0$ and $\sum_{b \in \mathcal{B}_0} \Delta_{v,x,b} > 0$. Let $\bar{\mathbf{w}} = \mathbf{w}' + (y/2) \mathbf{e}_{\mathcal{B}_0}$. Since each beamlet's intensity has either stayed the same or

increased, it is clear that

$$\theta_v < \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{w}_b$$

holds for every $v \in \mathcal{T}$, $x \in X$. For every $(v, x) \in \mathcal{T} \times X$ such that $\sum_{b \in \mathcal{B}_0} \Delta_{v,x,b} > 0$, we have

$$\begin{aligned} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{w}_b &= \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b + \sum_{b \in \mathcal{B}_0} \Delta_{v,x,b} \frac{y}{2} \\ &< \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b + \sum_{b \in \mathcal{B}_0} \Delta_{v,x,b} y \\ &\leq \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b + \sum_{b \in \mathcal{B}_0} \Delta_{v,x,b} \left(\frac{\gamma \theta_v - \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b}{\sum_{b \in \mathcal{B}_0} \Delta_{v,x,b}} \right) \\ &= \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b + \gamma \theta_v - \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b \\ &= \gamma \theta_v, \end{aligned}$$

where the first inequality follows from y being positive and the second inequality follows from the definition of y . For every $(v, x) \in \mathcal{T} \times X$ such that $\sum_{b \in \mathcal{B}_0} \Delta_{v,x,b} = 0$, we simply have

$$\begin{aligned} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{w}_b &= \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b + \sum_{b \in \mathcal{B}_0} \Delta_{v,x,b} \frac{y}{2} \\ &= \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w'_b \\ &< \gamma \theta_v. \end{aligned}$$

In either case, \bar{w} satisfies the condition of the proposition, for we have that

$$\theta_v < \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{w}_b < \gamma \theta_v$$

for every $v \in \mathcal{T}$ and $x \in X$ and that $\bar{w}_b > 0$ for every beamlet $b \in \mathcal{B}$. ■

Assumption 3 For every beamlet $b \in \mathcal{B}$ and every phase $x \in X$, there exists a voxel $v \in \mathcal{V}$ such that $\Delta_{v,x,b} > 0$.

This assumption is reasonable, because it simply asks that the patient geometry accounts for a sufficiently large amount of the healthy tissue around the tumour. This

condition is satisfied for the data used in Section 4.3.

The importance of Assumption 3 is that it implies that the robust problem has a bounded optimal solution set.

Proposition 4 *The set of optimal solutions to the robust problem (3.1) for any uncertainty set P is a bounded set.*

Proof: Let \mathbf{w}^* be an optimal solution of the robust problem (3.1), and let $b' \in \mathcal{B}$. We know already that $w_{b'}^* \geq 0$. Since $\sum_{x \in X} \bar{p}(x) = 1$, there must be a phase x' for which $\bar{p}(x') > 0$. By Assumption 3, there must be a voxel $v' \in \mathcal{V}$ such that $\Delta_{v',x',b'} > 0$. Let Z^* be the optimal objective function value of the robust problem. This value is finite because each component of every feasible intensity vector is bounded from below by 0. Then, we have that

$$Z^* = \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \bar{p}(x) w_b^* \geq \Delta_{v',x',b'} \bar{p}(x') w_{b'}^*.$$

Since $\Delta_{v',x',b'} \bar{p}(x') > 0$, it follows that

$$w_{b'}^* \leq \frac{Z^*}{\Delta_{v',x',b'} \bar{p}(x')}.$$

Since each optimal beamlet intensity w_b^* is bounded in this way, the set of optimal solutions of the robust problem (3.1) must be bounded. ■

A.2 An alternative robust problem

Proof of Lemma 1: We prove the result by showing that the feasible regions of the two problems are identical. Suppose that \mathbf{w} is feasible for (3.1). Then for every $\mathbf{p} \in P$ and $v \in \mathcal{T}$,

$$\theta_v \leq \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b \leq \gamma \theta_v.$$

Since $\mathbf{q}^1, \dots, \mathbf{q}^M \in P$,

$$\theta_v \leq \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} q^j(x) w_b \leq \gamma \theta_v$$

holds for every $j \in \{1, \dots, M\}$ and every $v \in \mathcal{T}$, which proves that \mathbf{w} is in the feasible region of (4.8).

For the reverse inclusion, suppose that \mathbf{w} is feasible for (4.8), and let \mathbf{p} be an arbitrary PMF in P . Since $P = \text{conv}\{\mathbf{q}^1, \dots, \mathbf{q}^M\}$, there exist scalars $\lambda_1, \dots, \lambda_M \geq 0$ such that $\sum_{j=1}^M \lambda_j = 1$ and

$$\sum_{j=1}^M \lambda_j \mathbf{q}^j = \mathbf{p}.$$

Because \mathbf{w} is feasible for (4.8), for each $j \in \{1, \dots, M\}$ and $v \in \mathcal{T}$ we have

$$\theta_v \leq \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} q^j(x) w_b \leq \gamma \theta_v,$$

or equivalently,

$$\lambda_j \theta_v \leq \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} \lambda_j q^j(x) w_b \leq \lambda_j \gamma \theta_v.$$

Adding up all three terms over j produces the desired inequality

$$\theta_v \leq \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} p(x) w_b \leq \gamma \theta_v. \quad \blacksquare$$

Before we prove Proposition 1, it is helpful to introduce some additional notation. Let the matrix $\mathbf{A}(\mathbf{q}^1, \dots, \mathbf{q}^M)$ be defined as

$$\mathbf{A}(\mathbf{q}^1, \dots, \mathbf{q}^M) = \begin{bmatrix} \mathbf{T}(\mathbf{q}^1, \dots, \mathbf{q}^M) \\ -\mathbf{T}(\mathbf{q}^1, \dots, \mathbf{q}^M) \\ -\mathbf{I} \\ \mathbf{0}^T \end{bmatrix},$$

where \mathbf{I} is the $|\mathcal{B}|$ -by- $|\mathcal{B}|$ identity matrix, $\mathbf{0}^T$ is a row vector of zeroes of length $|\mathcal{B}|$ and $\mathbf{T}(\mathbf{q}^1, \dots, \mathbf{q}^M)$ is a matrix defined (component-wise) as

$$T_{(v,j),b}(\mathbf{q}^1, \dots, \mathbf{q}^M) = \sum_{x \in X} \Delta_{v,x,b} q^j(x).$$

The rows of $\mathbf{T}(\mathbf{q}^1, \dots, \mathbf{q}^M)$ correspond to (v, j) pairs ($v \in \mathcal{T}, j \in \{1, \dots, M\}$) and the columns correspond to beamlets $b \in \mathcal{B}$.

Let the column vector \mathbf{c} be defined as

$$\mathbf{c} = \begin{bmatrix} -\gamma \mathbf{z} \\ \mathbf{z} \\ \mathbf{0} \\ 0 \end{bmatrix},$$

where $\mathbf{0}$ is a $|\mathcal{B}|$ -component column vector of zeroes, 0 is the scalar zero and \mathbf{z} is a column vector defined (component-wise) as

$$z_{(v,j)} = \theta_v.$$

To understand why \mathbf{A} and \mathbf{c} are defined in this way, consider the set

$$\mathcal{W} = \{\mathbf{w} \in \mathbb{R}^{|\mathcal{B}|} \mid \mathbf{A}(\mathbf{q}^1, \dots, \mathbf{q}^M)\mathbf{w} + \mathbf{c} \leq \mathbf{0}\},$$

where $\mathbf{0}$ is a $2M|\mathcal{T}| + |\mathcal{B}| + 1$ -component column vector of zeroes and the inequality is taken component-wise. The set \mathcal{W} is the set of all $\mathbf{w} \in \mathbb{R}^{|\mathcal{B}|}$ such that

$$\begin{aligned} \sum_{x \in X} \Delta_{v,x,b} q^j(x) w_b - \gamma \theta_v &\leq 0, \quad \forall j \in \{1, \dots, M\}, v \in \mathcal{T}, \\ - \sum_{x \in X} \Delta_{v,x,b} q^j(x) w_b + \theta_v &\leq 0, \quad \forall j \in \{1, \dots, M\}, v \in \mathcal{T}, \\ - w_b &\leq 0, \quad \forall b \in \mathcal{B}, \end{aligned}$$

which is precisely the set of feasible solutions to problem (4.8) with the defining PMFs $\mathbf{q}^1, \dots, \mathbf{q}^M$. In other words, the set of feasible \mathbf{w} can be written as a set of the form $\{\mathbf{w} \in \mathbb{R}^{|\mathcal{B}|} \mid g(\mathbf{w}) \leq 0\}$, where g is an affine function.

We first show that if the defining PMFs all converge to a PMF \mathbf{p}^* , then the affine functions defined using the matrix $\mathbf{A}(\cdot, \dots, \cdot)$ and the vector \mathbf{c} above also converge.

Proposition 5 *Let $(\mathbf{q}^{1,i})_{i=1}^\infty, \dots, (\mathbf{q}^{M,i})_{i=1}^\infty$ be sequences of PMFs in \mathcal{P} , each of which converges to \mathbf{p}^* . The sequence of affine functions $(f^i)_{i=1}^\infty$, where $f^i : \mathbb{R}^{|\mathcal{B}|} \rightarrow \mathbb{R}^{2M|\mathcal{T}|+|\mathcal{B}|+1}$, defined as*

$$f^i(\mathbf{w}) = \mathbf{A}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})\mathbf{w} + \mathbf{c},$$

converges pointwise to the affine function f defined as

$$f(\mathbf{w}) = \mathbf{A}(\mathbf{p}^*, \dots, \mathbf{p}^*)\mathbf{w} + \mathbf{c}.$$

Proof: For every $j \in \{1, \dots, M\}$ we have that as $i \rightarrow \infty$, $\mathbf{q}^{j,i} \rightarrow \mathbf{p}^*$. This means that for every (v, j) pair and beamlet b ,

$$T_{(v,j),b}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}) = \sum_{x \in X} \Delta_{v,x,b} q^{j,i}(x) \rightarrow \sum_{x \in X} \Delta_{v,x,b} p^*(x) = T_{(v,j),b}(\mathbf{p}^*, \dots, \mathbf{p}^*)$$

as $i \rightarrow \infty$. In other words, $\mathbf{T}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})$ converges component-wise to $\mathbf{T}(\mathbf{p}^*, \dots, \mathbf{p}^*)$. Since the only other matrices used to form $\mathbf{A}(\mathbf{q}^1, \dots, \mathbf{q}^M)$ are always constant ($-\mathbf{I}$ and $\mathbf{0}^T$), regardless of $\mathbf{q}^1, \dots, \mathbf{q}^M$, we can deduce that $\mathbf{A}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}) \rightarrow \mathbf{A}(\mathbf{p}^*, \dots, \mathbf{p}^*)$ component-wise as $i \rightarrow \infty$.

Let $\mathbf{w} \in \mathbb{R}^{|\mathcal{B}|}$. Let k be the index of a row of $\mathbf{A}(\cdot, \dots, \cdot)$; the k th coordinate function of f^i is then

$$f_k^i(\mathbf{w}) = \sum_{b \in \mathcal{B}} A_{k,b}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}) w_b + c_k.$$

We have already established that $A_{k,b}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}) \rightarrow A_{k,b}(\mathbf{p}^*, \dots, \mathbf{p}^*)$ as $i \rightarrow \infty$ for every $b \in \mathcal{B}$ and every row k of \mathbf{A} . Since the k th coordinate of f^i is just a linear combination of components of $\mathbf{A}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})$ plus c_k , it must converge to the same linear combination of the limits of those components plus c_k – in other words,

$$f_k^i(\mathbf{w}) = \sum_{b \in \mathcal{B}} A_{k,b}(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}) w_b + c_k \rightarrow \sum_{b \in \mathcal{B}} A_{k,b}(\mathbf{p}^*, \dots, \mathbf{p}^*) w_b + c_k = f_k(\mathbf{w}).$$

Since this holds for every coordinate function k , we must have that $f^i(\mathbf{w}) \rightarrow f(\mathbf{w})$ as $i \rightarrow \infty$, which is the required result. ■

Before we prove the next proposition, we need some additional definitions. Following Dantzig et al. [1967], we say that a sequence of subsets $(A_n)_{n=1}^\infty$ of a metric space X converges to a set A if $\underline{\lim}_{n \rightarrow \infty} A_n = \overline{\lim}_{n \rightarrow \infty} A_n = A$. This is the form of set convergence used in Proposition 1. The two limit operations are defined as

$$\underline{\lim}_{n \rightarrow \infty} A_n = \left\{ x \in X \mid x = \lim_{n \rightarrow \infty} x_n, \text{ where } x_n \in A_n \text{ for all but a finite number of } n \right\}$$

and

$$\overline{\lim}_{n \rightarrow \infty} A_n = \left\{ x \in X \mid x = \lim_{n \rightarrow \infty} x_{n_i}, \text{ where } \{n_i\} \text{ is an infinite subsequence in } \mathbb{Z}_+ \text{ and } x_{n_i} \in A_{n_i} \right\}.$$

We define the rank of an affine function $g(\mathbf{w}) = \mathbf{A}\mathbf{w} + \mathbf{c}$ as the rank of the matrix \mathbf{A} . If $g : \mathbb{R}^n \rightarrow \mathbb{R}^m$ and $I \subset \{1, \dots, m\}$, then the function g_I is the function consisting only of those coordinates of g corresponding to indices in I , i.e., $g_I(x) = \mathbf{A}_I x + \mathbf{c}_I$ where \mathbf{A}_I (\mathbf{c}_I) consists of the rows of \mathbf{A} (components of \mathbf{c}) corresponding to indices in I .

We can now prove Proposition 1.

Proof of Proposition 1: Let $(f^i)_{i=1}^{\infty}$ and f be defined as in the statement of Proposition 5. Let $H(f^i) = \{\mathbf{w} \in \mathbb{R}^{|\mathcal{B}|} \mid f^i(\mathbf{w}) \leq \mathbf{0}\}$ and $H(f) = \{\mathbf{w} \in \mathbb{R}^{|\mathcal{B}|} \mid f(\mathbf{w}) \leq \mathbf{0}\}$, where $\mathbf{0}$ is a $2M|\mathcal{T}| + |\mathcal{B}| + 1$ -component column vector of zeroes and the inequality is taken component-wise. We can immediately see that

$$H(f^i) = H(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})$$

and

$$H(f) = H(\mathbf{p}^*, \dots, \mathbf{p}^*).$$

The margin problem uncertainty set contains the singleton set $\{\mathbf{p}^*\}$, so the feasible region of the margin problem is contained in the feasible region of the nominal problem with respect to \mathbf{p}^* . The nominal problem with respect to \mathbf{p}^* is equivalent to the alternative robust problem (4.8) with all M defining PMFs set to \mathbf{p}^* . By Proposition 3, the margin problem has an interior solution, so the feasible region of the margin problem is non-empty. Therefore, $H(f)$ is also non-empty.

Now, define I as

$$I = \{k \mid f_k(\mathbf{w}) = 0 \text{ for all } \mathbf{w} \in H(f)\},$$

with k ranging over all the coordinates of f (row indices of the matrix $\mathbf{A}(\mathbf{p}^*, \dots, \mathbf{p}^*)$). The set I consists of only the index corresponding to the last row of the matrix $\mathbf{A}(\mathbf{p}^*, \dots, \mathbf{p}^*)$. This is because the \mathbf{w} of Proposition 3 satisfies $-\mathbf{I}\mathbf{w} < \mathbf{0}$ (from the condition on the w_b 's being strictly positive), $-\mathbf{T}(\mathbf{p}^*, \dots, \mathbf{p}^*)\mathbf{w} + \mathbf{z} < \mathbf{0}$ (because \mathbf{w} satisfies $\sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b > \theta_v$ for all $v \in \mathcal{T}, x \in X$) and $\mathbf{T}(\mathbf{p}^*, \dots, \mathbf{p}^*)\mathbf{w} - \gamma \mathbf{z} < \mathbf{0}$ (because \mathbf{w} satisfies $\sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b <$

$\gamma\theta_v$ for all $v \in \mathcal{T}, x \in X$). In contrast, the last coordinate of the affine function f is

$$\mathbf{0}^T \mathbf{w} + 0 = 0$$

regardless of \mathbf{w} . Therefore, the set I consists of only the index of this last coordinate.

This means that $\text{rank } f_I = \text{rank } [\mathbf{0}^T] = 0$ and for every i , $\text{rank } f_I^i = \text{rank } [\mathbf{0}^T] = 0$, and in particular that

$$\limsup_{i \rightarrow \infty} (\text{rank } f_I^i) \leq \text{rank } f_I. \quad (\text{A.2})$$

Invoking now Theorem II.2.2 of Dantzig et al. [1967] which requires (A.2), we have that either $\lim_{i \rightarrow \infty} H(f^i) = H(f)$ or $H(f^i)$ is empty for infinitely many i . The latter case is not possible by Proposition 3 (the \mathbf{w} of Proposition 3 belongs to every set $H(f^i)$). Therefore, we must have $\lim_{i \rightarrow \infty} H(f^i) = H(f)$ or equivalently, $\lim_{i \rightarrow \infty} H(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}) = H(\mathbf{p}^*, \dots, \mathbf{p}^*)$. ■

We now prove Proposition 2 (the convergence of optimal solution sets of the alternative robust problem).

Proof of Proposition 2: By Proposition 1, whenever $(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})$ converges to $(\mathbf{p}^*, \dots, \mathbf{p}^*)$ as $i \rightarrow \infty$, $H(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})$ converges to $H(\mathbf{p}^*, \dots, \mathbf{p}^*)$. For any $(\mathbf{q}^1, \dots, \mathbf{q}^M)$ -tuple, $H(\mathbf{q}^1, \dots, \mathbf{q}^M)$ is non-empty because it contains the \mathbf{w} of Proposition 3. The set $H(\mathbf{q}^1, \dots, \mathbf{q}^M)$ is also a closed and convex set (it is a polyhedron). The function Z is a linear function on $\mathbb{R}^{|\mathcal{B}|}$, and therefore continuous and quasiconvex. The set $M(Z \mid H(\mathbf{p}^*, \dots, \mathbf{p}^*))$ is non-empty because by the constraints in $H(\mathbf{p}^*, \dots, \mathbf{p}^*)$, each beamlet's intensity w_b is non-negative, so Z is always bounded from below by 0, and

$$\begin{aligned} & \text{minimize } Z(\mathbf{w}) \\ & \text{subject to } \mathbf{w} \in H(\mathbf{p}^*, \dots, \mathbf{p}^*) \end{aligned}$$

is a linear optimization problem. The set $M(Z \mid H(\mathbf{p}^*, \dots, \mathbf{p}^*))$ is a bounded set, because it is the set of optimal solutions of the nominal problem with respect to \mathbf{p}^* , which is bounded by Proposition 4. The set $U(M(Z \mid H(\mathbf{p}^*, \dots, \mathbf{p}^*)), \epsilon)$ is therefore the epsilon neighborhood of a bounded set, and so it must also be bounded. Since $U(M(Z \mid H(\mathbf{p}^*, \dots, \mathbf{p}^*)), \epsilon)$ is the union of open sets, it is open.

From these observations, the hypotheses of Theorem I.3.3 of Dantzig et al. [1967] are satisfied. Invoking it furnishes a neighborhood of $(\mathbf{p}^*, \dots, \mathbf{p}^*)$ such that whenever

$(\mathbf{q}^1, \dots, \mathbf{q}^M)$ is inside that neighborhood, $M(Z|H(\mathbf{q}^1, \dots, \mathbf{q}^M)) \subseteq U(M(Z|H(\mathbf{p}^*, \dots, \mathbf{p}^*)), \epsilon)$. Since each $(\mathbf{q}^{j,i})_{i=1}^\infty$ sequence converges to \mathbf{p}^* , there must be an $N \in \mathbb{Z}_+$ such that for $i > N$, $(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})$ is within that neighborhood, and thus $M(Z|H(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})) \subseteq U(M(Z|H(\mathbf{p}^*, \dots, \mathbf{p}^*)), \epsilon)$. ■

A.3 Convergence of optimal solutions sets and dose distributions under convex-convergent update algorithms

First we prove that the exponential smoothing and running average update algorithms are convex-convergent.

Proposition 6 *The exponential smoothing update algorithm with $\alpha \in (0, 1]$ and the running average update algorithm are convex-convergent update algorithms.*

Proof: The first defining property of a convex-convergent update algorithm is that for each $i \in \mathbb{Z}_+$, there exists an $\alpha_i \in [0, 1]$ such that $\ell^{i+1} = (1 - \alpha_i)\ell^i + \alpha_i \mathbf{p}^i$ and $\mathbf{u}^{i+1} = (1 - \alpha_i)\mathbf{u}^i + \alpha_i \mathbf{p}^i$. From equations (4.1), (4.2), (4.3) and (4.4) in Section 4.1.1, we know that both the exponential smoothing and running average updates satisfy this property: for exponential smoothing, we set $\alpha_i = \alpha$ for every $i \in \mathbb{Z}_+$, while for running average, we set $\alpha_i = i/(i + 1)$ for every $i \in \mathbb{Z}_+$.

The second defining property is that if $(\mathbf{p}^i)_{i=1}^\infty$ converges to \mathbf{p}^* , then so must $(\ell^i)_{i=1}^\infty$ and $(\mathbf{u}^i)_{i=1}^\infty$. To show the convergence property for the running average update, recall that for any real-valued sequence $(a_i)_{i=1}^\infty$ that converges to $a \in \mathbb{R}$, the sequence of arithmetic means $(\sigma_i)_{i=1}^\infty$ where $\sigma_i = (a_1 + \dots + a_i)/i$ also converges to a (see Exercise 14 of Chapter 3 in Rudin [1976]). The vector sequence $(\ell^i)_{i=1}^\infty$ is the sequence of arithmetic means of the vector sequence $(\ell^1, \mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots)$, and similarly $(\mathbf{u}^i)_{i=1}^\infty$ is the sequence of arithmetic means of the vector sequence $(\mathbf{u}^1, \mathbf{p}^1, \mathbf{p}^2, \mathbf{p}^3, \dots)$. For every $x \in X$, the scalar, real-valued sequence $(\ell^i(x))_{i=1}^\infty$ converges to $p^*(x)$ because $(\ell^1(x), p^1(x), p^2(x), p^3(x), \dots)$ converges to $p^*(x)$ and similarly $(u^i(x))_{i=1}^\infty$ converges to $p^*(x)$ because $(u^1(x), p^1(x), p^2(x), p^3(x), \dots)$ also converges to $p^*(x)$. Therefore, the lower and upper bound vectors converge to \mathbf{p}^* in every component ($x \in X$), which means that as $i \rightarrow \infty$, $\ell^i \rightarrow \mathbf{p}^*$ and $\mathbf{u}^i \rightarrow \mathbf{p}^*$ as vectors.

To show the convergence property for the exponential smoothing update, let us consider $(\ell^i)_{i=1}^\infty$; the proof for $(\mathbf{u}^i)_{i=1}^\infty$ follows in exactly the same way.

Let $\epsilon > 0$ be given. Since $\mathbf{p}^i \rightarrow \mathbf{p}^*$ as $i \rightarrow \infty$, there is an $N_1 \in \mathbb{Z}_+$ such that for all $i \geq N_1$, $\|\mathbf{p}^i - \mathbf{p}^*\| < \epsilon/2$. Also, since $0 < \alpha \leq 1$, there is an $N_2 \in \mathbb{Z}_+$ such that $N_2 > N_1$ and for all $i \geq N_2$, $(1 - \alpha)^{i-N_1} \|\ell^{N_1} - \mathbf{p}^*\| < \epsilon/2$.

For $i \geq N_2$, note that by iterating the definition of exponential smoothing, we can write ℓ^i as

$$\begin{aligned} \ell^i &= (1 - \alpha)^{i-N_1} \ell^{N_1} + (1 - \alpha)^{i-N_1-1} \alpha \mathbf{p}^{N_1} + (1 - \alpha)^{i-N_1-2} \alpha \mathbf{p}^{N_1+1} \\ &\quad + \cdots + \alpha(1 - \alpha) \mathbf{p}^{i-2} + \alpha \mathbf{p}^{i-1} \\ &= (1 - \alpha)^{i-N_1} \ell^{N_1} + \sum_{j=1}^{i-N_1} (1 - \alpha)^{i-N_1-j} \alpha \mathbf{p}^{N_1+j-1} \end{aligned} \quad (\text{A.3})$$

Note also that

$$\sum_{j=1}^{i-N_1} (1 - \alpha)^{i-N_1-j} \alpha = 1 - (1 - \alpha)^{i-N_1}. \quad (\text{A.4})$$

Putting equations (A.3) and (A.4) together, we have, for $i \geq N_2$,

$$\begin{aligned} \ell^i - \mathbf{p}^* &= (1 - \alpha)^{i-N_1} \ell^{N_1} + \sum_{j=1}^{i-N_1} (1 - \alpha)^{i-N_1-j} \alpha \mathbf{p}^{N_1+j-1} - \mathbf{p}^* \\ &= (1 - \alpha)^{i-N_1} \ell^{N_1} + \sum_{j=1}^{i-N_1} (1 - \alpha)^{i-N_1-j} \alpha \mathbf{p}^{N_1+j-1} - (1 - (1 - \alpha)^{i-N_1}) \mathbf{p}^* - (1 - \alpha)^{i-N_1} \mathbf{p}^* \\ &= (1 - \alpha)^{i-N_1} (\ell^{N_1} - \mathbf{p}^*) + \sum_{j=1}^{i-N_1} (1 - \alpha)^{i-N_1-j} \alpha \mathbf{p}^{N_1+j-1} - \sum_{j=1}^{i-N_1} (1 - \alpha)^{i-N_1-j} \alpha \mathbf{p}^* \\ &= (1 - \alpha)^{i-N_1} (\ell^{N_1} - \mathbf{p}^*) + \sum_{j=1}^{i-N_1} (1 - \alpha)^{i-N_1-j} \alpha (\mathbf{p}^{N_1+j-1} - \mathbf{p}^*). \end{aligned}$$

Taking norms, for $i \geq N_2$,

$$\begin{aligned} \|\ell^i - \mathbf{p}^*\| &= \left\| (1 - \alpha)^{i-N_1} (\ell^{N_1} - \mathbf{p}^*) + \sum_{j=1}^{i-N_1} (1 - \alpha)^{i-N_1-j} \alpha (\mathbf{p}^{N_1+j-1} - \mathbf{p}^*) \right\| \\ &\leq (1 - \alpha)^{i-N_1} \|\ell^{N_1} - \mathbf{p}^*\| + \sum_{j=1}^{i-N_1} (1 - \alpha)^{i-N_1-j} \alpha \|\mathbf{p}^{N_1+j-1} - \mathbf{p}^*\| \end{aligned}$$

$$\begin{aligned}
&< \frac{\epsilon}{2} + \sum_{j=1}^{i-N_1} (1-\alpha)^{i-N_1-j} \alpha \frac{\epsilon}{2} \\
&= \frac{\epsilon}{2} + (1 - (1-\alpha)^{i-N_1}) \frac{\epsilon}{2} \\
&\leq \frac{\epsilon}{2} + \frac{\epsilon}{2} \\
&= \epsilon,
\end{aligned}$$

where the first inequality follows from the triangle inequality, the second inequality follows from our choice of N_1 and N_2 , and the last inequality follows because $0 \leq (1 - (1 - \alpha)^{i-N_1}) \leq 1$. This proves that $\ell^i \rightarrow \mathbf{p}^*$ as $i \rightarrow \infty$. The proof that $\mathbf{u}^i \rightarrow \mathbf{p}^*$ follows in exactly the same way. ■

Before proving Theorem 1, we need to establish an auxilliary lemma. This lemma states that if we modify an uncertainty set P by taking the convex combination of ℓ and \mathbf{u} with an observed PMF $\tilde{\mathbf{p}}$, then the resulting uncertainty set is the same as one obtained by applying the same convex combination to the M defining PMFs $\mathbf{q}^1, \dots, \mathbf{q}^M$ whose convex hull is P .

Lemma 2 *Let $\alpha \in [0, 1]$. Let P be an uncertainty set defined by ℓ and \mathbf{u} , and let $\mathbf{q}^1, \dots, \mathbf{q}^M$ be M PMFs such that $P = \text{conv} \{\mathbf{q}^1, \dots, \mathbf{q}^M\}$. Let $\tilde{\mathbf{p}} \in \mathcal{P}$, and let \tilde{P} be the uncertainty set defined by*

$$\begin{aligned}
\tilde{\ell} &= (1 - \alpha)\ell + \alpha\tilde{\mathbf{p}}, \\
\tilde{\mathbf{u}} &= (1 - \alpha)\mathbf{u} + \alpha\tilde{\mathbf{p}}.
\end{aligned}$$

Then $\tilde{P} = \text{conv} \{\tilde{\mathbf{q}}^1, \dots, \tilde{\mathbf{q}}^M\}$, where

$$\tilde{\mathbf{q}}^j = (1 - \alpha)\mathbf{q}^j + \alpha\tilde{\mathbf{p}},$$

for every $j \in \{1, \dots, M\}$.

Proof: If $\alpha = 1$, then $\tilde{\ell} = \tilde{\mathbf{u}} = \tilde{\mathbf{p}}$, so $P = \{\tilde{\mathbf{p}}\}$. At the same time, $\tilde{\mathbf{q}}^j = \tilde{\mathbf{p}}$ for every j , so $\text{conv} \{\tilde{\mathbf{q}}^1, \dots, \tilde{\mathbf{q}}^M\} = \{\tilde{\mathbf{p}}\}$.

Now consider $\alpha < 1$. We will first show that $\tilde{P} \subseteq \text{conv} \{\tilde{\mathbf{q}}^1, \dots, \tilde{\mathbf{q}}^M\}$. Let $\mathbf{p}^2 \in \tilde{P}$, and define \mathbf{p}^1 as

$$\mathbf{p}^1 = \frac{1}{1 - \alpha}\mathbf{p}^2 - \frac{\alpha}{1 - \alpha}\tilde{\mathbf{p}}.$$

We claim that $\mathbf{p}^1 \in P$. For any $x \in X$,

$$\begin{aligned}
p^1(x) &= \frac{1}{1-\alpha} p^2(x) - \frac{\alpha}{1-\alpha} \tilde{p}(x) \\
&\geq \frac{1}{1-\alpha} \tilde{\ell}(x) - \frac{\alpha}{1-\alpha} \tilde{p}(x) \\
&= \frac{1}{1-\alpha} [(1-\alpha)\ell(x) + \alpha\tilde{p}(x)] - \frac{\alpha}{1-\alpha} \tilde{p}(x) \\
&= \ell(x) + \frac{\alpha}{1-\alpha} \tilde{p}(x) - \frac{\alpha}{1-\alpha} \tilde{p}(x) \\
&= \ell(x).
\end{aligned}$$

A similar argument can be used to establish that $p^1(x) \leq u(x)$ for each $x \in X$.

Since $\mathbf{p}^1 \in P$ and $P = \text{conv}\{\mathbf{q}^1, \dots, \mathbf{q}^M\}$, there exist $\lambda_1, \dots, \lambda_M \geq 0$ such that $\sum_{j=1}^M \lambda_j = 1$ and

$$\mathbf{p}^1 = \sum_{j=1}^M \lambda_j \mathbf{q}^j.$$

If we now form the sum $\sum_{j=1}^M \lambda_j \tilde{\mathbf{q}}^j$, we see that

$$\begin{aligned}
\sum_{j=1}^M \lambda_j \tilde{\mathbf{q}}^j &= \sum_{j=1}^M \lambda_j [(1-\alpha)\mathbf{q}^j + \alpha\tilde{\mathbf{p}}] \\
&= (1-\alpha) \sum_{j=1}^M \lambda_j \mathbf{q}^j + \alpha \sum_{j=1}^M \lambda_j \tilde{\mathbf{p}} \\
&= (1-\alpha)\mathbf{p}^1 + \alpha\tilde{\mathbf{p}} \\
&= (1-\alpha) \left[\frac{1}{1-\alpha} \mathbf{p}^2 - \frac{\alpha}{1-\alpha} \tilde{\mathbf{p}} \right] + \alpha\tilde{\mathbf{p}} \\
&= \mathbf{p}^2;
\end{aligned}$$

in other words, $\mathbf{p}^2 \in \text{conv}\{\tilde{\mathbf{q}}^1, \dots, \tilde{\mathbf{q}}^M\}$.

To prove that $\text{conv}\{\tilde{\mathbf{q}}^1, \dots, \tilde{\mathbf{q}}^M\} \subseteq \tilde{P}$, note that since $P = \text{conv}\{\mathbf{q}^1, \dots, \mathbf{q}^M\}$, each $\mathbf{q}^j \in P$. Then, for any $j \in \{1, \dots, M\}$ and any component $x \in X$,

$$\tilde{q}^j(x) = (1-\alpha)q^j(x) + \alpha\tilde{p}(x) \geq (1-\alpha)\ell(x) + \alpha\tilde{p}(x) = \tilde{\ell}(x)$$

and

$$\tilde{q}^j(x) = (1-\alpha)q^j(x) + \alpha\tilde{p}(x) \leq (1-\alpha)u(x) + \alpha\tilde{p}(x) = \tilde{u}(x).$$

As a result, each $\tilde{\mathbf{q}}^j$ belongs to \tilde{P} . Since \tilde{P} is a convex set, it will contain any convex combination of $\tilde{\mathbf{q}}^1, \dots, \tilde{\mathbf{q}}^M$. It therefore follows that $\text{conv} \{\tilde{\mathbf{q}}^1, \dots, \tilde{\mathbf{q}}^M\} \subseteq \tilde{P}$, and thus that $\tilde{P} = \text{conv} \{\tilde{\mathbf{q}}^1, \dots, \tilde{\mathbf{q}}^M\}$. ■

We use Lemma 2 in the proof of Theorem 1.

Proof of Theorem 1: The first lower and upper bound vectors ℓ^1 and \mathbf{u}^1 correspond to the uncertainty set P^1 , which is a bounded polyhedron. As such, it can be written as the convex hull of finitely many PMFs $\mathbf{q}^{1,1}, \dots, \mathbf{q}^{M,1}$. Since the update algorithm is convex-convergent, for every $i \in \mathbb{Z}_+$ there is an α_i such that

$$\begin{aligned}\ell^{i+1} &= (1 - \alpha_i)\ell^i + \alpha_i\mathbf{p}^i, \\ \mathbf{u}^{i+1} &= (1 - \alpha_i)\mathbf{u}^i + \alpha_i\mathbf{p}^i.\end{aligned}$$

For each $j \in \{1, \dots, M\}$ and each $i \in \mathbb{Z}_+$, define $\mathbf{q}^{j,i+1}$ recursively as

$$\mathbf{q}^{j,i+1} = (1 - \alpha_i)\mathbf{q}^{j,i} + \alpha_i\mathbf{p}^i.$$

Let P^i denote the uncertainty set specified by ℓ^i and \mathbf{u}^i . By repeatedly applying Lemma 2, we can see that $P^i = \text{conv} \{\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}\}$. By Lemma 1, problem (3.1) with the uncertainty set P^i and problem (4.8) with defining PMFs $\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i}$ are equivalent and have the same feasible region. Therefore, it follows that

$$\mathbf{w}^*(\ell^i, \mathbf{u}^i) = M(Z | H(\mathbf{q}^{1,i}, \dots, \mathbf{q}^{M,i})).$$

Also, as we noted immediately after Proposition 1, the nominal problem with respect to \mathbf{p}^* is equivalent to the alternative robust problem (4.8) with all M defining PMFs set to \mathbf{p}^* ; we therefore have that

$$\mathbf{w}^*(\mathbf{p}^*) = M(Z | H(\mathbf{p}^*, \dots, \mathbf{p}^*)).$$

By the definition of the sequences $(\mathbf{q}^{j,i})_{i=1}^\infty$ for each $j \in \{1, \dots, M\}$ and by the convergence property of convex-convergent update algorithms, $\mathbf{q}^{j,i} \rightarrow \mathbf{p}^*$ for each $j \in \{1, \dots, M\}$ as $i \rightarrow \infty$ (each $(\mathbf{q}^{j,i})_{i=1}^\infty$ sequence can be treated as a lower or upper bound vector sequence in the definition of a convex-convergent update algorithm). Applying

Proposition 2 with any $\epsilon > 0$, we obtain an $N \in \mathbb{Z}_+$ such that for every $i > N$,

$$\mathbf{w}^*(\ell^i, \mathbf{u}^i) \subseteq U(\mathbf{w}^*(\mathbf{p}^*), \epsilon),$$

as required. ■

Before we can prove Theorem 2 (the dose distribution from the adaptive robust method approaches the set \mathbf{D}), we need two additional lemmata.

Lemma 3 *Let $\mathbf{p}^i \in \mathcal{P}$ and $\mathbf{w}, \mathbf{y} \in \mathbb{R}^{|\mathcal{B}|}$. Then*

$$\|\Delta \mathbf{p}^i(\mathbf{w} - \mathbf{y})\| \leq |\mathcal{V}| \cdot \max \Delta_{v', x', b'} \cdot \|\mathbf{w} - \mathbf{y}\|,$$

where the maximum is taken over all voxels $v' \in \mathcal{V}$, motion states $x' \in X$ and beamlets $b' \in \mathcal{B}$.

Proof: We have

$$\begin{aligned} \|\Delta \mathbf{p}^i(\mathbf{w} - \mathbf{y})\| &= \sum_{v \in \mathcal{V}} \left| \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v, x, b} p^i(x) (w_b - y_b) \right| \\ &\leq \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v, x, b} p^i(x) |w_b - y_b| \\ &\leq \sum_{v \in \mathcal{V}} \sum_{x \in X} \sum_{b \in \mathcal{B}} \max \Delta_{v', x', b'} p^i(x) |w_b - y_b| \\ &= |\mathcal{V}| \cdot \max \Delta_{v', x', b'} \cdot \sum_{b \in \mathcal{B}} \left(\sum_{x \in X} p^i(x) \right) |w_b - y_b| \\ &= |\mathcal{V}| \cdot \max \Delta_{v', x', b'} \cdot \|\mathbf{w} - \mathbf{y}\|, \end{aligned}$$

as required. ■

Lemma 4 *Let $\mathbf{p}, \mathbf{q} \in \mathcal{P}$ and $\mathbf{w} \in \mathbb{R}^{|\mathcal{B}|}$, with each $w_b \geq 0$. Then*

$$\|\Delta(\mathbf{p} - \mathbf{q})\mathbf{w}\| \leq \max_{x' \in X} \sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v, x', b} w_b \cdot \|\mathbf{p} - \mathbf{q}\|.$$

Proof: We have

$$\|\Delta(\mathbf{p} - \mathbf{q})\mathbf{w}\| = \sum_{v \in \mathcal{V}} \left| \sum_{x \in X} \sum_{b \in \mathcal{B}} \Delta_{v, x, b} (p(x) - q(x)) w_b \right|$$

$$\begin{aligned}
&\leq \sum_{x \in X} \left(\sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v,x,b} w_b \right) |p(x) - q(x)| \\
&\leq \sum_{x \in X} \left(\max_{x' \in X} \sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v,x',b} w_b \right) |p(x) - q(x)| \\
&= \left(\max_{x' \in X} \sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v,x',b} w_b \right) \|\mathbf{p} - \mathbf{q}\|,
\end{aligned}$$

as required. ■

We now prove Theorem 2.

Proof of Theorem 2: Let $\epsilon > 0$ be given. The set of optimal solutions of problem (3.1) with $P = \{\mathbf{p}^*\}$, denoted $\mathbf{w}^*(\mathbf{p}^*)$, is closed (it is a polyhedron) and is bounded by Proposition 4. Since $\mathbf{w}^*(\mathbf{p}^*) \subseteq \mathbb{R}^{|\mathcal{B}|}$, $\mathbf{w}^*(\mathbf{p}^*)$ is compact.

The function $Y : \mathbb{R}^{|\mathcal{B}|} \rightarrow \mathbb{R}$ defined as

$$Y(\mathbf{w}) = \max_{x' \in X} \sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v,x',b} w_b$$

is a continuous function, so by the extreme value theorem, it attains a maximum on $\mathbf{w}^*(\mathbf{p}^*)$. Let \bar{Y} denote that maximum value.

Since $\mathbf{p}^i \rightarrow \mathbf{p}^*$ as $i \rightarrow \infty$, there must exist an $N_1 \in \mathbb{Z}_+$ such that, for $i > N_1$,

$$\|\mathbf{p}^i - \mathbf{p}^*\| < \frac{\epsilon}{3 \cdot \bar{Y}}.$$

By Theorem 1, there must also exist an $N_2 \in \mathbb{Z}_+$ such that, for $i > N_2$,

$$\mathbf{w}^*(\ell^i, \mathbf{u}^i) \subseteq U \left(\mathbf{w}^*(\mathbf{p}^*), \frac{\epsilon}{3 \cdot |\mathcal{V}| \cdot \max \Delta_{v',x',b'}} \right).$$

Set $N_3 = \max\{N_1, N_2\}$.

Define the functions $\Omega_i : \mathbb{R}^{|\mathcal{B}|} \rightarrow \mathbb{R}$ for each $i \leq N_3$ as

$$\Omega_i(\mathbf{w}) = \|\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^* \mathbf{w}\|.$$

Recall that for each i , $\mathbf{w}^i \in \mathbf{w}^*(\ell^i, \mathbf{u}^i)$ is the optimized intensity vector determined by the adaptive robust method for fraction i and \mathbf{w}^i/n is delivered to the patient in fraction i . Like Y , each Ω_i is also a continuous function, and since the set $\mathbf{w}^*(\mathbf{p}^*)$ is

compact, each Ω_i attains a maximum on $\mathbf{w}^*(\mathbf{p}^*)$. Let $\bar{\Omega}_i$ denote each of those maximum values. Since $1/n$ is a decreasing function on \mathbb{Z}_+ , there must be an $N_4 \in \mathbb{Z}_+$, $N_4 > N_3$, such that, for $n > N_4$,

$$\frac{1}{n} \sum_{i=1}^{N_3} \bar{\Omega}_i < \frac{\epsilon}{3}.$$

Now, let $n > N_4$. To prove our result, we will construct a dose vector inside \mathbf{D} whose epsilon neighborhood contains $1/n \cdot \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i$.

For each $i > N_3$, select $\mathbf{y}^i \in \mathbf{w}^*(\mathbf{p}^*)$ so that

$$\|\mathbf{w}^i - \mathbf{y}^i\| < \frac{\epsilon}{3 \cdot |\mathcal{V}| \cdot \max \Delta_{v',x',b'}}.$$

Such a \mathbf{y}^i exists because \mathbf{w}^i is contained inside the union of all open balls of size $\epsilon / (3 \cdot |\mathcal{V}| \cdot \max \Delta_{v',x',b'})$ centered at points in $\mathbf{w}^*(\mathbf{p}^*)$, so there must be at least one $\mathbf{y}^i \in \mathbf{w}^*(\mathbf{p}^*)$ such that

$$\mathbf{w}^i \in B\left(\mathbf{y}^i, \frac{\epsilon}{3 \cdot |\mathcal{V}| \cdot \max \Delta_{v',x',b'}}\right).$$

For $i \leq N_3$, select any $\mathbf{y}^i \in \mathbf{w}^*(\mathbf{p}^*)$.

Let $\mathbf{y} = 1/n \cdot \sum_{i=1}^n \mathbf{y}^i$. The beamlet intensity vector \mathbf{y} belongs to $\mathbf{w}^*(\mathbf{p}^*)$ because it is the convex combination of the beamlet intensity vectors $\mathbf{y}^1, \dots, \mathbf{y}^n$, which are all elements of the convex set $\mathbf{w}^*(\mathbf{p}^*)$. Since $\mathbf{y} \in \mathbf{w}^*(\mathbf{p}^*)$, the dose distribution \mathbf{d} defined as

$$\mathbf{d} = \Delta \mathbf{p}^* \mathbf{y} = \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^* \mathbf{y}^i$$

belongs to \mathbf{D} . We will now verify that the epsilon ball of \mathbf{d} indeed contains $1/n \cdot \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i$. We have

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i - \mathbf{d} \right\| &= \left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i - \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^* \mathbf{y}^i \right\| \\ &= \frac{1}{n} \left\| \sum_{i=1}^n (\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^* \mathbf{y}^i) \right\| \\ &\leq \frac{1}{n} \left\| \sum_{i=1}^{N_3} (\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^* \mathbf{y}^i) \right\| + \frac{1}{n} \left\| \sum_{i=N_3+1}^n (\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^* \mathbf{y}^i) \right\| \end{aligned}$$

$$\begin{aligned}
&= \frac{1}{n} \left\| \sum_{i=1}^{N_3} (\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^* \mathbf{y}^i) \right\| + \frac{1}{n} \left\| \sum_{i=N_3+1}^n (\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^i \mathbf{y}^i + \Delta \mathbf{p}^i \mathbf{y}^i - \Delta \mathbf{p}^* \mathbf{y}^i) \right\| \\
&\leq \frac{1}{n} \left\| \sum_{i=1}^{N_3} (\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^* \mathbf{y}^i) \right\| + \frac{1}{n} \left\| \sum_{i=N_3+1}^n (\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^i \mathbf{y}^i) \right\| \\
&\quad + \frac{1}{n} \left\| \sum_{i=N_3+1}^n (\Delta \mathbf{p}^i \mathbf{y}^i - \Delta \mathbf{p}^* \mathbf{y}^i) \right\| \\
&\leq \frac{1}{n} \sum_{i=1}^{N_3} \|\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^* \mathbf{y}^i\| + \frac{1}{n} \sum_{i=N_3+1}^n \|\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^i \mathbf{y}^i\| \\
&\quad + \frac{1}{n} \sum_{i=N_3+1}^n \|\Delta \mathbf{p}^i \mathbf{y}^i - \Delta \mathbf{p}^* \mathbf{y}^i\|,
\end{aligned}$$

where all of the inequalities follow by the triangle inequality. By Lemmata 3 and 4 and by our definitions of \bar{Y} and $\bar{\Omega}_i$, we have

$$\begin{aligned}
\left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i - \mathbf{d} \right\| &\leq \frac{1}{n} \sum_{i=1}^{N_3} \bar{\Omega}_i + \frac{1}{n} \sum_{i=N_3+1}^n |\mathcal{V}| \cdot \max \Delta_{v',x',b'} \cdot \|\mathbf{w}^i - \mathbf{y}^i\| \\
&\quad + \frac{1}{n} \sum_{i=N_3+1}^n \max_{x' \in X} \sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v,x',b} y_b^i \cdot \|\mathbf{p}^i - \mathbf{p}^*\| \\
&\leq \frac{1}{n} \sum_{i=1}^{N_3} \bar{\Omega}_i + \frac{1}{n} \sum_{i=N_3+1}^n |\mathcal{V}| \cdot \max \Delta_{v',x',b'} \cdot \|\mathbf{w}^i - \mathbf{y}^i\| + \frac{1}{n} \sum_{i=N_3+1}^n \bar{Y} \cdot \|\mathbf{p}^i - \mathbf{p}^*\|.
\end{aligned}$$

Since $n > N_4$, we now have that

$$\begin{aligned}
\left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i - \mathbf{d} \right\| &\leq \frac{1}{n} \sum_{i=1}^{N_3} \bar{\Omega}_i + \frac{1}{n} \sum_{i=N_3+1}^n |\mathcal{V}| \cdot \max \Delta_{v',x',b'} \cdot \|\mathbf{w}^i - \mathbf{y}^i\| + \frac{1}{n} \sum_{i=N_3+1}^n \bar{Y} \cdot \|\mathbf{p}^i - \mathbf{p}^*\| \\
&< \frac{\epsilon}{3} + \frac{(n - N_3)}{n} \cdot \frac{\epsilon}{3} + \frac{(n - N_3)}{n} \cdot \frac{\epsilon}{3} \\
&\leq \frac{\epsilon}{3} + \frac{\epsilon}{3} + \frac{\epsilon}{3} \\
&= \epsilon.
\end{aligned}$$

We have just shown that $\frac{1}{n} \cdot \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i \in B(\mathbf{d}, \epsilon)$. Since $B(\mathbf{d}, \epsilon) \subseteq U(\mathbf{D}, \epsilon)$, we have that

$$\frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i \in U(\mathbf{D}, \epsilon),$$

which proves the theorem. ■

Lastly, we prove the dose distribution convergence results for the daily and average prescient methods. For the purposes of this proof, we use $H(\mathbf{q})$ to denote the feasible region of the alternative robust problem (4.8) with only one defining PMF \mathbf{q} . In this case, $H(\mathbf{q})$ corresponds to the feasible region of the nominal problem with respect to \mathbf{q} .

Proof of Theorem 3: (a) (Daily prescient dose convergence.) First, observe that the robust problem (3.1) with $\ell = \mathbf{p}^i$ and $\mathbf{u} = \mathbf{p}^i$ contains only one PMF: \mathbf{p}^i . Therefore, the robust problem (3.1) with $\ell = \mathbf{u} = \mathbf{p}^i$ is equivalent to the alternative robust problem (4.8) with \mathbf{p}^i as its one and only defining PMF. For every $i \in \mathbb{Z}_+$, it then follows that

$$\mathbf{w}^*(\mathbf{p}^i, \mathbf{p}^i) = M(Z | H(\mathbf{p}^i)).$$

The feasible region of the nominal problem with respect to \mathbf{p}^* is equal to $H(\mathbf{p}^*)$, and we have that

$$\mathbf{w}^*(\mathbf{p}^*) = M(Z | H(\mathbf{p}^*)).$$

The PMF sequence $(\mathbf{p}^i)_{i=1}^\infty$ converges to \mathbf{p}^* , so we can invoke Proposition 2. Doing so we obtain that for every $\epsilon > 0$, there exists an $N \in \mathbb{Z}_+$ such that for all $i > N$,

$$M(Z | H(\mathbf{p}^i)) \subseteq U(M(Z | H(\mathbf{p}^*)), \epsilon),$$

or equivalently,

$$\mathbf{w}^*(\mathbf{p}^i, \mathbf{p}^i) \subseteq U(\mathbf{w}^*(\mathbf{p}^*), \epsilon).$$

From this point forward, the proof proceeds in exactly the same way as the proof of Theorem 2.

(b) (Average prescient dose convergence.) Just as in the proof of part (a), the robust problem (3.1) with $\ell = \mathbf{u} = 1/n \cdot \sum_{i=1}^n \mathbf{p}^i$ is equal to the alternative robust problem with $1/n \cdot \sum_{i=1}^n \mathbf{p}^i$ as its one and only defining PMF. Analogously to that proof, we have that

$$\mathbf{w}^* \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i, \frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) = M \left(Z \left| H \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) \right. \right),$$

and that

$$\mathbf{w}^*(\mathbf{p}^*) = M(Z | H(\mathbf{p}^*)).$$

The sequence of averages $(1/n \cdot \sum_{i=1}^n \mathbf{p}^i)_{n=1}^\infty$, converges to \mathbf{p}^* (again, we appeal to Exercise 14 of Chapter 3 in Rudin [1976], which shows that $1/n \cdot \sum_{i=1}^n p^i(x) \rightarrow p^*(x)$ as $n \rightarrow \infty$ if $p^i(x) \rightarrow p^*(x)$ as $i \rightarrow \infty$ for every $x \in X$).

Since $1/n \cdot \sum_{i=1}^n \mathbf{p}^i \rightarrow \mathbf{p}^*$ as $n \rightarrow \infty$, we can invoke Proposition 2. Doing so we obtain that for every $\epsilon > 0$, there exists an $N \in \mathbb{Z}_+$ such that for all $n > N$,

$$M \left(Z \left| H \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) \right. \right) \subseteq U \left(M(Z | H(\mathbf{p}^*)), \epsilon \right),$$

or equivalently,

$$\mathbf{w}^* \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i, \frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) \subseteq U \left(\mathbf{w}^*(\mathbf{p}^*), \epsilon \right).$$

From this point forward, the proof proceeds in the same way as the proof of Theorem 2 with suitable modifications; for completeness, we provide the full proof here.

Fix an $\epsilon > 0$. Just as in the proof of Theorem 2, we observe that $\mathbf{w}^*(\mathbf{p}^*)$ is compact. The function $Y : \mathbb{R}^{|\mathcal{B}|} \rightarrow \mathbb{R}$ defined as

$$Y(\mathbf{w}) = \max_{x' \in X} \sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v,x',b} w_b$$

is a continuous function, so by the extreme value theorem, it attains a maximum on $\mathbf{w}^*(\mathbf{p}^*)$. Let \bar{Y} denote that maximum value. Since $1/n \cdot \sum_{i=1}^n \mathbf{p}^i \rightarrow \mathbf{p}^*$ as $n \rightarrow \infty$, there must exist an $N_1 \in \mathbb{Z}_+$ such that, for $n > N_1$,

$$\left\| \frac{1}{n} \sum_{i=1}^n \mathbf{p}^i - \mathbf{p}^* \right\| < \frac{\epsilon}{2 \cdot \bar{Y}}.$$

By the result that we proved at the beginning of this proof, there is an $N_2 \in \mathbb{Z}_+$ such that for every $n > N_2$,

$$\mathbf{w}^* \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i, \frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) \subseteq U \left(\mathbf{w}^*(\mathbf{p}^*), \frac{\epsilon}{2 \cdot |\mathcal{V}| \cdot \max \Delta_{v',x',b'}} \right).$$

Set $N_3 = \max\{N_1, N_2\}$.

Suppose now that we have an $n > N_3$. To complete the proof, we will construct a dose vector \mathbf{d} inside \mathbf{D} whose epsilon ball contains $1/n \cdot \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^n$. Recalling that $\mathbf{w}^n \in \mathbf{w}^*(1/n \cdot \sum_{i=1}^n \mathbf{p}^i, 1/n \cdot \sum_{i=1}^n \mathbf{p}^i)$, select $\mathbf{y}^n \in \mathbf{w}^*(\mathbf{p}^*)$ such that

$$\|\mathbf{w}^n - \mathbf{y}^n\| < \frac{\epsilon}{2 \cdot |\mathcal{V}| \cdot \max \Delta_{v',x',b'}}.$$

This is possible because of the fact that

$$\mathbf{w}^n \in U\left(\mathbf{w}^*(\mathbf{p}^*), \frac{\epsilon}{2 \cdot |\mathcal{V}| \cdot \max \Delta_{v',x',b'}}\right),$$

which is made possible by our choice of N_3 . Define the dose distribution \mathbf{d} as $\mathbf{d} = \Delta \mathbf{p}^* \mathbf{y}^n$; since $\mathbf{y}^n \in \mathbf{w}^*(\mathbf{p}^*)$, clearly $\mathbf{d} \in \mathbf{D}$.

We now verify that $1/n \cdot \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^n$ belongs to the epsilon ball of \mathbf{d} . We have

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i - \mathbf{d} \right\| &= \left\| \Delta \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) \mathbf{w}^n - \Delta \mathbf{p}^* \mathbf{y}^n \right\| \\ &= \left\| \Delta \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) (\mathbf{w}^n - \mathbf{y}^n + \mathbf{y}^n) - \Delta \mathbf{p}^* \mathbf{y}^n \right\| \\ &= \left\| \Delta \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) (\mathbf{w}^n - \mathbf{y}^n) + \Delta \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) \mathbf{y}^n - \Delta \mathbf{p}^* \mathbf{y}^n \right\| \\ &= \left\| \Delta \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) (\mathbf{w}^n - \mathbf{y}^n) + \Delta \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i - \mathbf{p}^* \right) \mathbf{y}^n \right\|. \end{aligned}$$

Applying the triangle inequality to this last expression, we have

$$\left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i - \mathbf{d} \right\| \leq \left\| \Delta \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i \right) (\mathbf{w}^n - \mathbf{y}^n) \right\| + \left\| \Delta \left(\frac{1}{n} \sum_{i=1}^n \mathbf{p}^i - \mathbf{p}^* \right) \mathbf{y}^n \right\|.$$

By Lemmata 3 and 4 presented at the beginning of this section and by our definition of \bar{Y} , we then have

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^i - \mathbf{d} \right\| &\leq |\mathcal{V}| \cdot \max \Delta_{v',x',b'} \|\mathbf{w}^n - \mathbf{y}^n\| + \left(\max_{x' \in X} \sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v,x',by_b^n} \right) \cdot \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{p}^i - \mathbf{p}^* \right\| \\ &\leq |\mathcal{V}| \cdot \max \Delta_{v',x',b'} \|\mathbf{w}^n - \mathbf{y}^n\| + \bar{Y} \cdot \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{p}^i - \mathbf{p}^* \right\|. \end{aligned}$$

By our choice of N_1, N_2, N_3 and \mathbf{y}^n , we have that

$$\begin{aligned} \left\| \frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^n - \mathbf{d} \right\| &\leq |\mathcal{V}| \cdot \max \Delta_{v', x', b'} \|\mathbf{w}^n - \mathbf{y}^n\| + \bar{Y} \cdot \left\| \frac{1}{n} \sum_{i=1}^n \mathbf{p}^i - \mathbf{p}^* \right\| \\ &\leq \frac{\epsilon}{2} + \frac{\epsilon}{2} \\ &= \epsilon. \end{aligned}$$

We have just shown that $1/n \cdot \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^n \in B(\mathbf{d}, \epsilon)$. Since $B(\mathbf{d}, \epsilon) \subseteq U(\mathbf{D}, \epsilon)$, we have that

$$\frac{1}{n} \sum_{i=1}^n \Delta \mathbf{p}^i \mathbf{w}^n \in U(\mathbf{D}, \epsilon),$$

which is the required result. ■

Appendix B

Pathological PMF sequence example for Chapter 4

Consider a sequence of 30 PMFs where

$$\begin{aligned}\mathbf{p}^{2k+1} &= (0.05, 0.20, 0.50, 0.20, 0.05), & k = 0, \dots, 14, \\ \mathbf{p}^{2k+2} &= (0.30, 0.15, 0.10, 0.15, 0.30), & k = 0, \dots, 14.\end{aligned}$$

This PMF sequence oscillates between two PMFs over all 30 fractions, and these PMFs differ quite significantly from each other. Due to the oscillatory nature of this PMF sequence, we would expect that the lower and upper bounds from exponential smoothing would be constantly “out-of-sync” with the PMF that is realized on each day, resulting in tumour underdose by the end of the treatment.

We can verify this by plotting the realized PMFs and the lower and upper bound vectors from three exponential smoothing implementations for the last two fractions: Figure B.1 shows fraction 29, while Figure B.2 shows fraction 30. (For the $\alpha = 0.5$ and $\alpha = 0.9$ methods, the lower and upper bound vectors are virtually identical by the last fraction, which is why there is only one line for each of these methods on each of the plots.) We can very clearly see that in both fractions, the realized PMF is not in any of the uncertainty sets.

The consequences of trying to track the oscillating PMFs are visible in the dose statistics that are shown in Table B.1. In Figures B.3 and B.4 we also plot the minimum tumour dose as a percentage of 72Gy versus mean left lung dose as a percentage of the static margin dose for each of the treatments for this PMF sequence. The same




Figure B.1: Plot of observed PMF in fraction 29 versus lower/upper bound vectors for fraction 29 for three different exponential smoothing treatments.




Figure B.2: Plot of observed PMF in fraction 30 versus lower/upper bound vectors for fraction 30 for three different exponential smoothing treatments.

implementations (with the exception of $\alpha = 1$) considered for the first PMF sequence in Section 4.3 were considered for this sequence. The same patient anatomy and robust optimization parameters used for the real patient sequences were also used here. As we can see from Table B.1 and Figure B.3, the exponential smoothing update algorithm exhibits poorer performance with increasing α under this sequence. When $\alpha = 0$, the exponential smoothing and the static solutions are identical, so as α approaches 0, the exponential smoothing treatment will approach the static treatment with the same initial uncertainty set. For this particular sequence, it happens that the static nominal and robust treatments actually lead to high levels of tumour coverage, due to the fact that the nominal and robust uncertainty sets are very close to the average PMF of this sequence.

In contrast to exponential smoothing, the running average method is relatively insensitive to the oscillatory nature of this PMF sequence. The reason for this is that this particular PMF sequence has a well-behaved average. When the running average update algorithm is applied to this sequence, the lower and upper bound vectors tend to the average of the sequence, and the intensity vector delivered in each fraction converges to the intensity vector that protects against the average PMF. As we noted in our discussion of prescient solutions in Section 4.1.2, delivering the intensity vector that protects against the average PMF in every fraction results in no tumour underdose, and the running average treatment, for a large number of fractions, is delivering a very similar intensity vector to the average PMF intensity vector, which is why it is still able to ensure a high level of tumour coverage.

These results are not too concerning from a clinical standpoint, because a patient PMF sequence that changes so drastically from fraction to fraction is very unlikely to be observed in practice. In the context of the patient anatomy that we have used for this sequence and for the real patient sequences, a patient breathing according to this sequence would alternate between spending 10% of the time in full inhale or exhale (motion states 1 and 5 respectively) during one treatment session, and spending 60% of the time in those same breathing phases during the next treatment session. As long as the uncertainty set update algorithm retains enough memory to avoid simply tracking the latest PMF, these pathological cases will not significantly affect the quality of the final adaptive robust treatment.

Implementation	Min. tumour dose		Mean lung dose		Mean n. tissue dose	
	Gy	%	Gy	%	Gy	%
(S,N)	71.25	98.96	17.39	85.31	9.01	89.00
(ES(0.1),N)	70.58	98.02	17.37	85.22	9.01	88.96
(ES(0.5),N)	63.46	88.14	17.41	85.39	9.00	88.82
(ES(0.9),N)	54.02	75.03	17.37	85.21	8.94	88.25
(RA,N)	70.76	98.28	17.37	85.24	9.01	88.94
(S,R)	71.97	99.96	18.19	89.23	9.41	92.92
(ES(0.1),R)	71.86	99.80	17.76	87.11	9.16	90.48
(ES(0.5),R)	64.27	89.26	17.49	85.81	9.03	89.20
(ES(0.9),R)	53.89	74.85	17.40	85.37	8.96	88.43
(RA,R)	71.50	99.30	17.57	86.17	9.08	89.64
(S,M)	72.06	100.08	20.38	100.00	10.13	100.00
(ES(0.1),M)	72.04	100.06	18.34	89.97	9.42	92.97
(ES(0.5),M)	64.76	89.95	17.62	86.43	9.09	89.71
(ES(0.9),M)	54.29	75.40	17.50	85.83	8.99	88.75
(RA,M)	72.03	100.04	17.89	87.76	9.20	90.86
(DLYP)	72.00	100.00	17.39	85.31	8.91	88.01
(AVGP)	72.00	100.00	17.39	85.30	9.01	88.98

Table B.1: Dose statistics for the pathological PMF sequence.

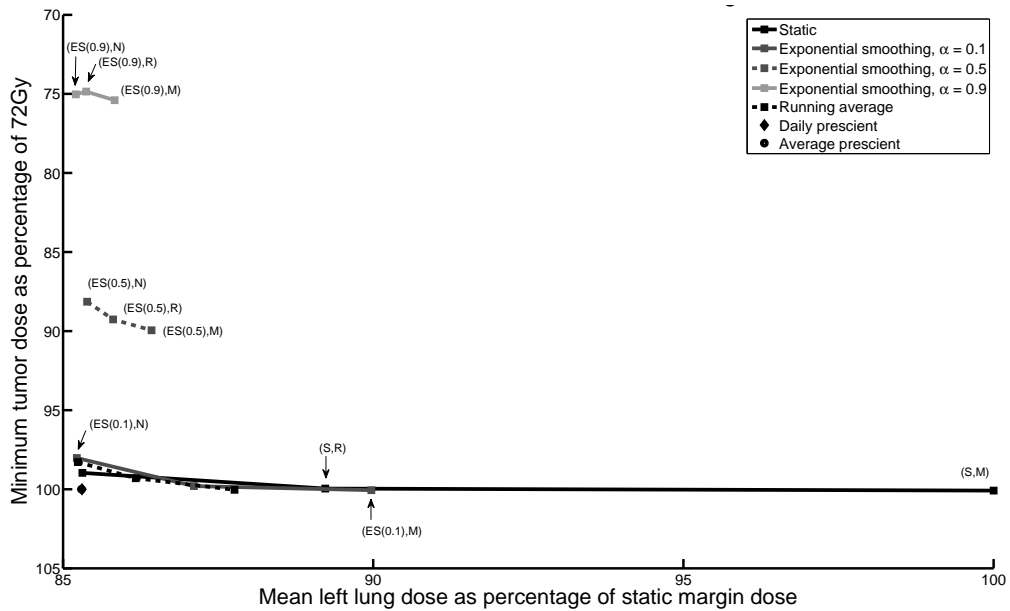


Figure B.3: Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the pathological PMF sequence (full).




Figure B.4: Plot of minimum tumour dose versus mean left lung dose for the different implementations applied to the pathological PMF sequence (zoomed-in).

Appendix C

Proofs for Chapter 5

To prove Theorem 4, we will make the same three assumptions (Assumptions 1, 2 and 3) made in Appendix A. The proof leverages a large part of the proof of Theorem 2 in Chapter 4; for completeness, we provide the full proof here.

Proof of Theorem 4: Let $\epsilon > 0$. By Proposition 4, the set $\mathbf{w}^*(\mathbf{p}^*)$, which is the set of optimal solutions to the robust problem (3.1) with $P = \{\mathbf{p}^*\}$, is a bounded set. Since it is also a closed set, it must therefore be a compact subset of $\mathbb{R}^{|\mathcal{B}|}$. The function $Y : \mathbb{R}^{|\mathcal{B}|} \rightarrow \mathbb{R}$ defined as

$$Y(\mathbf{w}) = \max_{x' \in X} \sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v,x',b} w_b$$

is a continuous function, so by the extreme value theorem, it attains a maximum on $\mathbf{w}^*(\mathbf{p}^*)$. Let \bar{Y} denote that maximum value.

Since $\mathbf{p}^i \rightarrow \mathbf{p}^*$ as $i \rightarrow \infty$, there must exist an $N_1 \in \mathbb{Z}_+$ such that, for $i > N_1$,

$$\|\mathbf{p}^i - \mathbf{p}^*\| < \frac{\epsilon}{2 \cdot \bar{Y}}.$$

By Theorem 1, there must also exist an $N_2 \in \mathbb{Z}_+$ such that, for $i > N_2$,

$$\mathbf{w}^*(\ell^i, \mathbf{u}^i) \subseteq U \left(\mathbf{w}^*(\mathbf{p}^*), \frac{\epsilon}{2 \cdot |\mathcal{V}| \cdot \max \Delta_{v',x',b'}} \right),$$

where the maximum, $\max \Delta_{v',x',b'}$, in the denominator of the radius is taken over all voxels $v' \in \mathcal{V}$, motion states $x' \in X$ and beamlets $b' \in \mathcal{B}$.

Set $N_3 = \max\{N_1, N_2\}$, and let $i > N_3$. To prove the result, we will construct a dose vector inside \mathbf{D} whose epsilon ball contains $\Delta \mathbf{p}^i \mathbf{w}^i$. Given $\mathbf{w}^i \in \mathbf{w}^*(\ell^i, \mathbf{u}^i)$, select

$\mathbf{y}^i \in \mathbf{w}^*(\mathbf{p}^*)$ such that $\|\mathbf{y}^i - \mathbf{w}^i\| < \epsilon/(2|\mathcal{V}| \max \Delta_{v',x',b'})$; this is possible by our use of Theorem 1 above. Define $\mathbf{d} = \Delta \mathbf{p}^* \mathbf{y}^i$. Since $\mathbf{y}^i \in \mathbf{w}^*(\mathbf{p}^*)$, clearly $\mathbf{d} \in \mathbf{D}$.

Now, we have that

$$\begin{aligned} \|\Delta \mathbf{p}^i \mathbf{w}^i - \mathbf{d}\| &= \|\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^i \mathbf{y}^i + \Delta \mathbf{p}^i \mathbf{y}^i - \Delta \mathbf{p}^* \mathbf{y}^i\| \\ &\leq \|\Delta \mathbf{p}^i \mathbf{w}^i - \Delta \mathbf{p}^i \mathbf{y}^i\| + \|\Delta \mathbf{p}^i \mathbf{y}^i - \Delta \mathbf{p}^* \mathbf{y}^i\| \\ &\leq |\mathcal{V}| \cdot \max \Delta_{v',x',b'} \cdot \|\mathbf{w}^i - \mathbf{y}^i\| + \max_{x' \in X} \sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v,x',b} y_b^i \cdot \|\mathbf{p}^i - \mathbf{p}^*\|, \end{aligned}$$


where the first inequality follows by the triangle inequality, and the second inequality follows by applying Lemmata 3 and 4. By our choice of N_1 , N_2 and N_3 and the definition of \bar{Y} , we have that

$$\begin{aligned} \|\Delta \mathbf{p}^i \mathbf{w}^i - \mathbf{d}\| &< |\mathcal{V}| \cdot \max \Delta_{v',x',b'} \cdot \frac{\epsilon}{2 \cdot |\mathcal{V}| \cdot \max \Delta_{v',x',b'}} + \max_{x' \in X} \sum_{v \in \mathcal{V}} \sum_{b \in \mathcal{B}} \Delta_{v,x',b} y_b^i \cdot \frac{\epsilon}{2 \cdot \bar{Y}} \\ &\leq |\mathcal{V}| \cdot \max \Delta_{v',x',b'} \cdot \frac{\epsilon}{2 \cdot |\mathcal{V}| \cdot \max \Delta_{v',x',b'}} + \bar{Y} \cdot \frac{\epsilon}{2 \cdot \bar{Y}} \\ &= \frac{\epsilon}{2} + \frac{\epsilon}{2} \\ &= \epsilon. \end{aligned}$$


This shows that $\Delta \mathbf{p}^i \mathbf{w}^i \in B(\mathbf{d}, \epsilon)$ for some $\mathbf{d} \in \mathbf{D}$, and thus that $\Delta \mathbf{p}^i \mathbf{w}^i \in U(\mathbf{D}, \epsilon)$. ■

Appendix D

Second PMF sequence results for Chapter 5




(a) Min. tumour dose vs. mean left lung dose.




(b) Max. tumour dose vs. mean left lung dose.

Figure D.1: Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0$ applied to the second PMF sequence.




(a) Min. tumour dose vs. mean left lung dose.




(b) Max. tumour dose vs. mean left lung dose.

Figure D.2: Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.1$ applied to the second PMF sequence.




(a) Min. tumour dose vs. mean left lung dose.




(b) Max. tumour dose vs. mean left lung dose.

Figure D.3: Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.5$ applied to the second PMF sequence.




(a) Min. tumour dose vs. mean left lung dose.




(b) Max. tumour dose vs. mean left lung dose.


Figure D.4: Plots of minimum tumour dose and maximum tumour dose versus mean left lung dose for implementations with $\alpha = 0.9$ applied to the second PMF sequence.




(a) Mean and max. underdose, (S, \cdot) implementations.




(b) Mean and max. overdose, (S, \cdot) implementations.




(c) Mean and max. underdose, ($R^{\pm}ES(0),\cdot$) implementations.




(d) Mean and max. overdose, ($R^{\pm}ES(0),\cdot$) implementations.




(e) Mean and max. underdose, ($R^{-}ES(0),\cdot$) implementations.



(f) Mean and max. overdose, ($R^{-}ES(0),\cdot$) implementations.




(g) Mean and max. underdose, ($R^{+}ES(0),\cdot$) implementations.




(h) Mean and max. overdose, ($R^{+}ES(0),\cdot$) implementations.


Figure D.5: Mean and max. underdose and overdose by fraction for implementations using the reactive, reactive⁻ or reactive⁺ exponential smoothing methods with $\alpha = 0$ or the static (non-adaptive, non-reactive) method, under the second PMF sequence.




(a) Mean and max. underdose, (ES(0.1),-) implementations.




(b) Mean and max. overdose, (ES(0.1),-) implementations.




(c) Mean and max. underdose, (R[±]ES(0.1),-) implementations.




(d) Mean and max. overdose, (R[±]ES(0.1),-) implementations.




(e) Mean and max. underdose, (R⁻ES(0.1),-) implementations.



(f) Mean and max. overdose, (R⁻ES(0.1),-) implementations.




(g) Mean and max. underdose, (R⁺ES(0.1),-) implementations.




(h) Mean and max. overdose, (R⁺ES(0.1),-) implementations.


Figure D.6: Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive[±], reactive⁻ or reactive⁺ exponential smoothing methods with $\alpha = 0.1$, under the second PMF sequence.




(a) Mean and max. underdose, (ES(0.5),-) implementations.




(b) Mean and max. overdose, (ES(0.5),-) implementations.




(c) Mean and max. underdose, (R⁺ES(0.5),-) implementations.




(d) Mean and max. overdose, (R⁺ES(0.5),-) implementations.




(e) Mean and max. underdose, (R⁻ES(0.5),-) implementations.



(f) Mean and max. overdose, (R⁻ES(0.5),-) implementations.




(g) Mean and max. underdose, (R⁺ES(0.5),-) implementations.




(h) Mean and max. overdose, (R⁺ES(0.5),-) implementations.


Figure D.7: Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive[±], reactive⁻ or reactive⁺ exponential smoothing methods with $\alpha = 0.5$, under the second PMF sequence.




(a) Mean and max. underdose, (ES(0.9),-) implementations.




(b) Mean and max. overdose, (ES(0.9),-) implementations.




(c) Mean and max. underdose, (R±ES(0.9),-) implementations.




(d) Mean and max. overdose, (R±ES(0.9),-) implementations.




(e) Mean and max. underdose, (R-ES(0.9),-) implementations.



(f) Mean and max. overdose, (R-ES(0.9),-) implementations.




(g) Mean and max. underdose, (R+ES(0.9),-) implementations.




(h) Mean and max. overdose, (R+ES(0.9),-) implementations.


Figure D.8: Mean and max. underdose and overdose by fraction for implementations using the non-reactive, reactive[±], reactive⁻ or reactive⁺ exponential smoothing methods with $\alpha = 0.9$, under the second PMF sequence.




(a) Mean and max. underdose, (DLYP) implementation.




(b) Mean and max. overdose, (DLYP) implementation.




(c) Mean and max. underdose, (R^\pm DLYP) implementation.




(d) Mean and max. overdose, (R^\pm DLYP) implementation.




(e) Mean and max. underdose, (R^- DLYP) implementation.



(f) Mean and max. overdose, (R^- DLYP) implementation.




(g) Mean and max. underdose, (R^+ DLYP) implementation.




(h) Mean and max. overdose, (R^+ DLYP) implementation.


Figure D.9: Mean and max. underdose and overdose by fraction for the non-reactive, reactive[±], reactive⁻ and reactive⁺ daily prescient methods under the second PMF sequence.




(a) Daily dose, (ES(0.1),R) implementation.




(b) Daily dose and target minimum and maximum doses, (ES(0.1),R) implementation.




(c) Daily dose, (R±ES(0.1),R) implementation.




(d) Daily dose and target minimum and maximum doses, (R±ES(0.1),R) implementation.




(e) Daily dose, (R-ES(0.1),R) implementation.



(f) Daily dose and target minimum and maximum doses, (R-ES(0.1),R) implementation.



(g) Daily dose, (R+ES(0.1),R) implementation.



(h) Daily dose and target minimum and maximum doses, (R+ES(0.1),R) implementation.

Figure D.10: Daily delivered dose and daily target minimum and maximum doses by fraction for the (ES(0.1),R), (R±ES(0.1),R), (R-ES(0.1),R) and (R+ES(0.1),R) implementations under the second PMF sequence. The lower error bar, median and upper error bar for each fraction correspond to the 5th, 50th and 95th percentiles of the appropriate distribution, respectively. On the left-hand plots, the dashed horizontal lines correspond to 2.4Gy and 2.64Gy.

Implementation	Min. tumour dose		Max. tumour dose		Mean lung dose		Mean n. tissue dose	
	Gy	%	Gy	%	Gy	%	Gy	%
(S,N)	64.37	89.40	85.82	108.36	17.18	85.11	8.99	88.92
(S,R)	71.99	99.98	79.06	99.83	17.99	89.12	9.39	92.90
(S,M)	72.05	100.07	78.42	99.02	20.18	100.00	10.11	100.00
(ES(0.1),N)	70.46	97.87	80.74	101.95	17.51	86.78	9.00	88.99
(ES(0.1),R)	71.98	99.98	79.13	99.91	17.68	87.59	9.12	90.20
(ES(0.1),M)	72.05	100.07	78.84	99.54	18.07	89.54	9.36	92.62
(ES(0.5),N)	71.61	99.46	79.49	100.37	17.55	86.96	8.98	88.87
(ES(0.5),R)	72.00	100.00	79.17	99.97	17.59	87.15	9.01	89.16
(ES(0.5),M)	72.01	100.01	79.10	99.87	17.69	87.67	9.06	89.65
(ES(0.9),N)	71.69	99.56	79.43	100.29	17.54	86.92	8.98	88.85
(ES(0.9),R)	71.96	99.94	79.20	100.00	17.58	87.08	9.00	89.03
(ES(0.9),M)	71.98	99.97	79.15	99.94	17.65	87.43	9.02	89.27
(DLYP)	72.00	100.00	79.20	100.00	17.55	86.97	8.98	88.84
(R [±] ES(0),N)	71.82	99.74	79.40	100.26	17.55	86.96	9.07	89.75
(R [±] ES(0),R)	71.99	99.99	79.20	100.00	17.55	86.96	9.10	90.02
(R [±] ES(0),M)	72.00	100.00	79.20	99.99	18.35	90.93	9.43	93.25
(R [±] ES(0.1),N)	71.97	99.96	79.20	100.00	17.56	87.02	8.99	88.91
(R [±] ES(0.1),R)	71.98	99.97	79.20	100.01	17.59	87.13	9.00	89.06
(R [±] ES(0.1),M)	72.00	100.00	79.20	100.00	17.68	87.59	9.09	89.95
(R [±] ES(0.5),N)	71.96	99.94	79.20	100.00	17.57	87.03	8.98	88.87
(R [±] ES(0.5),R)	71.96	99.94	79.21	100.02	17.58	87.08	8.99	88.92
(R [±] ES(0.5),M)	71.96	99.94	79.20	100.01	17.68	87.61	9.02	89.19
(R [±] ES(0.9),N)	71.94	99.91	79.23	100.03	17.56	87.02	8.98	88.86
(R [±] ES(0.9),R)	71.94	99.92	79.22	100.03	17.57	87.06	8.99	88.89
(R [±] ES(0.9),M)	71.94	99.91	79.22	100.03	17.64	87.39	9.00	89.04
(R [±] DLYP)	72.00	100.00	79.20	100.00	17.56	87.00	8.98	88.83

Table D.1: Dose statistics for the second PMF sequence under the non-reactive and reactive[±] methods.

Implementation	Min. tumour dose		Max. tumour dose		Mean lung dose		Mean n. tissue dose	
	Gy	%	Gy	%	Gy	%	Gy	%
(R ⁻ ES(0),N)	71.83	99.76	83.73	105.72	17.55	86.94	9.17	90.74
(R ⁻ ES(0),R)	72.02	100.03	79.20	100.00	17.95	88.93	9.33	92.33
(R ⁻ ES(0),M)	72.12	100.16	79.19	99.99	19.18	95.03	9.87	97.59
(R ⁻ ES(0.1),N)	71.99	99.99	80.41	101.53	17.57	87.03	9.03	89.35
(R ⁻ ES(0.1),R)	72.02	100.02	79.20	100.00	17.65	87.45	9.10	90.00
(R ⁻ ES(0.1),M)	72.04	100.06	79.20	100.00	17.98	89.07	9.30	92.02
(R ⁻ ES(0.5),N)	71.99	99.99	79.53	100.41	17.57	87.05	9.00	88.99
(R ⁻ ES(0.5),R)	72.01	100.01	79.20	100.01	17.60	87.18	9.02	89.24
(R ⁻ ES(0.5),M)	72.02	100.03	79.20	100.00	17.69	87.65	9.08	89.79
(R ⁻ ES(0.9),N)	71.99	99.98	79.48	100.35	17.56	86.98	8.99	88.95
(R ⁻ ES(0.9),R)	72.00	100.00	79.29	100.11	17.58	87.12	9.00	89.06
(R ⁻ ES(0.9),M)	72.01	100.01	79.24	100.05	17.64	87.42	9.03	89.27
(R ⁻ DLYP)	72.00	100.00	79.20	100.00	17.55	86.95	8.98	88.83
(R ⁺ ES(0),N)	67.79	94.15	79.36	100.20	17.25	85.45	8.94	88.44
(R ⁺ ES(0),R)	71.98	99.98	79.04	99.80	17.57	87.06	9.15	90.54
(R ⁺ ES(0),M)	72.00	100.00	78.09	98.60	18.81	93.19	9.59	94.89
(R ⁺ ES(0.1),N)	70.63	98.09	79.20	100.00	17.51	86.75	8.96	88.59
(R ⁺ ES(0.1),R)	71.97	99.95	79.14	99.92	17.56	87.02	9.03	89.29
(R ⁺ ES(0.1),M)	72.00	100.00	78.74	99.42	17.77	88.03	9.16	90.57
(R ⁺ ES(0.5),N)	71.59	99.43	79.20	100.00	17.55	86.93	8.97	88.76
(R ⁺ ES(0.5),R)	71.94	99.91	79.14	99.93	17.58	87.10	9.00	89.01
(R ⁺ ES(0.5),M)	71.82	99.75	79.07	99.83	17.67	87.54	9.03	89.33
(R ⁺ ES(0.9),N)	71.66	99.53	79.20	100.00	17.55	86.95	8.97	88.77
(R ⁺ ES(0.9),R)	71.87	99.83	79.17	99.96	17.57	87.04	8.99	88.93
(R ⁺ ES(0.9),M)	71.90	99.86	79.13	99.92	17.64	87.38	9.01	89.11
(R ⁺ DLYP)	72.00	100.00	79.20	100.00	17.57	87.06	8.98	88.84

Table D.2: Dose statistics for the second PMF sequence under the reactive⁻ and reactive⁺ methods.

Bibliography

- American Cancer Society. *Cancer Facts and Figures 2010*. American Cancer Society, Atlanta, 2010.
- American Cancer Society. *Cancer Facts and Figures 2011*. American Cancer Society, Atlanta, 2011.
- P. Anand, A.B. Kunnumakara, C. Sundaram, K.B. Harikumar, S.T. Tharakan, O.S. Lai, B. Sung, and B.B. Aggarwal. Cancer is a preventable disease that requires major lifestyle changes. *Pharmaceutical research*, 25(9):2097–2116, 2008.
- A. Ben-Tal and A. Nemirovski. Robust truss topology design via semidefinite programming. *SIAM Journal on Optimization*, 7(4):991–1016, 1997.
- A. Ben-Tal and A. Nemirovski. Robust solutions of uncertain linear programs. *Operations Research Letters*, 25(1):1–14, 1999.
- A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.
- A. Ben-Tal, B. Golany, A. Nemirovski, and J.P. Vial. Retailer-supplier flexible commitments contracts: A robust optimization approach. *Manufacturing & Service Operations Management*, 7(3):248–271, 2005.
- A. Ben-Tal, L. El Ghaoui, and A.S. Nemirovski. *Robust optimization*. Princeton University Press, 2009.
- D. Bertsimas and C. Caramanis. Adaptability via sampling. In *Proceedings of the 46th IEEE Conference on Decision and Control*, pages 4717–4722, 2007.
- D. Bertsimas and C. Caramanis. Finite adaptability in multistage linear optimization. *IEEE Transactions on Automatic Control*, 55(12):2751–2766, 2010.

- D. Bertsimas and M. Sim. The price of robustness. *Operations Research*, 52(1):35–53, 2004.
- D. Bertsimas and A. Thiele. A robust optimization approach to inventory theory. *Operations Research*, 54(1):150–168, 2006.
- D. Bertsimas, D.B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.
- M. Birkner, D. Yan, M. Alber, J. Liang, and F. Nüsslin. Adapting inverse planning to patient and organ geometrical variation: algorithm and implementation. *Medical Physics*, 30:2822–2831, 2003.
- T. Bortfeld, K. Jokivarsi, M. Goitein, J. Kung, and S.B. Jiang. Effects of intra-fraction motion on IMRT dose delivery: statistical analysis and simulation. *Physics in Medicine and Biology*, 47:2203–2220, 2002.
- T. Bortfeld, S.B. Jiang, and E. Rietzel. Effects of motion on the total dose distribution. *Seminars in Radiation Oncology*, 14(1):41–51, 2004.
- T. Bortfeld, T.C.Y. Chan, A. Trofimov, and J.N. Tsitsiklis. Robust management of motion uncertainty in intensity-modulated radiation therapy. *Operations Research*, 56(6):1461–1473, 2008.
- Canadian Cancer Society’s Steering Committee. *Canadian Cancer Statistics 2010*. Canadian Cancer Society, Toronto, 2010.
- Canadian Cancer Society’s Steering Committee. *Canadian Cancer Statistics 2011*. Canadian Cancer Society, Toronto, 2011.
- T.C.Y. Chan, T. Bortfeld, and J.N. Tsitsiklis. A robust approach to IMRT optimization. *Physics in Medicine and Biology*, 51:2567–2583, 2006.
- M. Chu, Y. Zinchenko, S.G. Henderson, and M.B. Sharpe. Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty. *Physics in Medicine and Biology*, 50:5463–5477, 2005.
- J.D. Cox, N. Azarnia, R.W. Byhardt, K.H. Shin, B. Emami, and T.F. Pajak. A randomized phase I/II trial of hyperfractionated radiation therapy with total doses of 60.0

- Gy to 79.2 Gy: possible survival benefit with greater than or equal to 69.6 Gy in favorable patients with Radiation Therapy Oncology Group stage III non-small-cell lung carcinoma: report of Radiation Therapy Oncology Group 83-11. *Journal of Clinical Oncology*, 8(9):1543–1555, 1990.
- G.B. Dantzig, J. Folkman, and N. Shapiro. On the continuity of the minimum set of a continuous function. *Journal of Mathematical Analysis and Applications*, 17(3):519–548, 1967.
- A. de la Zerda, B. Armbruster, and L. Xing. Formulating adaptive radiation therapy (ART) treatment planning into a closed-loop control framework. *Physics in Medicine and Biology*, 52:4137–4153, 2007.
- G. Deng and M.C. Ferris. Neuro-dynamic programming for fractionated radiotherapy planning. In C.J.S. Alves, P.M. Pardalos, and L.N. Vicente, editors, *Optimization in Medicine*, pages 47–70. Springer Verlag, 2008.
- M.A. Ebert. Viability of the EUD and TCP concepts as reliable dose indicators. *Physics in Medicine and Biology*, 45:441–457, 2000.
- M.C. Ferris and M.M. Voelker. Fractionation in radiation treatment planning. *Mathematical Programming*, 101(2):387–413, 2004.
- A. Fredriksson, A. Forsgren, and B. Hårdemark. Minimax optimization for handling range and setup uncertainties in proton therapy. *Medical Physics*, 38:1672–1684, 2011.
- D.P. Gierga, J. Brewer, G.C. Sharp, M. Betke, C.G. Willett, and G.T.Y. Chen. The correlation between internal and external markers for abdominal tumors: implications for respiratory gating. *International Journal of Radiation Oncology, Biology, Physics*, 61(5):1551–1558, 2005.
- H. Keller, M.A. Ritter, and T.R. Mackie. Optimal stochastic correction strategies for rigid-body target motion. *International Journal of Radiation Oncology, Biology, Physics*, 55(1):261–270, 2003.
- F.M. Kong, R.K. Ten Haken, M.J. Schipper, M.A. Sullivan, M. Chen, C. Lopez, G.P. Kalemkerian, and J.A. Hayman. High-dose radiation improved local tumor control and overall survival in patients with inoperable/unresectable non-small-cell lung

- cancer: long-term results of a radiation dose escalation study. *International Journal of Radiation Oncology, Biology, Physics*, 63(2):324–333, 2005.
- K.L. Lam, R.K.T. Haken, D. Litzenberg, J.M. Balter, and S.M. Pollock. An application of Bayesian statistical methods to adaptive radiotherapy. *Physics in Medicine and Biology*, 50:3849–3858, 2005.
- J.G. Li and L. Xing. Inverse planning incorporating organ motion. *Medical Physics*, 27:1573–1578, 2000.
- J. Löf, B.K. Lind, and A. Brahme. An adaptive control algorithm for optimization of intensity modulated radiotherapy considering uncertainties in beam profiles, patient set-up and internal organ motion. *Physics in Medicine and Biology*, 43:1605–1628, 1998.
- W. Lu, G.H. Olivera, Q. Chen, K.J. Ruchala, J. Haimerl, S.L. Meeks, K.M. Langen, and P.A. Kupelian. Deformable registration of the planning image (kVCT) and the daily images (MVCT) for adaptive radiation therapy. *Physics in Medicine and Biology*, 51:4357–4374, 2006.
- A.E. Lujan, E.W. Larsen, J.M. Balter, and R.K. Ten Haken. A method for incorporating organ motion due to breathing into 3D dose calculations. *Medical Physics*, 26:715–720, 1999.
- A.E. Lujan, J.M. Balter, and R.K. Ten Haken. A method for incorporating organ motion due to breathing into 3D dose calculations in the liver: Sensitivity to variations in motion. *Medical Physics*, 30:2643–2649, 2003.
- L.K. Mell, A.K. Mehrotra, and A.J. Mundt. Intensity-modulated radiation therapy use in the US, 2004. *Cancer*, 104(6):1296–1303, 2005.
- A. Mestrovic, M.P. Milete, A. Nichol, B.G. Clark, and K. Otto. Direct aperture optimization for online adaptive radiation therapy. *Medical Physics*, 34:1631, 2007.
- R. Mohan, X. Zhang, H. Wang, Y. Kang, X. Wang, H. Liu, K.K. Ang, D. Kuban, and L. Dong. 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, 2005.

- K. Ohara, T. Okumura, M. Akisada, T. Inada, T. Mori, H. Yokota, and M.J.B. Calaguas. Irradiation synchronized with respiration gate. *International Journal of Radiation Oncology, Biology, Physics*, 17(4):853–857, 1989.
- A. Ólafsson and S.J. Wright. Efficient schemes for robust IMRT treatment planning. *Physics in Medicine and Biology*, 51:5621–5642, 2006.
- J. Osterman, W. Weaver, J. Slater, and P. Glass. Failure-proof method for intrinsic field subtraction. *Manhattan Journal of Physics*, 1(4):7–10, 1959.
- C.A. Perez, M. Bauer, S. Edelstein, B.W. Gillespie, and R. Birch. Impact of tumor control on survival in carcinoma of the lung treated with irradiation. *International Journal of Radiation Oncology, Biology, Physics*, 12(4):539–547, 1986.
- H. Rehbinder, C. Forsgren, and J. Löf. Adaptive radiation therapy for compensation of errors in patient setup and treatment delivery. *Medical Physics*, 31:3363–3371, 2004.
- R. Rengan, K.E. Rosenzweig, E. Venkatraman, L.A. Koutcher, J.L. Fox, R. Nayak, H. Amols, E. Yorke, A. Jackson, C.C. Ling, and S.A. Leibel. Improved local control with higher doses of radiation in large-volume stage iii non-small-cell lung cancer. *International Journal of Radiation Oncology, Biology, Physics*, 60(3):741–747, 2004.
- H.E. Romeijn and J.F. Dempsey. Intensity modulated radiation therapy treatment plan optimization. *Top*, 16(2):215–243, 2008.
- H.E. Romeijn, R.K. Ahuja, J.F. Dempsey, and A. Kumar. A New Linear Programming Approach to Radiation Therapy Treatment Planning Problems. *Operations Research*, 54(2):201, 2006.
- W. Rudin. *Principles of Mathematical Analysis*. McGraw-Hill Publishing Co., 1976.
- D. Schulz-Ertner and H. Tsujii. Particle radiation therapy using proton and heavier ion beams. *Journal of Clinical Oncology*, 25(8):953–964, 2007.
- K. Sheng, J. Cai, J. Brookeman, J. Molloy, J. Christopher, and P. Read. A computer simulated phantom study of tomotherapy dose optimization based on probability density functions (pdf) and potential errors caused by low reproducibility of pdf. *Medical Physics*, 33:3321–3326, 2006.

- M. Sir, M.A. Epelman, and S.M. Pollock. Stochastic programming for off-line adaptive radiotherapy. *Annals of Operations Research*, pages 1–31, 2010. URL <http://dx.doi.org/10.1007/s10479-010-0779-x>.
- M.Y. Sir. *Optimization of radiotherapy considering uncertainties caused by daily setup procedures and organ motion*. PhD thesis, University of Michigan, 2007.
- J.J. Sonke, L. Zijp, P. Remeijer, and M. van Herk. Respiratory correlated cone beam ct. *Medical Physics*, 32:1176–1186, 2005.
- S.V. Spirou and C.S. Chui. Generation of arbitrary intensity profiles by dynamic jaws or multileaf collimators. *Medical Physics*, 21:1031–1041, 1994.
- H.D. Thames and J.H. Hendry. *Fractionation in radiotherapy*. Taylor and Francis London, 1987.
- C. Thieke, T. Bortfeld, and K.H. Küfer. Characterization of dose distributions through the max and mean dose concept. *Acta Oncologica*, 41(2):158–161, 2002.
- L. Toschi, F. Cappuzzo, and PA Janne. Evolution and future perspectives in the treatment of locally advanced non-small cell lung cancer. *Annals of Oncology*, 18 (Supplement 9), 2007.
- A. Trofimov, E. Rietzel, H.M. Lu, B. Martin, S. Jiang, G.T.Y. Chen, and T. Bortfeld. Temporo-spatial IMRT optimization: concepts, implementation and initial results. *Physics in Medicine and Biology*, 50:2779–2798, 2005.
- Y. Tsunashima, T. Sakae, Y. Shioyama, K. Kagei, T. Terunuma, A. Nohtomi, and Y. Akine. Correlation between the respiratory waveform measured using a respiratory sensor and 3d tumor motion in gated radiotherapy. *International Journal of Radiation Oncology, Biology, Physics*, 60(3):951–958, 2004.
- J. Unkelbach and U. Oelfke. Inclusion of organ movements in IMRT treatment planning via inverse planning based on probability distributions. *Physics in Medicine and Biology*, 49:4005–4029, 2004.
- J. Unkelbach and U. Oelfke. Incorporating organ movements in inverse planning: assessing dose uncertainties by Bayesian inference. *Physics in Medicine and Biology*, 50:121–139, 2005.

- J. Unkelbach, T.C.Y. Chan, and T. Bortfeld. Accounting for range uncertainties in the optimization of intensity modulated proton therapy. *Physics in Medicine and Biology*, 52:2755–2773, 2007.
- M. van Herk. Errors and margins in radiotherapy. *Seminars in Radiation Oncology*, 14(1):52–64, 2004.
- M. van Herk, P. Remeijer, C. Rasch, and J.V. Lebesque. The probability of correct target dosage: dose-population histograms for deriving treatment margins in radiotherapy. *International Journal of Radiation Oncology, Biology, Physics*, 47(4):1121–1135, 2000.
- S. Webb. Adapting IMRT delivery fraction-by-fraction to cater for variable intrafraction motion. *Physics in Medicine and Biology*, 53:1–21, 2008.
- Q. Wu, J. Liang, and D. Yan. Application of dose compensation in image-guided radiotherapy of prostate cancer. *Physics in Medicine and Biology*, 51:1405–1419, 2006.
- Q.J. Wu, D. Thongphiew, Z. Wang, B. Mathayomchan, V. Chankong, S. Yoo, W.R. Lee, and F.F. Yin. On-line re-optimization of prostate IMRT plans for adaptive radiation therapy. *Physics in Medicine and Biology*, 53:673–691, 2008.
- D. Yan, F. Vicini, J. Wong, and A. Martinez. Adaptive radiation therapy. *Physics in Medicine and Biology*, 42:123–132, 1997a.
- D. Yan, J. Wong, F. Vicini, J. Michalski, C. Pan, A. Frazier, E. Horwitz, and A. Martinez. Adaptive modification of treatment planning to minimize the deleterious effects of treatment setup errors. *International Journal of Radiation Oncology, Biology, Physics*, 38(1):197–206, 1997b.
- D. Yan, E. Ziaja, D. Jaffray, J. Wong, D. Brabbins, F. Vicini, and A. Martinez. The use of adaptive radiation therapy to reduce setup error: a prospective clinical study. *International Journal of Radiation Oncology, Biology, Physics*, 41(3):715–720, 1998.
- C.X. Yu, D.A. Jaffray, and J.W. Wong. The effects of intra-fraction organ motion on the delivery of dynamic intensity modulation. *Physics in Medicine and Biology*, 43:91–104, 1998.