

# **A framework for the robust delivery of respiratory motion adaptive arc radiotherapy**

by

**Eric Jessie Christiansen**

A thesis submitted to the Faculty of Graduate and Postdoctoral Affairs in  
partial fulfillment of the requirements for the degree of

**Doctor of Philosophy**

in

**Physics**

with specialization in

**Medical Physics**

Carleton University

Ottawa, Ontario, Canada

September 2020

© 2020

Eric Jessie Christiansen

# Abstract

Respiratory motion during radiotherapy can distort the highly conformal delivery of dose to lung tumours, compromising treatment outcomes and elevating dose to healthy tissues. A new optimization and delivery framework is proposed to track and adapt for the tumour motion during arc delivery in real time. An aperture library, consisting of a radiation field aperture for each combination of beam angle and tumour position, is optimized before delivery. The tracked tumour position would then be used to select the appropriate aperture from the library, ensuring that the tumour is under continuous irradiation throughout delivery. The instantaneous delivered dose therefore depends on the respiratory motion trajectory realized by the patient. Robust optimization is applied to optimize the aperture library, such that the expectation value of the dose, which considers all possible respiratory trajectories, gives the desired dose to the tumour with minimal dose to healthy tissues.

A Markov chain is used to model the respiratory motion, assigning each trajectory a probability. This model is trained using 5 min of patient motion data acquired in another study, and tested on the subsequent 5 min of data.

Statistical tests confirm that the model is capable of adequately estimating the distribution of the patient's respiratory motion from fractions of a second to the order of a minute into delivery.

The optimization framework is developed in matRad, an open-source treatment planning software. A novel dose calculation and optimization method is implemented, which accurately accounts for the continuous aperture motion undergone during arc delivery and tumour motion adaptation.

Plans optimized under this framework are compared to treatments planned with other respiratory motion compensation methods, with focus on the quality of plans and their robustness with respect to variations in the motion trajectory. Robustness is evaluated by simulating trajectories and the corresponding dose from the probabilistic model. The proposed framework is found to generate plans which have more or slightly less robustness to variations depending on the method being compared, with an increase in robustness resulting in more consistent treatment outcomes. Compared to those methods which are more robust, the proposed framework delivers a lower dose to healthy organs.

Great are the works of the Lord,  
studied by all who delight in them.

(Psalm 111:2, New Revised Standard Version)

# Acknowledgments

First I must thank my supervisors, Dr. Emily Heath and Dr. Tong Xu, for their guidance, encouragement, and support throughout my doctoral studies. This work would not have been possible without the many discussions, suggestions, and second (and third) looks at my code that we had together.

I wish to express my thanks to Dr. Mark Bangert, Dr. Niklas Wahl, and the rest of the matRad team at the German Cancer Research Center for their support and assistance with the matRad toolkit. I am especially grateful for their invitation to Heidelberg to present my work and share it with the global matRad community.

I would like to thank Dr. Jason Belec, at The Ottawa Hospital, for his assistance with radiotherapy treatment planning and contouring protocols. I am also grateful to Dr. Frédéric Tessier, at National Research Council Canada, for his support with the Monte Carlo code VMC++, and for the many late hours we spent together debugging the software.

I would like to acknowledge the many sources of funding I have received, so that I could focus on this work for the past five years: the Kiwanis Club of Ottawa Medical Foundation & Dr. Kanta Marwah Scholarship in Medical

Physics; the Ontario Graduate Scholarship; the John Lyndhurst Kingston Memorial Scholarship; the Carleton University Department of Physics Scholarship; and the Jasbir Singh Sidhu Memorial Book Fund. I would also like to thank the Carleton University Faculty of Graduate and Postdoctoral affairs and the Canadian Organization of Medical Physicists for their financial support in travel and attendance to scientific conferences.

Finally, I wish to thank my friends and family for their love and support, in particular to my parents and sister. I would not be here without them. To my wife, Kyra, I extend my endless gratitude for being everything I need, and more.

# Preface

This thesis represents the research work that I performed throughout my doctoral program at Carleton University. All work was supervised by Dr. Emily Heath and Dr. Tong Xu. All code was developed and implemented in the open-source treatment planning system matRad [1], written in the MATLAB<sup>®</sup> [2] programming language. The source code will be made available with the matRad distribution on GitHub at <https://github.com/e0404/matRad>.

Chapter 2 has been modified for inclusion in this thesis from a research article, published as:

E. Christiansen, E. Heath, and T. Xu, “Continuous aperture dose calculation and optimization for volumetric modulated arc therapy,” *Physics in Medicine & Biology*, vol. 63, no. 21, p. 21NT01, 2018,

© Institute of Physics and Engineering in Medicine, reproduced with permission, all rights reserved. The Version of Record can be found online at <https://doi.org/10.1088/1361-6560/aae65e>. Parts of this chapter were also presented at scientific conferences:

E. Christiansen, M. Bangert, T. Xu, and E. Heath, “Implementation of VMAT in matRad, an open source treatment planning toolkit,” in *Sixty-Third Annual Scientific Meeting of the Canadian Organization of Medical Physicists, Medical Physics*, vol. 44, no. 8, 2017, p. 4394,

E. Christiansen, E. Heath, and T. Xu, “A Continuous Aperture Approach for VMAT Plan Optimization,” in *Scientific Abstracts and Sessions, Medical Physics*, vol. 45, no. 6, 2018, p. e412.

Two separate articles representing the work in Chapters 3 and 4, respectively, are currently in preparation to be submitted to a scientific journal for publication. Parts of these chapters were presented at scientific conferences:

E. Christiansen, E. Heath, and T. Xu, “Markov chain modeling of respiratory motion,” in *Sixty-Fifth Annual Scientific Meeting of Canadian Organization of Medical Physicists, Medical Physics*, vol. 46, no. 11, 2019, p. 5395,

E. Christiansen, E. Heath, and T. Xu, “A framework for robust optimization of 4D-VMAT,” in *Sixty-Fifth Annual Scientific Meeting of Canadian Organization of Medical Physicists, Medical Physics*, vol. 46, no. 11, 2019, p. 5372.

# Table of Contents

Abstract	ii
Dedication	iv
Acknowledgments	v
Preface	vii
Table of Contents	ix
List of Tables	xiv
List of Figures	xvi
List of Abbreviations	xxi
<b>1 Introduction</b>	<b>1</b>
1.1 The use of radiation for the treatment of cancer . . . . .	1
1.2 External beam radiation therapy . . . . .	4
1.2.1 EBRT modalities . . . . .	4

1.2.2	EBRT planning . . . . .	9
1.2.3	EBRT optimization . . . . .	11
1.3	Intrafraction motion and its effects on treatment . . . . .	14
1.4	Motion mitigation methods in EBRT . . . . .	16
1.5	Proposed solution and thesis overview . . . . .	20

**2 Continuous aperture dose calculation and optimization for**

<b>VMAT</b>		<b>24</b>
2.1	Introduction . . . . .	24
2.2	Materials and methods . . . . .	28
2.2.1	Continuous aperture VMAT . . . . .	28
2.2.1.1	Fluence and dose calculation . . . . .	28
2.2.1.2	Optimization . . . . .	36
2.2.2	Discrete aperture VMAT approximation . . . . .	40
2.2.2.1	Fluence and dose calculation . . . . .	40
2.2.2.2	Optimization . . . . .	41
2.2.3	Dose angle interval study . . . . .	42
2.2.4	Treatment planning study . . . . .	44
2.3	Results . . . . .	45
2.3.1	Dose angle interval study . . . . .	45
2.3.2	Treatment planning study . . . . .	48
2.4	Discussion . . . . .	56
2.5	Conclusions . . . . .	59

<b>3</b>	<b>Markov chain respiratory motion model</b>	<b>61</b>
3.1	Introduction . . . . .	61
3.2	Summary of Markov chains . . . . .	65
3.2.1	The stationary distribution . . . . .	68
3.2.2	Maximum likelihood estimation of Markov chains . . . . .	69
3.2.3	Monte Carlo simulation of Markov chains . . . . .	72
3.3	Materials and methods . . . . .	73
3.3.1	Respiratory Markov chain model and parameters . . . . .	73
3.3.2	Respiratory data acquisition and processing . . . . .	75
3.3.3	Model evaluation and selection . . . . .	78
3.4	Results . . . . .	89
3.4.1	Fraction A . . . . .	89
3.4.2	Fraction B . . . . .	94
3.4.3	Fraction model intercomparisons . . . . .	102
3.5	Discussion . . . . .	105
3.6	Conclusions . . . . .	109
<b>4</b>	<b>4D-VMAT optimization and delivery framework</b>	<b>110</b>
4.1	Introduction . . . . .	110
4.2	Materials and methods . . . . .	114
4.2.1	4D-VMAT framework . . . . .	114
4.2.1.1	Dose calculation . . . . .	119
4.2.1.2	Probabilistic optimization . . . . .	128
4.2.2	Calculation of the 4D dose-influence matrix . . . . .	130

4.2.3	4D digital phantom . . . . .	131
4.2.4	4D optimization and parameters . . . . .	133
4.2.5	Validation and verification . . . . .	137
4.2.5.1	Dose expectation value and variance . . . . .	138
4.2.5.2	Dependence on dose calculation angular and number of dose calculation geometries . . . . .	140
4.2.5.3	Fluence calculation angle interval . . . . .	145
4.2.6	Comparison of 4D methods . . . . .	148
4.3	Results . . . . .	152
4.3.1	Validation and verification . . . . .	152
4.3.1.1	Expectation value and variance . . . . .	152
4.3.1.2	Dependence on dose calculation angular and number of dose calculation geometries . . . . .	152
4.3.1.3	Fluence calculation angle interval . . . . .	157
4.3.2	Comparison of 4D methods . . . . .	157
4.3.2.1	Patient 1: conventional fractionation . . . . .	157
4.3.2.2	Patient 1: SBRT . . . . .	164
4.3.2.3	Other patients . . . . .	170
4.3.2.4	Fraction model intercomparisons . . . . .	172
4.4	Discussion . . . . .	178
4.5	Conclusions . . . . .	187
<b>5</b>	<b>Conclusions and future work</b>	<b>188</b>
5.1	Future work . . . . .	193

List of References	198
Appendix A Derivation of bixel-based fluence in continuous aperture VMAT	217
Appendix B Derivation of the gradient in continuous aperture VMAT	229
Appendix C Discrete expression for $\bar{d}_i$ , the mean dose in 4D-VMAT	237
Appendix D Discrete expression for $\overline{d_i^2}$ , the mean squared dose in 4D-VMAT	244

# List of Tables

2.1	Minimum and maximum values for VMAT machine delivery constraints. . . . .	39
2.2	Intervals for the three gantry angle sets used in the VMAT planning study for each anatomical site considered. . . . .	45
2.3	TG119 case: results of the discrete aperture and continuous aperture VMAT optimized plan, as well as the dose from the former plan recalculated using continuous aperture fluence calculation. . . . .	50
2.4	Prostate case: results of the discrete aperture and continuous aperture VMAT optimized plan, as well as the dose from the former plan recalculated using continuous aperture fluence calculation. . . . .	51
2.5	Head & neck case: results of the discrete aperture and continuous aperture VMAT optimized plan, as well as the dose from the former plan recalculated using continuous aperture fluence calculation. . . . .	53

3.1	Trial Markov chain time interval factor values $r$ used to search for the optimal Markov chain respiratory motion model. . . .	80
4.1	Tumour parameters for the simulated patient phantoms investigated in the 4D optimization study. . . . .	134
4.2	Results of the dose distribution comparison for the dose angle interval study in lung. . . . .	154
4.3	Results of the dose distribution comparison for the study investigating the number of patient geometries. . . . .	156
4.4	CTV DVH objectives and thresholds for patient 1, conventionally fractionated, for each optimization method considered. . .	159
4.5	CTV DVH objectives and thresholds for patient 1, SBRT, for each optimization method considered. . . . .	166

# List of Figures

1.1	Photograph of a radiotherapy linac. . . . .	5
1.2	Close-up photograph of an MLC. . . . .	7
1.3	Dose distributions achievable with an MLC, using either 3D-CRT or IMRT. . . . .	8
1.4	Ideal and realistic DVHs for both an OAR and a target. . . .	12
2.1	Discretization of the patient and radiation field in matRad dose calculation and plan optimization. . . . .	30
2.2	Definitions of aperture and intensity variables for continuous aperture and discrete aperture VMAT optimization over an angular arc. . . . .	32
2.3	Definition of the user-specified angles $\theta_{\text{FMO}}$ , $\theta_{\text{DAO}}$ , and $\theta_{\text{dose}}$ . . .	37
2.4	TG119 case: comparison between dose recalculated using discrete aperture and continuous aperture calculation schemes at 4 angular intervals $\Delta\theta_{\text{dose}}$ . . . . .	46
2.5	Head & neck case: comparison between dose recalculated using discrete aperture and continuous aperture calculation schemes at 4 angular intervals $\Delta\theta_{\text{dose}}$ . . . . .	47

3.1	Transition diagram for a simple Markov chain. . . . .	66
3.2	Sample of the respiratory motion data used to train the model, consisting of the principal component of the motion of an optical surface marker. . . . .	77
3.3	Position binning and breathing phase segmentation of the filtered respiratory data. . . . .	79
3.4	The position histogram for time $\Delta t = 1$ s constructed from the testing respiratory data. . . . .	82
3.5	The position bin histograms actually observed and expected from the Markov chain respiratory motion model, for a particular time after the triggering position bin is observed. . . . .	85
3.6	Distribution of $\chi_{30s}^2$ estimated from a Monte Carlo simulation of the Markov chain respiratory motion model. . . . .	87
3.7	Distribution of $\chi_{\Sigma}^2$ estimated from a Monte Carlo simulation of the Markov chain respiratory motion model (fraction A). . . . .	90
3.8	Image of $p_{\Sigma}$ as a function of the number of bins recording the current elapsed time in the current breathing phase $N_T$ , and the Markov chain time interval factor $r$ . . . . .	91
3.9	Image of $p_{\Sigma}$ as a function of the number of position sub-bins per bin $N_x$ , and the Markov chain time interval factor $r$ . . . . .	92
3.10	Image of $p_{\Sigma}$ as a function of the number of bins recording the current elapsed time in the current breathing phase $N_T$ , and the number of velocity bins $N_v$ . . . . .	93

3.11	Graphs of $\chi_{\text{obs},\Delta t}^2$ and the corresponding $p_{\Delta t}$ , both as functions of $\Delta t$ (fraction A). . . . .	95
3.12	Images of the observed and expected position bin histograms, both as functions of $\Delta t$ (fraction A). . . . .	96
3.13	Observed and simulated sequence of respiratory motion as a function of time (fraction A). . . . .	97
3.14	Graphs of $\chi_{\text{obs},\Delta t}^2$ and the corresponding $p_{\Delta t}$ , both as functions of $\Delta t$ (fraction B). . . . .	99
3.15	Images of the observed and expected position bin histograms, both as functions of $\Delta t$ (fraction B). . . . .	100
3.16	Observed and simulated sequence of respiratory motion as a function of time (fraction B). . . . .	101
3.17	Graphs of $\chi_{\text{obs},\Delta t}^2$ and the corresponding $p_{\Delta t}$ , both as functions of $\Delta t$ (fraction B data tested against fraction A model). . . . .	103
3.18	Graphs of $\chi_{\text{obs},\Delta t}^2$ and the corresponding $p_{\Delta t}$ , both as functions of $\Delta t$ (fraction A data tested against fraction B model). . . . .	104
3.19	Stationary position bin distributions for the fraction A and B models. . . . .	106
4.1	Schematic of aperture library, consisting of apertures and dose rates defined at control points in the space of gantry angle $\theta$ and tumour position $\omega$ . . . . .	116

4.2	Dose deposited from a single bixel, irradiating two patient geometries: the patient at maximum exhale and maximum inhale. This latter dose is also shown deformed to apply to the anatomical voxels in the reference geometry of maximum exhale.	122
4.3	Allowed trajectory segments $\omega = \Omega^{k,l \rightarrow l'}(\Theta^{-1}(\theta))$ in the 4D-VMAT dose calculation model.	125
4.4	Schematic of the set of fluence calculation angles $\{\theta_{\text{flu}}\}$ in the space of gantry angle $\theta$ and tumour position bin $l$ .	127
4.5	Transverse slice of the XCAT-generated reference CT geometry for patient 1.	133
4.6	Trajectory of the embedded tumour in the XCAT phantom for patient 1.	135
4.7	Histogram of the normalized dose difference $\epsilon_i$ , scaled to display the probability density. Also shown for comparison is a graph of the probability density function of the standard normal distribution.	153
4.8	DVHs for patient 1, conventionally fractionated, optimized with CO and PO.	160
4.9	DVHs for patient 1, conventionally fractionated, optimized with STO and DAD.	162
4.10	DVHs for patient 1, conventionally fractionated, optimized with MKO and 3D-CTV.	163
4.11	DVHs for patient 1, conventionally fractionated, optimized with CO and PO for 1 fraction only.	165

4.12 DVHs for patient 1, SBRT, optimized with CO, for 5 fractions and 1 fraction only. . . . .	167
4.13 DVHs for patient 1, SBRT, optimized with STO and DAD. . .	169
4.14 DVHs for patient 1, SBRT, optimized with MKO and 3D-CTV.	171
4.15 DVHs for patient 2, conventionally fractionated, optimized with PO and DAD. . . . .	173
4.16 DVHs for patient 3, conventionally fractionated, optimized with CO and MKO. . . . .	174
4.17 DVHs for patient 4, conventionally fractionated, optimized with CO and MKO. . . . .	175
4.18 DVHs for patient 1, conventionally fractionated, optimized with CO, using different respiratory motion models for optimization and delivery simulation. . . . .	177
4.19 DVHs for patient 1, conventionally fractionated, optimized with STO and MKO, with delivery simulated using the fraction B respiratory model. . . . .	179
4.20 Value of the objective function $F(d)$ as a function of itera- tion number throughout the DAO step of plan optimization. Conventionally-fractionated CO and PO plans for patient 1 are shown. . . . .	180
C.1 Allowed trajectory segments $\omega = \Omega^{k,l \rightarrow l'}(\Theta^{-1}(\theta))$ in the 4D- VMAT dose calculation model. . . . .	240

# List of Abbreviations

The following list is in alphabetical order.

---

Abbreviation	Definition
3D-CRT	three-dimensional conformal radiation therapy
4D-CT	four-dimensional computed tomography
4D-VMAT	four-dimensional volumetric modulated arc therapy
ALERT-RA	aperture library-enabled real-time robust adaptation
ART	adaptive radiation therapy
bixel	beam element
CO	conventional optimization
CSU	combined statistical uncertainty
CT	computed tomography
CTV	clinical target volume
DAO	direct aperture optimization
DICOM	Digital Imaging and Communications in Medicine
DVF	deformation vector field
DVH	dose-volume histogram
EBRT	external beam radiation therapy

---

Table continued on following page.

---

Abbreviation	Definition
FMO	fluence map optimization
HU	Hounsfield unit
IMRT	intensity-modulated radiation therapy
ITV	internal target volume
linac	linear accelerator
MKO	motion kernel optimization
MLC	multi-leaf collimator
MLE	maximum likelihood estimation
MRI	magnetic resonance imaging
MU	monitor unit
OAR	organ at risk
PDVH	percentile dose-volume histogram
PO	probabilistic optimization
PTV	planning target volume
radiotherapy	radiation therapy
RCSU	relative combined statistical uncertainty
SBRT	stereotactic body radiation therapy
VMAT	volumetric modulated arc therapy
voxel	volume element

---

# Chapter 1

## Introduction

### 1.1 The use of radiation for the treatment of cancer

In Canada, approximately 226 000 new cases of cancer are expected to be diagnosed in 2020 [8]. Lung cancer is the most commonly diagnosed cancer among the general population, with 29 800 new cases expected in this same year. It is also one of the most fatal, with a five-year survival rate of 19% [9]. In 2007 32% of all patients with lung cancer in Ontario received radiation therapy (radiotherapy) for the treatment of cancer [10]. This percentage is also observed in the United States: 32% of patients with non-small cell lung cancer received radiotherapy in the time period 1999 to 2008 [11].

Ionizing radiation consists of particles such as photons, electrons, neutrons, and protons, with sufficient energy to either directly (photons, electrons, and protons) or indirectly (neutrons) ionize atoms, releasing energetic electrons.

When this occurs in a medium, these electrons travel throughout the medium, imparting their energy along the way to other atoms in turn; the high-energy charged particles, such as protons and electrons, can also directly deposit energy themselves. The amount of energy deposited to the medium in this way is quantified as the dose  $d$ , defined as the energy deposited in an infinitesimal volume divided by the mass of that volume. The SI unit of dose is the gray, Gy, defined as  $1 \text{ Gy} = 1 \text{ J/kg}$ . In living matter, the primary effect of ionizing radiation is to induce strand breaking in the chromosomal DNA of cells, which can lead to eventual cell death [12, 13].

A radiotherapy treatment consists of treating a patient with high-energy fields of ionizing radiation, used to deliver a highly-focused dose to a cancerous tumour with the aim of reducing its size or destroying it. Although the ultimate target of the radiation is the tumour, current clinical treatment methods unavoidably subject the surrounding normal tissues to varying amounts of radiation. The impact of radiation on living matter is generally more acute on cancerous cells than on healthy ones due to the rapidly proliferating nature of cancer cells [14]. High doses of radiation can, however, cause disorders to healthy tissues which depend on the organ subjected to radiation and on the particular amount, and can include skin reddening, reduced effectiveness of salivary glands, radiation pneumonitis, and induction of second cancers [15, 16]. The goal of radiotherapy is therefore to deliver a high curative dose to the tumour, while minimizing the dose to the surrounding organs at risk (OARs).

There are three main methods of delivering ionizing radiation to a patient: brachytherapy, nuclear medicine therapy, and external beam radiation therapy (EBRT). In both brachytherapy and nuclear medicine, the source of radiation consists of radioactive material internal to the patient, possibly placed on the skin, in the case of brachytherapy. Brachytherapy uses encapsulated or sealed sources, typically in either the form of permanently implanted seeds or as sources which are temporarily placed in or near the tumour via catheters in the patient. Nuclear medicine therapy, on the other hand, uses radioactive nuclei attached to drugs, which when injected into the patient are preferentially taken up by the tumour.

EBRT instead uses an external source of energetic particles to treat the patient; this source can consist of radioactive material, an X-ray tube, or, most commonly, a linear accelerator (linac). The linac accelerates electrons to MeV energies towards a metal target, typically tungsten, inside which the electrons undergo quick deceleration, losing their energy to bremsstrahlung photons. These photons, also with energy in the MeV range, are then used to irradiate the patient. Alternatively, the high-energy electrons may be used as a radiation source, in which case a scattering foil is used in place of the tungsten target to increase the size of the beam exiting from the linac. Photon-based EBRT is the most common form of radiotherapy for the treatment of lung cancer [17], which is the focus of this thesis, and will therefore be examined more closely in the following sections.

## 1.2 External beam radiation therapy

### 1.2.1 EBRT modalities

Modern EBRT typically employs a high-energy X-ray linac mounted on a gantry, used to treat a patient lying horizontally on a treatment couch, as seen in Figure 1.1. The gantry is able to rotate in the transverse plane of the patient, either during delivery or in between delivery of individual treatment beams. A multi-leaf collimator (MLC) is generally employed when using X-ray based linacs to modulate the fluence of the photon field exiting the linac. As shown in Figure 1.2, the MLC consists of two banks of parallel leaves, usually tungsten, which move independently in and out of the beam path.

Different modalities of EBRT make use of the MLC in various ways. For example, in three-dimensional conformal radiation therapy (3D-CRT), the individual leaves of the MLC are set to conform the beam aperture to the shape of the tumour at various gantry angles around the patient. Intensity-modulated radiation therapy (IMRT), on the other hand, uses multiple apertures at each gantry angle, none necessarily conformal, to modulate the fluence output by the linac. The modulation is done with the goal of delivering a dose distribution which is more conformal to the tumour, and with therefore less dose given to normal tissues; an illustration of this is given in Figure 1.3. The radiation can be delivered either while the leaves are stationary (static IMRT) or in motion (dynamic IMRT). Static IMRT is mechanically simpler,



Figure 1.1: Photograph of a radiotherapy linac. The linac itself is situated vertically, pointing down towards the couch. The entire gantry can rotate around the patient, who would lie horizontally on the adjustable couch during treatment. The photograph “[Linear Accelerator - Elekta Compact model](#)” by [User:Narenfox](#) is licensed under [CC BY-SA 4.0](#). No modifications to the original work have been made.

while dynamic IMRT delivers a larger total amount of fluence to the patient due to the modulation required, thereby increasing the volume of patient receiving a low dose. The continuous leaf motion used in dynamic IMRT, on the other hand, allows for a more conformal dose distribution [18].

A full treatment is delivered over the course of several days to weeks in discrete fractions; a possible fractionation scheme for lung tumours, for example, is to deliver 2 Gy per fraction for 33 daily fractions. This is done to allow the non-cancerous cells some time to recover from radiation-induced DNA damage [19]; during this time however the tumour cells also undergo repopulation. The fractionation scheme is therefore chosen to strike a balance between the damage, and subsequent repair and repopulation, of both normal tissues and tumour cells [20]. Stereotactic body radiation therapy (SBRT) is a method of radiotherapy which uses image guidance, such as radiography, to decrease the uncertainty of the internal position of the tumour. The dose to the patient may then be sufficiently localized to the tumour with minimal dose to OARs such that a treatment course with fewer fractions can be used [21]; for example, 60 Gy in 5 fractions. The higher dose per fraction has the advantage of increasing the biological effectiveness of the treatment while keeping a similarly low toxicity to normal tissues [22].

Volumetric modulated arc therapy [23, 24] (VMAT) is a form of dynamic IMRT in which the gantry continuously rotates around the patient while the beam intensity is modulated using dynamic MLC leaf motion. Because of the continuous delivery, gantry rotation, and leaf movement, VMAT plans tend

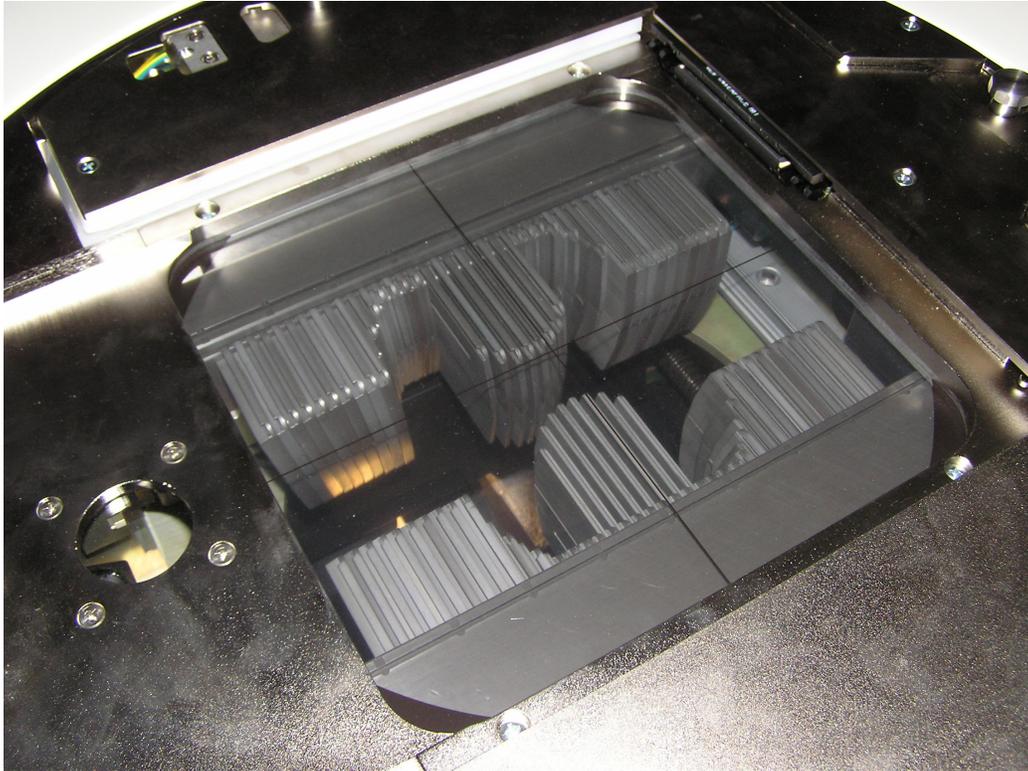
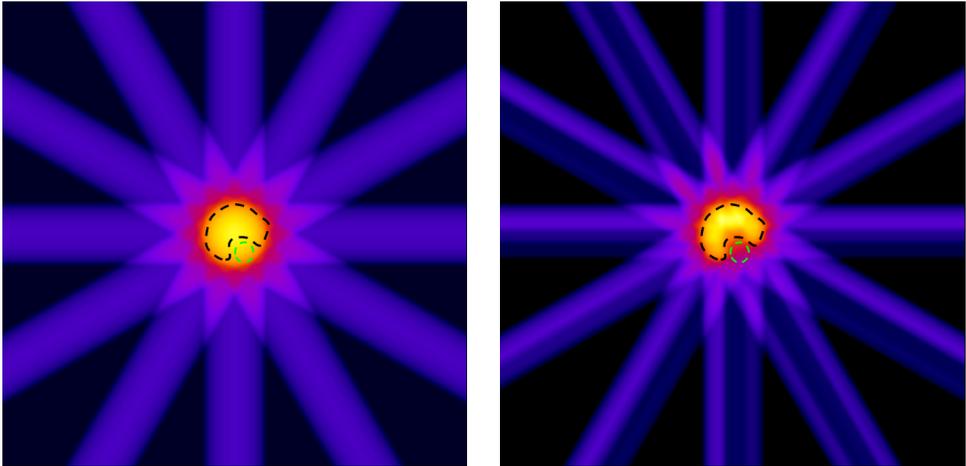


Figure 1.2: Close-up photograph of an MLC, which would be attached to the bottom end of the linac shown in Figure 1.1. Each leaf on the two opposed banks can move in and out of the beam path independently, thereby modulating the intensity of the photon field. The photograph “[Multi leaf collimator](#)” by [User:Egg](#) is licensed under [CC BY-SA 2.5](#). No modifications to the original work have been made.



(a) 3D-CRT dose distribution.

(b) IMRT dose distribution.

Figure 1.3: Dose distributions achievable with an MLC, using either 3D-CRT or IMRT. A horseshoe-shaped tumour is outlined in black, with an OAR inside the bend of the tumour outlined in green. A total of 12 beams are delivered in each case, at equal angular intervals. In 3D-CRT, the MLC aperture conforms to the shape of the tumour at each angle, resulting in a very large dose to the OAR. In IMRT, the MLC modulates the photon field at the different gantry angles, diminishing the fluence, and resulting dose, to the OAR. Images courtesy of Dr. Tong Xu.

to have shorter treatment times compared to IMRT plans with similar plan quality [25].

### 1.2.2 EBRT planning

Whichever radiotherapy treatment modality has been selected, many of the steps taken to plan the delivery remain the same. A computed tomography (CT) image of the target and surrounding tissues is commonly acquired for the express purpose of treatment planning and dose calculation. Following this, the target and relevant OARs are delineated on the CT image, to record the dose delivered to each volume. Four different volumes are used to define the target to be treated, to account for microscopic spread of the tumour, as well as position variation and uncertainty. The first is the gross tumour volume (GTV), which is defined as the extent of the tumour which can be identified using, for example, diagnostic imaging techniques or palpation. The clinical target volume (CTV) includes the GTV, as well as a margin to include microscopic disease spread which cannot be imaged or palpated. Third, the internal target volume (ITV) accounts for internal motion of the CTV. Finally, the planning target volume (PTV) allows for setup uncertainties while positioning the patient for delivery [26]. In this thesis, these uncertainties are not accounted for, therefore only the CTV and ITV are considered.

The treatment objectives are then determined, which include a desired dose to be delivered to the PTV as well as threshold doses to OARs set to a specific level to prevent radiation-induced complications [15, 16]. A plan is

constructed to fulfill these objectives, consisting of the set of MLC apertures and linac output to be delivered at the chosen gantry angles. For the simpler treatment modalities of 3D-CRT and static IMRT, these can be selected by hand by an experienced treatment planner. Dynamic IMRT and VMAT, however, have many more variables which must be individually set to yield an optimal plan. The method of computer-aided inverse plan optimization, which will be discussed in Section 1.2.3, is therefore applied to plan the delivery of these more complex modalities; in practice, static IMRT plans are also optimized in this way.

The finished treatment plan contains a series of MLC apertures and linac output levels at different beam delivery angles. The monitor unit (MU) is used as the unit of linac output; under reference conditions (a  $10\text{ cm} \times 10\text{ cm}$  open field is commonly used, with a source to measuring distance of 100 cm), 100 MU is defined to deliver a dose of 1 Gy to a specific depth of water. Many dose calculation methods exist to calculate the dose delivered throughout the patient, with varying degrees of accuracy and efficiency. Monte Carlo algorithms, which simulate the transport and energy deposition of energetic particles throughout the patient, have been shown to be the most accurate dose calculation method for lung cancer [27].

Before treatment, the plan is evaluated to assess if the desired prescription dose will be delivered, as well as the extent to which dose to OARs is minimized. The dose-volume histogram (DVH) is a tool used to evaluate the quality of treatment plans: for each volume which has been delineated, the DVH

indicates the percentage of that volume which receives a given dose or greater. Example DVHs for both a target and OAR are shown in Figure 1.4. For OARs, a DVH which shows a lower volume at any dose is indicative of a plan with greater quality; in practice, threshold dose-volume points which are informed by clinical studies are used to prevent specific complications [15, 16]. The ideal DVH for the target, on the other hand, is one in which as large a volume as possible receives the prescription dose, with a sharp decrease in the volume receiving dose higher than the prescription.

Once a plan has been constructed or optimized with acceptable quality, the plan is delivered to a phantom which simulates the patient, and the actual delivered dose is measured. This is done to verify that the planned dose distribution is an accurate representation of the delivered dose. After this quality assurance step is complete, delivery of the treatment on the patient can commence.

### 1.2.3 EBRT optimization

In inverse plan optimization, an objective function  $F(d)$  is defined such that a lower value of the objective function indicates a plan with greater quality. Typically each volume of interest, such as the delineated OARs and the target, contribute one or more terms to the objective function. Each term should increase as the dose inside the respective volume deviates from the desired dose. This can be accomplished in the target, for example, using a squared

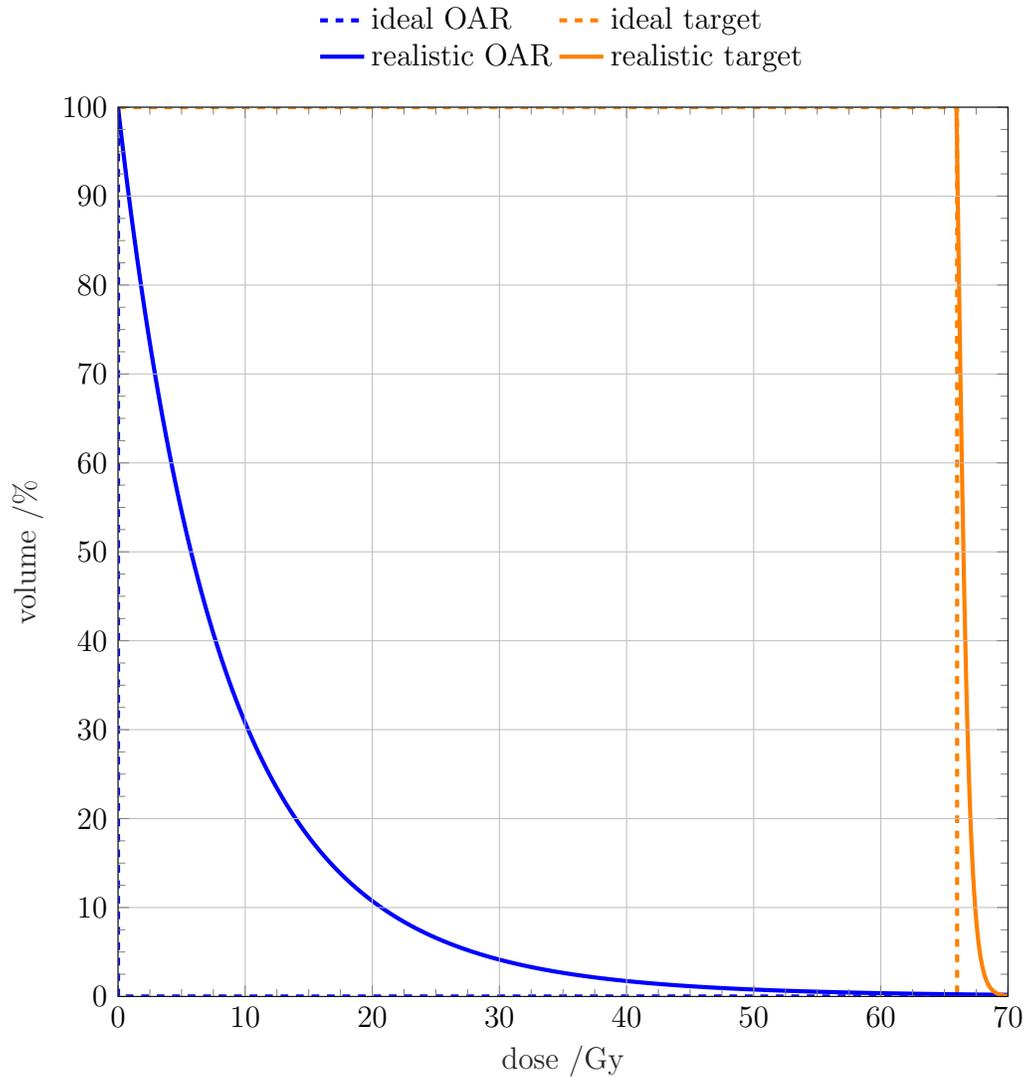


Figure 1.4: Ideal and realistic DVHs for both a target and an OAR. Ideally the 100 % of the target volume should receive exactly the prescription dose (in this case, 66 Gy); in practice there will be some degree of dose inhomogeneity. On the other hand, the OAR would ideally receive a dose of 0 Gy; in reality, OARs in or near the path of the beam will inevitably receive some amount of dose.

deviation term:

$$F_T = w_T \frac{1}{N_T} \sum_{i \in T} (d_i - d_T)^2. \quad (1.1)$$

The sum here is over all volume elements (voxels)  $i$  inside the target  $T$ .  $d_i$  then is the dose deposited in voxel  $i$ , which is compared to  $d_T$ , the desired dose to the target (usually taken to be the prescription dose). The sum is divided by the number of voxels inside the target  $N_T$  to calculate a mean square deviation, which is then multiplied by a weighting or importance factor  $w_T$ .  $f_T$  is seen to decrease as the dose to voxels inside the target comes closer to  $d_T$ . The form of objective function in Equation (1.1) is not necessarily appropriate for OARs, since for these volumes the objective function should not be penalized if the dose is lower than a particular threshold. A Heaviside step function  $H(x)$ , defined to be 0 when  $x < 0$  and 1 when  $x \geq 0$ , is commonly used in the definition of the objective function term for OARs, thus:

$$F_O = w_O \frac{1}{N_O} \sum_{i \in O} H(d_i - d_O) (d_i - d_O)^2. \quad (1.2)$$

This ensures that for OAR  $O$ , only voxels which have dose larger than the threshold dose  $d_O$  are counted when calculating  $F_O$ . The final objective function then is taken to be the sum of all volume-specific terms.

Once the form and parameters of each term are set by the treatment planner, optimization methods such as gradient descent or simulated annealing can be applied to minimize the objective function [28]. The final result of the optimization is a set of MLC apertures and MU values which delivers

the prescription dose to the target, and minimal dose to OARs. Machine constraints, such as a minimum and maximum MLC leaf speed, gantry rotation speed, and MU rate, can be incorporated into the optimizer such that the final plan is able to be delivered.

### **1.3 Intrafraction motion and its effects on treatment**

Organ and tumour motion which occurs during a radiotherapy treatment fraction, known as intrafraction motion, can come from many sources. These include those which occur on longer time scales, such as bladder filling, peristalsis, and tumour drift, and also sources which act at a faster frequency, such as respiratory and cardiac motion. The focus of this thesis is the problem of motion due to respiration, specifically as it applies to lung tumours.

The main muscles responsible for respiration are the diaphragm and those in the chest wall. During inspiration, the diaphragm is contracted and descends inferiorly and anteriorly, while the chest wall expands superiorly and anteriorly. In this way the volume of the thoracic cavity is increased, thereby pulling air into the lungs due to the difference in pressure from outside the body. For relaxed breathing, exhalation is passive: the lung and chest walls are elastic, and therefore return to their pre-inhalation positions at the end of exhalation [29]. This entire process occurs with a period on the order of 3 s to 5 s in relaxed breathing.

The lung, lung tumour, as well as adjacent organs such as the pancreas, liver, breast, prostate, and kidneys, are all subject to varying degrees of motion due to the respiratory process. Lung tumours in particular can exhibit motion with a range from a few millimetres to as high as 34 mm [30]. The motion in the superior-inferior direction is generally the largest, but there does exist motion in the anterior-posterior, as well as the left-right directions. Further generalized statements on patterns in the tumour motion due to respiration is not possible due to the observed large variation in this motion from one patient to another [29]. This is due to factors such as the type of breathing (quiet or deep, for example), as well as on the position of the tumour within the lung.

The use of the ITV, defined in Section 1.2 as a volume which encloses the tumour throughout its range of motion, constitutes a traditional method to compensate for the respiratory motion in radiotherapy: no matter the breathing pattern exhibited by the patient, the tumour will receive the prescription dose. Further solutions will be discussed in Section 1.4.

If however no allowance or compensation is made for respiratory motion during the treatment planning process, the conformal dose to the target which was achieved during treatment planning will be blurred by the motion of the tumour [31]. Specifically, the tumour would pass in and out of the high-dose region, thereby lowering the dose to the target, and increasing the dose to the healthy lung volume. Thus the quality of the actual delivered plan is diminished, with the potential for compromised treatment outcomes

and an increased chance of radiation-induced complications. For dynamic IMRT modalities in which the instantaneous high- and low-dose regions are continuously changing, the interplay between the tumour and MLC motion can further intensify this effect [32].

## 1.4 Motion mitigation methods in EBRT

There exist many ways to account or compensate for respiratory motion which have been proposed or are in current clinical use. The conventional approach, described in Section 1.2, is to deliver the prescription dose to the ITV, thereby accounting for the uncertainty in the tumour location due to respiratory motion. The ITV can be generated using a four-dimensional CT (4D-CT) image of the patient, which consists of a set three-dimensional CT image taken at various points throughout the respiratory cycle. The 4D-CT therefore captures the respiratory motion of the tumour as well as the various organs which surround the lungs. The ITV is unnecessarily large, since it assumes that the tumour will occupy each of its possible locations 100% of the time; in reality, part of this volume will inevitably be composed of normal tissue. This can be mitigated by instead using a probabilistic margin derived to ensure the CTV receives the desired coverage, within specified random and systematic errors (for example, due to respiratory motion) [33]. This has been shown to reduce the amount of normal lung tissue receiving a high dose, as compared to treatment using an ITV [34]. Another solution,

which in this thesis is referred to as motion kernel optimization (MKO), is possible if the probability distribution of the target position is known. In this method the CTV is effectively convolved with a motion kernel, weighted by this probability distribution, to represent the average motion of the target throughout delivery. Plan optimization is then performed with the goal of delivering the prescription dose to this blurred target volume, resulting in a lower dose to OARs as compared to ITV optimization [35].

With respiratory gating methods, radiation is delivered to the patient only during a specific portion of the breathing cycle. This has the advantage of reducing the amount of motion observed while the beam is on, although some residual motion may still persist within the gating window [36]. The smaller the gating window used, the less residual motion, however this comes at a cost of increased treatment times. A real-time tumour tracking system is required to perform gating. Systems which track an external surrogate signal, such as markers placed on the patient's abdomen as used by the Real-time Position Management system (Varian Medical Systems, Palo Alto, United States of America), are generally non-invasive. A surrogate signal can however become desynchronized with the actual tumour motion, resulting in a discrepancy between when the tumour position lies within the gating window and the delivery of radiation [37]. Fiducial markers implanted inside or near to the tumour, which can then be viewed using radiography, can instead be used to improve the accuracy of gated delivery [38]. This however requires an invasive implantation procedure which may not be suitable for all patients. A

system called PeTrack which uses positron-emitting markers implanted in the tumour has also been proposed for real-time tracking in radiotherapy [39], with successful usage in gating [40]. Magnetic resonance imaging (MRI), having excellent soft-tissue contrast and requiring no implantation, has recently been applied in real-time tumour and anatomy tracking for use in gating [41].

Breath hold techniques are a form of respiratory gating in which the patient is instructed to hold their breath before the treatment beam is to be delivered. The breath may be held at the end of a deep inspiration [42] or at some other point during the cycle; the total duration of the breath hold is patient-dependent, and may last from 10s to 30s [29]. The hold may be performed voluntarily, in which case verbal coaching is often used in conjunction with a respiratory monitor system. Alternatively the hold can be actively aided with the use of a device to maintain a reproducible breath hold [43]. Like respiratory gating, this technique has the disadvantage of longer treatment times. If an active-breathing control device is not used, a lack of breath-hold stability and reproducibility can prevent the tumour from reliably falling within the required treatment delivery window.

Compression and immobilization techniques have been used to reduce the range of respiratory motion during radiotherapy, while still allowing limited breathing [44]. The compression is achieved using a plate attached to a body frame which is pressed against the patient's abdomen. The range of the motion has been observed to reduce from 8 mm to 20 mm without compression down to 2 mm to 11 mm with compression [45]. This technique does however

cause patient discomfort throughout the treatment delivery. Furthermore, the compression is not entirely reproducible from one fraction to another, necessitating the use of pre-treatment verification imaging [29].

There also exist several methods which restrict the delivered dose to the moving target throughout its range of motion by continuously repositioning the MLC aperture; a real-time tumour tracking system is required for these methods. This may take the form of implanted tumour fiducial [46] or positron-emitting [39] markers, MRI [47], or the use of a surrogate signal which is correlated to the real-time tumour position. As with respiratory gating, this correlation may vary throughout the treatment, although it can be periodically updated using radiography images of the tumour [48, 49]. This motion information can be used to simply deform the MLC aperture to follow the tumour motion, with no other changes [50, 51].

Alternatively, with real-time tracking comes the potential to exploit the temporal degree of freedom in optimization. With such techniques, called 4D optimization, the optimization algorithm can, for example, selectively increase the amount of dose delivered in instances for which the target is farther away from a dose-sensitive OAR than in others [35]. This has been applied to VMAT treatment planning [52], with a decreased dose to OARs as compared to a VMAT plan delivered using an ITV. In this four-dimensional VMAT (4D-VMAT) method, however, the patient's breathing motion during delivery is assumed to be known exactly when performing treatment plan optimization, and to remain constant between fractions; in other words, the

relationship between the angle of the continuously rotating gantry and the tumour position remains fixed. There is no accounting or accommodation for if the patient’s breathing varies, for example, by becoming faster, slower, deeper, or shallower than anticipated. This could have a significant effect on the resulting plan quality, since an MLC aperture optimized to direct radiation to the tumour at a particular gantry angle and position may miss its target entirely if the motion changes from what was planned.

## 1.5 Proposed solution and thesis overview

In this thesis, the aperture library-enabled real-time robust adaptation (ALERT-RA) framework is introduced. As its name suggests, the goal of this framework is to optimize and deliver 4D-VMAT with real-time respiratory tracking and corresponding adaptation. Possible variations in the realized tumour trajectory are explicitly taken into account by optimizing an aperture library, consisting of an MLC aperture and MU linac output for each combination of gantry angle and tumour position. Throughout the course of delivery, the real-time tumour position as well as the current gantry angle would be used to select the appropriate aperture to be formed by the MLC, ensuring that the tumour is continuously irradiated by the linac. Any of the tracking techniques mentioned previously could in principle be used, although more accurate methods, such as MRI or PeTrack, would be preferred.

The 4D-VMAT plan optimization and delivery framework was developed

and implemented in the open-source treatment planning system matRad [1], written in the MATLAB® [2] programming language. This treatment planning system is capable of optimization dynamic IMRT and proton delivery, however, VMAT treatment planning and optimization had previously not been supported. The implementation of a non-motion adaptive VMAT optimization method is described in Chapter 2. A novel dose calculation and optimization method for VMAT which accurately accounts for the continuous MLC leaf travel during delivery is also introduced, and compared to traditional discrete methods of calculation [3].

Robust optimization is a technique which accounts for the random nature of variables which have an effect on the delivery of dose to the patient, with the aim of reproducibly delivering the prescription dose to the target and the desired low doses to OARs [53]. Examples of these random variables include setup errors, interfraction motion, and intrafraction motion such as respiration. Robust optimization was therefore applied to optimize the aperture library required for the ALERT-RA framework, with the aim of yielding high-quality 4D-VMAT plans which are robust to changes in the respiratory motion. In this case the observed tumour trajectory, having an effect on the sequence of delivered MLC apertures and subsequent dose, is taken to be the random variable.

Robust optimization is often performed by optimizing either the plan quality for the random variable with the worst outcome [54], or the expectation

value of the plan quality [55]. In this thesis, the latter approach is taken:

$$\min \bar{F},$$

$$\bar{F} = \int \Pr [\Omega] F (d [\Omega]) \mathcal{D}\Omega. \quad (1.3)$$

Here,  $d [\Omega]$  is the dose due to a particular respiratory trajectory  $\Omega$ , while  $\Pr [\Omega]$  is the probability to observe that trajectory. A necessary component of this framework therefore is a patient-specific probabilistic model of the tumour motion which can calculate the probability of each trajectory. A method to construct a Markov chain respiratory motion model which can estimate these probabilities is introduced in Chapter 3. The model is trained on real patient respiratory data acquired during radiotherapy delivery. The accuracy of the model-calculated probability distribution of the tumour position at various times during delivery is evaluated by statistical comparison with data from the same patient at a later time during delivery.

The design, implementation, validation, and evaluation of this 4D-VMAT plan optimization and delivery framework is done in Chapter 4. The resulting plans are compared to plans optimized using no motion compensation, MKO, as well as two motion tracking techniques: simple deformation of the MLC aperture to the motion of the tumour [51], and 4D-VMAT optimized under the assumption of a single fixed trajectory [52]. Particular attention is paid to the robustness of the plan quality, including adequate coverage of the target by the prescription dose, as well as maintaining a low dose to OARs.

Finally, concluding remarks are given in Chapter 5, with discussion on possible avenues of future work after this study.

## Chapter 2

# Continuous aperture dose calculation and optimization for VMAT

In this chapter a summary of the implementation of non-motion adaptive VMAT in matRad is given. A new dose calculation and optimization method is also introduced and developed, which accurately take into account the continuous motion of the leaves forming the MLC apertures. This optimization method is then compared against traditional discrete aperture calculation methods.

### 2.1 Introduction

As mentioned in Section 1.2, VMAT [23, 56, 24, 57, 58] is a form of EBRT in which the linac gantry rotates continuously while modulating the beam intensity using dynamic MLC leaf motion. Because of the continuous delivery, gantry rotation, and leaf movement, VMAT plans tend to have shorter

treatment times compared to IMRT plans with similar plan quality [25]. A VMAT treatment plan generally consists of defining MLC apertures and delivered MU at a series of angles around the patient. The machine will then transition smoothly between these states as it delivers the plan. A treatment plan which results in a desirable dose distribution (i.e., prescription dose in the target, low dose in the OARs) is obtained by optimizing an objective function which penalizes undesired aspects in the dose distribution. The machine states at each gantry angle are variables in this optimization. A physically deliverable plan is ensured by performing a constrained optimization, using the MLC leaf speed, gantry rotation speed, and MU rate as constraints.

A variety of approaches to the VMAT plan optimization problem exist, including direct aperture optimization (DAO) using either deterministic [57, 58] or stochastic [56, 24] methods. These approaches differ mainly in how the optimization problem is initialized, and in how the variables are updated on each iteration of the optimization algorithm. The deterministic methods both use a three-step approach, consisting of: fluence map optimization (FMO), leaf sequencing, followed by DAO based on gradient descent. This sequence of steps will be fully explained in Section 2.2.1.2. The stochastic methods, on the other hand, initialize the optimization problem using the aperture shapes conforming to the beam's eye view of the target at each angle. During optimization, the aperture shapes and corresponding MU are randomly varied, with changes kept if the objective function is reduced. The stochastic optimization method proposed by Otto [24] starts with a coarse

sampling of the gantry angles around the patient; as optimization continues, new samples are progressively included.

Although VMAT delivery features continuous gantry rotation and leaf motion, most treatment planning systems perform dose calculation under an assumption of discrete apertures changing instantaneously from one discrete angle to the next [25]. The calculated dose will differ from the delivered dose due to two sources of error. The first source of error is that the dose is calculated at discrete angles, whereas in the actual delivery fluence is delivered continuously as a function of gantry angle. This will be referred to in this thesis as the discrete dose angle approximation. Depending on the angular interval and the density and medium variations present in the patient geometry, this can lead to errors in the calculated dose.

The second source of error is that the fluence delivered as the MLC leaves sweep between two apertures is not correctly calculated; instead, the fluence calculation is treated as though the entire fluence at a given angle is delivered by a single static aperture. This is referred to in this thesis as the discrete aperture approximation. This can lead to errors in the final patient dose when apertures at two adjacent angles are significantly different from one another. The discrete aperture approximation will then underestimate the dose in the volume of the patient corresponding to the transition region between the two aperture shapes. This is especially problematic if there is an OAR in this volume, as the delivered dose in the OAR will be higher than what was planned.

Many authors have proposed various solutions to the problems described above. For example, VMAT plans may be recalculated using aperture interpolation at a finer gantry angle resolution after [59] or during [57, 58, 60] optimization. The former approach may reveal a reduction in plan quality compared to what was originally optimized, which is problematic for treatment planners desiring to achieve a given plan quality, while the latter approach increases optimization time and memory requirements. Alternatively, the fluence delivered by the continuously moving leaves may be calculated analytically, which will be referred to in this thesis as the continuous aperture fluence calculation. Some authors have incorporated the continuous aperture fluence calculation into different steps in the VMAT optimization process, either as a post-optimization step [61] or during a more restrictive sliding window VMAT optimization [62].

In this study, a continuous aperture VMAT optimization algorithm was implemented following the three-step approach. The DAO is done with the continuous aperture fluence calculation, which correctly calculates the dose delivered as leaves sweep between two angles. The benefit of this algorithm over simply performing a post-optimization calculation is determined by comparing final optimized plans from both discrete aperture and continuous aperture VMAT optimizations, and then recalculating the dose for the discrete aperture VMAT optimized plan using continuous aperture fluence calculation.

Note that the continuous aperture fluence calculation does not affect errors introduced by the use of discrete dose angles in the dose calculation.

Therefore the two sources of error present in discrete aperture VMAT discussed above, those due to the discrete dose angle approximation and the discrete aperture fluence approximation, are investigated separately using a series of recalculation schemes. The results of this study will help determine the angular interval at which the errors introduced by the discrete dose angle approximation become insignificant.

## 2.2 Materials and methods

### 2.2.1 Continuous aperture VMAT

To accurately calculate the dose delivered during VMAT, taking into account the continuous MLC leaf motion and gantry rotation, an analytical calculation method for continuous aperture VMAT dose calculation and optimization was developed, described in the following section.

#### 2.2.1.1 Fluence and dose calculation

All plan optimization and dose calculation was implemented in matRad [1], an open-source treatment planning system written in MATLAB<sup>®</sup> [2]. matRad is capable of importing patient imaging (Digital Imaging and Communications in Medicine, DICOM) data, as well as performing dose calculation and plan optimization for photons, protons, and carbon ions. matRad uses the dose-influence matrix to calculate the delivered dose: the patient is subdivided into voxels  $\{i\}$ , and the radiation fields at several discrete angles around the

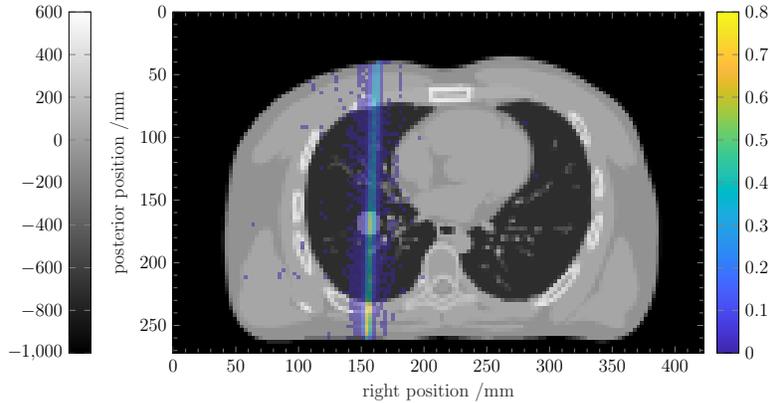
patient  $\{\theta_k\}$  are subdivided into beam elements (bixels)  $\{j\}$ , as illustrated in Figure 2.1. The dose to voxel  $i$  is calculated thus:

$$d_i = \sum_{j,k} D_{ij\theta_k} \phi_{j\theta_k}. \quad (2.1)$$

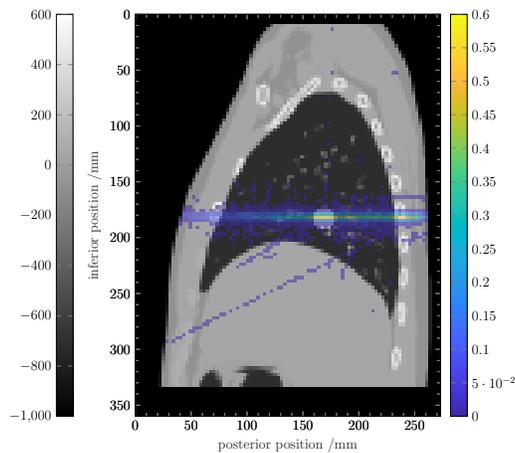
Here,  $D_{ij\theta_k}$  is the dose contribution to voxel  $i$  from bixel  $j$  at angle  $\theta_k$  at unit bixel fluence, and  $\phi_{j\theta_k}$  is the fluence at bixel  $j$  at angle  $\theta_k$ . This dose depends on the patient geometry, including the density and material composition of each voxel, as well as the characteristics of the incident radiation beam. The dose-influence matrix  $D_{ij\theta_k}$  is calculated in matRad using a singular value decomposed pencil beam algorithm [63], and all bixels are assumed to deliver dose independently from the others. A 5 mm  $\times$  5 mm bixel size was used for the dose calculations in this study, corresponding to the leaf width in the Elekta Agility™ MLC (Stockholm, Sweden).

Equation (2.1) calculates dose correctly if the fluence is delivered at discrete angles around the patient, for example in IMRT delivery. In a VMAT plan, however, fluence is delivered while the gantry rotates continuously around the patient, and the MLC leaves move and dynamically modulate the field. In this continuous delivery, the sum over discrete angles  $\theta_k$  is converted to an integral over continuous angles  $\theta$ :

$$d_i = \int \sum_j D_{ij}(\theta) \frac{d\phi_j}{d\theta}(\theta) d\theta. \quad (2.2)$$



(a) Transverse slice.



(b) Sagittal slice.

Figure 2.1: Discretization of the patient and radiation field in matRad dose calculation and plan optimization. Transverse (Figure 2.1a) and sagittal (Figure 2.1b) slices of a CT show the patient subdivided into voxels  $\{i\}$ . The CT images are in greyscale, and expressed in Hounsfield units (HU). The dose (colour, in units of Gy per 100 MU) due to a single bixel  $j$  from beam angle  $\theta_k$  (incident from the posterior side) is shown superimposed on the CT images. The CT images were generated using XCAT [64], and dose was calculated using VMC++ [65, 66] (see Sections 4.2.3 and 4.2.2 respectively for more information).

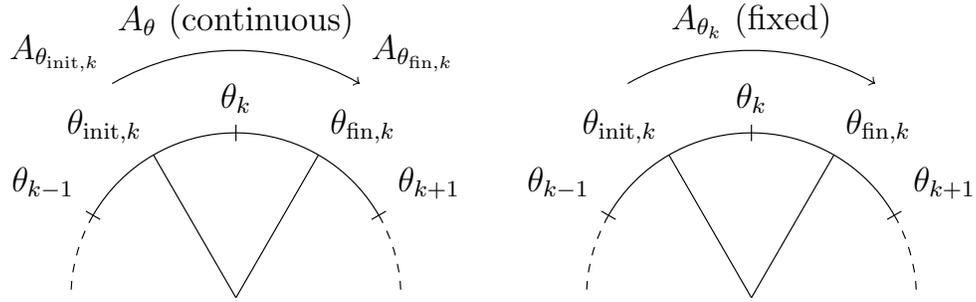
Here, the  $D_{ij}(\theta)$  is the dose-influence matrix for a continuous gantry angle  $\theta$ , and  $\frac{d\phi_j}{d\theta}(\theta)$  is the rate of change of the accumulated bixel fluence with respect to  $\theta$ .

If the  $D_{ij}(\theta)$  do not vary significantly over a certain angular interval  $\Delta\theta$ , the integral in Equation (2.2) can be approximated by the discrete sum in Equation (2.1), with the elements in the discrete set  $\{\theta_k\}$  separated by the same interval  $\Delta\theta$ ; this is the discrete dose angle approximation. The bixel fluence at a particular angle  $\theta_k$  is given by the integral of the differential fluence along an arc section centred at  $\theta_k$  subtending an angle  $\Delta\theta$ , from  $\theta_{\text{init},k} = \theta_k - \Delta\theta/2$ ,  $\theta_{\text{fin},k} = \theta_k + \Delta\theta/2$ :

$$\phi_{j\theta_k} = \int_{\theta_{\text{init},k}}^{\theta_{\text{fin},k}} \frac{d\phi_j}{d\theta}(\theta) d\theta. \quad (2.3)$$

In the continuous aperture fluence calculation method, it is assumed that the linac intensity rate  $\dot{\psi}_{\theta_k}$  (directly proportional to the MU rate delivered by the linac) and the speed of each leaf are constant over the arc section from  $\theta_{\text{init},k}$  to  $\theta_{\text{fin},k}$  (the speed of each leaf may be different).

The fluence of a particular bixel  $j$  accumulated during the arc section around  $\theta_k$  depends on the motion of the pair (left and right) of MLC leaves in the row to which the bixel belongs, and also on the linac intensity delivered during this arc section. This is illustrated by Figure 2.2a. The calculation of the bixel fluence can be simplified by splitting it into two components (the bixel subscript  $j$  and angle subscript  $\theta_k$  are omitted for simplicity):  $\phi_L$ ,



(a) Continuous aperture VMAT.

(b) Discrete aperture VMAT.

Figure 2.2: Definitions of aperture and intensity variables for continuous aperture (Figure 2.2a) and discrete aperture (Figure 2.2b) VMAT optimization over an angular arc  $[\theta_{\text{init},k}, \theta_{\text{fin},k}]$ . In continuous aperture VMAT,  $A_{\theta_{\text{init},k}}$  ( $A_{\theta_{\text{fin},k}}$ ) refers to the aperture at angle  $\theta_{\text{init},k}$  ( $\theta_{\text{fin},k}$ ); the aperture  $A_\theta$  varies continuously between these two angles and is continuous at the limits between adjacent arcs  $[\theta_{\text{init},k}, \theta_{\text{fin},k}]$ . In discrete aperture VMAT, the symbol  $A_{\theta_k}$  refers to the aperture at angle  $\theta_k$ , the collection of left and right MLC leaf position variables, which remain fixed over the arc  $[\theta_{\text{init},k}, \theta_{\text{fin},k}]$  and change instantaneously from one arc to the next. The linac intensity rate  $\dot{\psi}_{\theta_k}$  is constant over the interval for both continuous aperture and discrete aperture VMAT, leading to the variable  $\psi_{\theta_k}$ , the total linac intensity delivered over the arc  $[\theta_{\text{init},k}, \theta_{\text{fin},k}]$ .

defined to be the component of the bixel fluence released by the left leaf in the absence of the right leaf; and  $\phi_R$ , defined to be the component of the bixel fluence blocked by the right leaf in the absence of the left leaf. Then  $\phi = \phi_L - \phi_R$ , with:

$$\phi_L = \int_{\theta_{\text{init}}}^{\theta_{\text{fin}}} \frac{d\phi_L}{d\theta}(\theta) d\theta, \quad (2.4a)$$

$$\phi_R = \int_{\theta_{\text{init}}}^{\theta_{\text{fin}}} \frac{d\phi_R}{d\theta}(\theta) d\theta. \quad (2.4b)$$

Assuming without loss of generality that both left and right leaves travel from left to right at constant speed (the speeds of the two leaves may be different), let the initial and final positions of the left and right leaves be  $L_{\text{init}}$ ,  $L_{\text{fin}}$ ,  $R_{\text{init}}$ , and  $R_{\text{fin}}$ .

Begin with the evaluation of  $\phi_L$ , which will only depend on the motion of the left leaf.

Changing the coordinates to  $x$ , the position of the left leaf, the integral becomes:

$$\phi_L = \int_{L_{\text{init}}}^{L_{\text{fin}}} \frac{d\phi_L}{dx}(x) dx. \quad (2.5)$$

$\frac{d\phi_L}{dx}(x)$  can be related to  $\frac{d\phi_L}{dt}(x)$  using the chain rule:

$$\begin{aligned} \frac{d\phi_L}{dx}(x) &= \frac{d\phi_L}{dt}(x) \frac{dt}{dx}(x) \\ &= \frac{1}{v_L} \frac{d\phi_L}{dt}(x) \\ &= \frac{\Delta t}{L_{\text{fin}} - L_{\text{init}}} \frac{d\phi_L}{dt}(x), \end{aligned} \quad (2.6)$$

where  $v_L$  is the constant speed of the left leaf in the arc section, and  $\Delta t$  is the time taken for the gantry to rotate through the arc section.

$\frac{d\phi_L}{dt}(x)$  is the rate of fluence released to the bixel by the left leaf when the left leaf is at position  $x$ . At any instant in time, this is given by the constant linac intensity rate  $\dot{\psi}$  multiplied by the fraction of the bixel not blocked by the left leaf ( $l$  and  $r$  denote the locations of the left and right boundaries of the bixel, respectively):

$$\frac{d\phi_L}{dt}(x) = \begin{cases} \dot{\psi} & \text{for } x < l \\ \dot{\psi} \frac{r-x}{r-l} & \text{for } l \leq x < r \\ 0 & \text{for } x \geq r. \end{cases} \quad (2.7)$$

Inserting Equation (2.7) into Equation (2.6) and substituting  $\psi = \dot{\psi}\Delta t$ , where  $\psi$  is the total linac intensity delivered in the arc section, yields:

$$\frac{d\phi_L}{dx}(x) = \begin{cases} \psi \frac{1}{L_{\text{fin}} - L_{\text{init}}} & \text{for } x < l \\ \psi \frac{r-x}{(L_{\text{fin}} - L_{\text{init}})(r-l)} & \text{for } l \leq x < r \\ 0 & \text{for } x \geq r. \end{cases} \quad (2.8)$$

This expression for the differential fluence released to the bixel by the left leaf is integrated along the leaf trajectory from  $L_{\text{init}}$  to  $L_{\text{fin}}$ , according to Equation (2.5). There are six possible cases for this integral, depending on the initial and final left leaf positions  $L_{\text{init}}$  and  $L_{\text{fin}}$  relative to the left and

right edges  $l$  and  $r$  of the bixel.

The simplest of these cases occurs for bixels which are fully crossed by both the right and left leaves in its row. For the left leaf, this is represented symbolically by  $L_{\text{init}} < l$  &  $r < L_{\text{fin}}$ . Then the integral of Equation (2.5) is evaluated thus:

$$\begin{aligned}
\phi_L &= \int_{L_{\text{init}}}^{L_{\text{fin}}} \frac{d\phi_L}{dx}(x) dx \\
&= \int_{L_{\text{init}}}^l \psi \frac{1}{L_{\text{fin}} - L_{\text{init}}} dx + \int_l^r \psi \frac{r - x}{(L_{\text{fin}} - L_{\text{init}})(r - l)} dx + \int_r^{L_{\text{fin}}} 0 dx \\
&= \psi \frac{(r + l)/2 - L_{\text{init}}}{L_{\text{fin}} - L_{\text{init}}}.
\end{aligned} \tag{2.9}$$

The derivation of  $\phi_R$  follows a similar procedure as that for  $\phi_L$ ; the changes amount to making the following substitutions:  $L_{\text{init}} \rightarrow R_{\text{init}}$ , and  $L_{\text{fin}} \rightarrow R_{\text{fin}}$ .

Thus the final expression for the fluence of the bixel in this case is:

$$\begin{aligned}
\phi &= \phi_L - \phi_R \\
&= \psi \left( \frac{(r + l)/2 - L_{\text{init}}}{L_{\text{fin}} - L_{\text{init}}} - \frac{(r + l)/2 - R_{\text{init}}}{R_{\text{fin}} - R_{\text{init}}} \right).
\end{aligned} \tag{2.10}$$

A derivation of all of the other cases may be found in Appendix A.

Note that  $L_{\text{init}}$ ,  $R_{\text{init}}$ ,  $L_{\text{fin}}$ , and  $R_{\text{fin}}$  of the MLC leaves are defined by the two apertures  $A_{\theta_{\text{init},k}}$  and  $A_{\theta_{\text{fin},k}}$ , shown in Figure 2.2a. Within the above framework, the bixel fluence can now be calculated analytically when the apertures change continuously from  $A_{\theta_{\text{init},k}}$  to  $A_{\theta_{\text{fin},k}}$ . It is these aperture variables, as well as the linac intensity during each arc,  $\psi_{\theta_k}$ , which are used

as variables during the plan optimization described below.

### 2.2.1.2 Optimization

Aside from the implementation of VMAT which is presently being discussed, matRad is capable of IMRT inverse planning, with both FMO and DAO available. The optimization itself is accomplished using the IPOPT algorithm [67]. matRad interfaces with IPOPT using callback functions which are used to evaluate the objective function and the gradient of the objective function with respect to the optimization variables. Constrained optimization is also possible with IPOPT. In this case, additional callback functions are required that evaluate the constraints and the Jacobian of the constraint functions with respect to the variables.

The three-step optimization sequence for VMAT [57, 58] was implemented in matRad. Briefly, the process starts with FMO, performed on a coarse set of angles. Next, the resulting bixel fluence maps are sequenced into a set of apertures [68] which are distributed over a smaller angular interval. A final refinement of the aperture shapes and intensities is then carried out using DAO. Before optimization, three sets of gantry angles are specified by the user: FMO initialized angles  $\{\theta_{\text{FMO}}\}$ , DAO optimization angles  $\{\theta_{\text{DAO}}\}$ , and dose calculation angles  $\{\theta_{\text{dose}}\}$  ( $\{\theta_{\text{FMO}}\} \subseteq \{\theta_{\text{DAO}}\} \subseteq \{\theta_{\text{dose}}\}$ ). The relation of these angles is illustrated in Figure 2.3.

In the FMO step, the fluence of each bixel in fields in the  $\{\theta_{\text{FMO}}\}$  angle set is independently optimized with the goal of minimizing a user-defined objective

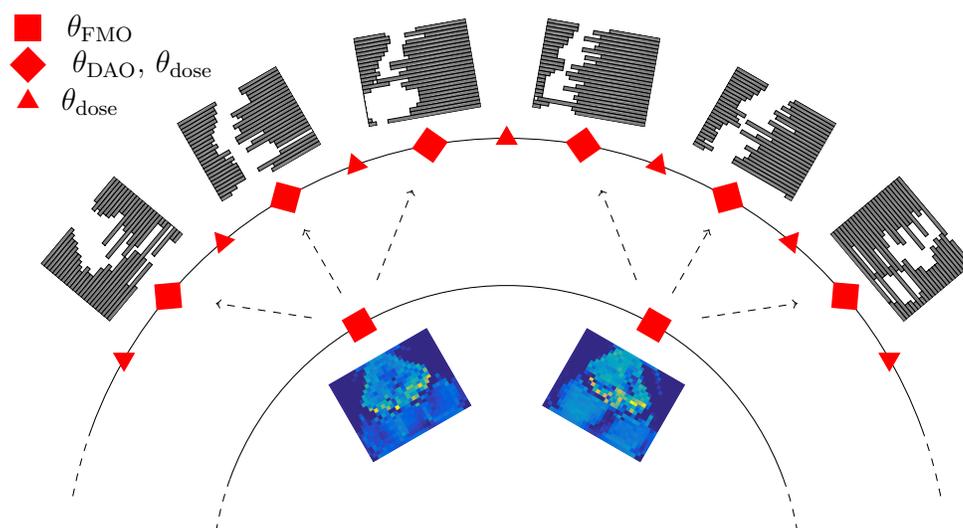


Figure 2.3: Definition of the user-specified angles  $\theta_{\text{FMO}}$ ,  $\theta_{\text{DAO}}$ , and  $\theta_{\text{dose}}$  (legend in upper-right; note that every  $\theta_{\text{DAO}}$  angle is also a  $\theta_{\text{dose}}$  angle). The arcs indicate the path of the gantry around the patient. Optimized fluence maps are shown at each  $\theta_{\text{FMO}}$ ; these are converted into aperture shapes and intensities and spread to the three neighbouring  $\theta_{\text{DAO}}$  angles (denoted by the dashed arrows). Apertures are distributed in alternating leaf travel direction between successive  $\theta_{\text{FMO}}$  angles. Aperture shapes and intensities are interpolated at the  $\theta_{\text{dose}}$  angles.

function  $F(d)$ . Square dose deviation terms, including under and over dosage, were used in the definition of the objective function. The derivatives of the objective function with respect to each bixel fluence, required for the optimization, are calculated by matRad [1].

After FMO, each fluence map belonging to a particular  $\theta_{\text{FMO}}$  is converted into aperture shapes and aperture intensities  $\psi$  using a sliding window sequencing algorithm [68], which are then distributed to the closest optimization angles  $\theta_{\text{DAO}}$  neighbouring  $\theta_{\text{FMO}}$ . As illustrated in Figure 2.3, apertures are distributed so that they sweep from right to left for the first  $\theta_{\text{FMO}}$ , then in alternating directions for the subsequent  $\theta_{\text{FMO}}$  angles; this is done to minimize the leaf travel between successive  $\theta_{\text{DAO}}$  angles. The sequencing algorithm will generally output more apertures than the number of neighbouring  $\theta_{\text{DAO}}$  angles. To ensure that the fluence map is best reproduced, the sequenced apertures having the largest intensity-area product are kept with their aperture intensities rescaled to maintain the original summed intensity-area product.

This set of aperture shapes and intensities is used as the initialization for the DAO step. The aperture intensity delivered in the arc section surrounding each  $\theta_{\text{DAO}}$ , as well as the leaf positions at the beginning and end of these arc sections, are used as optimization variables in the DAO step (see Figure 2.2a). From these variables, the fluence and resulting dose can be directly calculated according to the continuous aperture method derived in the preceding section and in Appendix A. Upon each iteration of the DAO algorithm, the leaf positions and aperture intensities at the  $\{\theta_{\text{dose}}\}$  angles are interpolated to

Table 2.1: Minimum and maximum values for VMAT machine delivery constraints.

Constraint	Min value	Max value
Leaf speed / (cm/s)	0	6
Gantry rotation speed / ( $^{\circ}$ /s)	0	6
MU rate / (MU/min)	75	600

maintain the dose calculation accuracy. Because the continuous aperture fluence calculation method already takes into account the linear motion of leaves, this interpolation will not, and is not intended to, improve the bixel fluence accuracy. Instead, it is done only to improve the angular sampling of the dose-influence matrices.

The DAO step includes constraints placed on MLC leaf speed, gantry rotation speed, and MU rate. Callback functions for these constraints are defined for use by IPOPT. In addition, the Jacobian for these constraint functions with respect to the optimization variables is calculated. Minimum and maximum values used for the delivery constraints are listed in Table 2.1. After optimization, total delivery time for each plan is estimated by decreasing the time difference between two  $\theta_{\text{DAO}}$  angles such that the most restrictive constraint is at its limiting value.

The derivatives of the objective function with respect to the leaf positions and the aperture intensities are required for the DAO step. Built on the analytical framework for continuous aperture fluence and dose calculation

described in the previous section, a full derivation of these derivatives can be found in Appendix B.

During the DAO step, the variables corresponding to the aperture intensities were scaled in order to force the linac intensity derivatives and leaf position derivatives to be of the same order of magnitude [69], greatly improving the convergence of the objective function.

## 2.2.2 Discrete aperture VMAT approximation

Most current VMAT treatment planning systems perform dose calculation and optimization under an assumption of discrete apertures changing instantaneously from one discrete angle to the next [25]. This conventional discrete aperture method was implemented in matRad to assess the improvements in dose calculation accuracy and differences in plan quality, if any, by the proposed continuous aperture scheme.

### 2.2.2.1 Fluence and dose calculation

The dose-influence matrix method of dose calculation found in Equation (2.1) is also used when calculating dose in discrete aperture VMAT, however, the discrete aperture approximation is used. In this approximation, the aperture shapes are assumed to not vary over the angular arc  $[\theta_{\text{init},k}, \theta_{\text{fin},k}]$  shown in Figure 2.2b; instead, the MLC leaf positions instantaneously change from one such interval to the next.

This results in the well-known expression for the bixel influence in discrete

aperture VMAT [25]:

$$\phi_{j\theta_k} = \psi_{\theta_k} f_{j\theta_k}. \quad (2.11)$$

Here,  $\psi_{\theta_k}$  is the total linac intensity delivered over the arc section surrounding  $\theta_k$ , and  $f_{j\theta_k}$  is the fractional portion of bixel  $j$  not blocked by the leaves, which depends only on the positions of the tips of the left and right leaves at angle  $\theta_k$  in the row corresponding to bixel  $j$ .

### 2.2.2.2 Optimization

The optimization of VMAT plans with discrete aperture dose calculation follows closely the sequence of continuous aperture VMAT optimization described in Section 2.2.1. First, FMO is performed on the coarse set of gantry angles  $\{\theta_{\text{FMO}}\}$ , and the resulting fluence maps are converted into aperture shapes, which are spread to the  $\{\theta_{\text{DAO}}\}$  set of angles.

Finally, DAO is performed, but with a different set of optimization variables than in continuous aperture optimization. The aperture intensity delivered in the arc section surrounding each  $\{\theta_{\text{DAO}}\}$  is optimized; since the aperture is fixed over each of these arc sections, however, only the leaf positions at each  $\{\theta_{\text{DAO}}\}$  are optimized (see Figure 2.2b). Similar to the continuous aperture VMAT case, dose calculation is done at the  $\{\theta_{\text{dose}}\}$  angles, with leaf positions and aperture intensities interpolated from the variables at the  $\{\theta_{\text{DAO}}\}$  angles.

Because of the new expressions for the fluence and the accompanying change of variables, the derivatives with respect to the leaf positions are

changed; expressions for these derivatives in the discrete aperture approximation have been previously published by Wild et al. [69].

A similar scaling of the variables corresponding to the aperture intensities was done as in the discrete aperture VMAT optimization, in order to force the intensity and leaf position derivatives to be of the same order of magnitude [69]. Furthermore, the same objective function and delivery constraints were used in the continuous aperture VMAT optimization as in the discrete aperture case.

### 2.2.3 Dose angle interval study

The following describes the steps taken to determine the validity of the discrete dose angle and discrete aperture approximations at a given dose angle interval, as well as the ability of the continuous aperture dose calculation method to correct the errors introduced by the discrete aperture approximation.

VMAT plans were optimized and calculated on three patient datasets, all originating from the CORT dataset [70]: the TG119 horseshoe-shaped target [71], a prostate patient, and a head & neck patient. These datasets are included with the matRad distribution, and include the patient CT as well as structure sets comprising the target(s) and OARs.

Plans for each of the three test cases were optimized using the discrete aperture VMAT algorithm described in Section 2.2.2 with angular intervals of  $\Delta\theta_{\text{FMO}} = 28^\circ$ ,  $\Delta\theta_{\text{DAO}} = 4^\circ$ , and  $\Delta\theta_{\text{dose}} = 4^\circ$ . The dose for each plan was then recalculated at successively smaller intervals  $\Delta\theta_{\text{dose}}$  using one of two schemes.

In scheme A, the original discrete apertures are interpolated at a finer  $\Delta\theta_{\text{dose}}$  angular interval, using the same discrete aperture fluence calculation method. For smaller angular intervals  $\Delta\theta_{\text{dose}}$ , successive apertures will tend to be more similar, gradually reducing the dose discrepancy introduced by the discrete aperture approximation; the dose-influence matrices at adjacent angles will also tend to be more similar, reducing the effect introduced by the discrete dose angle approximation. Therefore any dose differences in the higher resolution plans compared to the original at  $\Delta\theta_{\text{dose}} = 4^\circ$  will be due to the breakdown of one or both of the assumptions in the discrete dose angle and discrete aperture approximations: that is, that the dose-influence matrix is constant in an angular neighbourhood surrounding the  $\{\theta_{\text{dose}}\}$  angles, and that the leaves move instantaneously from one angular interval to the next.

In scheme B, the leaf positions are interpolated at the new  $\{\theta_{\text{dose}}\}$  angles, and the continuous aperture method is used to calculate the fluence of each bixel. This scheme tests the validity of the discrete dose angle approximation only, since the discrete aperture approximation is not used.

Each plan was recalculated at angular intervals of  $\Delta\theta_{\text{dose}} = 0.5^\circ, 1^\circ, 2^\circ,$  and  $4^\circ$  using the two dose calculation schemes. The dose distributions were compared to the reference dose, which was selected to be the one resulting from scheme B recalculated at  $\Delta\theta_{\text{dose}} = 0.5^\circ$ , this scheme being the most accurate with respect to the actual delivered dose.

## 2.2.4 Treatment planning study

The VMAT optimization was performed on the same three patient datasets mentioned in Section 2.2.3. Angular intervals for the three gantry angle sets are listed in Table 2.2; these values were informed by the results of the dose angle interval study, Section 2.3.1, in particular Figures 2.4 and 2.5. The head and neck case used a  $\Delta\theta_{\text{dose}}$  interval of  $2^\circ$  and a  $\Delta\theta_{\text{DAO}}$  interval of  $4^\circ$ . This means that during DAO, apertures were optimized at  $4^\circ$ , while dose was calculated to a interval of  $2^\circ$ , with aperture interpolation being performed at the non-optimized angles.

Plans were first optimized using the discrete aperture method with the goal of attaining DVH objectives based on the particular anatomical site; each case had at least one target with a prescription dose of 2 Gy per fraction (both the prostate and head & neck cases featured multiple targets). For the TG119 case, these objectives originated from the TG119 report [71], while for the two patient cases, treatment planning guidelines at The Ottawa Hospital were used. All objectives are reported in Tables 2.3 to 2.5.

After optimization, these discrete aperture VMAT plans were recalculated using the continuous aperture dose calculation. The resulting effect on plan quality of using the more accurate dose calculation was assessed.

Finally, the effectiveness of the continuous aperture VMAT optimization algorithm was tested by re-optimizing the same three cases with the same target and OAR objective doses used in the discrete aperture plans.

Table 2.2: Intervals for the three gantry angle sets used in the VMAT planning study for each anatomical site considered. Refer to Figure 2.3 for the definitions of each type of angle.

Anatomical site	$\Delta\theta_{\text{FMO}}/^\circ$	$\Delta\theta_{\text{DAO}}/^\circ$	$\Delta\theta_{\text{dose}}/^\circ$
TG119	28	4	4
Prostate	28	4	4
Head & neck	28	4	2

## 2.3 Results

### 2.3.1 Dose angle interval study

Comparisons between the dose distributions  $d_i$  for discrete aperture VMAT plans recalculated using the two schemes and the corresponding reference dose distribution  $d_i^{\text{ref}}$  can be found in Figures 2.4 and 2.5 (prostate results not shown). The reference dose is defined to be the distribution resulting from scheme B recalculated at  $\Delta\theta_{\text{dose}} = 0.5^\circ$ . Each figure shows a semi-log plot of the percentage of the volume of interest (vertical axis) for which the relative local dose difference  $\Delta d_i(\%) = 100\% \cdot \frac{d_i - d_i^{\text{ref}}}{d_i^{\text{ref}}}$  exceeds a certain amount (horizontal axis). The volume of interest is defined to be the union of volumes corresponding to the targets and the OARs which contribute a term in the objective function; these are specified in Tables 2.3 to 2.5.

The dose recalculated using scheme A (discrete aperture fluence calculation) approaches the reference dose as the angular interval decreases: the best 99% of voxels have maximum dose differences that decrease from 47.7%,

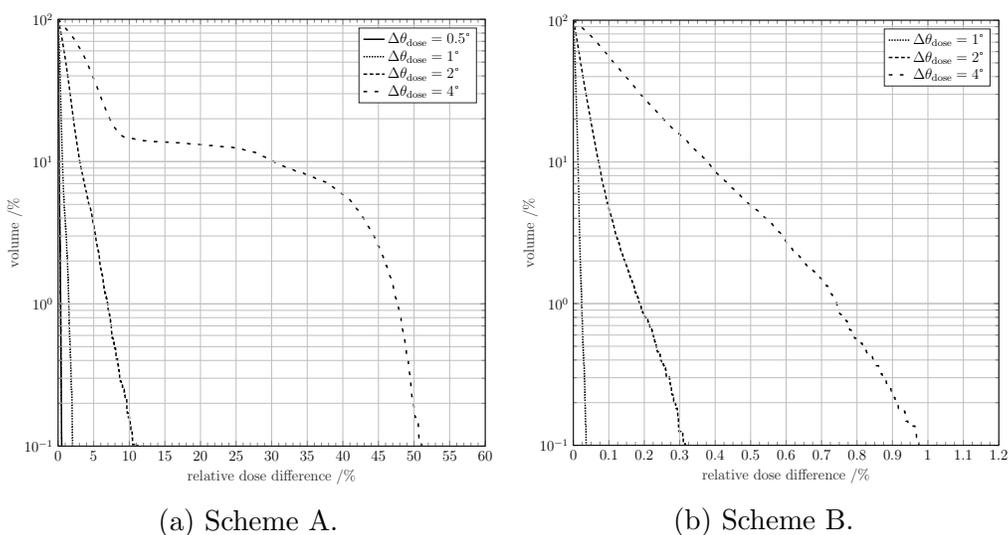


Figure 2.4: TG119 case: comparison between dose recalculated using scheme A (discrete aperture) and B (continuous aperture) at 4 angular intervals  $\Delta\theta_{\text{dose}}$  and the reference dose (scheme B at  $\Delta\theta_{\text{dose}} = 0.5^\circ$ ) for the discrete aperture VMAT plan optimized at  $\Delta\theta_{\text{dose}} = 4^\circ$ . The graph shows the percentage of the volume of interest (target and OAR) for which the relative local dose difference exceeds a certain amount.

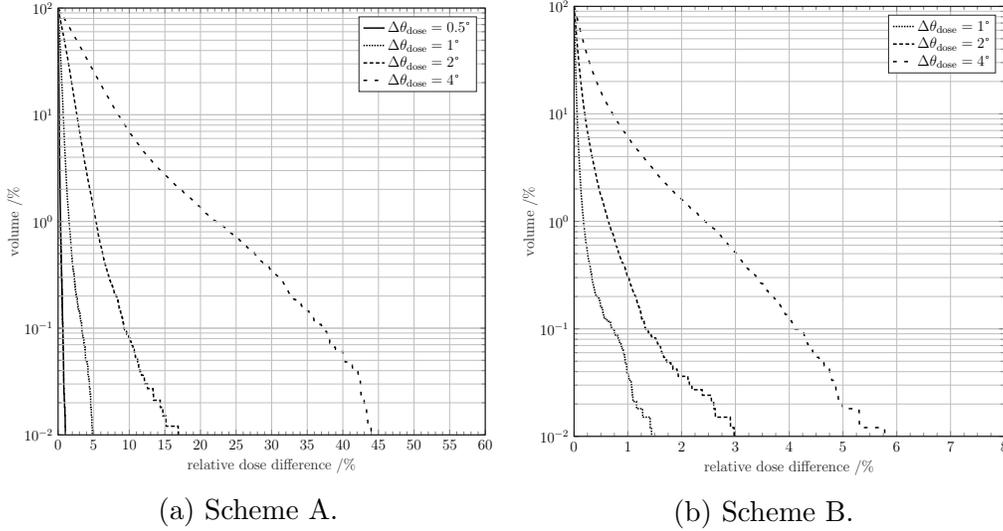


Figure 2.5: Head & neck case: comparison between dose recalculated using scheme A (discrete aperture) and B (continuous aperture) at 4 angular intervals  $\Delta\theta_{\text{dose}}$  and the reference dose (scheme B at  $\Delta\theta_{\text{dose}} = 0.5^\circ$ ) for the discrete aperture VMAT plan optimized at  $\Delta\theta_{\text{dose}} = 4^\circ$ . The graph shows the percentage of the volume of interest (targets and OARs) for which the relative local dose difference exceeds a certain amount.

10.9 %, and 22.3 % at  $\Delta\theta_{\text{dose}} = 4^\circ$  to 0.4 %, 0.2 %, and 0.3 % at  $\Delta\theta_{\text{dose}} = 0.5^\circ$ , for the TG119, prostate, and head & neck plans, respectively. This excellent agreement between schemes A and B at the fine  $\Delta\theta_{\text{dose}}$  angular interval indicates that the discrete aperture and continuous aperture dose calculation methods converge in the limit of smaller angular interval.

The dose recalculated using scheme B yields small dose differences relative to the reference, even at  $\Delta\theta_{\text{dose}} = 4^\circ$ . Since this scheme uses the continuous aperture dose calculation, these differences are purely due to the changes in the dose-influence matrix. The maximum dose differences between scheme B at  $\Delta\theta_{\text{dose}} = 4^\circ$  and the reference dose are 0.7 %, 1.0 %, and 2.4 % for the

best 99% of voxels in the TG119, prostate, and head & neck plans. The dose differences are smallest in the TG119 plan and largest in the head & neck plan, which can be related to the heterogeneity in each plan’s CT image, and thus to the variation of the dose-influence matrix as a function of angle.

Figures 2.4 and 2.5 show that in order to achieve  $< 1\%$  dose difference in the best 99% of voxels using the discrete aperture approximation, it is necessary to calculate dose to  $\Delta\theta_{\text{dose}} = 0.5^\circ$ . To achieve similar agreement using continuous aperture, one can use  $4^\circ$  for TG119 and prostate, and  $2^\circ$  for head & neck, representing a reduction in the required number of  $\theta_{\text{dose}}$  by a factor of 4 or 8.

### 2.3.2 Treatment planning study

The results of the treatment planning study are reported in Tables 2.3 to 2.5, which compare the investigated DVH objectives that were extracted from the three dose distributions: discrete aperture VMAT optimization, the discrete aperture plan recalculated using continuous aperture fluence calculation, and continuous aperture VMAT optimization. Throughout these tables, the symbol  $d_{p\%}$  is used to signify the smallest dose which is received by at least  $p\%$  of a particular volume, while  $d_{\text{mean}}$  represents the mean dose to that volume.  $V_d$ , on the other hand, represents the percentage of the volume which receives a dose at least as high as  $d$ .

All objectives in all three cases were met for both the discrete aperture and continuous aperture optimized plans. A comparison of the plan objective

values between these two approaches for each case reveals that they yield dose distributions with similar plan qualities overall. The plan quality of the discrete aperture optimized plans decreases noticeably when recalculated using the continuous aperture fluence calculation.

For example, in the TG119 case, continuous aperture recalculation of the discrete aperture VMAT plan showed that the 50 Gy target was simultaneously under- and overdosed: the  $D_{95\%}$  prescription dose was not met, and the uniformity threshold of  $D_{10\%} < 55$  Gy was marginally exceeded. Furthermore, the DVH objective of  $D_{10\%} < 25$  Gy in the core OAR was exceeded by 5 Gy. The estimated delivery time for the discrete aperture VMAT plan was 1.8 min, and for the continuous aperture plan was 2.3 min.

The high- and low-dose thresholds for the 76 Gy PTV in the prostate case were not met by the discrete aperture VMAT recalculated plan. All other objectives were elevated as compared to the discrete aperture VMAT optimized plan, but not in excess of the corresponding thresholds. The estimated delivery time for the discrete aperture VMAT plan was 1.8 min, and for the continuous aperture plan was 2.1 min.

Table 2.3: TG119 case: results of the discrete aperture and continuous aperture VMAT optimized plan, as well as the dose from the former plan recalculated using continuous aperture fluence calculation. Each row corresponds to a different DVH objective specified in the TG119 report. A quantity in bold indicates that the objective for a particular plan exceeds the corresponding threshold.

Volume	DVH objective	Threshold	Discrete aperture optimized	Discrete aperture recalculated	Continuous aperture optimized
50 Gy PTV	$D_{95\%}/\text{Gy}$	>50	50.0	<b>49.6</b>	50.0
	$D_{10\%}/\text{Gy}$	<55	52.1	<b>55.1</b>	52.4
Core	$D_{10\%}/\text{Gy}$	<25	21.7	<b>30.1</b>	21.7

Table 2.4: Prostate case: results of the discrete aperture and continuous aperture VMAT optimized plan, as well as the dose from the former plan recalculated using continuous aperture fluence calculation. Each row corresponds to a different DVH objective used as guidelines at The Ottawa Hospital. A quantity in bold indicates that the objective for a particular plan exceeds the corresponding threshold.

Volume	DVH objective	Threshold	Discrete aperture optimized	Discrete aperture recalculated	Continuous aperture optimized
76 Gy PTV	$D_{95\%}/\text{Gy}$	>76	76.0	<b>75.1</b>	76.0
	$D_{2\%}/\text{Gy}$	<83.6	82.5	<b>84.2</b>	82.4
46 Gy PTV	$D_{95\%}/\text{Gy}$	>46	47.3	46.8	47.4
Bladder	$V_{65\text{ Gy}}/\%$	<50	20.0	20.3	20.5
	$V_{70\text{ Gy}}/\%$	<35	18.0	18.0	18.0
	$V_{75\text{ Gy}}/\%$	<25	15.7	15.8	15.8
	$V_{80\text{ Gy}}/\%$	<15	8.3	12.0	7.8
Rectum	$V_{60\text{ Gy}}/\%$	<50	19.2	26.6	16.7
	$V_{65\text{ Gy}}/\%$	<35	12.1	17.9	11.0
	$V_{70\text{ Gy}}/\%$	<25	7.9	11.6	8.0
	$V_{75\text{ Gy}}/\%$	<15	5.5	7.3	5.5

Table continued on following page.

Table 2.4: (continued)

Volume	DVH objective	Threshold	Discrete aperture optimized	Discrete aperture recalculated	Continuous aperture optimized
Penile bulb	$D_{\text{mean}}/\text{Gy}$	$< 5.25$	4.6	4.6	4.5
Left femur	$V_{50\text{Gy}}/\%$	$< 5$	0.1	0.3	0.1
Right femur	$V_{50\text{Gy}}/\%$	$< 5$	0	0	0.2

Table 2.5: Head & neck case: results of the discrete aperture and continuous aperture VMAT optimized plan, as well as the dose from the former plan recalculated using continuous aperture fluence calculation. Each row corresponds to a different DVH objective specified used as guidelines at The Ottawa Hospital. A quantity in bold indicates that the objective for a particular plan exceeds the corresponding threshold.

Volume	DVH objective	Threshold	Discrete aperture optimized	Discrete aperture recalculated	Continuous aperture optimized
70 Gy PTV	$D_{99\%}/\text{Gy}$	>66.5	69.3	<b>65.8</b>	69.8
	$D_{95\%}/\text{Gy}$	>70	70.0	<b>68.7</b>	70.4
	$D_{20\%}/\text{Gy}$	<77	72.0	75.1	72.1
	$D_{\max}/\text{Gy}$	<80.5	74.7	<b>80.6</b>	74.8
	$D_{\text{mean}}/\text{Gy}$	<73.5	71.3	72.9	71.6
63 Gy PTV	$D_{99\%}/\text{Gy}$	>59.8	59.9	<b>52.9</b>	59.8
	$D_{95\%}/\text{Gy}$	>63	65.5	<b>60.4</b>	66.2
Spinal cord	$D_{0.1 \text{ cm}^3}/\text{Gy}$	<45	42.6	<b>45.9</b>	41.8
	$D_{\max}/\text{Gy}$	<48	42.9	46.0	41.9
Brainstem	$D_{0.1 \text{ cm}^3}/\text{Gy}$	<50	46.6	46.7	45.1
	$D_{\max}/\text{Gy}$	<54	47.3	48.3	46.4

Table continued on following page.

Table 2.5: (continued)

Volume	DVH objective	Threshold	Discrete aperture optimized	Discrete aperture recalculated	Continuous aperture optimized
Optic chiasm	$D_{\max}/\text{Gy}$	<50	3.9	3.7	3.7
Left parotid	$D_{50\%}/\text{Gy}$	<30	18.4	23.5	17.6
	$D_{\text{mean}}/\text{Gy}$	<26	24.5	<b>28.1</b>	23.9
Right parotid	$D_{50\%}/\text{Gy}$	<30	16.6	20.0	16.2
	$D_{\text{mean}}/\text{Gy}$	<26	21.9	25.1	21.4
TM joint	$D_{0.1\text{ cm}^3}/\text{Gy}$	<70	63.1	56.0	63.6
Larynx	$D_{67\%}/\text{Gy}$	<50	40.7	44.9	31.9
	$D_{\text{mean}}/\text{Gy}$	<45	42.5	<b>47.5</b>	34.9
Lips	$D_{\text{mean}}/\text{Gy}$	<20	17.1	<b>20.3</b>	18.0
	$D_{\max}/\text{Gy}$	<30	23.2	<b>30.1</b>	27.5
Optic nerve	$D_{\max}/\text{Gy}$	<50	3.9	3.8	3.8

After recalculating the discrete aperture VMAT plan using continuous aperture fluence calculation in the head & neck case, neither the  $D_{95\%}$  nor the  $D_{99\%}$  dose thresholds for the 70 Gy and 63 Gy PTVs were attained. Furthermore, the  $D_{0.1\text{cm}^3}$  objective of 45 Gy for the spinal cord was exceeded, along with the  $D_{\text{mean}}$  thresholds for the left parotid gland and the larynx. Finally, both the  $D_{\text{mean}}$  and  $D_{\text{max}}$  thresholds for the lips were marginally exceeded. The other objectives were elevated as compared to the discrete aperture VMAT optimized plan, but not in excess of the corresponding thresholds. The estimated delivery times for the discrete aperture VMAT plan was 2.4 min, and for the continuous aperture plan was 3.6 min.

These results show that the discrete aperture VMAT optimization overestimates the actual plan quality delivered by the linac: in all three cases, certain DVH objectives were not within the specified thresholds when dose was recalculated using the continuous aperture fluence calculation, with the targets receiving less than the prescription dose and the OARs receiving a higher dose. Instead of delivering fluence in discrete apertures, fluence is spread between two apertures as the leaves sweep from one position to the next. This spreads the dose deposition over a greater volume, shifting dose away from the high-dose targets to the low-dose OARs. Note that if the total number of MU was rescaled to meet the required target coverage, the dose in the OARs would be increased, leading to a greater excess of DVH objective thresholds.

The desired plan quality is able to be achieved using the continuous

aperture VMAT optimization, which does not require dose recalculation. Furthermore, the delivery times of the discrete aperture and continuous aperture VMAT plans are comparable in all three cases, and did not exceed 4 min.

## 2.4 Discussion

The results of the angle interval study show that significant dose errors can appear in discrete aperture VMAT plans optimized at a  $4^\circ$  interval, which are revealed when recalculated at a smaller  $0.5^\circ$  interval, agreeing with findings reported by other authors [24, 59]. The results from the scheme B interpolation imply that these errors mainly stem from the discrete aperture approximation, not the discrete dose angle approximation, and therefore can be corrected by using the continuous aperture fluence calculation. Indeed, a much coarser gantry angle spacing may be used for a similar dose accuracy if the continuous aperture fluence calculation is employed instead of the discrete aperture method (by a factor of 4 or 8, depending on the patient geometry and heterogeneity). The smaller  $\Delta\theta_{\text{dose}}$  of  $2^\circ$  required for the head & neck case (versus  $4^\circ$  for the TG119 and prostate) is due to the increased heterogeneity in this anatomical site, causing a breakdown of the discrete dose angle approximation.

This reduction in the angular interval has the advantage of lesser computational time and memory requirements to calculate and store the dose-influence

matrix. For example, in performing a discrete aperture VMAT optimization for the TG119 case with a  $\Delta\theta_{\text{dose}}$  interval of  $2^\circ$  (dosimetric results not shown), the dose-influence matrix was calculated in 90 min, occupying 41.3 GB in system memory; the total optimization time for both FMO and DAO was 206 min. When the optimization was done using the continuous aperture calculation at a larger  $\Delta\theta_{\text{dose}}$  interval of  $4^\circ$  (which will still result in a much more accurate dose calculation, according to Figure 2.4), the dose-influence matrix was calculated in 36 min, occupying 20.7 GB in system memory, and the total optimization time was only 89 min. All dose calculation and plan optimization in this thesis was performed with an Intel<sup>®</sup> Xeon<sup>®</sup> Processor E5-1620 v3 with a frequency of 3.50 GHz, with 64 GB of DDR4 memory with frequency 2133 MHz. Although absolute numbers will depend on the specific hardware and software implementation, dose calculation method, and optimization targets, these results show that the continuous aperture algorithm can generate a VMAT plan (initial dose calculation and optimization) in about 40% of the time and using half of the system memory compared to the discrete aperture approximation.

As shown in the treatment planning study, the errors introduced by the discrete aperture approximation do have a detrimental effect on treatment plan quality. Comparing plan quality from the optimized dose to the dose recalculated using the continuous aperture fluence calculation, a decreased dose coverage and dose uniformity in the targets, as well as an increased dose in the OARs is observed. This is problematic for treatment planners, since

the plan quality at the end of the optimization will not necessarily be equal to the delivered plan quality. If, however, optimization is performed with the continuous aperture fluence calculation, the results of the treatment planning study show that the desired plan quality can be achieved using the more accurate calculation method.

It is difficult to directly compare plan quality and delivery time results from this algorithm to the direct leaf trajectory optimization approach proposed by Papp and Unkelbach [62], since that study used different patient geometries and prescription doses. Nevertheless, in that approach, the plan quality improves significantly when a longer treatment time is selected, until the best quality is reached at 3 min to 4 min total time. The longer treatment time provides the MLC leaves more time to modulate the field, thus leading to a better plan. The possibilities for modulation are limited, however, by the restriction that leaves must travel unidirectionally in a given  $18^\circ$  arc section before changing direction only at the edge of the field.

In contrast, the proposed continuous aperture VMAT optimization algorithm is not hindered by this restriction; furthermore, the MU rate is allowed to vary. These features have the potential to shorten the treatment time, which was 2 min to 4 min, depending on the case. Decreasing the maximum leaf speed from 6 cm/s to 3 cm/s (as used by Papp and Unkelbach [62]) increases estimated the estimated treatment time for the TG119 and prostate cases to 2.7 min, and for the head & neck case to 5.4 min, which is comparable to the 3 min to 4 min times reported by Papp and Unkelbach [62]. Overall,

this comparison is inconclusive because of the difference in patient geometry and the fact that the direct leaf trajectory study used smaller prescription doses than this study.

## 2.5 Conclusions

The dosimetric errors introduced by the discrete dose angle and discrete aperture approximations typically used in VMAT treatment planning were each tested. The discrete aperture approximation was found to have the largest contribution, and could be eliminated using the continuous aperture fluence calculation. Errors introduced by the discrete dose angle approximation were minimized by an informed selection of the dose calculation angular interval. These dosimetric errors were shown to translate into a decrease in the actual delivered quality of a discrete aperture VMAT plan. The desired plan quality was able to be attained by incorporating the continuous aperture fluence calculation into the VMAT optimization procedure. Plan optimization was done with variable but constrained MLC leaf speed, gantry rotation speed, and MU rate, with no restrictions such as sliding window placed on the leaf travel directions. Using an appropriate dose calculation angular interval with the continuous aperture VMAT optimization, the treatment planner can be assured that, as far as the continuous nature of VMAT is concerned, the final optimized dose will be an accurate representation of the delivered dose. The open-source and modular nature of matRad aided greatly in the

implementation and analysis of both the discrete aperture and continuous aperture VMAT optimization algorithms.

## Chapter 3

# Markov chain respiratory motion model

In this chapter the probabilistic respiratory motion model, a central component of the robust 4D-VMAT optimization and delivery framework to be developed in Chapter 4, is constructed. This Markov chain model is trained using real patient respiratory data acquired during radiotherapy, and evaluated by comparing model predictions of the tumour position probability distribution with data from the same patient at a later time during delivery.

### 3.1 Introduction

As discussed in Section 1.3, if it is not compensated for or eliminated entirely, intrafraction respiratory patient motion can distort the conformal dose distribution to the tumour achieved during plan optimizations described in Chapter 2. Any method which seeks to mitigate these effects by compensating

the delivered radiotherapy plan for intrafraction motion, such as gating or tracking, must be accompanied by some model of the patient motion. Since the breathing pattern of each patient is unique to some extent [29], this model must be specific to each patient.

Respiratory motion models generally serve one of two functions: to predict the future motion of the patient from a set of past observations [72, 73], or to quantify the probability distribution of, for example, the position of the tumour throughout the delivery [35, 74]. Typically, prediction-type models are used in conjunction with real-time motion measurement techniques to minimize the effects of latency on the radiotherapy plan adaptations. The distribution-type models, on the other hand, are employed in the treatment planning stage, as the probability distribution of the tumour motion is used to calculate, and optimize to the desired treatment goals, the expectation value of the dose before any treatment fraction has been delivered. Since the latter is one of the goals of this thesis, and in particular the treatment planning framework developed in Chapter 4, the subject of this chapter is the construction of a distribution-type respiratory motion model.

Distribution-type models themselves may come in two varieties: time-independent, or time-dependent. A time-independent distribution is typically constructed by simply calculating the proportion of time that the tumour spends at each position throughout delivery (see Section 3.2.1). Time-dependent distributions on the other hand account for the fact that if the

tumour position is monitored, then at the beginning of treatment the distribution is that of a Dirac delta function, and the distribution at future times is likely to be centred around a periodic trajectory.

Perhaps the simplest time-dependent distribution-type model used in radiotherapy planning is the assumption that the tumour follows a single periodic trajectory, for example that of a sinusoidal or modified sinusoidal function [75, 76]. This singular motion trajectory is a special case of a time-dependent probability distribution in which the distribution at any given time point is a Dirac delta function centred at the position defined by the trajectory. This can be seen as a manifestation of the nominal formulation of the “continuum of robustness” introduced by Chan et al. [74], as applied to 4D radiotherapy where there is no uncertainty in the motion of the tumour. This distribution cannot account for any possible variation in the breathing pattern which will occur during delivery, reducing its applicability.

On the other end of the continuum of robustness is the margin formulation, which is characterized by complete uncertainty in the motion of the tumour [74]. This is commonly realized by assuming that the target volume is simultaneously at all possible positions of its range of motion throughout the treatment, thus creating the ITV. In between these two extremes are so-called robust formulations, in which there exists some uncertainty in the tumour motion. This uncertainty is characterized by a non-Delta function time-dependent distribution of the tumour position, which will be the subject of this chapter.

Typically the form of the respiratory motion model used is chosen to fulfill the requirements of a particular 4D optimization and delivery method. As described in Section 1.5, the 4D-VMAT optimization method introduced in this thesis considers all possible trajectories when performing dose calculation and optimization of an aperture library. The tumour is then tracked during delivery, and the aperture corresponding to the current position is selected from the planned library. As will be shown in Section 4.2.1, the patient's respiratory motion is treated as a random process evolving in time, and the expectation value and variance of the dose over all trajectories are used during the library optimization.

A central component of both the expectation value and variance calculations is the calculation of the probability of observing a particular trajectory. This probability must be calculated from a patient-specific probabilistic breathing model. In particular, two items are required, from the derivation shown in Section 4.2.1.1 and Appendix C: the position probability distribution at any time after delivery has commenced; and the probability to transition from one position to another in any arbitrary interval of time. As mentioned previously, a time-dependent distribution-type respiratory motion model is therefore required.

The aim of the work in this chapter then was to construct and evaluate such a respiratory motion model which can accurately calculate the probability distribution of position as a function of time. For this, a Markov chain type model was selected, which consists of a series of states and a matrix encoding

the probabilities to transition between two states. This class of stochastic models is well-suited to both describe this periodic motion seen in respiration as well as calculate the required probabilities. First a general description and summary of Markov chains is given, before defining the states of the Markov chain from a series of adjustable and optimizable parameters. These states were determined from the position of the tumour, its velocity, the breathing phase of the patient (inhaling, exhaling, end of exhale), and by the elapsed time in the current breathing phase. After constructing the Markov chain, the accuracy of the respiratory motion model is verified by performing statistical comparisons of observed data and corresponding model calculations.

## 3.2 Summary of Markov chains

The following section is included to introduce Markov chains and summarize some important properties and results which will be used throughout this chapter. This is based on two sets of lecture notes, written separately by Tolver [77] and Noé et al. [78].

A discrete-time Markov chain is a stochastic model of a system which can be in one of several states observed at discrete points of time separated by a constant time interval  $\tau$ . The defining characteristic of a Markov chain is that the probability distribution of states at a certain time  $t + \tau$  depends only on the known state of the system at the previous time point  $t$ . In other words, Markov chains have no memory of their past behaviour.

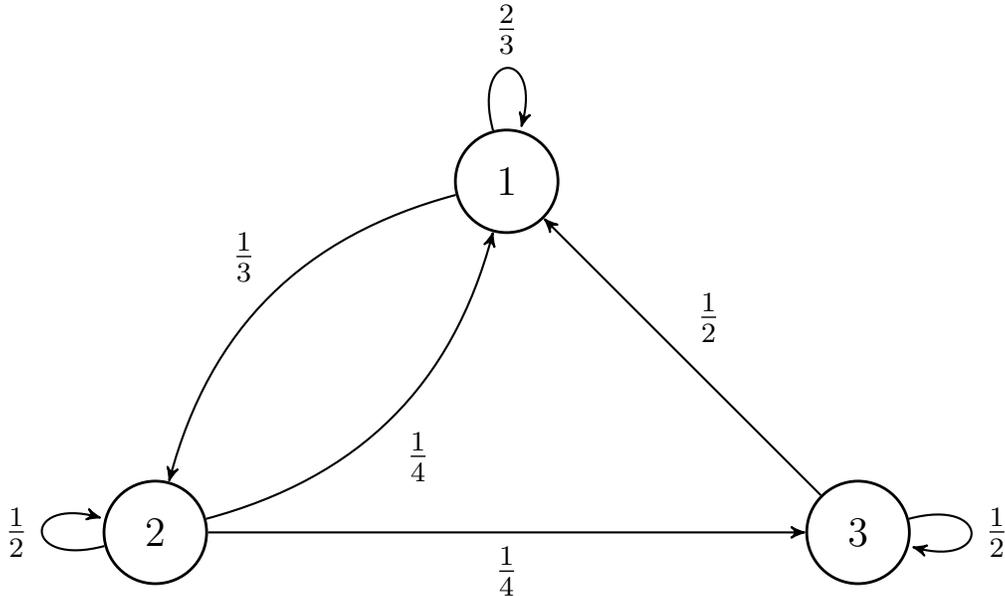


Figure 3.1: Transition diagram for a simple Markov chain. The set of states is  $S = \{1, 2, 3\}$ , and the possible transitions between these states are displayed as arrows, with associated probabilities. The corresponding transition matrix is  $\mathbf{P} = \begin{pmatrix} \frac{2}{3} & \frac{1}{3} & 0 \\ \frac{1}{4} & \frac{1}{2} & \frac{1}{4} \\ \frac{1}{2} & 0 & \frac{1}{2} \end{pmatrix}$ .

The Markov chain consists of a system  $\mathcal{S}$ , the set of  $N_S$  states  $S$  describing that system, and an  $N_S \times N_S$  transition matrix  $\mathbf{P}$  which governs the time evolution of the system. The  $(\alpha, \beta)$  element of  $\mathbf{P}$ ,  $(\mathbf{P})_{\alpha, \beta}$ , is defined as the probability for the system to transition from state  $\alpha$  to  $\beta$  in the time interval  $\tau$ . A simple example of a Markov chain is shown in Figure 3.1.

The transition matrix has two properties which can be derived from the definition just given. Since the matrix elements describe probabilities,  $(\mathbf{P})_{\alpha, \beta}$  is limited to real values between 0 and 1. Second, the probability to transition from a state  $\alpha$  to any state, including  $\alpha$ , after time  $\tau$  must be unity. This

probability can be calculated from the sum of  $(\mathbf{P})_{\alpha,\beta}$  over all final states  $\beta$ , therefore:

$$\sum_{\beta \in S} (\mathbf{P})_{\alpha,\beta} = 1, \forall \alpha \in S. \quad (3.1)$$

A square matrix which satisfies these properties is known in mathematics as a stochastic matrix.

If the probability distribution of states at time  $t$  is known,  $P_\alpha(t)$ , the distribution at time  $t + \tau$  can be calculated using the transition matrix. The probability of an initial state  $\alpha$  is multiplied by the probability to transition from  $\alpha$  to the final state  $\beta$ , before summing over all initial states (note that unless otherwise noted, summation over all states  $\alpha \in S$  is implied):

$$P_\beta(t + \tau) = \sum_{\alpha} P_\alpha(t) (\mathbf{P})_{\alpha,\beta}. \quad (3.2)$$

The probability distribution can be represented as a  $1 \times N_S$  row vector  $\vec{P}(t)$  whose elements are the  $P_\alpha(t)$ . Equation (3.2) can then be simplified as the  $\beta$  component of the following vector equation, where the  $*$  operator represents left-multiplication:

$$\vec{P}(t + \tau) = \vec{P}(t) * \mathbf{P}. \quad (3.3)$$

Since the starting time  $t$  is arbitrary, the probability distribution at any time  $t + n\tau$ ,  $n \in \mathbb{N}$ , can be calculated by induction as:

$$\vec{P}(t + n\tau) = \vec{P}(t) * \mathbf{P}^n, \quad (3.4)$$

where  $\mathbf{P}^n$  is the  $n$ th power of the transition matrix. This equation describes how the probability distribution of states evolves from the initial distribution as the number of time intervals  $n$  increases.

### 3.2.1 The stationary distribution

A stationary probability distribution of states is one which is independent of time under the actions of the Markov chain. It can be shown that if the Markov chain is irreducible (all of the states are accessible from each other) and recurrent (for any initial state  $\alpha$ , the probability to return to  $\alpha$  is 1), the stationary distribution exists and is unique for a particular Markov chain. The stationary distribution  $\pi_\alpha$  can be found using Equation (3.3), imposing that  $\vec{\pi} = \vec{P}(t + \tau) = \vec{P}(t)$ :

$$\vec{\pi} = \vec{\pi} * \mathbf{P}. \quad (3.5)$$

Thus the stationary distribution  $\pi_\alpha$  are the elements of the left eigenvector of  $\mathbf{P}$  with eigenvalue 1.

Given any initial distribution  $\vec{P}(0)$  and the conditions of irreducibility and recurrence on the Markov chain, the distribution  $\vec{P}(n\tau)$  will converge to the stationary distribution  $\vec{\pi}$  in the limit  $n \rightarrow \infty$  [78, 77]. The rate of convergence can be determined by examining the discrepancy between  $\vec{P}(n\tau)$

and  $\vec{\pi}$ ; in this study, the  $L^2$ -norm is used to quantify this error:

$$\begin{aligned} E(n\tau) &= \left\| \vec{P}(n\tau) - \vec{\pi} \right\| \\ &= \sqrt{\sum_{\alpha} (P_{\alpha}(n\tau) - \pi_{\alpha})^2}. \end{aligned} \quad (3.6)$$

The error can be shown to decay exponentially in  $n$  with rate  $k = -1/\ln \lambda_2$ , where  $\lambda_2$  is the second-largest eigenvalue of  $\mathbf{P}$  (the largest being 1) [78]. In this study, the convergence time is defined as the minimal time required to achieve an error smaller than 0.01.

### 3.2.2 Maximum likelihood estimation of Markov chains

In most circumstances the transition matrix  $\mathbf{P}$  is unknown. Instead, the state of the system is observed at discrete times separated by the interval  $\tau$ , resulting in the sequence  $\alpha_{\text{obs}}(n\tau) \in S$ ,  $n \in \{0, \dots, N\}$ . The observed data is assumed to have originated from a Markov chain with transition matrix  $\mathbf{P}$ , which can thus be estimated from the data.

Maximum likelihood estimation (MLE) is a common method to generate a model from observed data, in which the estimate maximizes the probability of observing the given data. MLE can be applied to Markov chains in the following way to estimate the elements for the transition matrix  $\mathbf{P}$  [78]. First the transition count matrix,  $\mathbf{Z} \in \mathbb{N}^{N_S \times N_S}$ , is constructed by letting  $(\mathbf{Z})_{\alpha,\beta}$

be the number of transitions observed in the sequence  $\alpha_{\text{obs}}(n\tau)$  from state  $\alpha$  at time  $(n-1)\tau$  to state  $\beta$  at time  $n\tau$ , summed over all  $n \in \{1, \dots, N\}$ . Then the maximum likelihood estimate of the transition matrix is equal to the intuitive guess given by the normalized transition count matrix:

$$(\mathbf{P})_{\alpha,\beta} = \frac{(\mathbf{Z})_{\alpha,\beta}}{\sum_{\gamma} (\mathbf{Z})_{\alpha,\gamma}}. \quad (3.7)$$

Here and in the following set of equations, the symbols  $\gamma$ ,  $\eta$ ,  $\mu$ , and  $\nu$  are used as dummy summation indices.

The stationary distribution of this Markov chain is approximately given by:

$$\pi_{\alpha} = \frac{\sum_{\eta} (\mathbf{Z})_{\alpha,\eta}}{\sum_{\mu,\nu} (\mathbf{Z})_{\mu,\nu}}. \quad (3.8)$$

This can be verified by inserting the proposed stationary distribution into the right side of the eigenvalue equation of Equation (3.5) and expanding the \* product:

$$\begin{aligned} (\vec{\pi} * \mathbf{P})_{\beta} &= \sum_{\alpha} \pi_{\alpha} (\mathbf{P})_{\alpha,\beta} \\ &= \sum_{\alpha} \frac{\left( \sum_{\eta} (\mathbf{Z})_{\alpha,\eta} \right) (\mathbf{Z})_{\alpha,\beta}}{\left( \sum_{\mu,\nu} (\mathbf{Z})_{\mu,\nu} \right) \left( \sum_{\gamma} (\mathbf{Z})_{\alpha,\gamma} \right)} \\ &= \frac{\sum_{\alpha} (\mathbf{Z})_{\alpha,\beta}}{\sum_{\mu,\nu} (\mathbf{Z})_{\mu,\nu}}. \end{aligned} \quad (3.9)$$

The  $\vec{\pi}$  in Equation (3.8) is the stationary distribution if and only if it is equal

to Equation (3.9). Setting aside the change in component from  $\alpha$  to  $\beta$ , these two expressions differ only in the summation index of the transition count matrix  $(\mathbf{Z})_{\alpha,\beta}$  in the numerator. The summation for  $\pi_\beta$  counts the number of observed transitions from state  $\beta$  to any other state, while the summation for  $(\vec{\pi} * \mathbf{P})_\beta$  counts the number of transitions to  $\beta$  from any other state. Since the system must have entered a state  $\beta$  before it can leave it, these two values will be equal, except for possibly the very first and last initial states in observed in the sequence  $\alpha_{\text{obs}}(n\tau)$ . The denominator for both expressions is equal to the total number of observed transitions from and to any state.

This stationary distribution corresponds exactly to the time-independent probability distribution that is commonly used to model respiratory motion. The time-independent distribution is equal to the probability that the system will be in a particular state, independent of time. The maximum likelihood estimate for this distribution is equal to the number of times a state is observed,  $N_\alpha$ , divided by the total number of observations:

$$P_\alpha^{\text{TI}} = \frac{N_\alpha}{\sum_\beta N_\beta}. \quad (3.10)$$

Using similar reasoning as before,  $N_\alpha$  is equal to the total number of observed transitions from  $\alpha$  to any other state. Thus  $N_\alpha = \sum_\beta (\mathbf{Z})_{\alpha,\beta}$ , and therefore the time-independent distribution from Equation (3.10) is equal to the stationary distribution from Equation (3.8). Thus it has been shown that the probability distribution of states in the Markov chain converges to the time-independent

distribution as the number of time intervals  $n \rightarrow \infty$ .

### 3.2.3 Monte Carlo simulation of Markov chains

The distribution of the Markov chain states themselves, while it can be easily calculated using Equation (3.4), may not itself be of interest directly; instead, the distribution of some function of the states may be required. Monte Carlo simulations are often used to estimate distributions of functions of random variables which are difficult to calculate using analytical methods, but for which the underlying random variables follow known distributions. This method can be applied to Markov chains, as the sequence of states  $\alpha(n\tau)$ ,  $n \in \{0, \dots, N\}$  can be readily simulated using the transition matrix  $\mathbf{P}$ .

To do this, the initial state  $\alpha(0)$  must first be randomly sampled from the initial distribution  $P_\alpha(0)$ . The rest of the sequence can be sampled sequentially, from  $n = 1$  to  $n = N$ : if the state at time  $(n - 1)\tau$  is known to be  $\alpha((n - 1)\tau)$ , then the state at time  $n\tau$  can be sampled from the transition matrix  $(\mathbf{P})_{\alpha(n\tau), \alpha((n-1)\tau)}$ , where the second element is held fixed. The distribution of functions of the Markov states can then be estimated by calculating the desired function for a particular sample of the state sequence of the system, and repeating the entire process until a desired level of uncertainty has been reached.

## 3.3 Materials and methods

### 3.3.1 Respiratory Markov chain model and parameters

For the purposes of this study, the system  $\mathcal{S}$  described by the Markov chain model is the patient's breathing pattern, and the Markov state space  $S$  is defined by the combination of the position and velocity of the principal component of motion, the breathing phase, and the current elapsed time in the current breathing phase. The principal component is defined as the component of motion along the axis upon which the entire motion exhibits the greatest variance [79]; for tumour motion due to respiration, this axis lies nearly in the sagittal plane. The position was included since the goal is to have a model which accurately calculates the position distribution. Meanwhile, the inclusion of the velocity helps to distinguish instances in which the patient undergoes faster or slower breathing motion.

Similarly, the breathing phase differentiates between the three broader respiratory phases of inhaling, exhaling, and end-of-exhale, each of which is characterized by specific motion characteristics, as shown in Figure 3.3. Finally, current elapsed time in the current breathing phase was included in the Markov state definition to model a smoother transition between these breathing phases. Specifically, as described in Section 3.2, Markov chains do not have memory: the transition probability depends only on the current state of the system. However, the longer the patient has been in the current breathing phase, the more likely they will transition to next phase. Thus the

current elapsed time in the current breathing phase is introduced to model such behaviour.

The set of possible states  $S$  was defined by the following tunable parameters, each of which takes on positive integer values. First,  $N_l$  is the number of position bins  $l$ , which defines the position resolution of dose calculation in the proposed 4D-VMAT optimization method (see Section 4.2.1.1). Position sub-bins were introduced to differentiate a state which has just entered a particular position bin from a state which is about to leave that bin;  $N_x$  is the number of such position sub-bins per bin.  $N_v$  is the number of velocity bins, and  $N_T$  is the number of bins recording the current elapsed time in the current breathing phase.

Each state  $\alpha$  then corresponds to a unique combination of position bin, and sub-bin, velocity bin, breathing phase, and elapsed time in the current breathing phase. Lookup tables from  $\alpha$  to each of these physical observables were generated, so that probabilities and simulations could be converted from the base Markov states to the observables. For example,  $l_\alpha$  is the position bin corresponding to state  $\alpha$ .

The final variable parameter of the Markov chain is the discrete time interval  $\tau$ , which represents the base time between which transition probabilities can be calculated. For small enough values of  $\tau$ , very few distinct transitions will be observed, limiting the model's use. On the other hand, a respiratory motion model with a large value of  $\tau$  (for example, on the order of the breathing period) would omit or pass over much (or all) of the actual

respiratory motion. Each of the parameters described above were tuned according to the methods presented in Section 3.3.3.

### 3.3.2 Respiratory data acquisition and processing

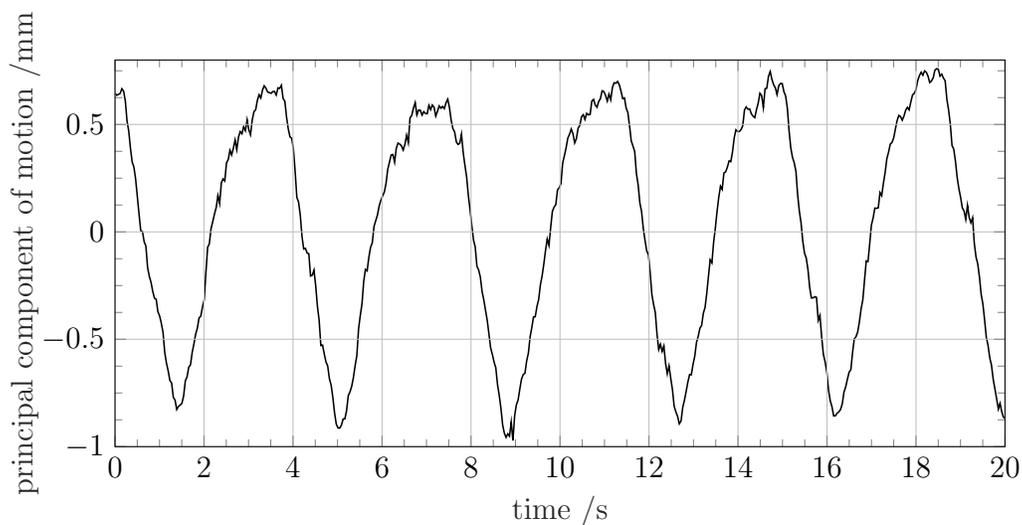
Individual datasets from a database of clinical patient respiratory data acquired during 102 CyberKnife treatment fractions delivered on 32 unique patients at Georgetown University Hospital [80] were used to construct and test the motion model; the database can be found at the following URL: [https://signals.rob.uni-luebeck.de/index.php?title=Signals\\_@\\_ROB](https://signals.rob.uni-luebeck.de/index.php?title=Signals_@_ROB). The data consists of the position of optical surface markers on the patient, sampled at a frequency of  $f_s = 26$  Hz. For the purposes of this study, two treatment fractions delivered on the same patient were selected to develop two unique models; the fractions and respective models are individually referred to as A and B. This was done to compare the models trained on data gathered during two separate fractions. The following procedure of the data processing and model construction was applied separately to the respiratory data from each fraction.

First, a 10 min interval with no baseline drift was selected to train and test the model: the first 5 min were used for training, and the following 5 min for testing. This interval was chosen as it is on the order of VMAT delivery times, therefore the range of patient motion occurring in this time interval will be representative of motion during delivery. The principal component of the three-dimensional marker motion, which was included with the full

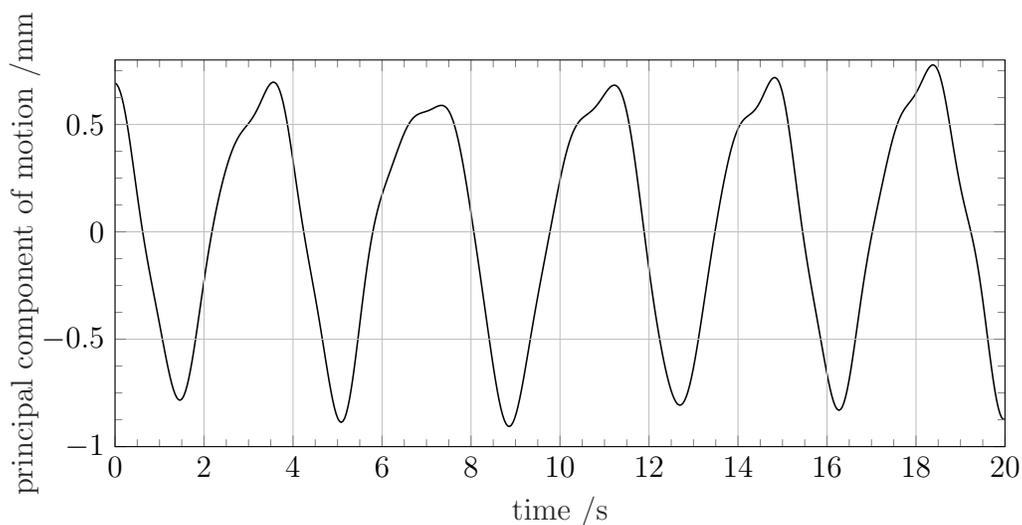
data, was extracted and used throughout this study as a surrogate for the actual tumour position (see Section 3.5 for more discussion on this). The time interval for the Markov chain was taken to be  $\tau = r\tau_s$ , with  $r$  a variable integer factor and  $\tau_s = \frac{1}{f_s}$ . The principal component signal was filtered, using a low-pass Chebyshev Type I infinite impulse response of order 8 [81]. The cutoff frequency of this filter was  $f_0 = \frac{f_s}{2} \cdot \frac{0.8}{r}$ , with  $\frac{f_s}{2}$  being the Nyquist frequency of the original signal.

This was done when training the model to filter out components with frequency larger than the Nyquist frequency corresponding to the Markov chain time interval  $\tau$ . Examples of the unfiltered and filtered respiratory signals are shown in Figure 3.2. The Markov chain time interval factor  $r$  was varied in order to find the model which fit best to the testing data. To ensure that the testing conditions remained constant while varying the model parameters, the testing data used to evaluate the model was not filtered.

The velocity of the principal component was then calculated on the filtered signal using the finite difference estimate at the sample points. Finally, the respiratory motion signals were segmented into breathing phases, as shown in Figure 3.3, which reflect a broad type of respiratory motion: inhaling, exhaling, end-of-exhale, or irregular breathing [82]. In the data used in this study, inhalation is characterized by a steep decrease in the principal component, exhalation by a steep increase in the principal component, and end of exhale by a shallow increase. Each respiratory motion data point was classified to one of the states defining the Markov chain system according to



(a) Nonfiltered respiratory signal.



(b) Filtered respiratory signal.

Figure 3.2: Sample of the respiratory motion data used to train the model, consisting of the principal component of the motion of an optical surface marker. Both the original nonfiltered and filtered data are shown, the latter of which is used for training. A low-pass Chebyshev Type I infinite impulse response of order 8, with a cutoff frequency of 1.1531 Hz ( $r = 9$ ), was used.

its position, velocity, breathing phase, and current elapsed time in the current breathing phase. The position binning and breathing phase segmentation is also illustrated in Figure 3.3. The final result is an observed sequence of states  $\alpha_{\text{obs}}(n\tau_s) \in \mathcal{S}$ ,  $n \in \{0, \dots, N\}$ , with the time interval given by  $\tau_s = \frac{1}{f_s}$ ,  $f_s = 26$  Hz, and  $N\tau_s = 5$  min.

The Markov chain transition matrix  $\mathbf{P}$  was estimated from the training data using MLE, as described in Section 3.2.2. Any state not reached at all by the observed sequence  $\alpha_{\text{obs}}(n\tau_s)$  was eliminated from the state space to guarantee irreducibility and recurrence of the Markov chain. Since the time interval for the Markov chain is  $\tau = r\tau_s$ , the transition count matrix  $\mathbf{Z}$  was defined by summing transitions observed in the sequence  $\alpha_{\text{obs}}(n\tau_s)$  from time  $(n-r)\tau_s$  to  $n\tau_s$  for all  $n \leq r$ , rather than the original definition which counted transitions from time  $(n-1)\tau_s$  to  $n\tau_s$ . The final transition matrix  $\mathbf{P}$ , which was calculated using Equation (3.7), thus describes the probability to transition between states in the time interval  $\tau = r\tau_s$ , as desired.

### 3.3.3 Model evaluation and selection

The optimal combination of parameter values  $N_x$ ,  $N_v$ ,  $N_T$ , and  $r$  was determined for both models A and B separately by varying the value of each, constructing a corresponding Markov chain model for the given set, and selecting the combination whose model fit best to the testing data. As will be discussed in Section 4.2.1.1, in the application of 4D treatment plan optimization, the ideal respiratory motion model is one which can accurately

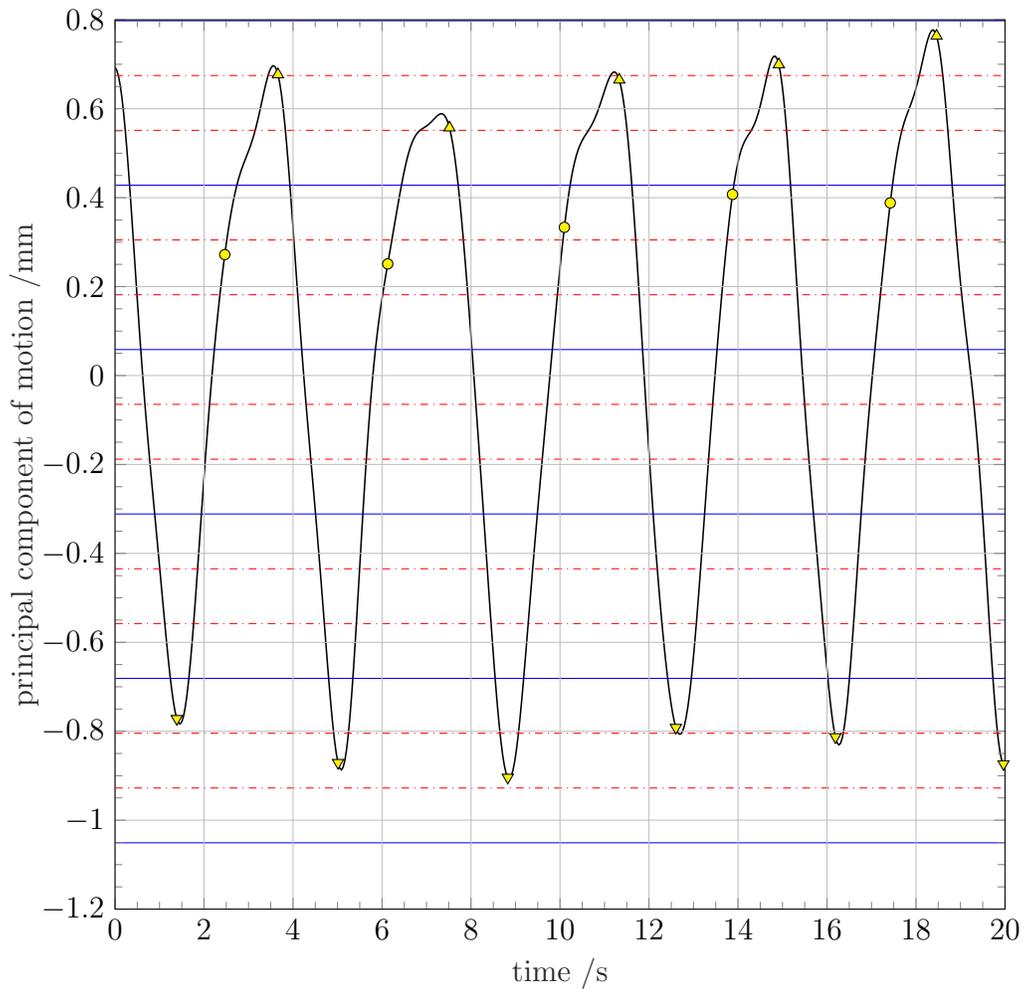


Figure 3.3: Position binning and breathing phase segmentation of the filtered respiratory data. The  $N_l = 5$  position bins and the  $N_x = 3$  sub-bins are demarcated by solid blue lines and dashed red lines, respectively. The bin number increases from  $l = 1$  to  $l = 5$  from the top of the graph to the bottom. The yellow marks indicate the beginning of a breathing phase: inhale (upward pointing triangle), exhale (downward pointing triangle), and end-of-exhale (circle). The data was also binned according to the velocity and the elapsed time in the current breathing phase (not shown).

Table 3.1: Trial Markov chain time interval factor values  $r$  used to search for the optimal Markov chain respiratory motion model. The corresponding Markov chain time interval,  $\tau = r\tau_s$  is shown, along with the inverse  $\frac{1}{\tau}$ .

Markov chain time interval factor value $r$	Markov chain time interval $\tau/s$	$\frac{1}{\tau}/\text{Hz}$
1	0.0385	25.9439
2	0.0771	12.9719
3	0.1156	8.6480
4	0.1542	6.4860
9	0.3869	2.8827
26	1.0022	0.9978

calculate the probability distribution of the position bin at any arbitrary time after treatment has begun.  $N_x$ ,  $N_v$ , and  $N_T$  were varied from integer values 1 to 10, while  $r$  was given six trial values shown in Table 3.1. The number of position bins  $N_l$ , on the other hand, was kept fixed at 5, since this is the number of position bins which was used for the 4D dose calculation and plan optimization (see Sections 4.2.5.2 and 4.3.1.2).

To quantify the probability distribution of the position bin, the position histogram was defined by tallying the position bin of the respiratory motion at some time  $\Delta t$  after every data point which falls in the triggering position bin (taken to be bin  $l = 1$ , maximum exhale). This process of constructing the observed position histogram  $O_{\Delta t}^l$  from the 5 min of testing respiratory data can be seen in Figure 3.4. Note that any instances of the triggering bin

after the critical time  $t_c = 5 \text{ min} - \Delta t$  will have no contribution, since the time at which the position bin would be tallied exceeds the total duration of 5 min.

The position histogram expected to be observed, given the Markov chain model, can be calculated as the probability of observing position bin  $l$  at time  $\Delta t$  after the triggering position bin, multiplied by the expected number of occurrences of the triggering bin  $l = 1$  in a sequence of respiratory motion of duration  $5 \text{ min} - \Delta t$  ( $N_{5 \text{ min} - \Delta t}^1$ ):

$$E_{\Delta t}^l = \Pr(L(t + \Delta t) = l | L(t) = 1) \cdot N_{5 \text{ min} - \Delta t}^1. \quad (3.11)$$

Here,  $t$  is arbitrary and can be set to 0 for simplicity, while  $L(t)$  is the sequence of observed position bins. The required duration of the sequence is  $5 \text{ min} - \Delta t$ , reflecting the critical time noted when discussing the observed histogram. The expression  $\Pr(L(t + \Delta t) = l | L(t) = 1)$  represents the probability to observe position bin  $l$  at time  $t + \Delta t$ , given that the position bin at time  $t$  is 1.

The probability  $\Pr(L(\Delta t) = l | L(0) = 1)$  can be calculated from the transition matrix  $\mathbf{P}$  using Equation (3.4) with some adjustments. First, the condition  $L(0) = 1$  determines the initial distribution of states, under the additional assumption that the states  $\alpha$  with  $l_\alpha = 1$  are uniformly distributed. Specifically, the elements of the initial probability distribution  $\vec{P}(0)$  are given

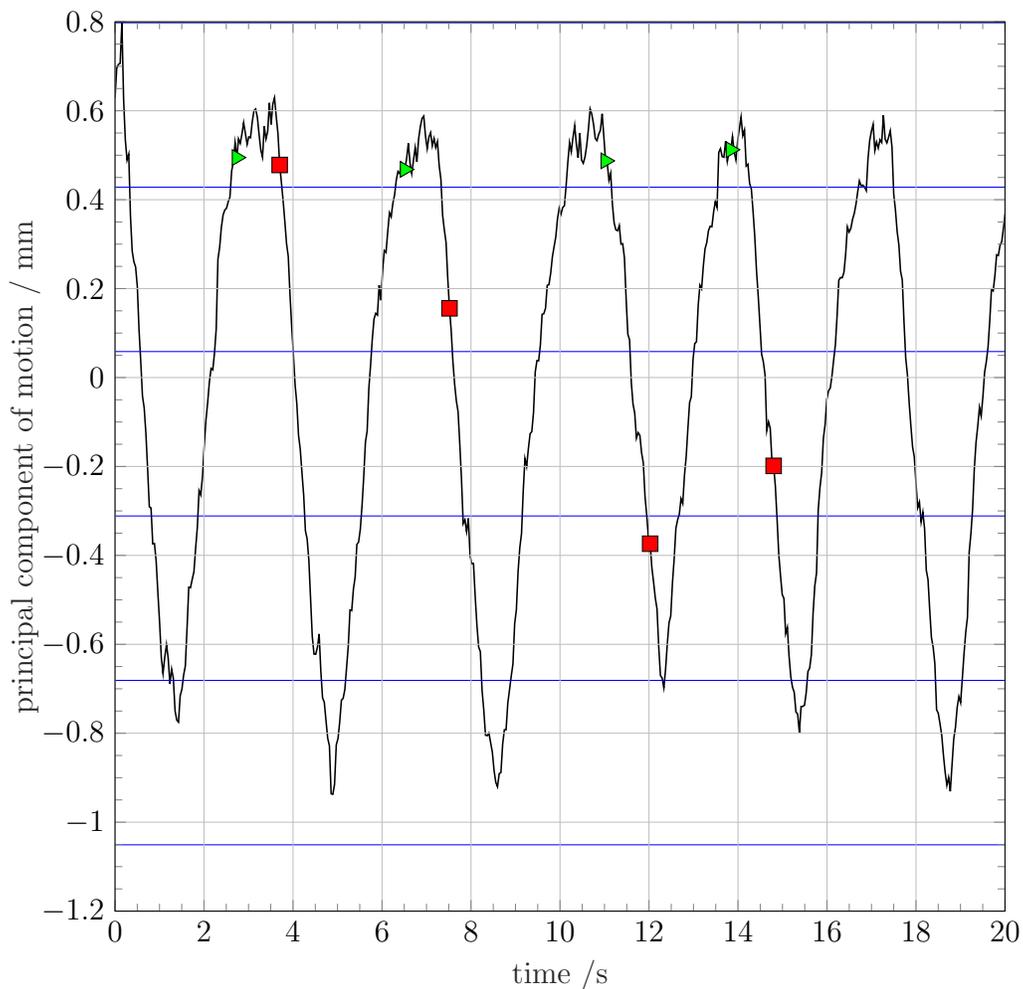


Figure 3.4: The position histogram for time  $\Delta t = 1$  s constructed from the testing respiratory data (black line). The blue lines indicate the boundaries of each position bin, as in Figure 3.3. For each time point at which the triggering position bin ( $l = 1$ , maximum exhale) is observed in the respiratory data, the position bin  $l$  which is observed  $\Delta t = 1$  s later is tallied. Four examples of the triggering points are shown indicated by the green triangles, along with the corresponding delayed points indicated by the red squares, thus tallying the bins  $l = 1, 2, 4, 3$  (in order).

by:

$$P_\alpha(0) = \begin{cases} \frac{1}{N_{l_\alpha=1}} & \text{for } l_\alpha = 1 \\ 0 & \text{for } l_\alpha \neq 1, \end{cases} \quad (3.12)$$

where  $N_{l_\alpha=1}$  is the total number of states  $\alpha$  for which  $l_\alpha = 1$ . Therefore the probability distribution at time  $\Delta t$  is given by:

$$\vec{P}(\Delta t) = \vec{P}(0) * \mathbf{P}^n, \quad (3.13)$$

with  $n = \frac{\Delta t}{\tau}$ . The fraction  $\frac{\Delta t}{\tau}$  is not an integer for arbitrary values of  $\Delta t$ . Non-integer powers of matrices are not well-defined for non-invertible matrices [83]; in general, the transition matrix  $\mathbf{P}$  is not guaranteed to be invertible, therefore the number of intervals  $n$  of length  $\tau$  is taken to be the whole integer nearest to this fraction. The model therefore makes the assumption that any changes in probability within time intervals smaller than  $\tau$  are not significant; this assumption will be tested for various values of  $\tau$  by performing this particular model evaluation. The probability to observe position bin  $l$  is finally obtained by summing the probability of observing any state  $\beta$  whose corresponding position bin  $l_\beta$  matches:

$$\Pr(L(\Delta t) = l | L(0) = 1) = \sum_{\beta | l_\beta = l} P_\beta(\Delta t). \quad (3.14)$$

The expected number of trajectory data points which fall in the triggering position bin,  $N_{5_{\min-\Delta t}}^1$  from Equation (3.11), cannot be directly computed

from the transition matrix  $\mathbf{P}$ . Instead, the Monte Carlo method, as described in Section 3.2.3, was used to estimate this quantity. A total of 1000 sample sequences of duration  $5 \text{ min} - \Delta t$  were generated from the Markov chain model, whereupon each sequence of states was converted to a sequence of position bins using the lookup table  $l_\alpha$ . The number of occurrences of position bin 1 was counted for each sample, and the mean across all 1000 samples was assigned to the quantity  $N_{5 \text{ min} - \Delta t}^1$ . Example observed and expected position histograms are shown in Figure 3.5.

The observed and expected position histograms,  $O_{\Delta t}^l$  and  $E_{\Delta t}^l$ , were compared using the  $\chi^2$  statistic (obs here stands for observed):

$$\chi_{\text{obs}, \Delta t}^2 = \sum_l \frac{(O_{\Delta t}^l - E_{\Delta t}^l)^2}{E_{\Delta t}^l}. \quad (3.15)$$

Test statistics such as  $\chi^2$  can be used to determine the goodness of fit of a statistical model with observational data. For example, the  $p$ -value is defined as the probability of obtaining a test statistic at least as extreme as the one actually observed, under the hypothesis that the given model is correct. In this case, since a larger value of  $\chi^2$  indicates greater disagreement of the model with the data, the  $p$ -value for a given  $\chi_{\text{obs}, \Delta t}^2$  can be calculated from the distribution of  $\chi_{\Delta t}^2$  as:

$$p_{\Delta t} = \Pr(\chi_{\Delta t}^2 \geq \chi_{\text{obs}, \Delta t}^2). \quad (3.16)$$

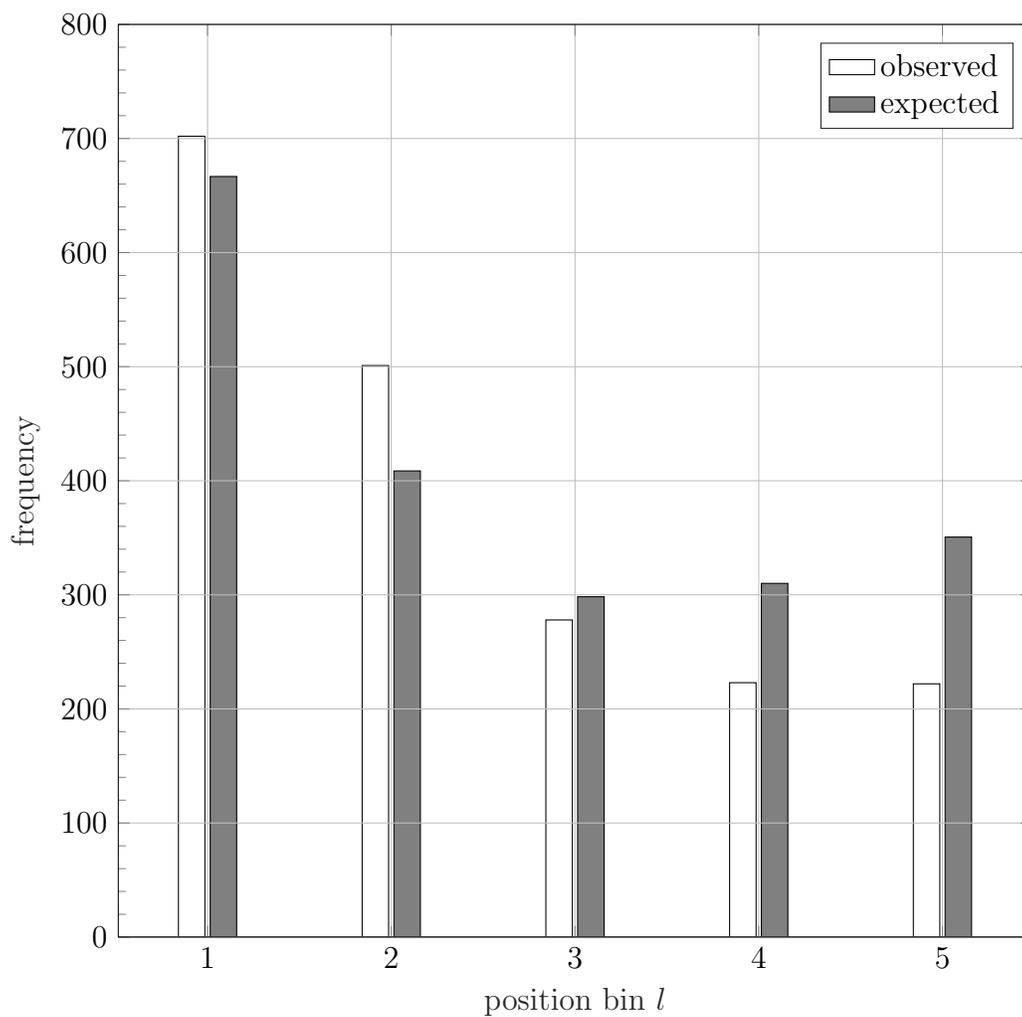


Figure 3.5: The position bin histograms actually observed,  $O_{\Delta t}^l$ , and expected,  $E_{\Delta t}^l$ , from the Markov chain respiratory motion model, for a particular time of  $\Delta t = 30$  s after the triggering bin is observed. The model used to calculate the expected histogram was constructed with the optimal combination of parameters for fraction A, as found in Section 3.4.1.

$\chi_{\Delta t}^2$  follows the actual  $\chi^2$ -distribution with  $N_l$  degrees of freedom if the observed data  $O_{\Delta t}^l$  are independent and Gaussian distributed with both mean and variance  $E_{\Delta t}^l$  (i.e., Poisson variables with a large mean). However, neither of these conditions are guaranteed to be met in this case; instead, the actual distribution of  $\chi_{\Delta t}^2$ , and thus  $p_{\Delta t}$ , can be estimated using Monte Carlo simulations. Just as how  $N_{5\text{min}-\Delta t}^1$  was estimated, a total of 1000 sample sequences of the position bin, each of duration 5 min, were generated from the Markov chain model. The position histogram  $O_{\text{MC},\Delta t}^l$  (MC for Monte Carlo) was constructed for each sequence, from which  $\chi_{\text{MC},\Delta t}^2$  could be calculated using Equation (3.15). The model hypothesis distribution of  $\chi_{\Delta t}^2$  was then estimated from these 1000 samples, from which  $p_{\Delta t}$  could be calculated as the fraction of samples which had  $\chi_{\text{MC},\Delta t}^2 \geq \chi_{\text{obs},\Delta t}^2$ . By way of example, the distribution of  $\chi_{30\text{s}}^2$  for the particular combination of parameters in Section 3.4.1, as estimated from the Monte Carlo samples, is shown in Figure 3.6.

The observed and expected position histograms, along with the corresponding  $\chi_{\text{obs},\Delta t}^2$  and  $p_{\Delta t}$  values, were calculated for a range of values of the time interval  $\Delta t$ , corresponding to transition times required to compute probabilities for the 4D treatment plan optimization. In this study, these were taken to be integer multiples of the time to update the MLC leaf positions, 40 ms (see Section 4.2.5.3), up to a maximum value given by the convergence time  $\Delta t_{\text{converge}}$  of the position distribution.  $E_{\Delta t}^l$  follows the same stationary

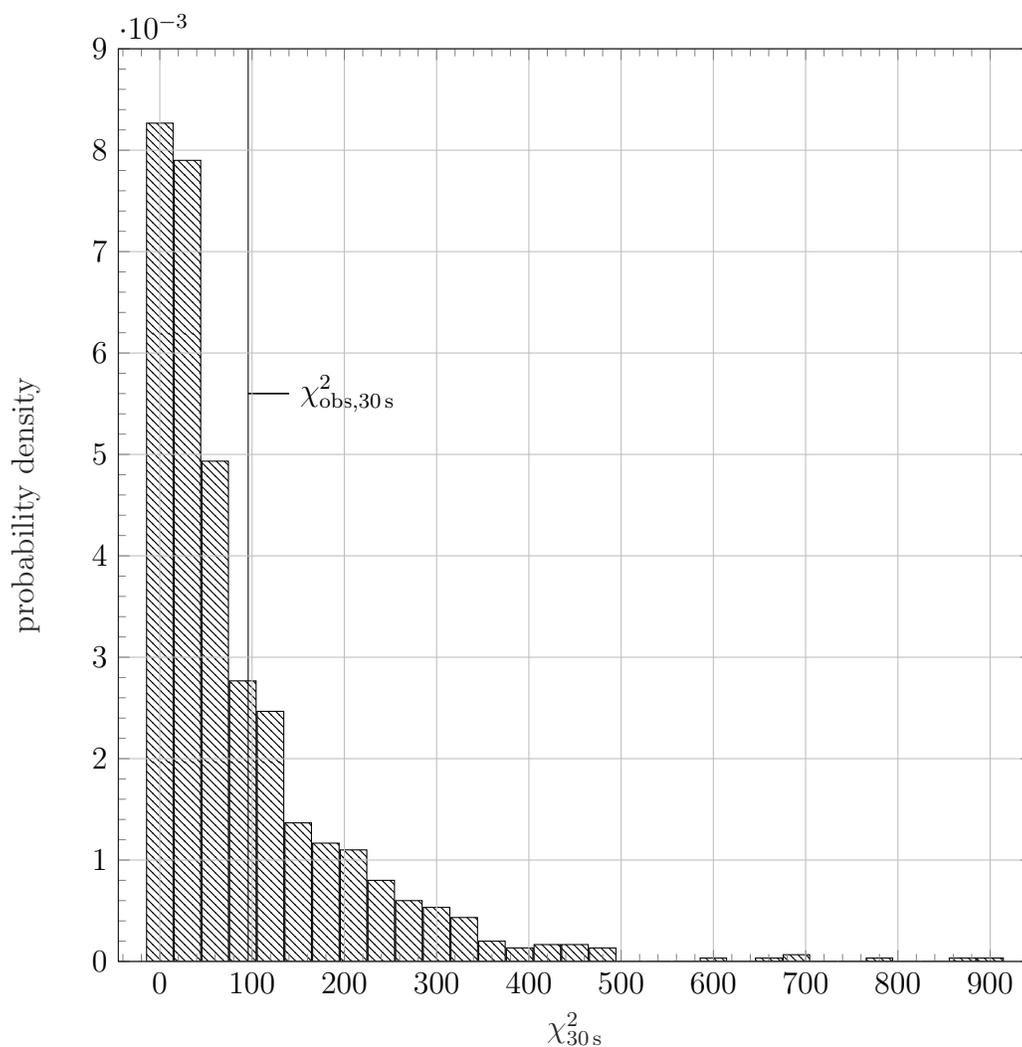


Figure 3.6: Distribution of  $\chi_{30s}^2$  estimated from a Monte Carlo simulation of the Markov chain respiratory motion model. The model used to perform the simulations was constructed with the optimal combination of parameters for fraction A, as found in Section 3.4.1. The distribution is shown as a histogram, scaled to represent the estimated probability density function. The vertical line indicates the actual value  $\chi_{obs,30s}^2 = 95.774$ , corresponding to the histograms in Figure 3.5. The corresponding  $p$ -value,  $p_{30s} = 0.345$ , is equal to the area to the right of the line according to Equation (3.16).

distribution for any transition time greater than  $\Delta t_{\text{converge}}$ , thus it is unnecessary to continue to calculate  $\chi_{\text{obs},\Delta t}^2$ . This convergence time was found to vary with the parameters used, therefore the maximal value of  $\Delta t_{\text{converge}} = 90$  s (for fraction A models) or  $\Delta t_{\text{converge}} = 110$  s (for fraction B models), determined according to Section 3.2.1, was used throughout this study.

To quantify the overall agreement between the observed and expected position bin histograms, the  $\chi_{\text{obs},\Delta t}^2$  were summed over all values of  $\Delta t$ :

$$\chi_{\text{obs},\Sigma}^2 = \sum_{\Delta t} \chi_{\text{obs},\Delta t}^2. \quad (3.17)$$

The corresponding  $p$ -value was estimated from the distribution of  $\chi_{\text{MC},\Sigma}^2 = \sum_{\Delta t} \chi_{\text{MC},\Delta t}^2$ , as the fraction of Monte Carlo samples for which the  $\chi_{\text{MC},\Sigma}^2$  was larger than the observed  $\chi_{\text{obs},\Sigma}^2$ . The optimal combination of parameter values  $N_x$ ,  $N_v$ ,  $N_T$ , and  $r$  was deemed to be the one which yielded the largest value of  $p_\Sigma$ .

In addition to comparing the model and testing data from the same fraction, intercomparisons were performed to determine if the models trained on the data originating from the two fractions were interchangeable. Specifically, the predictions of the fraction A model were compared to the observations in the fraction B testing data, and vice versa.

## 3.4 Results

### 3.4.1 Fraction A

The largest value of  $p_\Sigma$  for the fraction A data was determined to be equal to 0.653, at the combination of parameters given by  $N_x = 1$ ,  $N_v = 2$ ,  $N_T = 6$ , and  $r = 9$  (Markov chain time interval  $\tau = 386.9$  ms). The full distribution of  $\chi_\Sigma^2$  estimated from the Monte Carlo sampling described in Section 3.3.3 is shown in Figure 3.7. The large value of  $p_\Sigma$  demonstrates the overall excellent agreement between the position histograms observed in the testing data and expected by the model over the range of values  $\Delta t$ .

Select colour images of  $p_\Sigma$ , taken to be two-dimensional slices in the four-dimensional parameter space of  $N_x$ ,  $N_v$ ,  $N_T$ , and  $r$ , are shown in Figures 3.8 to 3.10. Each slice intersects the point of the largest  $p_\Sigma$  value,  $N_x = 1$ ,  $N_v = 2$ ,  $N_T = 6$ , and  $r = 9$ .

The agreement of the position bin histograms at particular values of  $\Delta t$  can be seen in the graphs of  $\chi_{\text{obs},\Delta t}^2$  and  $p_{\Delta t}$  as a function of  $\Delta t$  in Figure 3.11. For values of  $\Delta t$  from 0 s to 50 s, both  $\chi_{\text{obs},\Delta t}^2$  and  $p_{\Delta t}$  are seen to oscillate from values indicating good to bad agreement. This is due to the differences in time resolution between the observed and expected histograms: while  $O_{\Delta t}^l$  can vary at the 40 ms resolution of the  $\Delta t$ ,  $E_{\Delta t}^l$  is constant within each Markov chain time interval,  $\tau = 386.9$  ms. Thus there can be significant variations in the agreement between these two histograms within time intervals on the order of 386.9 ms, as seen. Despite this, the observed and expected histograms agree

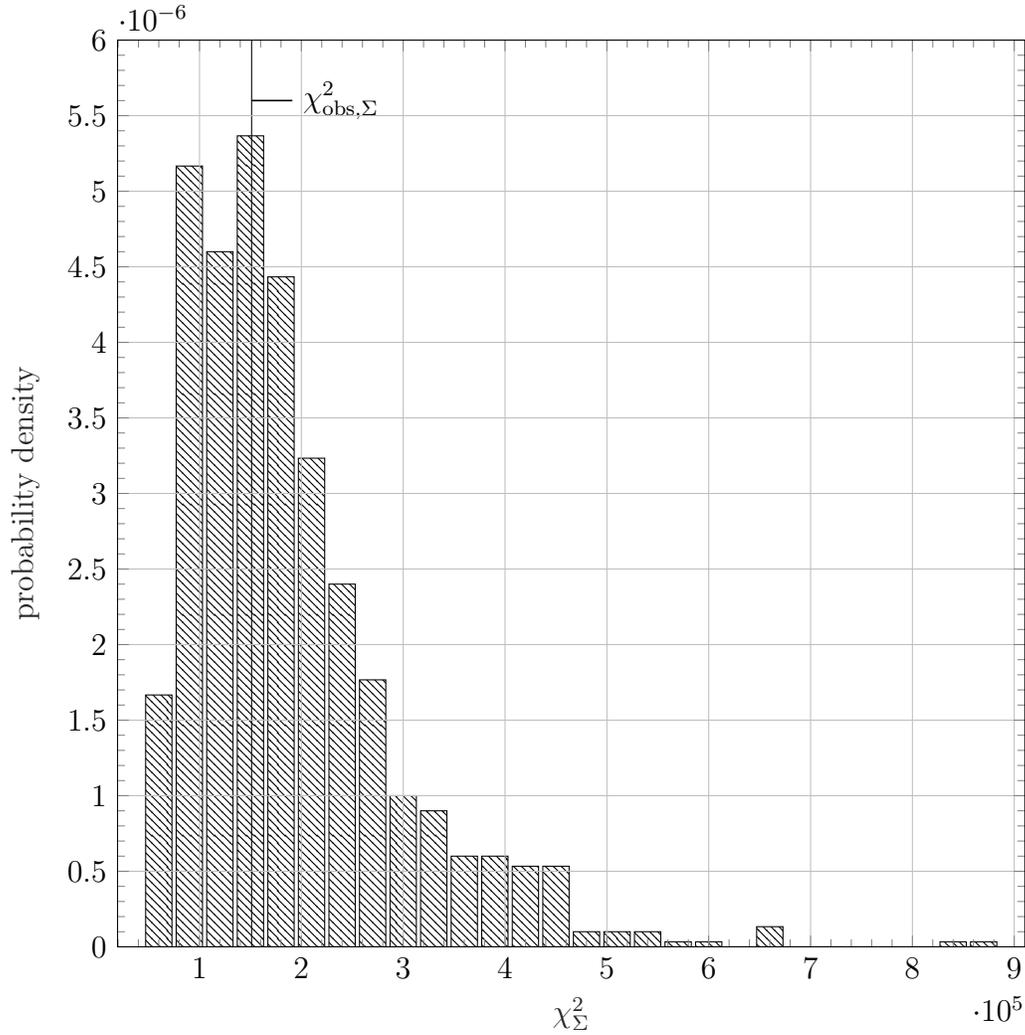


Figure 3.7: Distribution of  $\chi^2_\Sigma$  estimated from a Monte Carlo simulation of the Markov chain respiratory motion model for fraction A. The model used to perform the simulations was constructed with the optimal combination of parameters, as found in Section 3.4.1. The distribution is shown as a histogram, scaled to represent the estimated probability density function. The vertical line indicates the actual value  $\chi^2_{\text{obs},\Sigma} = 151\,148$ . The corresponding  $p$ -value,  $p_\Sigma = 0.653$ , is equal to the area to the right of the line.

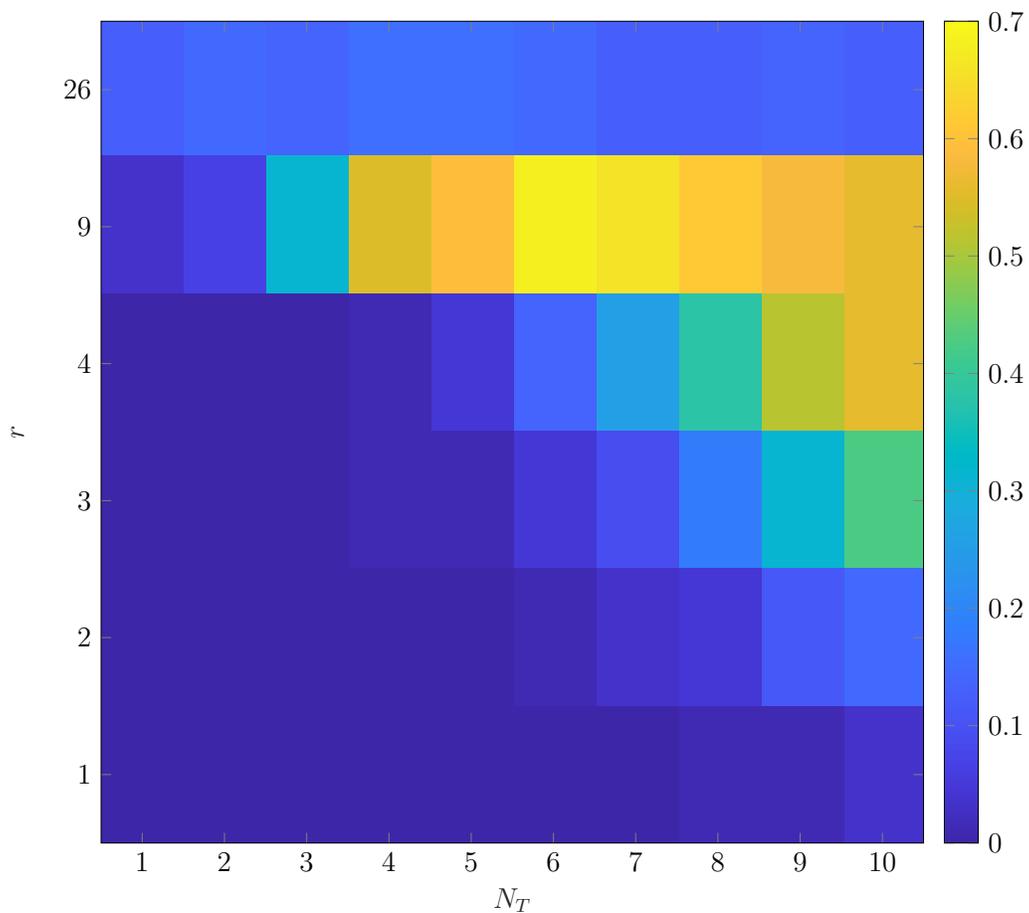


Figure 3.8: Image of  $p_\Sigma$  as a function of the number of bins recording the current elapsed time in the current breathing phase  $N_T$ , and the Markov chain time interval factor  $r$ . The remaining parameters are fixed at  $N_x = 1$ ,  $N_v = 2$ .

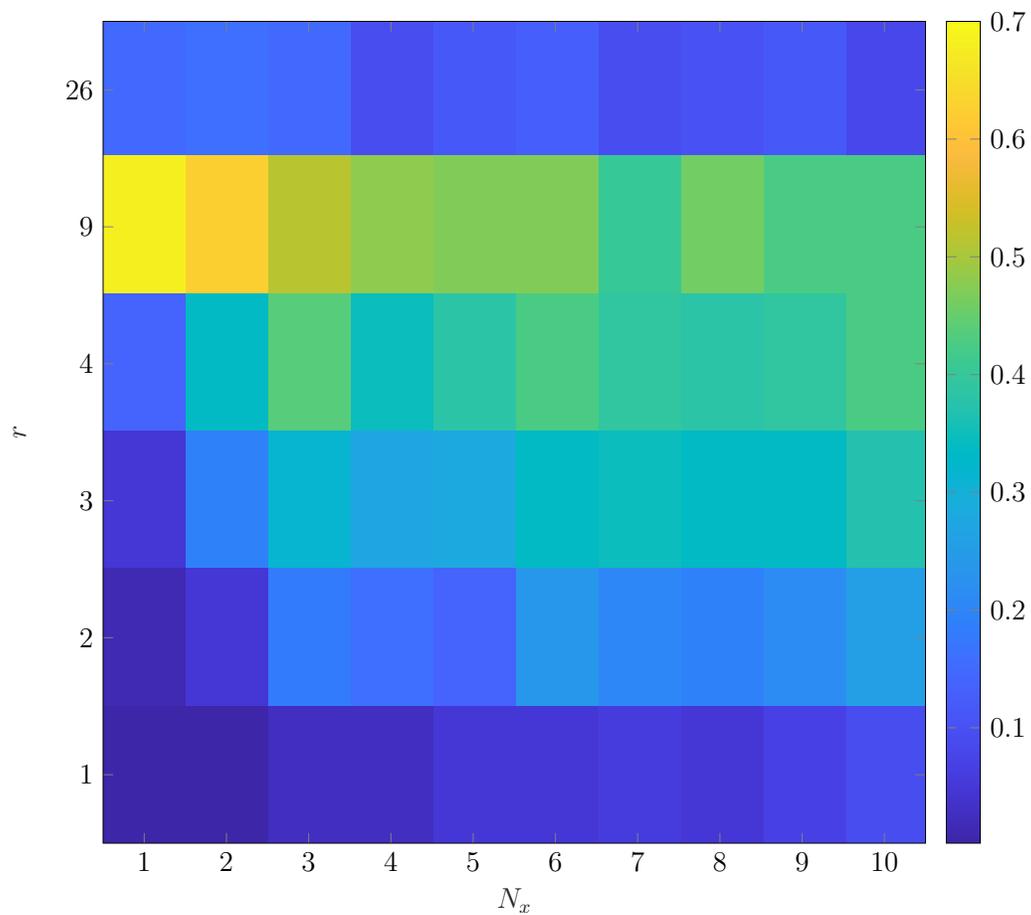


Figure 3.9: Image of  $p_\Sigma$  as a function of the number of position sub-bins per bin  $N_x$ , and the Markov chain time interval factor  $r$ . The remaining parameters are fixed at  $N_v = 2$ ,  $N_T = 6$ .

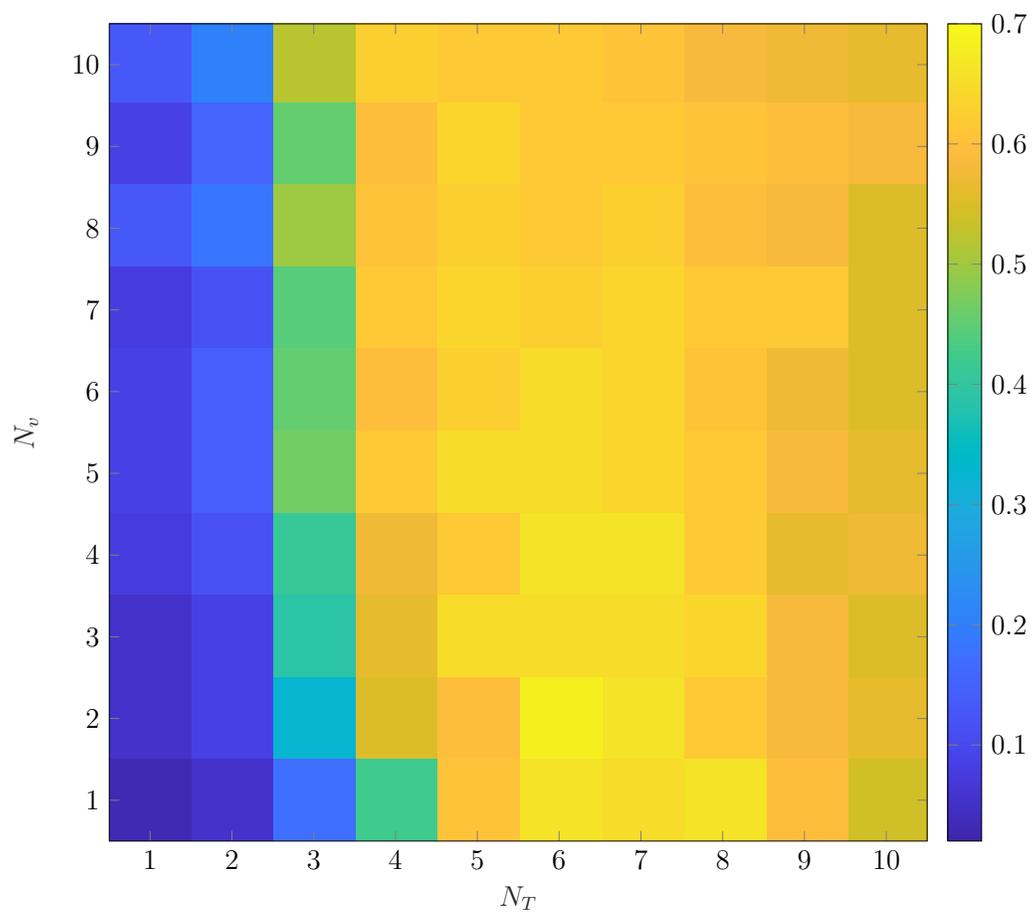


Figure 3.10: Image of  $p_\Sigma$  as a function of the number of bins recording the current elapsed time in the current breathing phase  $N_T$ , and the number of velocity bins  $N_v$ . The remaining parameters are fixed at  $N_x = 1$ ,  $r = 9$ .

well with each other when considering the entire range of  $\Delta t$ ; in particular, the agreement is excellent past  $\Delta t = 50$  s, with  $p_\Sigma$  staying above about 0.4.

The histograms are also shown directly as colour images for comparison in Figure 3.12. The periodic behaviour of the respiratory motion modelled by the Markov chain is shown in these images, as the peak of the distribution of the position bin is shown to oscillate with a similar period from  $l = 1$  to 5 for both observed and expected histograms. Further, the convergence of the position distribution to the stationary distribution as described in Section 3.2.1 is seen in both histograms.

A 20 s sample sequence generated from the Markov chain is shown in comparison to the observed sequence in Figure 3.13. Both sequences are expressed as the principal component as a function of time. For the sample sequence, this was accomplished by first transforming the Markov state sequence  $\alpha(n\tau)$  to the respective position sub-bin  $x(n\tau)$  using the lookup table  $x_\alpha$ , and then converting each sub-bin in the sequence into the value of the principal component which it represents.

### 3.4.2 Fraction B

The largest value of  $p_\Sigma$  for the fraction B data was determined to be equal to 0.126, at the combination of parameters given by  $N_x = 10$ ,  $N_v = 10$ ,  $N_T = 10$ , and  $r = 4$  (Markov chain time interval  $\tau = 154.2$  ms). The smaller value of  $p_\Sigma$  indicates that the overall agreement between the position histograms observed in the testing data and expected by the model for fraction B is

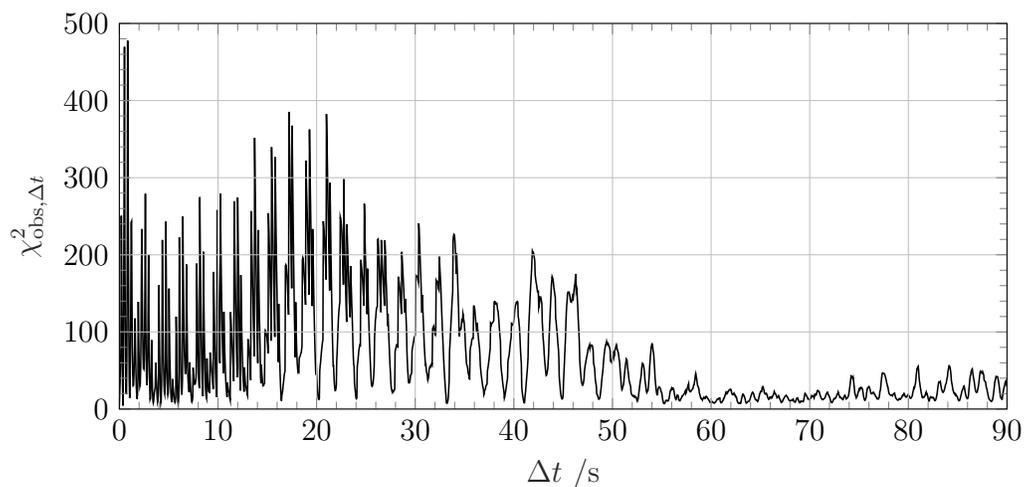
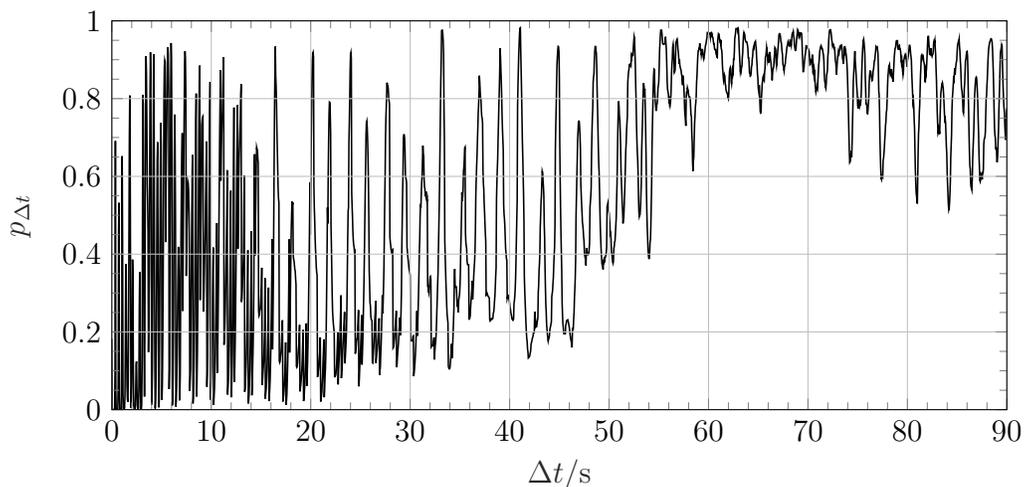
(a) Graph of  $\chi_{\text{obs}, \Delta t}^2$ .(b) Graph of  $p_{\Delta t}$ .

Figure 3.11: Graphs of  $\chi_{\text{obs}, \Delta t}^2$  and the corresponding  $p_{\Delta t}$ , both as functions of  $\Delta t$ . The model used to calculate these two was constructed with the optimal combination of parameters for fraction A, as found in Section 3.4.1.

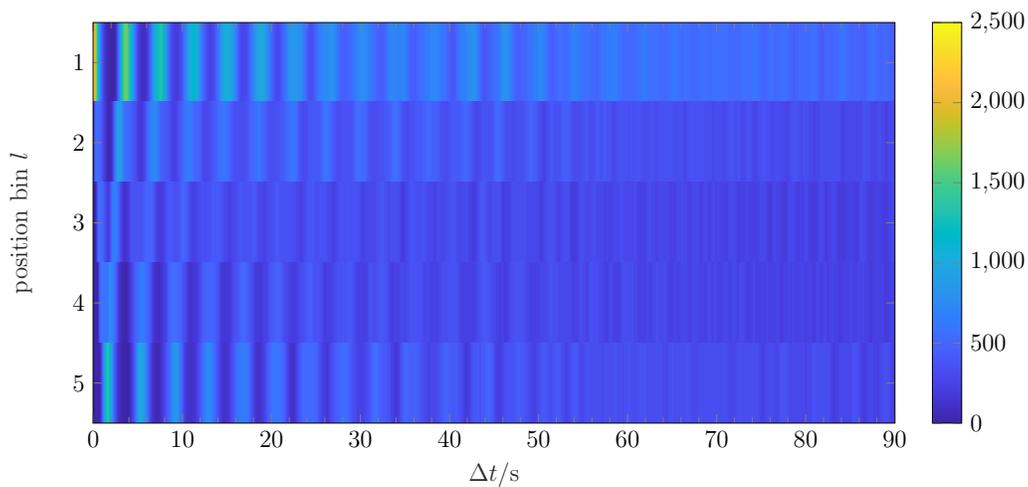
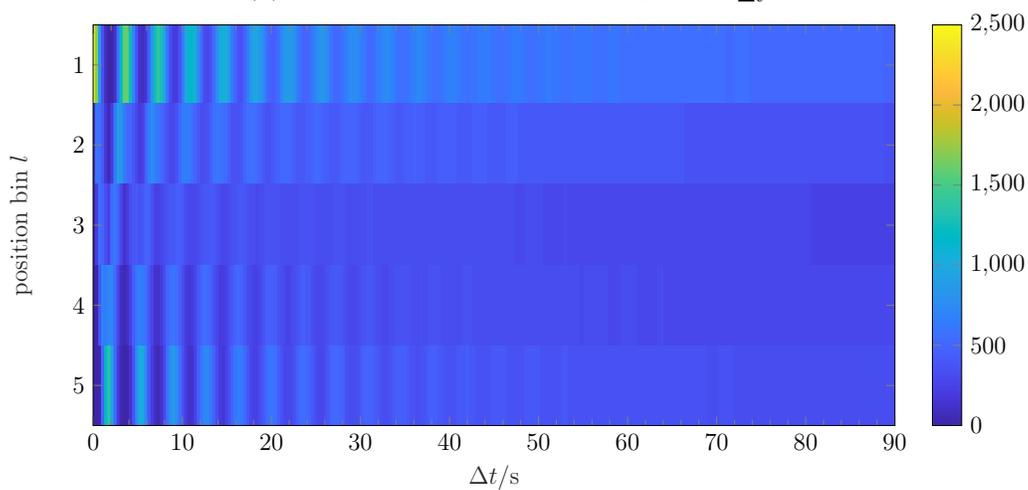
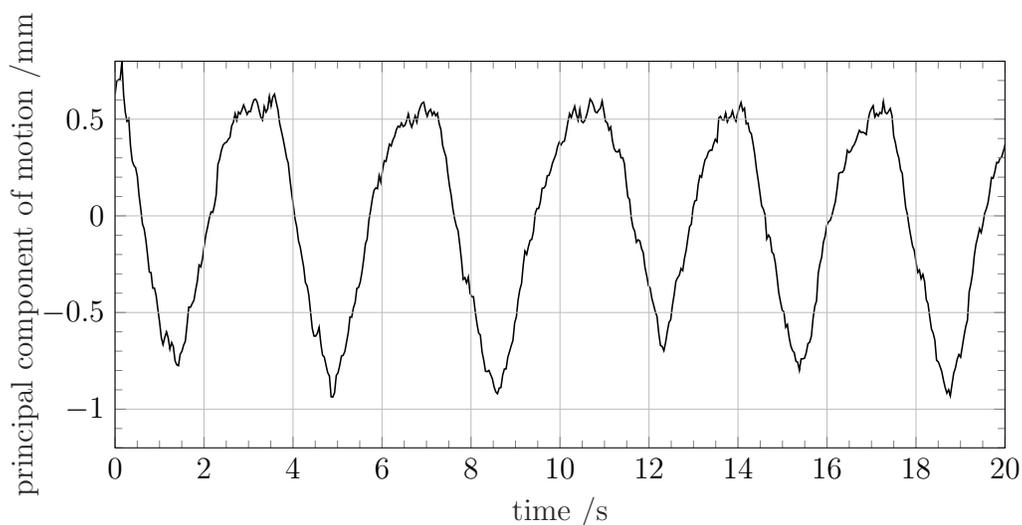
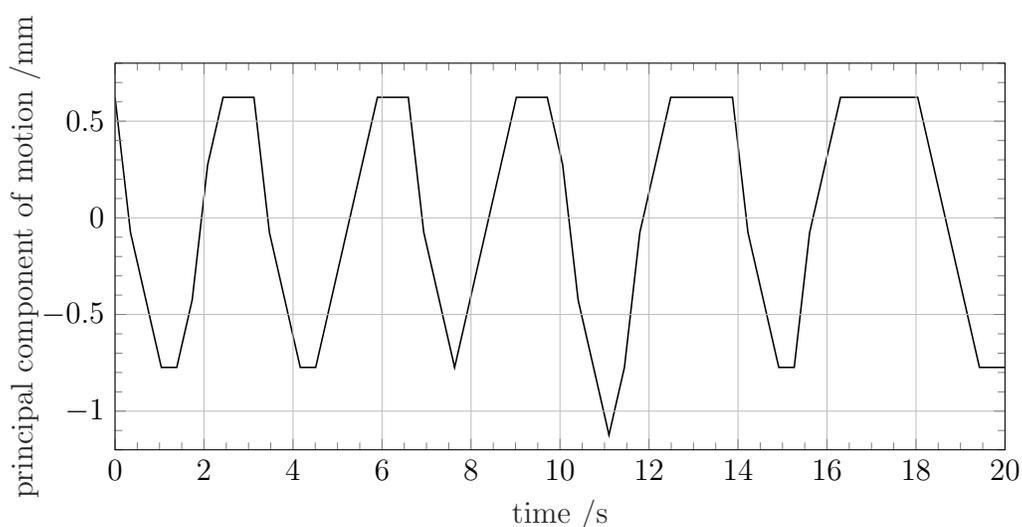
(a) Observed position bin histogram  $O_{\Delta t}^l$ .(b) Expected position bin histogram  $E_{\Delta t}^l$ .

Figure 3.12: Images of the observed and expected position bin histograms, both as functions of  $\Delta t$ . The model used to calculate these two was constructed with the optimal combination of parameters for fraction A, as found in Section 3.4.1.



(a) Observed sequence.



(b) Sample sequence.

Figure 3.13: Observed and simulated sequence of respiratory motion as a function of time. Both sequences are represented as the principal component of motion; this was done for the Markov sample using the position sub-bin lookup table  $x_\alpha$ , and then converting the position sub-bin index to the corresponding value of the principal component of motion. The model used to generate the sample was constructed with the optimal combination of parameters for fraction A, as found in Section 3.4.1, thus the reason for the low apparent resolution in the graph of the sample sequences: there are only  $N_l \cdot N_x = 5$  total position sub-bins in this particular model.

lessened compared to the corresponding agreement for fraction A, although can still be considered to be a good fit to the data.

This is further investigated in the graphs of  $\chi_{\text{obs},\Delta t}^2$  and  $p_{\Delta t}$  as a function of  $\Delta t$ , shown in Figure 3.14. As seen in the fraction A model, for values of  $\Delta t$  from 0 s to 50 s, both  $\chi_{\text{obs},\Delta t}^2$  and  $p_{\Delta t}$  are seen to oscillate from values indicating good to bad agreement. However, for values of  $\Delta t$  exceeding 50 s  $p_{\Sigma}$  never decreases below 0.25, indicating good overall agreement as the model converges to the stationary distribution.

Both the observed and expected histograms are shown in Figure 3.15. As in the corresponding images for fraction A model, both the periodic behaviour of the respiratory motion, as well as the convergence of the position distribution to the stationary distribution, are correctly modelled by the Markov chain. It should be noted however that the model prediction for the convergence is much quicker than what is actually observed, as the image in Figure 3.15b is smoother than that of Figure 3.15a. This is reflected in the much higher values of  $\chi_{\text{obs},\Delta t}^2$  for this fraction model than those observed for fraction A, especially in the first 50 s.

A 20 s sample sequence generated from the Markov chain is shown in comparison to the observed sequence in Figure 3.16, both expressed as the principal component as a function of time. The sample sequence shown here is smoother than that generated from the fraction A model in Figure 3.13, both because of the greater number of position sub-bins and the smaller Markov chain interval  $\tau$ .

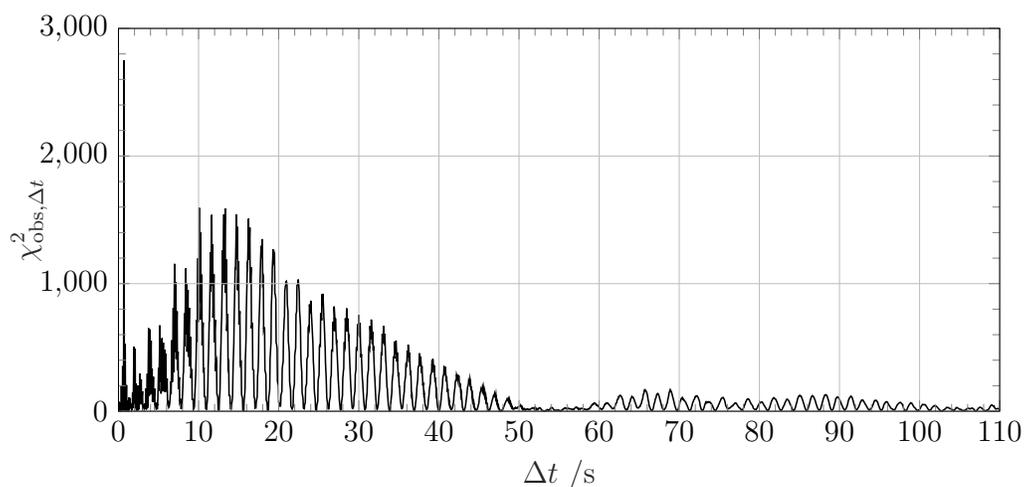
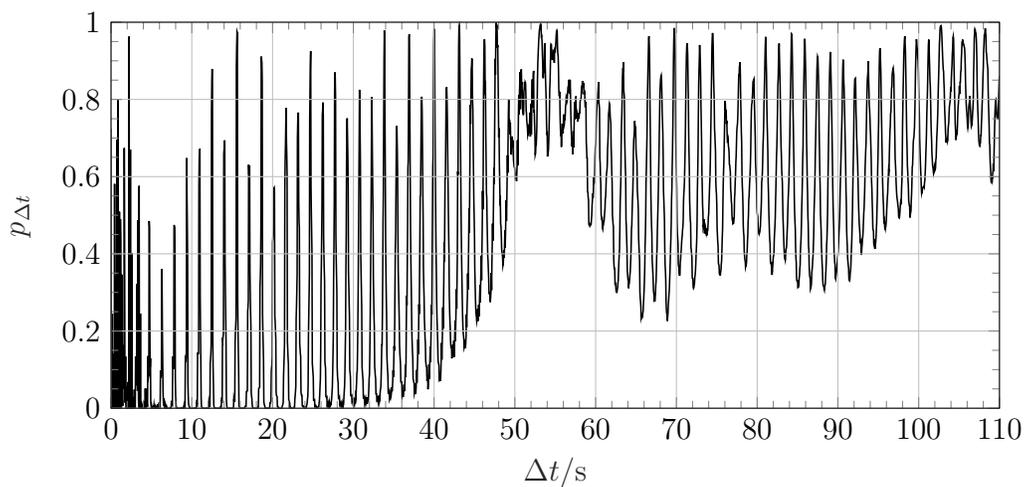
(a) Graph of  $\chi_{obs, \Delta t}^2$ .(b) Graph of  $p_{\Delta t}$ .

Figure 3.14: Graphs of  $\chi_{obs, \Delta t}^2$  and the corresponding  $p_{\Delta t}$ , both as functions of  $\Delta t$ . The model used to calculate these two was constructed with the optimal combination of parameters for fraction B, as found in Section 3.4.2.

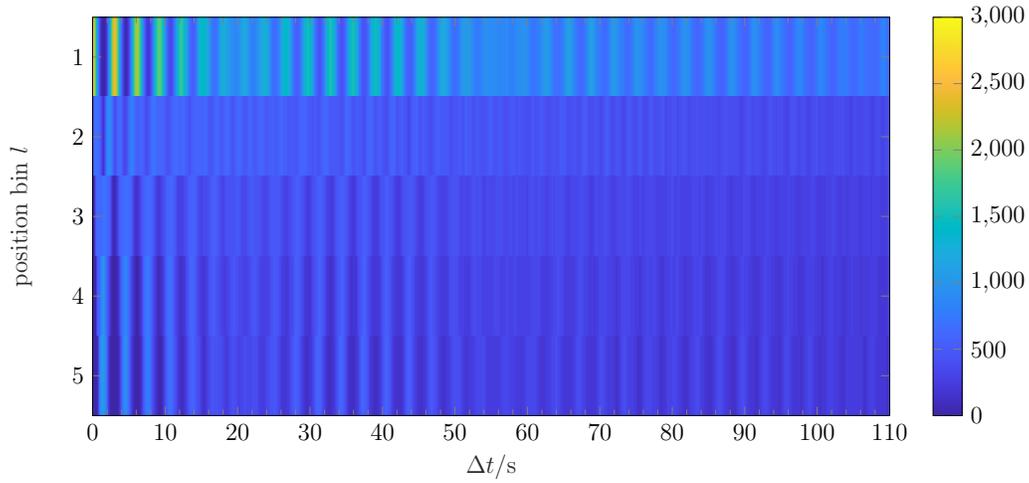
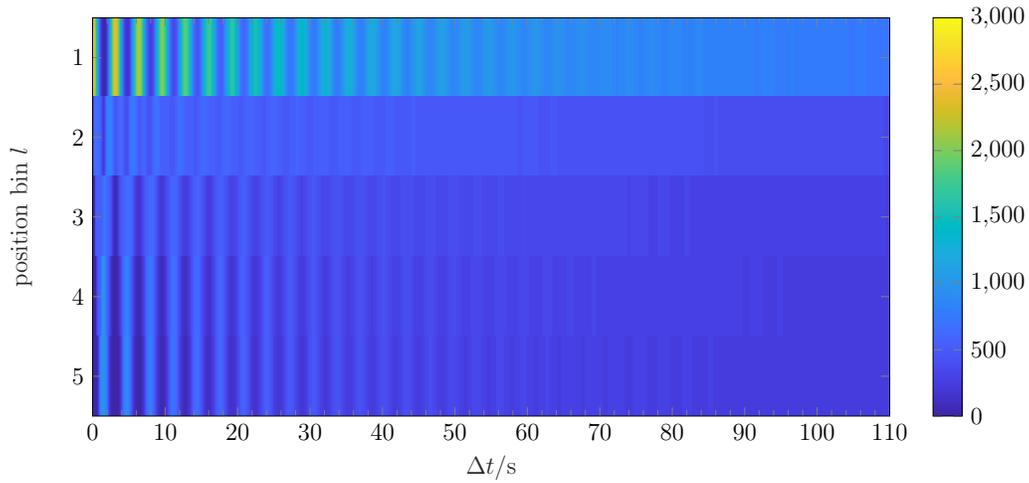
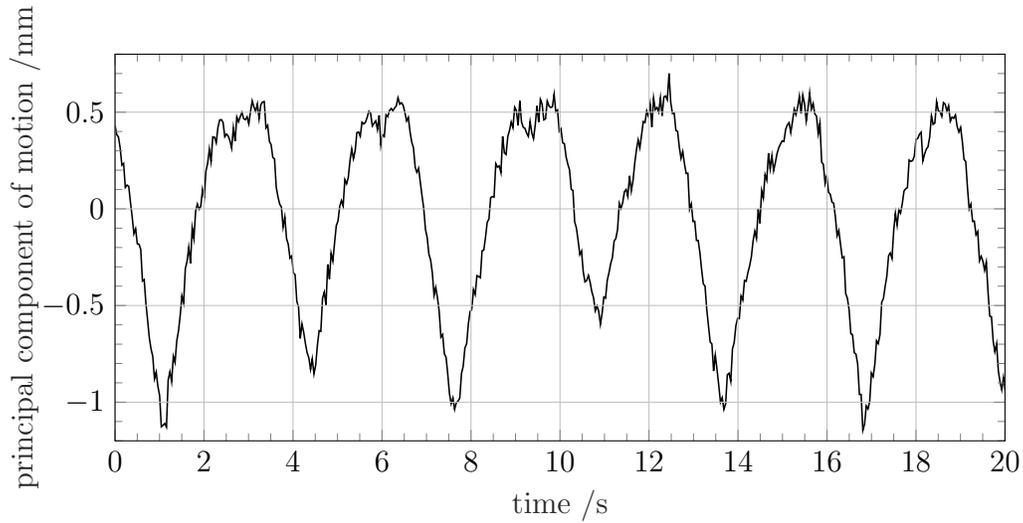
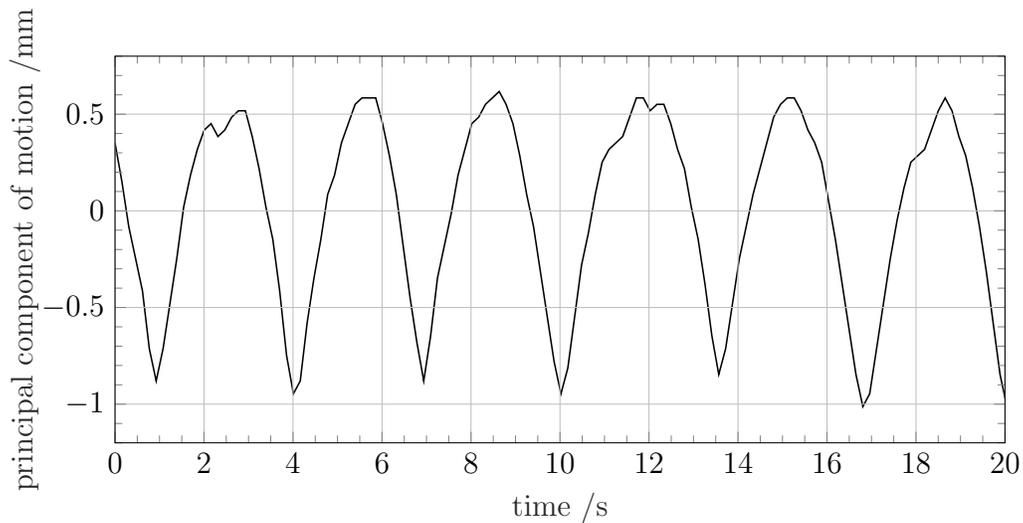
(a) Observed position bin histogram  $O_{\Delta t}^l$ .(b) Expected position bin histogram  $E_{\Delta t}^l$ .

Figure 3.15: Images of the observed and expected position bin histograms, both as functions of  $\Delta t$ . The model used to calculate these two was constructed with the optimal combination of parameters for fraction B, as found in Section 3.4.2.



(a) Observed sequence.



(b) Sample sequence.

Figure 3.16: Observed and simulated sequence of respiratory motion as a function of time. Both sequences are represented as the principal component of motion; this was done for the Markov sample using the position sub-bin lookup table  $x_\alpha$ , and then converting the position sub-bin index to the corresponding value of the principal component of motion. The model used to generate the sample was constructed with the optimal combination of parameters for fraction B, as found in Section 3.4.2.

### 3.4.3 Fraction model intercomparisons

The 5 min of fraction B testing data was compared to the fraction A model, with  $\chi_{\text{obs},\Delta t}^2$  and  $p_{\Delta t}$  results shown in Figure 3.17. This was done by effectively comparing the observed fraction B histogram in Figure 3.15a to the expected fraction A histogram in Figure 3.12b. The agreement between these two fractions is especially poor for values of  $\Delta t$  less than 10 s, and only somewhat improves as  $\Delta t$  increases, with  $p_{\Delta t}$  oscillating between a maximal value of 0.28 and 0. The summed  $\chi_{\text{obs},\Sigma}^2$  was found to be equal to 2 794 742, a far greater value than any observed in the Monte Carlo simulations of model A (Figure 3.7). The corresponding value of  $p_{\Sigma} = 0$  indicates that the observed data from fraction B is wholly incompatible with the model trained on the fraction A data.

Conversely, the 5 min of fraction A testing data was compared to the fraction B model, with similarly poor agreement seen in the graphs of  $\chi_{\text{obs},\Delta t}^2$  and  $p_{\Delta t}$  (Figure 3.18). The summed  $\chi_{\text{obs},\Sigma}^2$  was found to be equal to 1 474 278, with a corresponding value of  $p_{\Sigma} = 0.001$ . These results also indicate that the observed data from fraction A is incompatible with the model trained on the fraction B data.

The stationary probability distributions of the position bin for both respiratory models are shown in Figure 3.19. These were calculated from the stationary distributions of the Markov chain states,  $\pi_{\alpha}$ , by accumulating the probabilities from each state  $\alpha$  into their respective position bin using the lookup tables  $l_{\alpha}$ , defined in Section 3.3.1. At larger values of  $\Delta t$ , the model

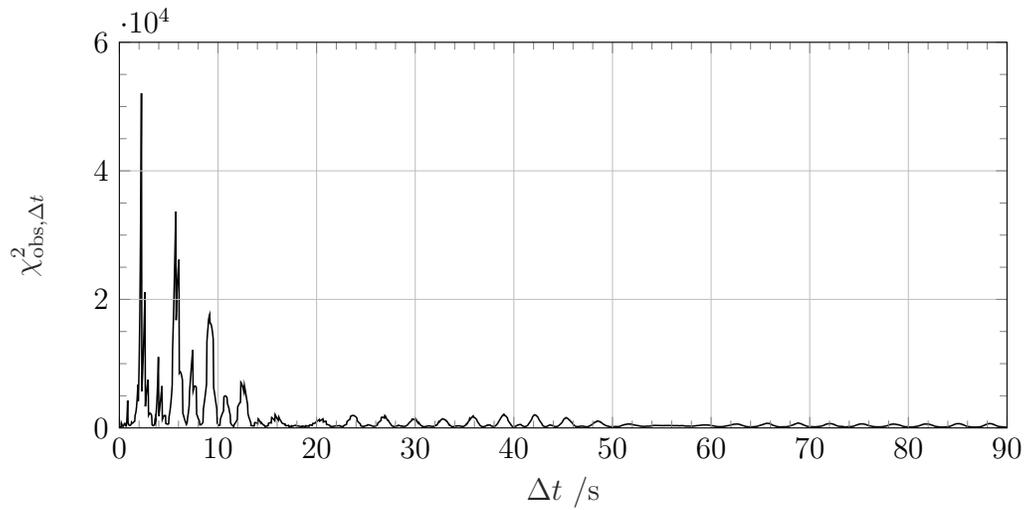
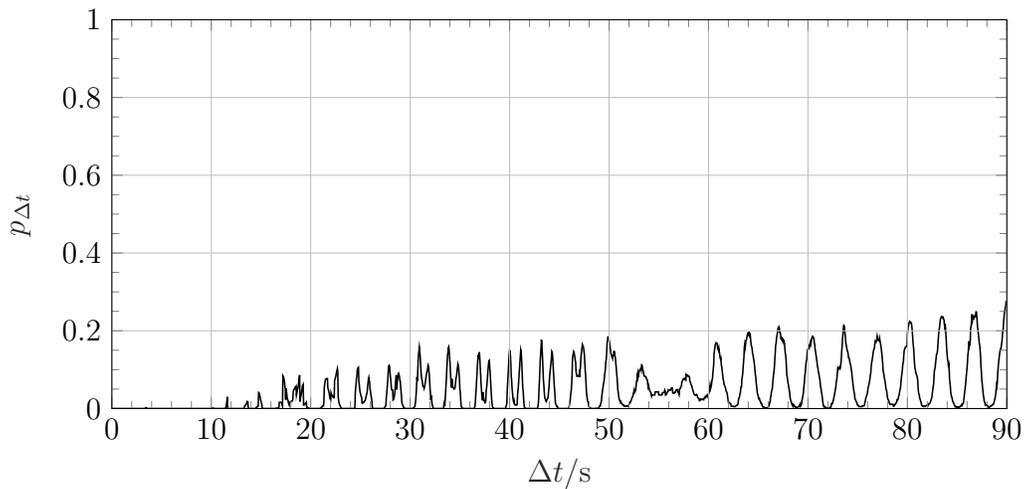
(a) Graph of  $\chi_{\text{obs}, \Delta t}^2$ .(b) Graph of  $p_{\Delta t}$ .

Figure 3.17: Graphs of  $\chi_{\text{obs}, \Delta t}^2$  and the corresponding  $p_{\Delta t}$ , both as functions of  $\Delta t$ . The model used to calculate these two was constructed with the optimal combination of parameters for fraction A, as found in Section 3.4.1, but the observation data was taken from fraction B.

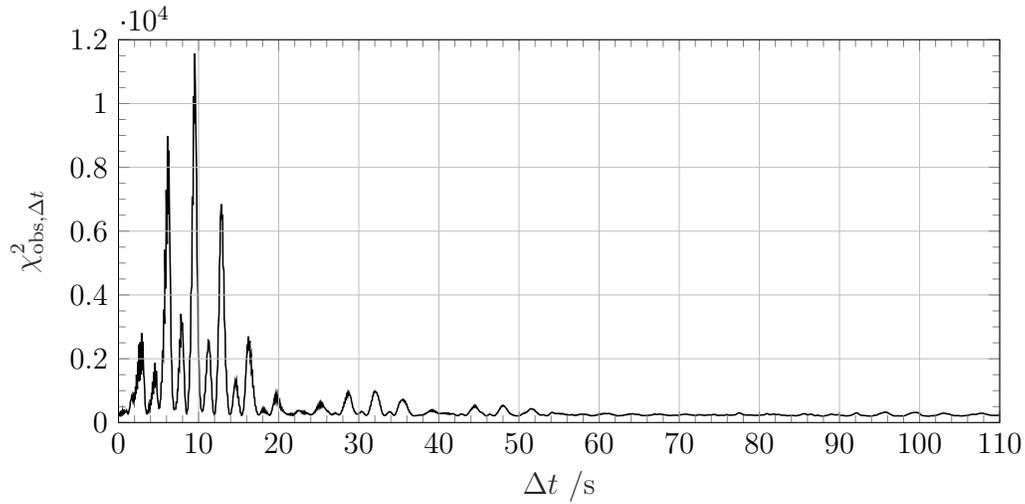
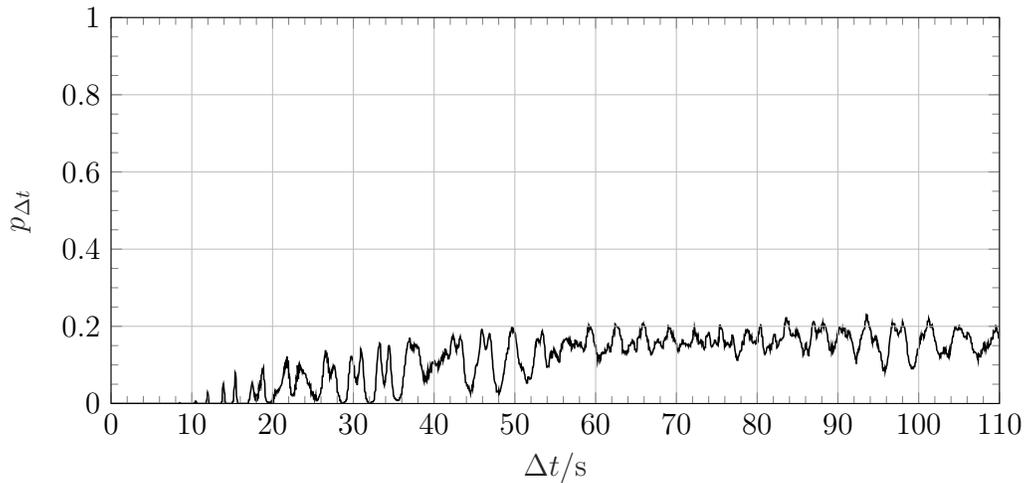
(a) Graph of  $\chi_{\text{obs}, \Delta t}^2$ .(b) Graph of  $p_{\Delta t}$ .

Figure 3.18: Graphs of  $\chi_{\text{obs}, \Delta t}^2$  and the corresponding  $p_{\Delta t}$ , both as functions of  $\Delta t$ . The model used to calculate these two was constructed with the optimal combination of parameters for fraction B, as found in Section 3.4.2, but the observation data was taken from fraction A.

trained from the fraction B data has a notably higher probability to be in position bin  $l = 1$ , at the end of exhale, whereas the stationary distribution for the fraction A model is more uniform.

### 3.5 Discussion

The excellent agreement between the position histograms observed in the 5 min of testing data and the corresponding model calculations demonstrates that the Markov chain adequately models the respiratory motion of actual patients, from fractions of a second to 90 s (or 110 s) into delivery. It is important to note that the Markov chain model is only aware of state transitions occurring at the time interval  $\tau$  (386.9 ms or 154.2 ms, depending on the fraction), and that the behaviour at longer time scales emergent from these underlying transitions is not directly controlled by the model. Considering this, it is remarkable that the models are able to reproduce the data from each respective fraction. This is made possible by the definition of the various components making up each Markov state, as described in Section 3.3.1.

For example, the images of  $p_\Sigma$  for the fraction A model in Figures 3.8 and 3.9 show that the Markov chain time interval factor  $r = 9$ , and the corresponding Markov chain time interval  $\tau = 386.9$  ms, is strongly preferred over others. Figure 3.8 also shows that as  $r$ , and subsequently  $\tau$ , decreases, the optimal number of bins recording the elapsed time in the current breathing phase increases. This reflects the fact that a model with a smaller value of  $\tau$

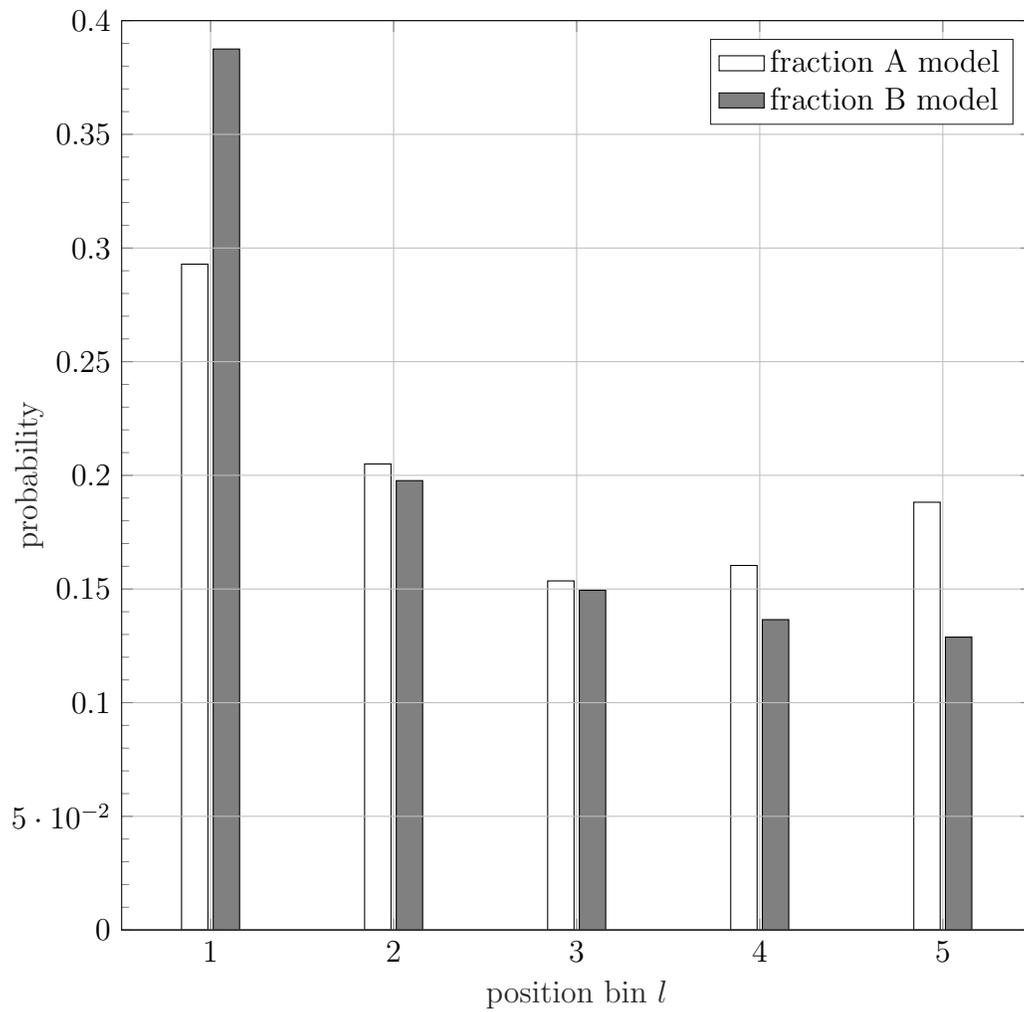


Figure 3.19: Stationary position bin distributions for the fraction A and B models.

spends more time points in any given breathing phase; a higher number of bins  $N_T$  is therefore required to differentiate these states.

The image in Figure 3.9 demonstrates that the value used for  $N_x$ , the number of position sub-bins per bin, is relatively unimportant, save when the value  $r = 9$  for the Markov chain time interval factor is chosen; in this case, the lowest values of  $N_x$  yield considerably higher values of  $p_\Sigma$ . Finally, the agreement of the model to the testing data is seen to be completely independent of  $N_v$ , the number of velocity bins, in Figure 3.10.

Although the sequences simulated from each model, shown in Figures 3.13 and 3.16, have lost many of the higher frequency components seen in their respective observed sequences, other features are emulated, such as the overall period, longer pauses in the exhale region, and a variation between shallower and deeper inhalations. These are considered to be more important than the high-frequency behaviour in the observed sequences for the purposes of 4D-VMAT treatment planning, since only five position bins, corresponding to distinct patient geometries, are used for dose calculation (see Sections 4.2.5.2 and 4.3.1.2). On the other hand, the greater number of position sub-bins in model fraction B can allow for the modelling of more complex patterns of respiratory motion, such as the dips in the peaks in the exhaling breathing phases seen in Figure 3.16.

The respiratory data used to train and test the model for a patient requiring such a respiratory motion model would need to be acquired before radiotherapy treatment, just as are the CT simulation images. As discussed

in Section 3.3.2, the data used in this study consists of the measured position of optical surface markers on the patient, which is a surrogate for the actual motion of the tumour; specifically, the principal component was used to determine the tumour motion. This surrogate data would then need to be related to the actual position of the tumour, for example by acquiring radiography images and the surrogate signal simultaneously as the patient is breathing [48, 49]. Alternatively, the tumour motion could be measured directly using, for example, positron-emitting markers implanted in the tumour [39], or a magnetic resonance linac [47].

As described in Section 3.3.2, the model was trained on the 5 min of respiratory data immediately preceding the testing data. Thus the training and testing datasets are distinct, yet related as they originate from the same patient, and indeed the same treatment fraction. The agreement is significantly worsened when intercomparing a model and testing data from two different fractions, as seen in Section 3.4.3. Indeed, the two fractions A and B seem to be wholly incompatible with one another, when only looking at the position histograms. The stationary position histograms of each fraction model, shown in Figure 3.19, are also significantly different from one another. Finally, the average period of trajectories simulated from the fraction A model is longer than the period for the fraction B model, at 3.6 s compared to 3.1 s.

This poses the problem of determining if a respiratory motion model trained on data acquired before the treatment, either for plan optimization or even just before delivery, is representative of patient breathing during

treatment. Since the focus of this thesis is on radiotherapy treatment planning, this problem can be reformulated as investigating the effect on the dose delivered to a patient if treatment planning and delivery are performed with different respiratory motion models. This is treated in Section 4.3.2.4, following the development of the proposed 4D-VMAT optimization framework.

### 3.6 Conclusions

A method was introduced to construct patient-specific Markov chain respiratory motion models for the purposes of probability estimation. The states of the Markov chain model were defined by the position of a tracking marker, its velocity, the breathing phase of the patient (inhaling, exhaling, end of exhale), and by the elapsed time in the current breathing phase. Such a model was trained and tested using patient respiratory data collected during radiotherapy treatments, and was found to adequately estimate the distribution of the patient's respiratory motion from fractions of a second to the order of a minute into delivery. Ultimately, the Markov chain motion model enables the calculation of the probability of individual respiratory trajectories, which is required for robust 4D treatment plan optimization.

## Chapter 4

# 4D-VMAT optimization and delivery framework

In this chapter the 4D-VMAT optimization and delivery framework ALERT-RA, first introduced in Section 1.5, is fully developed, validated, and compared to other motion adaptive VMAT methods. In particular the robustness of the plan quality with respect to the patient’s respiratory motion is evaluated for each method.

### 4.1 Introduction

Inverse planning of radiotherapy plans is a process whereby computer optimization algorithms are used to find the treatment plan parameters which maximize the quality of the delivered treatment plan. As described in Chapter 2, in VMAT the plan consists of specifying the shapes of the apertures formed by the MLC at a discrete set of angles around the patient [23, 24].

The quality of a radiotherapy plan generally increases as the dose to the target becomes closer and more uniform with respect to the prescribed dose, and as the dose to OARs decreases below clinically-guided threshold.

In Section 1.3 it was described how intrafraction motion, particularly due to respiration, can lead to shifting and blurring of the high-quality dose distributions obtained during optimization, decreasing the dose to the tumour and increasing the dose to OARs. This can potentially lead to a lower probability of tumour control, and an increased risk of radiation-induced side-effects.

Many approaches have been developed to limit or compensate for this motion. Some of the methods implemented at the delivery stage include gating, breath-hold, and compression techniques, as discussed in Section 1.4. Techniques also exist which compensate for intrafraction motion in the planning stage. First a 4D-CT is acquired, which is composed of a set of CT images at different points in the patient's respiratory motion, recording the tumour and organ displacement as a function of respiratory phase. A typical clinical solution is to define the ITV which encloses the CTV throughout its entire range of motion, as seen on the 4D-CT [26]. The radiotherapy plan is then optimized to deliver the prescription dose to this volume, ensuring that regardless of the tumour motion, the required dose is always achieved. This larger treated volume results however in a greater dose to surrounding OARs, increasing the chance and severity of radiation-induced complications. An alternative solution to the ITV, called MKO, was introduced in Section 1.4. In

this method the dose-influence matrix, which, as will be seen Section 4.2.1.1, depends on the changing tumour position and corresponding patient geometry, is weighted and summed over all tumour positions according to the probability distribution of the patient motion. Plan optimization is then conducted using this effective dose-influence matrix [35]. A full explanation of this is given in Section 4.2.6.

There are also several methods which consider the fraction of the respiratory cycle that the tumour spends in each position along its trajectory, thereby restricting the delivered dose to the moving CTV and obviating the need for an ITV. In addition to the decreased dose to OARs, these techniques, which fall under the umbrella term of optimized tracking, have the advantage of an extra temporal degree of freedom for optimization. Specifically, the optimization algorithm can selectively increase the amount of dose delivered in instances for which the target is farther away from a dose-sensitive OAR than in others [35].

Two particular 4D-VMAT optimization methods will be investigated more closely in this study: single trajectory optimization (STO) [52] and direct aperture deformation (DAD) [51]. In STO it is assumed that the patient's breathing motion during delivery is known at the optimization stage, and does not change from one fraction to another. A one-to-one correspondence between the tumour position and gantry angle is fixed before the apertures belonging to this sequence are optimized. If the patient breathes at a faster or slower rate than anticipated, the entire delivery may be slowed down or

sped up as required, assuming that the machine speed constraints are met. However, with these methods there is no clear way to modify the plan if the patient's breath becomes more shallow or deep than expected. These deviations of the breathing pattern at delivery, if not considered during optimization, may have significant effects on the resulting plan quality.

DAD, on the other hand, optimizes apertures on a single reference tumour position, as in 3D-VMAT. Each aperture is then deformed to follow the real-time motion of the target as the gantry rotates around the patient. This technique does not however account for the varying target and OAR geometry, unlike other optimized tracking methods. Treatment plans using DAD may therefore deliver a higher dose to OARs and a lower or more non-uniform dose to the target.

Introduced in Section 1.5, robust optimization is a technique which accounts for the random nature of variables which have an effect on the delivery of dose to the patient, with the aim of reproducibly delivering the prescription dose to the target and the desired low doses to OARs [53]. This is typically done by optimizing the expectation value of the plan quality. In this study, robust optimization is applied to the uncertain nature of the patient respiratory trajectory, with the aim of yielding high-quality 4D-VMAT plans which are robust to changes in the trajectory. The resulting plans are then compared with plans generated using MKO, STO, and DAD. Specifically, the overall quality of the plans are compared, as well as the robustness of the plan quality to variations in the patient respiratory trajectory.

The 4D-VMAT approach proposed here uses an aperture library, consisting of an MLC aperture and MU linac output for each combination of gantry angle and tumour position. Throughout the course of delivery, the real-time tumour position as well as the current gantry angle would be used to select the appropriate aperture to be formed by the MLC. The use of robust optimization to generate the aperture library would ensure that no matter the respiratory motion trajectory observed, the tumour is irradiated with the desired dose, and a sufficiently low dose to OARs is given. The robust optimization framework introduced in this chapter requires a model which can calculate the probability of arbitrary respiratory trajectories. The Markov chain respiratory motion model which was introduced, constructed, and evaluated in Chapter 3 will be used for this very purpose.

## 4.2 Materials and methods

### 4.2.1 4D-VMAT framework

The foundation of the proposed ALERT-RA framework consists of a library of MLC apertures and MU rates which are defined on a grid of control points in the space of gantry angle  $\theta$  and tumour position  $\omega$  (being the continuous version of the discretized tumour position bin  $l$  used throughout Chapter 3); Figure 4.1 shows a schematic of the library. This is an extension of the optimization framework for 3D-VMAT presented in Chapter 2; correspondingly, the MLC leaves are considered to transition smoothly from the apertures

defined at one control point to the next. The three sets of gantry angles defined in Section 2.2.1.2 for 3D-VMAT, FMO initialized angles  $\{\theta_{\text{FMO}}\}$ , DAO optimization angles  $\{\theta_{\text{DAO}}\}$ , and dose calculation angles  $\{\theta_{\text{dose}}\}$ , are also used in 4D-VMAT optimization method. Thus the control points of the aperture library in the  $\theta$  direction correspond to the beginning and end of the arc section surrounding each  $\theta_{\text{DAO}}$ , as was done in that section. In the  $\omega$  direction, the control points are defined by a discrete set of patient geometries used for dose calculation (see Sections 4.2.1.1 and 4.2.3).

It is assumed that during delivery, the position of the tumour is tracked in real time, and at each control point the aperture corresponding to the current position and gantry angle is selected from the library and delivered. The aperture library must be optimized so that a high-quality dose distribution  $d$  is consistently delivered to the patient, independent of the patient motion undergone during treatment delivery. As indicated in Section 1.2.3 and used throughout Chapter 2, in inverse treatment planning an objective function  $F(d)$  is defined to quantify the overall quality of the delivered plan, with consideration to the dose to the target and to OARs;  $F(d)$  generally decreases for increased plan quality, therefore the goal of the optimization is to minimize the objective function. In general the patient will undergo a different unknown respiratory trajectory each treatment fraction. Since the dose delivered in each fraction  $d$  depends on the trajectory (specific calculations are performed in Section 4.2.1.1), the delivered plan quality for each fraction, quantified by the objective function, is a random variable.



Traditionally only a single dose scenario, here identified with the respiratory trajectory, is considered when calculating dose in the optimization. This approach fails when it is unknown which scenario, and therefore which dose, is realized during delivery. Robust treatment planning enables the optimization of the aperture library with the aim of minimizing the objective function  $F(d)$  for all dose scenarios simultaneously, by treating the dose scenario as a random variable. Specifically, in conventional robust optimization (hereafter referred to as conventional optimization, or CO), the expectation value of the dose  $\bar{d}$  over all scenarios is computed, and the objective function evaluated at this dose is minimized:

$$\begin{aligned} & \min F(\bar{d}), \\ \bar{d} &= \int \Pr[\Omega] d[\Omega] \mathcal{D}\Omega. \end{aligned} \quad (4.1)$$

Here,  $\Omega$  is a label for the respiratory trajectory undergone by the patient during delivery,  $d[\Omega]$  is the dose delivered to the patient due to a particular trajectory  $\Omega$ , and  $\Pr[\Omega]$  is the probability of observing that trajectory.

Probabilistic optimization (PO), on the other hand, seeks to minimize the expectation value of the objective function evaluated at each dose scenario:

$$\begin{aligned} & \min \bar{F}, \\ \bar{F} &= \int \Pr[\Omega] F(d[\Omega]) \mathcal{D}\Omega. \end{aligned} \quad (4.2)$$

Previous studies [55, 84, 85] have shown that plans optimized under PO are more robust to variations in the underlying random variables (in this case, the respiratory trajectory  $\Omega$ ) than those optimized with CO. Since plan quality robustness to variations in the tumour trajectory, in addition to plan quality, is a desirable outcome of the treatment planning, PO was performed to optimize the aperture library. Plans were also optimized with CO to compare plan quality and robustness between the two methods.

The novelty of this 4D-VMAT optimization and delivery framework lies in Equations (4.1) and (4.2), which combines two aspects. The first is a method of dose calculation which correctly accounts for the interaction between the tumour motion trajectory and the corresponding delivery adaptation. Secondly, in this method the trajectories themselves are considered to be random variables with definite probability. This in contrast to MKO, for which the tumour position is assumed to have a static distribution in time.

An expression for the expectation value of the dose delivered during a single 4D-VMAT fraction, Equation (4.1), is derived in Section 4.2.1.1. The derivation starts with a calculation of the dose in 3D-VMAT, moves on to the dose delivered during a single trajectory  $\Omega$ , before finally calculating the expectation value over all trajectories. Following this, the expectation value of the objective function for PO, Equation (4.1), is calculated in Section 4.2.1.2.

#### 4.2.1.1 Dose calculation

In the 3D-VMAT method introduced in Chapter 2, the dose-influence matrix  $D_{ij}(\theta)$  is defined as the dose imparted to voxel  $i$  if bixel  $j$  at gantry angle  $\theta$  has unit fluence (see Figure 2.1 for an illustration of this). Following the treatment of continuous aperture dose calculation described in Section 2.2.1.1, the differential fluence delivered by each bixel  $j$  at angle  $\theta$  can be calculated given the trajectories of the right and left leaves of MLC row  $J$  to which  $j$  belongs:  $L_J = L_J(\theta)$  and  $R_J = R_J(\theta)$ . Practically,  $L_J$  and  $R_J$  is defined at a discrete set of angles  $\{\theta\}$  and interpolated between these control points. The differential fluence is given the symbol  $\frac{d\phi_j}{d\theta}(L_J, R_J) d\theta$ , and thus the dose delivered to voxel  $i$  during a 3D-VMAT delivery is (summation over bixel index  $j$  is implied):

$$d_i = \int D_{ij}(\theta) \frac{d\phi_j}{d\theta}(L_J, R_J) d\theta. \quad (4.3)$$

As will be seen, an integral over time  $t$  is required in 4D-VMAT dose calculations, therefore it is useful to include the integral at this point in 3D-VMAT. The insertion of this integral should not change the numerical result from Equation (4.3), since it is already accurate. Furthermore, when doing calculations in 4D-VMAT it is required to select the particular time  $t$  at which the gantry reaches angle  $\theta$ , given the gantry trajectory  $\theta = \Theta(t)$ . Taking these two facts into account, the Dirac delta function  $\delta(t - \Theta^{-1}(\theta))$  must accompany the integral over time, resulting in the following expression for

the dose in 3D-VMAT:

$$d_i = \iint \delta(t - \Theta^{-1}(\theta)) D_{ij}(\theta) \frac{d\phi_j}{d\theta}(L_J, R_J) d\theta dt. \quad (4.4)$$

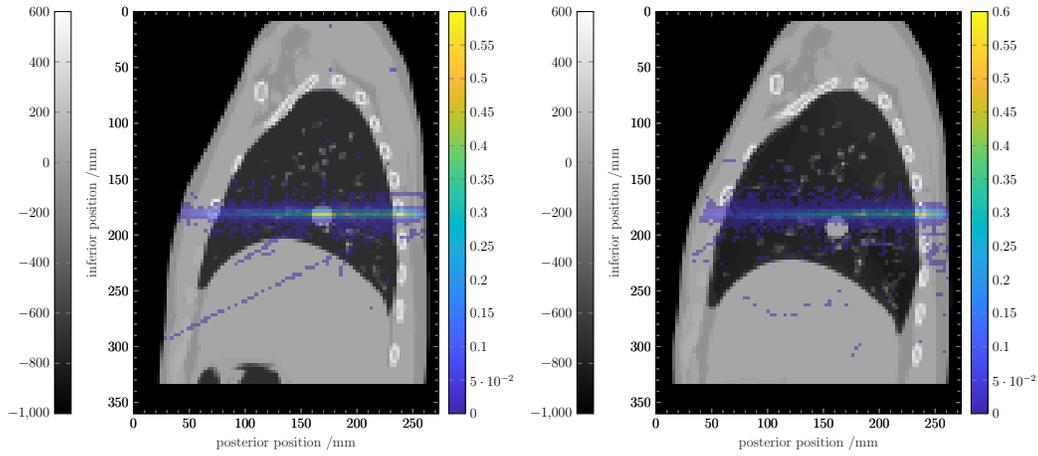
In 4D-VMAT dose is planned to be delivered while the patient is undergoing respiratory motion. Since each tumour position  $\omega$  will correspond to a different patient geometry, as captured in the 4D-CT, the dose-influence matrix is in general a function of both the tumour position and gantry angle:  $\tilde{D}_{ij}(\theta, \omega)$ . Here the tilde over  $D$  indicates that the dose calculated on each patient geometry is deformed backed to a reference geometry, taken to be maximum exhale. This is so that the voxel indices  $\{i\}$  always refer to the same anatomical voxels independent of tumour position  $\omega$ , which is required for consistent dose accumulation [35].

To illustrate this, images of the dose from a single bixel to two patient geometries are shown in Figure 4.2, along with the corresponding deformed dose. Figure 4.2a shows the dose from a bixel aligned with the centre of the tumour at maximum exhale. In Figure 4.2b, the dose due to that same bixel is shown, but this time incident on the patient geometry at maximum inhale: instead of the centre, the superior edge of the tumour is now irradiated by the beam. Finally, Figure 4.2c shows this same dose, now deformed back to the first geometry, that of maximum exhale. In this image, the superior edge of the tumour is still clearly irradiated by the beam; indeed, the entire beamline in the lung is now curved upward, reflecting the relative motion between the

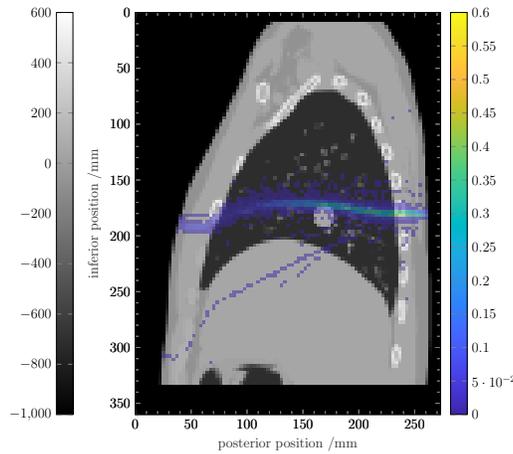
two geometries.

This dose deformation process can be carried out using one of two methods, both of which require the set of vectors defining the motion of each anatomical voxel from the reference geometry to the deformed geometry, known as the deformation vector field (DVF). Practically, the DVF can be obtained from a 4D-CT of the patient by performing deformable image registration to determine the correspondence between the different geometries [86]. In the energy mapping method, the dose is calculated on the deformed geometry and then mapped back to the reference geometry using the DVF; in principle dose calculation can be done using any method [87, 88]. The voxel warping method, on the other hand is well-suited to Monte Carlo dose calculation: in this method, particle transport is performed in the deformed geometry, while the energy deposited by the particles is scored in the reference geometry [66]. This is the method used in this study, as implemented in the Monte Carlo code VMC++; further details are given in Section 4.2.2. Note that both of these methods also account the effect of the geometry deformation on varying voxel density, and thus on particle transport and energy deposition.

The total dose delivered in a 4D-VMAT plan will depend on the respiratory trajectory experienced by the patient, represented by the tumour position as a function of time:  $\omega = \Omega(t)$ . To calculate this dose, an integral over position  $\omega$  is inserted in Equation (4.4), with an accompanying delta function to pick



(a) Dose to maximum exhale geometry (b) Dose to maximum inhale geometry. (reference).



(c) Dose to maximum inhale geometry, deformed to the reference geometry.

Figure 4.2: Dose deposited from a single bixel, irradiating two patient geometries: the patient at maximum exhale (Figure 4.2a), and the patient at maximum inhale (Figure 4.2b). The dose (colour, in units of Gy per 100 MU) is superimposed on a sagittal slice of the CT image (greyscale, in HU) for each of these geometries. The beam is entering the patient from the right side of the image, the patient’s posterior side. Figure 4.2c shows the same dose as in Figure 4.2b, but deformed to apply to the anatomical voxels in the reference geometry of maximum exhale, on which it is superimposed. The CT images were generated using XCAT [64], and dose was calculated using VMC++ [65, 66]. The CT image for the inhale geometry in Figure 4.2b was further modified to ensure local preservation of mass [89].

out the particular position reached at time  $t$ :

$$\begin{aligned} d_i [\Omega] &= \iiint \delta(\omega - \Omega(t)) \delta(t - \Theta^{-1}(\theta)) \tilde{D}_{ij}(\theta, \omega) \frac{d\phi_j}{d\theta}(L_J, R_J) d\theta dt d\omega \\ &= \iint \delta(\omega - \Omega(\Theta^{-1}(\theta))) \tilde{D}_{ij}(\theta, \omega) \frac{d\phi_j}{d\theta}(L_J, R_J) d\theta d\omega. \end{aligned} \quad (4.5)$$

The 4D dose-influence matrix is used, and in the second line, the integral over time has been performed, eliminating one of the delta functions. Here the left and right leaf positions are functions of  $\Theta$ ,  $\Omega$ ,  $\theta$ , and  $\omega$ . The dependence on  $\Theta$  and  $\Omega$  is made explicit here because practically  $L_J$  and  $R_J$  will be defined at a discrete set of points in  $(\theta, \omega)$  space, with the leaf positions taking on interpolated values between these control points depending on the trajectories  $\Theta$  and  $\Omega$  which are followed. In other words, since the sequence of apertures from the aperture library is selected based these trajectories, the fluence calculation also then depends on the respiratory trajectory, taking into account the interplay of tumour motion and MLC leaf movement.

It is now possible to calculate the expectation value of the dose delivered in a 4D-VMAT plan. Given a model which can calculate the probability  $\Pr[\Omega]$  of observing a respiratory trajectory  $\Omega$ , the expectation value is simply:

$$\begin{aligned} \bar{d}_i &= \int \Pr[\Omega] d_i[\Omega] \mathcal{D}\Omega \\ &= \iiint \Pr[\Omega] \delta(\omega - \Omega(\Theta^{-1}(\theta))) \tilde{D}_{ij}(\theta, \omega) \frac{d\phi_j}{d\theta}(L_J, R_J) d\theta d\omega \mathcal{D}\Omega. \end{aligned} \quad (4.6)$$

Here the integration is over all possible trajectories  $\Omega$ . In the second line, Equation (4.5) has been substituted in for  $d_i[\Omega]$ . A full derivation of the conversion of this continuous expression for dose into a discrete formulation which can be computed is performed in Appendix C.

The final result of the derivation is a series of equations (Equations (C.2) and (C.8)):

$$\bar{d}_i = \sum_l \sum_k \tilde{D}_{ij}^{kl} \phi_j^{kl}, \quad (4.7)$$

$$\begin{aligned} \phi_j^{kl} = \sum_{l'} \left[ \Pr(\Omega^{k,l \rightarrow l'}) \int_{\theta^k}^{\theta^{l'}} \frac{d\phi_j}{d\theta} (L_J^{k,l \rightarrow l'}, R_J^{k,l \rightarrow l'}) d\theta \right. \\ \left. + \Pr(\Omega^{k,l' \rightarrow l}) \int_{\theta^k}^{\theta^{l'}} \frac{d\phi_j}{d\theta} (L_J^{k,l' \rightarrow l}, R_J^{k,l' \rightarrow l}) d\theta \right]. \end{aligned} \quad (4.8)$$

Here the space of gantry angle and tumour position has been subdivided into a grid of discrete points  $\{(\theta^k, \omega^l)\}$ , with respective intervals of  $\Delta\theta$  and  $\Delta\omega$ .  $\tilde{D}_{ij}^{kl}$  is the dose-influence matrix centred at each such point,  $\tilde{D}_{ij}(\theta^k, \omega^l)$ , and is multiplied by the effective fluence for bixel  $j$  at this point,  $\phi_j^{kl}$ . This expression features the probability to observe a particular trajectory segment  $\Omega^{k,l \rightarrow l'}$ , in which the tumour position changes from  $\omega^l$  to  $\omega^{l'}$  as the gantry rotates between angles  $\theta^k \pm \Delta\theta/2$  (as well as the trajectory segment  $\Omega^{k,l' \rightarrow l}$ ). An example of such a trajectory segment is shown graphically in Figure 4.3. The effective fluence  $\phi_j^{kl}$  can therefore be understood intuitively as the probability to observe a trajectory segment transitioning either from or to the tumour position  $l$  ( $\Omega^{k,l \rightarrow l'}$  or  $\Omega^{k,l' \rightarrow l}$ , respectively), multiplied by the fluence delivered

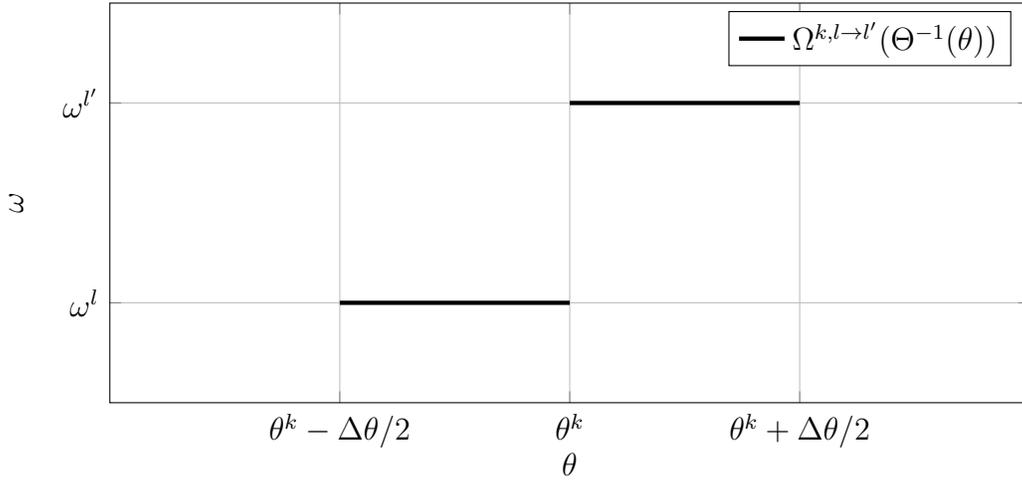


Figure 4.3: Allowed trajectory segments  $\omega = \Omega^{k,l \rightarrow l'}(\Theta^{-1}(\theta))$  in the 4D-VMAT dose calculation model.

during that trajectory segment (encoded in the integrals), summed over all tumour positions  $l'$ .

The probability to observe  $\Omega^{k,l \rightarrow l'}$  can be calculated from the definition of conditional probability, as the probability of the tumour to arrive at position  $\omega^l$  when the gantry arrives at angle  $\theta^k - \Delta\theta/2$ , multiplied by the probability for the tumour to transition from  $\omega^l$  to  $\omega^{l'}$  as the gantry angle increases by  $\Delta\theta$ . Therefore a suitable respiratory motion model, such as the one introduced in Chapter 3, must accurately calculate the probability distribution of the tumour position at the time points corresponding to the gantry angles  $\theta^k$ . Note that in this model, the probabilities of observing specific states  $\alpha$ , which do not necessarily correspond to the tumour position bins  $l$ , are calculated directly. The lookup tables  $l_\alpha$ , as defined in Section 3.3.1, are then used to accumulate the probabilities from each state into their respective position

bin.

The grid  $\{(\theta^k, \omega^l)\}$  is selected to fulfill two conditions:  $\tilde{D}_{ij}(\theta, \omega)$  should be relatively constant over the intervals of length  $\Delta\theta$  and  $\Delta\omega$  surrounding each point  $(\theta^k, \omega^l)$ ; and the gantry rotation times corresponding to the angular intervals  $\Delta\theta$  are small enough such that as the tumour position transitions from  $\omega^l$  to  $\omega^{l'}$  in this arc, no significant dose calculation errors are made by ignoring the dose deposited to intermediary patient geometries.

The set of dose calculation angles,  $\{\theta_{\text{dose}}\}$ , as well as the number of deformed geometries corresponding to independent tumour positions  $\omega^l$ , are both defined to fulfill the first condition (see Section 4.2.5.2 for more on this topic). For the second condition, an additional set of gantry angles,  $\{\theta_{\text{flu}}\}$ , is introduced, with a sufficiently small interval  $\Delta\theta_{\text{flu}}$ . The fluence released by the MLC during the trajectory segment  $\Omega^{k,l \rightarrow l'}$  is calculated on this set of gantry angles, as is the corresponding probability. This fluence is then assigned to the nearest dose calculation angle, to be then multiplied into the corresponding portion of the dose-influence matrix. A schematic of the  $\{\theta_{\text{flu}}\}$  angle set is shown in Figure 4.4.

In practice the interval between fluence calculation angles,  $\Delta\theta_{\text{flu}}$ , was selected to correspond to the Markov chain time interval  $\tau$  using a fixed gantry rotation speed  $\dot{\theta}$ , thus:

$$\Delta\theta_{\text{flu}} = \dot{\theta} \cdot \tau. \quad (4.9)$$

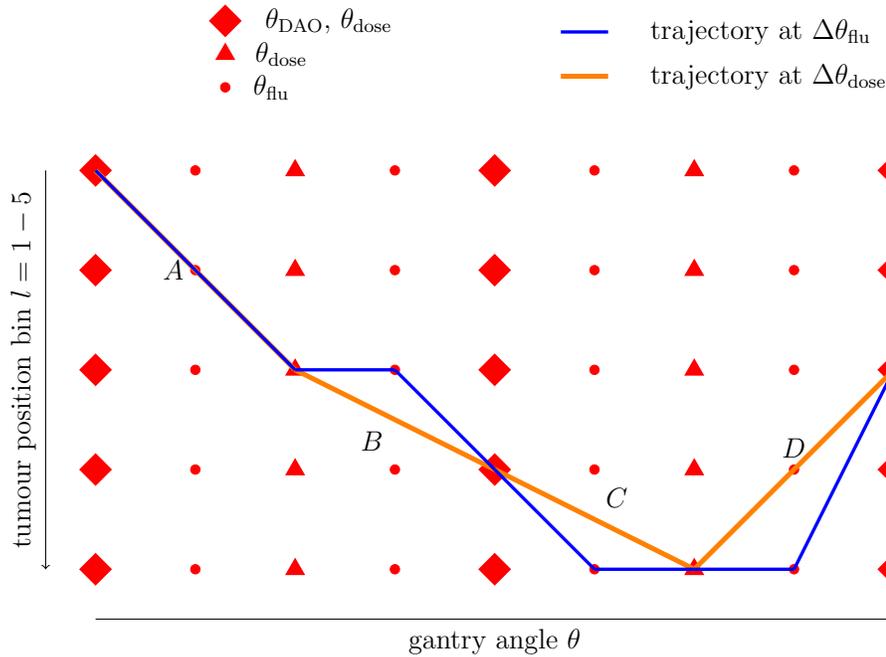


Figure 4.4: Schematic of the set of fluence calculation angles  $\{\theta_{\text{flu}}\}$  in the space of gantry angle  $\theta$  and tumour position bin  $l$ . The dose calculation  $\{\theta_{\text{dose}}\}$  and DAO optimization angles  $\{\theta_{\text{DAO}}\}$  are also shown. Apertures at the  $\theta_{\text{flu}}$  and  $\theta_{\text{dose}}$  angles are interpolated from the apertures at the two adjacent  $\theta_{\text{DAO}}$  angles with the same tumour position. The blue line shows the trajectory of a tumour at the  $\Delta\theta_{\text{flu}}$  spacing, while the orange line shows this same trajectory resampled at the larger  $\Delta\theta_{\text{dose}}$  angle spacing. In segment *A*, the orange trajectory deposits dose only to position bins  $l = 1$  and  $l = 3$ , while the blue trajectory correctly deposits dose to the intermediate bin  $l = 2$  in addition to the two endpoints. In addition to this, the MLC aperture is updated at the intermediate  $\theta_{\text{flu}}$  angle. This scenario is replicated similarly for segments *B*, *C*, and *D*.

This choice is validated in Sections 4.2.5.3 and 4.3.1.3. Assuming that there are  $k$  angular intervals of length  $\Delta\theta_{\text{flu}}$  before the arc corresponding to gantry angle  $\theta^k$ , the probability  $\Pr(\Omega^{k,l \rightarrow l'})$  required can be calculated from the Markov chain model as the probability to observe tumour position  $\omega^l$  at time  $\Delta t = k\tau$ , multiplied by the probability to transition from  $\omega^l$  to  $\omega^{l'}$  in a time interval  $\Delta t = \tau$  (see Equation (C.7)). These two component probabilities can be calculated using a similar set of equations as Equations (3.12) to (3.14).

#### 4.2.1.2 Probabilistic optimization

As stated in Section 4.2.1, the goal of PO is to minimize the expectation value of the objective function,  $\bar{F}$ . If a quadratic difference form of the objective function  $F(d)$  is used, as is typical, then it has been shown that the expectation value of the objective function for a treatment course of  $N_{\text{frac}}$  fractions is equal to the objective function evaluated at the expectation value of the dose plus a variance term which includes the variance of the dose of each voxel [55]:

$$\bar{F} = F(\bar{d}) + \frac{V}{N_{\text{frac}}}. \quad (4.10)$$

Both the expectation value and variance of the dose are here evaluated for a single fraction only. The variance term  $V$  is a linear combination of the dose variance of each voxel  $i$ ,  $\sigma_{d_i}^2$ . In this study, only voxels in the target volume are included in the variance term  $V$  [84].

Equation (4.10) then has a simple interpretation: the expectation value of

the objective function is equal to the sum of two terms: the intuitive guess of  $F(\bar{d})$ , and an additional term encoding the variation of the dose due to a single fraction from its expectation value. This latter term is inversely proportional to the number of fractions in the treatment, which is a reflection of the fact that the random fluctuations of the dose in any one fraction become less significant to the dose summed over  $N_{\text{frac}}$  fractions as this number increases.

The expectation value of the dose can be calculated using Equation (4.6). The variance of the dose to voxel  $i$  is equal to the difference of the expectation value of the squared dose and the square of the expectation value,  $\sigma_{d_i}^2 = \overline{d_i^2} - \bar{d}_i^2$ . The second term can be calculated by squaring Equation (4.6), while the first is given by:

$$\overline{d_i^2} = \int \text{Pr} [\Omega] d_i^2 [\Omega] \mathcal{D}\Omega, \quad (4.11)$$

To continue this calculation, two instances of  $d_i [\Omega]$  from Equation (4.5) must be inserted into this equation, each with different integration and summation variables  $\theta$ ,  $\omega$ , and  $j$ :

$$\begin{aligned} \overline{d_i^2} = & \iiint \iiint \text{Pr} [\Omega] \delta (\omega_1 - \Omega (\Theta^{-1} (\theta_1))) \tilde{D}_{ij_1} (\theta_1, \omega_1) \frac{d\phi_{j_1}}{d\theta_1} \\ & \cdot \delta (\omega_2 - \Omega (\Theta^{-1} (\theta_2))) \tilde{D}_{ij_2} (\theta_2, \omega_2) \frac{d\phi_{j_2}}{d\theta_2} d\theta_1 d\omega_1 d\theta_2 d\omega_2 \mathcal{D}\Omega. \end{aligned} \quad (4.12)$$

The explicit dependence of the differential fluence  $\frac{d\phi_j}{d\theta}$  on the leaf positions  $(L_J, R_J)$  has been removed to clean up the notation. The conversion of this complicated continuous expression into a discrete formulation which can be

computed is performed in Appendix D.

## 4.2.2 Calculation of the 4D dose-influence matrix

All optimization and dose calculation was performed within the open-source treatment planning system matRad [1]. The Monte Carlo code VMC++ [65] and the accompanying deformable geometry library [66] were used for the calculation of the deformed dose-influence matrices  $\tilde{D}_{ij}^{kl}$ . A Monte Carlo calculation engine was chosen here over the pencil beam algorithm used throughout Chapter 2 for its improved dose calculation accuracy in low-density volumes (i.e., the lung). The existing interface between the matRad and VMC++ code was expanded to support the deformable geometry library, and to enable the use of a phase space particle input source (previously only spectrum sources were supported). A phase space file taken from a BEAMnrc model of a 6 MV Elekta Infinity™ linac with Agility™ MLC (Stockholm, Sweden) linac was used throughout this study [90].

The deformable geometry library code requires as input both a reference patient geometry, as well as the DVF, the set of vectors defining the motion of each anatomical voxel from the reference geometry to the deformed geometry. The simulation of the energetic particle transport is performed in the deformed geometry, while the dose is accumulated in the original voxels in the reference geometry, as described in Section 4.2.1.1. The acquisition of both the reference geometry as well as the DVF for each deformed geometry is described in Section 4.2.3.

A dose angle interval of  $\Delta\theta_{\text{dose}} = 4^\circ$  was selected following the method introduced and results presented in Sections 4.2.5.2 and 4.3.1.2, respectively.

### 4.2.3 4D digital phantom

The XCAT numerical phantom [64] was used to generate the CT phantom at the reference patient geometry (maximum exhale), as well as the DVFs required for calculating dose to the deformed patient geometries. A total of 5 such deformed geometries were considered (see Sections 4.2.5.2 and 4.3.1.2 for a discussion on this), with each geometry chosen to be representative of one of 5 tumour position bins. These equal-width position bins were constructed from the trajectory of the tumour position, as shown in Figure 4.6. The geometry corresponding to the mean tumour position in each bin was then selected as the representative geometry for that bin. Aside from the fact that inhaling is faster than exhaling, the XCAT phantom makes no distinction between the inhaling and exhaling portions of the trajectory. That is, the phantom does not exhibit hysteresis, which would occur if the tumour travelled a different path through the lung for these two portions. For this reason, only the geometries corresponding to the inhaling part of the trajectory are required, as seen in Figure 4.6.

Since dose is accumulated on the reference geometry, all OAR and target volumes were contoured on this image following clinical guidelines at The Ottawa Hospital for lung tumour treatment planning. The CTV was obtained by performing a 5 mm radial expansion around the visible tumour (GTV).

In order to calculate the dose to the deformed geometry, the DVFs used as input to VMC++ must not lead to any instances of voxel collapse. This occurs if a voxel vertex in the reference geometry is pushed through one of the voxel faces in the deformed geometry, inverting the voxel orientation; particle transport cannot be done in this case. Any instances of voxel collapse in the XCAT DVFs were eliminated by applying a moving average filter to the DVF. This presents a simplified yet more complete method compared to one proposed by Eiben et al. [91], which separated the phantom into two regions, the inner region (internal organs such as the lungs and heart), and the outer region on which these organs slide (chest wall, ribs, etc.). In this method, voxel collapse is eliminated in each of these two regions separately, but persists along the boundary between these two regions, which precludes the use of the resulting DVF in calculating dose to deformable geometries with VMC++.

A total of four patient cases, all simulated in XCAT as described above, were considered for the optimization study, with varying tumour size and location, as listed in Table 4.1. A CT image for patient number 1 is shown in Figure 4.5, along with contours of the target volumes and OARs. The same DVFs were used for each of the patient cases; the motion of the tumour was taken as that of the surrounding lung tissue. A graph of the tumour trajectory for patient 1 is shown in Figure 4.6. Two tumour sizes were included to investigate the effect of the size with respect to the motion amplitude on the plan quality and robustness. The different positions were selected to vary the

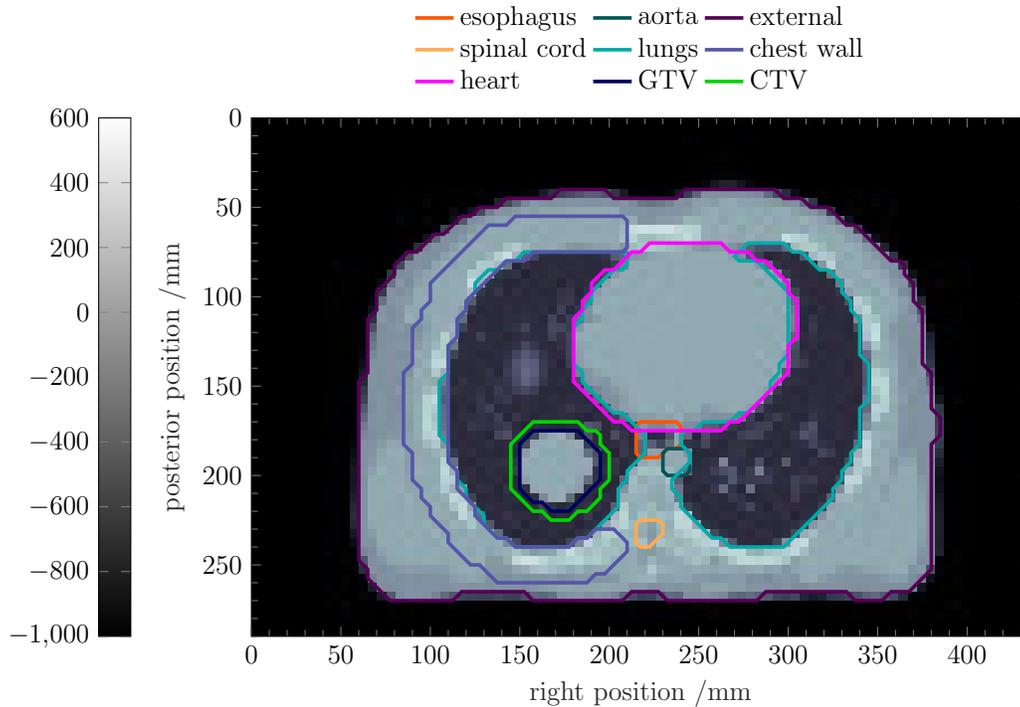


Figure 4.5: Transverse slice of the XCAT-generated reference CT geometry for patient 1, at the position of the centre of the tumour. Also shown are the contours for the target volumes (GTV and CTV), and the OARs. The chest wall is defined as a 2 cm expansion away from the ipsilateral lung of the tumour, while the external contour encloses the entire body. The trachea was also contoured, but is not visible on this slice.

resulting tumour motion amplitude, as well as to place the tumour nearer to different OARs.

#### 4.2.4 4D optimization and parameters

As mentioned in Section 4.2.2, all plan optimization was implemented and performed within the open-source treatment planning system matRad [1].

Table 4.1: Tumour parameters for the simulated patient phantoms investigated in the 4D optimization study. For each patient a spherical tumour with varying diameter is inserted at different positions in the lung. The resulting motion amplitude at the centre of the tumour is also shown.

Patient number	Tumour diameter /mm	Anatomical location	Motion amplitude /mm
1	40	Right lung, near the diaphragm	11.5
2	20	Right lung, near the diaphragm	10.4
3	20	Right lung, near the heart	8.5
4	20	Left lung, near the spinal cord	3.3

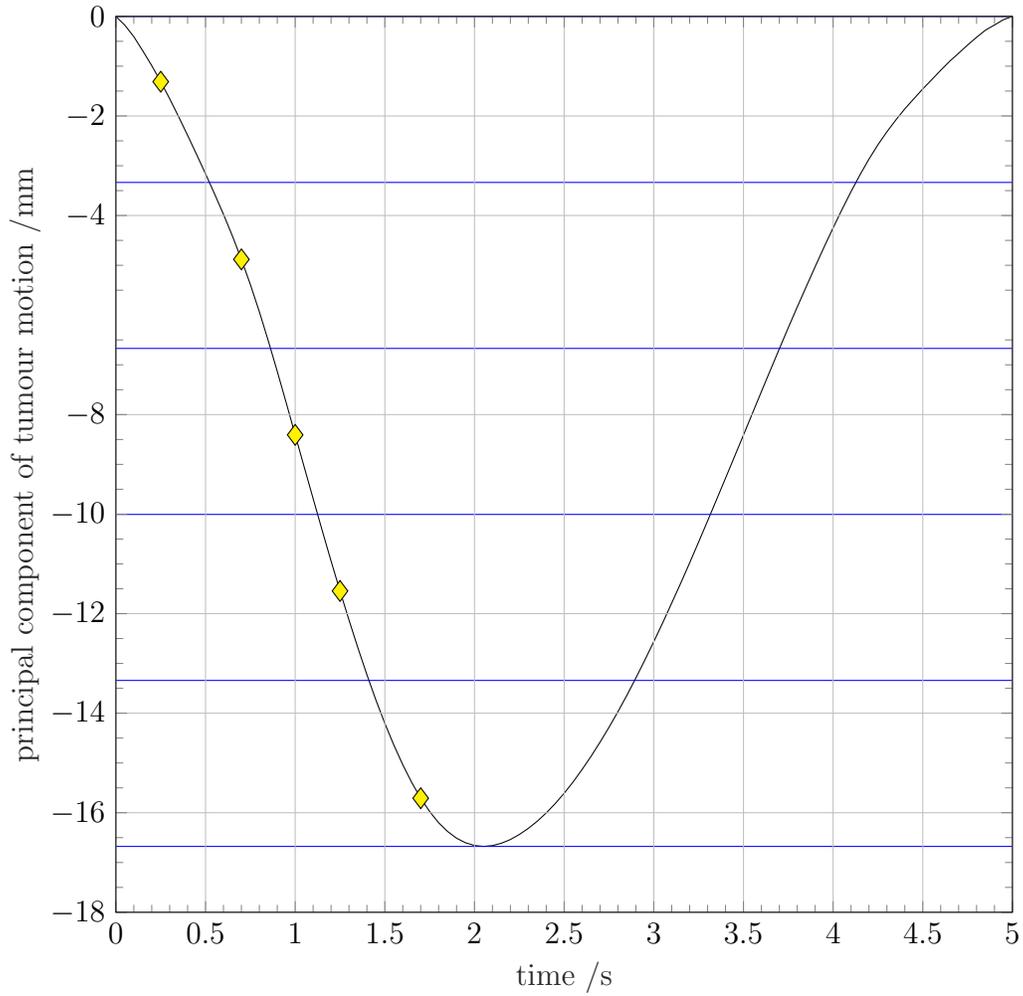


Figure 4.6: Trajectory of the embedded tumour in the XCAT phantom for patient 1. The principal component of the full three-dimensional motion is shown in black. A total of 100 geometries were sampled and sorted into 5 bins, whose borders are indicated by the blue lines. The mean tumour position for each bin, considering both exhaling and inhaling portions of the trajectory, is represented by the yellow diamonds. The corresponding patient geometry for each of these positions is then selected to represent the bin.

The aperture libraries for the robust 4D-VMAT optimizations were initialized by first performing fluence map optimization (FMO) on the reference CT geometry, at an angular interval of  $\Delta\theta_{\text{FMO}} = 28^\circ$ . This was followed by a leaf sequencing step, in which the fluence maps were sequenced into a finite set of apertures [68] surrounding the patient at an angular interval of  $\Delta\theta_{\text{DAO}} = 8^\circ$ . These apertures were then deformed from the reference geometry to each geometry in the aperture library by tracking the relative motion of the target. Specifically, the average of the DVF along each bixel ray was computed, counting only voxels inside the target, as described by Feng et al. [92].

Finally, DAO was used to optimize the aperture library. The adaptive nature of the delivery was taken into account during the optimization using the framework described in Section 4.2.1.1; in particular, the expectation value of the dose when considering all respiratory trajectories was used, as shown in Equation (4.8). During DAO, constraints were placed on the MLC leaf speed and MU rate, with minimum and maximum values listed in Table 2.1.

It can be seen from the description in Section 4.2.1.2 that the plan robustness to the variation of the patient's respiratory motion pattern increases as the number of fractions  $N_{\text{frac}}$  of a treatment also increases. To investigate this effect, two fractionation schemes were considered: a conventional treatment course with a prescription dose of 66 Gy delivered in  $N_{\text{frac}} = 33$  fractions, and an SBRT course with 60 Gy delivered in  $N_{\text{frac}} = 5$  fractions. The CTV was used as the treatment volume receiving the prescription dose for all of the optimization methods considered in this study, described in Section 4.2.6.

The plan quality for both the target and OARs was assessed by comparing DVH metrics to treatment planning guidelines currently in use at The Ottawa Hospital.

The respiratory model trained with the patient fraction A data in Section 3.4.1 was used to optimize all treatment plans, as well as to evaluate the plan robustness, using a method described in Section 4.2.6. The gantry speed was fixed throughout delivery to maintain the correspondence between the fluence gantry angle interval  $\Delta\theta_{\text{flu}}$  and the Markov chain time interval  $\tau$ , following Equation (4.9). For the conventionally-fractionated plans, a delivery time of 95 s was used, with a corresponding interval of  $\Delta\theta_{\text{flu}} = 1.3^\circ$  ( $\tau = 386.9$  ms). For the SBRT plan, a longer delivery time of 410 s was used to accommodate for the larger dose per fraction. This delivery time corresponds to an interval of  $\Delta\theta_{\text{flu}} = 0.3^\circ$ , the selection of which was validated for the conventionally-fractionated plan as described in Section 4.2.5.3. Following the discussion in that passage and in Section 4.2.1.1, the errors introduced by a discrete set of fluence calculation angles are related to the length of time required for the gantry to rotate from one angle to the next: for example, a longer time between angles would result in larger tumour and aperture transitions made during this time.

#### 4.2.5 Validation and verification

The 4D-VMAT optimization framework ALERT-RA was implemented within the open-source treatment planning system matRad [1]. This implementation

was validated using a series of tests, designed to verify the consistency of the calculated expectation value and variance of the dose, and to justify both the number of deformed patient geometries used for dose calculation and the angular interval  $\Delta\theta_{\text{flu}}$ .

All validation was performed using conventionally-fractionated VMAT plans optimized on patient 1. The first and third studies, the details of which are described in Sections 4.2.5.1 and 4.2.5.3, respectively, were performed on an aperture library taken before 4D-VMAT plan optimization was done; that is, immediately after the leaf sequencing step. The dose calculation grid study, described in Section 4.2.5.2, used a 3D-VMAT plan optimized to deliver dose to the CTV in the reference patient geometry.

#### 4.2.5.1 Dose expectation value and variance

The expectation value and variance of the dose to voxel  $i$  for a single fraction,  $\bar{d}_i$  and  $\sigma_{d_i}^2$ , are both calculated using the equations developed throughout Sections 4.2.1.1 and 4.2.1.2 and Appendices C and D. These quantities are calculated using a probabilistic Markov chain model of the patient's respiratory motion. To verify the accuracy of these equations and their implementation, the analytical calculations of  $\bar{d}_i$  and  $\sigma_{d_i}^2$  were compared to a random sampling of these same variables.

Specifically, the dose to an individual fraction was simulated by first sampling a respiratory trajectory  $\omega = \Omega(t)$  from the motion model, following the method of Section 3.2.3, then calculating the dose due to the sampled

trajectory using Equation (4.5). In practice the dose is calculated using Equations (4.7) and (4.8), with the probability of each trajectory segment  $\Omega^{k,l \rightarrow l'}$  equal to 0, except for those trajectory segments which make up the actual sample trajectory of interest, whose probabilities are set to 1. This was performed for a total of  $N_{\text{sample}} = 50\,000$  samples; the dose to voxel  $i$  for the  $n$ th sample is given the symbol  $d_{i,n}$ .

The sample mean of the dose to voxel  $i$  was calculated from the individual sample doses:

$$\tilde{d}_i = \frac{1}{N_{\text{sample}}} \sum_{n=1}^{N_{\text{sample}}} d_{i,n}. \quad (4.13)$$

This value can be compared to its analytical counterpart  $\bar{d}_i$  by making use of the central limit theorem [93]: as  $N_{\text{sample}} \rightarrow \infty$  the distribution of the random variable  $\tilde{d}_i$  converges to a normal distribution with mean  $\bar{d}_i$  and variance  $\frac{\sigma_{d_i}^2}{N_{\text{sample}}}$ , independent of the distribution of the dose to an individual sample  $d_{i,n}$ . As stated earlier,  $\bar{d}_i$  and  $\sigma_{d_i}^2$  correspond to the analytical calculations of the expectation value and variance of the dose to voxel  $i$ . Equivalently, the normalized dose difference

$$\epsilon_i = \frac{\tilde{d}_i - \bar{d}_i}{\sqrt{\sigma_{d_i}^2 / N_{\text{sample}}}} \quad (4.14)$$

converges to the standard normal distribution, i.e. with mean 0 and variance 1. Note that this is true under the hypothesis that the calculated expectation value and variance accurately represent their true values.

This hypothesis can be tested for all  $N_{\text{vox}}$  voxels collectively by calculating the sum of the squares of the  $\epsilon_i$ :

$$\chi^2 = \sum_{i=1}^{N_{\text{vox}}} \epsilon_i^2. \quad (4.15)$$

This should follow the  $\chi^2$ -distribution with  $N_{\text{vox}}$  degrees of freedom [93], thus a  $p$ -value can be calculated which represents the probability of obtaining a  $\chi^2$  value at least as extreme as the one observed, given that the analytical values of the expectation value and variance of the dose are correct. The results of this statistical analysis are shown in Section 4.3.1.1.

#### 4.2.5.2 Dependence on dose calculation angular and number of dose calculation geometries

As described in Section 4.2.1.1, the grid of points  $\{(\theta^k, \omega^l)\}$  is selected to fulfill two conditions, one of which is that  $\tilde{D}_{ij}(\theta, \omega)$  should be relatively constant over the intervals of length  $\Delta\theta$  and  $\Delta\omega$  surrounding each point  $(\theta^k, \omega^l)$ . As mentioned in Section 2.4, the smaller these intervals are, the more computer memory and calculation time are required;  $\Delta\theta$  and  $\Delta\omega$  should therefore be set at the largest value which fulfills this condition.

These intervals can be selected by following the general method of Section 2.2.3; that is, by recalculating the dose due to a particular plan at successively smaller dose angle intervals  $\Delta\theta_{\text{dose}}$  (or with increasingly more geometries corresponding to independent tumour positions  $\omega^l$ ) until the dose

has converged. In Section 2.2.3, convergence is quantified by comparing the dose calculated using a particular dose calculation angular interval to the reference dose ( $\Delta\theta_{\text{dose}} = 0.5^\circ$ ). Specifically,  $\Delta\theta_{\text{dose}}$  is deemed to be sufficiently small when 99% of the voxels in the volume of interest have less than 1% dose difference with respect to the reference dose.

When using a Monte Carlo dose calculation engine, the dose to each voxel is a random quantity; therefore it is important to separate dose differences due to the underlying statistical uncertainty and the sought-after systematic differences, which are to be minimized. The method introduced by Kawrakow and Fippel [94] does this by building a statistical model of the dose difference to each voxel, factoring in the type A uncertainty. First, the relative dose difference to each voxel  $i$  is calculated as (note that throughout this section, the resolution in the tumour position direction may be omitted, but the analysis is meant to apply to both directions  $\theta$  and  $\omega$ ):

$$x_i = \frac{D_i(\Delta\theta_{\text{dose}}) - D_i(\text{ref})}{\sqrt{\sigma_i^2(\Delta\theta_{\text{dose}}) + \sigma_i^2(\text{ref})}}. \quad (4.16)$$

Here,  $D_i(\Delta\theta_{\text{dose}})$  and  $D_i(\text{ref})$  refer to the mean of the dose calculated using the  $\Delta\theta_{\text{dose}}$  angular interval and the reference interval, respectively. The  $\sigma_i^2$  represent the statistical variance of these respective quantities due to the Monte Carlo calculation uncertainty only; that is, they have nothing whatsoever to do with the variance due to the uncertainty in the patient's respiratory trajectory. The denominator in Equation (4.16) is therefore the

combined statistical uncertainty (CSU) of the numerator.

Assuming no correlations between the two dose distributions, as well as the absence of any systematic differences,  $x_i$  should follow the standard normal distribution. The first assumption could be eliminated, and the second somewhat relaxed, by fitting the distribution of  $x_i$  instead to the following parametrized distribution:

$$f(x) = \frac{1}{\sqrt{2\pi\sigma^2}} \left( \alpha_1 \exp\left(-\frac{(x - \delta_1)^2}{2\sigma^2}\right) + \alpha_2 \exp\left(-\frac{(x - \delta_2)^2}{2\sigma^2}\right) + (1 - \alpha_1 - \alpha_2) \exp\left(-\frac{x^2}{2\sigma^2}\right) \right). \quad (4.17)$$

The third term, the normal distribution centred at  $x = 0$ , represents the assumption of no systematic differences. This term has its weight reduced by two factors  $\alpha_1$  and  $\alpha_2$ , both restricted to values between 0 and 1, such that  $\alpha_1 + \alpha_2 \leq 1$ . Each factor multiplies a normal distribution shifted by  $\delta_1$  and  $\delta_2$ , respectively; these distributions therefore represent possible systematic differences between the two dose. Finally, the parameter  $\sigma^2 \in (0, 1]$  is introduced to account for possible correlations between  $D_i(\Delta\theta_{\text{dose}})$  and  $D_i(\text{ref})$  due to the use of the same radiation source model [95].

Larger values of the parameters  $\alpha_1$ ,  $\alpha_2$ ,  $\delta_1$ , and  $\delta_2$  indicate greater systematic differences between the two doses. These parameters were determined for each  $\Delta\theta_{\text{dose}}$  using a maximum likelihood fit of  $x_i$  to  $f(x)$ . The correlation

parameter  $\sigma^2$  in Equation (4.17) was determined before these fits by comparing two dose distributions calculated using the same dose calculation angular interval. A modified version of the distribution in Equation (4.17) was fit to the resulting  $x_i$ , with the systematic parameters restricted to 0 (effectively a normal distribution with unknown variance  $\sigma^2$ ). Statistical independence was ensured for all of the calculations described here by varying the seed input for the random number generator in each dose calculation, and by randomly varying the starting particle in the phase space file.

This method was used to determine the convergence first for the dose angle interval  $\Delta\theta_{\text{dose}}$ ; the reference case for this study was taken to be  $\Delta\theta_{\text{dose}} = 0.5^\circ$ , as in Section 2.2.3. Dose was recalculated based on a 3D-VMAT plan optimized to the reference patient geometry at successively smaller values of  $\Delta\theta_{\text{dose}}$ , always using the continuous aperture method derived in Section 2.2.1.1. For this particular portion of the study, the dose was recalculated on this same reference geometry: no motion effects were included. Only voxels with dose greater than 50% of the maximum dose were included in the fit of  $f(x)$ . Convergence was deemed to have been reached when both  $\delta_1$  and  $\delta_2$  were small enough such that the maximum relative systematic dose difference per voxel is less than 1%. These parameters are defined the following way:

$$\Delta_1 = \text{RCSU}_{\text{max}} \cdot |\delta_1|, \quad (4.18a)$$

$$\Delta_2 = \text{RCSU}_{\text{max}} \cdot |\delta_2|. \quad (4.18b)$$

Since in the definition of  $x_i$  in Equation (4.16) the dose difference has been normalized by the CSU, each  $\delta$  parameter should be multiplied by the CSU per voxel to obtain a value for the systematic difference in units of dose. A factor of  $\sigma$  must also be multiplied into the CSU to offset its inclusion in the definitions of  $\delta_1$  and  $\delta_2$  through Equation (4.17); in this thesis, any reported values of the CSU include this factor of  $\sigma$ . The resulting quantities,  $\text{CSU}_i \cdot \delta_1$  and  $\text{CSU}_i \cdot \delta_2$ , can then be divided by the reference dose to voxel  $i$ , giving a relative systematic dose difference per voxel  $i$ . This can be accomplished by multiplying each  $\delta$  parameter by the relative combined statistical uncertainty (RCSU) instead of the CSU. The maximum absolute value of these quantities across all voxels are taken to be the relative systematic dose differences,  $\Delta_1$  and  $\Delta_2$ .

Following this, the number of patient geometries corresponding to independent tumour positions  $\omega^l$  was increased from 1 to 10 to determine convergence, with the dose calculated on 10 geometries used as the reference. The patient geometries were selected such that the corresponding tumour positions were equal to the mean position for each position bin, as described in Section 4.2.3 and shown in Figure 4.6. Dose was again recalculated based on a 3D-VMAT plan optimized to the reference patient geometry; the dose angle interval of convergence (Section 4.3.1.2),  $\Delta\theta_{\text{dose}} = 4^\circ$  was used throughout these calculations. Since the focus of this study is limited to the patient geometries and not on the probabilistic motion model, a very simple model was used. The same 3D-VMAT plan was delivered to each geometry independently and

then averaged, with the probability for each tumour position bin  $l$  given as the fraction of time spent in that bin (see Figure 4.6).

#### 4.2.5.3 Fluence calculation angle interval

As described in Section 4.2.1.1, an additional set of gantry angles,  $\{\theta_{\text{flu}}\}$ , is introduced to ensure that when calculating the dose delivered in the trajectory segment  $\Omega^{k,l \rightarrow l'}$ , no significant dose calculation errors are made by ignoring the dose deposited to patient geometries between those corresponding to tumour positions  $\omega^l$  and  $\omega^{l'}$ . These errors are caused by simplifying the corresponding tracked MLC motion from a trajectory which might have pass through many other control points in the aperture library as the tumour transitions between  $\omega^l$  and  $\omega^{l'}$ , and also by only depositing dose to the two corresponding geometries, as illustrated in Figure 4.4. The interval between fluence calculation angles,  $\Delta\theta_{\text{flu}}$ , was selected to correspond to the Markov chain time interval  $\tau$  of the respiratory motion model, using a fixed gantry rotation speed  $\dot{\theta}$ .

For the model optimized for patient fraction A in Section 3.4.1, a time interval of  $\tau = 386.9$  ms was used. Using a delivery time of 95 s for a full  $360^\circ$  rotation gives a corresponding angular interval of  $\Delta\theta_{\text{flu}} = 1.3^\circ$ . This model therefore assumes that the tumour moves instantaneously from one discrete position to the next every 386.9 ms (or  $1.3^\circ$ ), with no dose deposition to intermediary patient geometries. The dose calculation errors introduced by this assumption can be assessed using the methods of Sections 2.2.3

and 4.2.5.2: that is, by comparing the dose calculated at  $\Delta\theta_{\text{flu}} = 1.3^\circ$  to a reference dose distribution calculated at a smaller angular interval  $\Delta\theta_{\text{flu}}$ .

In this case, the reference dose distribution was selected at  $\Delta\theta_{\text{flu}} = 0.1465^\circ$ , corresponding to a time interval of 40 ms. This was chosen to coincide with the time to update the Elekta Agility MLC, which has been proposed for real-time tumour tracking and treatment plan adaptation [96, 97]. The first of the two sources of error described above, associated with a simplified leaf trajectory, is therefore eliminated entirely, since for this particular MLC the leaves cannot be updated any faster. The second source of error can be initially dealt with heuristically by considering that a typical breathing period is 3 s to 5 s; a time interval of 40 ms therefore represents at most a  $\sim 1\%$  portion of the breathing period. 5 unique deformed patient geometries are used to represent the entire respiratory cycle for the purposes of dose calculation (see Sections 4.2.5.2 and 4.3.1.2), therefore 10 total geometries are passed through as the patient inhales and then exhales. Assuming for the moment that approximately equal time is spent in each patient geometry, a time interval representing a  $\sim 10\%$  portion of the breathing period would be sufficiently small to eliminate the possibility of skipped patient geometries. Thus the selection of 40 ms far exceeds this requirement. It is shown later, in Section 4.3.1.3, that the dose distribution converges for values of  $\Delta\theta_{\text{flu}}$  far greater than  $0.1465^\circ$  (indeed, there is minimal systematic difference between  $\Delta\theta_{\text{flu}} = 1.3^\circ$  and  $\Delta\theta_{\text{flu}} = 0.1465^\circ$ ), confirming the preceding argument.

For values of  $\Delta\theta_{\text{flu}}$  with a corresponding time interval shorter than the

$\tau = 386.9$  ms used in the Markov chain model, the expectation value of the dose cannot be calculated analytically. This is because when calculating the probability to transition between two tumour position bins, the transition time is limited to integer multiples of  $\tau$ . In this case, since the reference dose uses time intervals of 40 ms, a direct calculation of the expectation value of the dose is not possible. Instead, a Monte Carlo sampling of the Markov chain model was used to calculate the sample mean, as described in Section 4.2.5.1 and using Equation (4.13), as well as the sample variance of the mean:

$$s_{d_i}^2 = \frac{1}{N_{\text{sample}} (N_{\text{sample}} - 1)} \sum_{n=1}^{N_{\text{sample}}} (d_{i,n} - \tilde{d}_i)^2. \quad (4.19)$$

A total of  $N_{\text{sample}} = 50\,000$  were used to calculate these quantities. For each sample, a sequence of position bins  $\omega^l$  is sampled at the time interval  $\tau = 386.9$  ms. To calculate the MLC leaf movement, the resulting fluence, and delivered dose to corresponding patient geometry at the required time interval of 40 ms, each sequence was resampled to 40 ms before dose was calculated. This was done by using linear interpolation, and then rounding the result to the nearest integer corresponding to the bin number  $l$ .

Following this, the analysis described in Section 4.2.5.2 was applied to compare the trial dose distribution calculated using  $\Delta\theta_{\text{flu}} = 1.3^\circ$  (386.9 ms),  $D_i(\Delta\theta_{\text{flu}})$ , to the reference stochastic distribution calculated at  $\Delta\theta_{\text{flu}} = 0.1465^\circ$  (40 ms),  $D_i(\text{ref})$ . The variable  $x_i$  was calculated using a modified version of

Equation (4.16):

$$x_i = \frac{D_i(\Delta\theta_{\text{flu}}) - D_i(\text{ref})}{\sqrt{\sigma_i^2(\Delta\theta_{\text{flu}}) + \sigma_i^2(\text{ref})}}. \quad (4.20)$$

In this case,  $\sigma_i^2(\text{ref})$  refers to the variance of the mean of the dose to voxel  $i$  with respect to the uncertainty in the patient's respiratory trajectory for the reference distribution, and is given by Equation (4.19).  $\sigma_i^2(\Delta\theta_{\text{flu}})$ , on the other hand, is set to 0, since the mean for this dose distribution was calculated analytically, and is therefore perfectly reproducible. The distribution of  $x_i$  was then fit to Equation (4.17), and the parameters  $\alpha_1$ ,  $\alpha_2$ ,  $\delta_1$ , and  $\delta_2$  extracted.

#### 4.2.6 Comparison of 4D methods

The proposed methods of robust optimization, CO and PO, were tested by comparing the quality and robustness of plans optimized among a total of six different methods. Specifically, plan quality was analyzed using DVHs of the expectation value of the total dose across  $N_{\text{frac}}$  fractions. Since the dose, and the resulting DVH, is subject to uncertainty, percentile dose-volume histograms (PDVHs) were calculated to determine the robustness of plans to changes in the patient's respiratory trajectory. Whereas the DVH indicates the percentage volume which receives a dose greater than some dose  $D$ , for a given percentage  $p$ , the PDVH indicates the percentage volume which receives dose greater than  $D$  with probability  $p$  [98]. The PDVHs were generated from the DVHs of a total of 100 simulated treatment courses, each consisting of the sum of the doses from  $N_{\text{frac}}$  simulated fractions. The dose to an

individual fraction was simulated using the same procedure as described in Section 4.2.5.1, by sampling a respiratory trajectory  $\Omega$  from the motion model and calculating the respective dose  $d_i[\Omega]$  using Equation (4.5). The PDVHs therefore simulate the actual dose delivery of the plan, using a particular motion model. Unless otherwise noted, trajectories were sampled starting with an initial position bin of  $l = 1$  and in the end of exhale breathing phase.

The CO and PO robust optimization methods described in Section 4.2.1 were both compared. The Markov chain respiratory motion model introduced in Chapter 3 was used as the probabilistic respiratory motion model, using the optimal combination of parameters for fraction A found in Section 3.4.1. As briefly mentioned in Section 3.5, fraction model intercomparisons were also performed to determine the effect of changing the respiratory motion model from optimization to delivery. This was done by calculating PDVHs with trajectories sampled from the fraction B model found in Section 3.4.2, rather than the fraction A model used for optimization.

Two variations of 3D-VMAT optimization were also tested: 3D-CTV and MKO. This latter optimization method assumes that the position of the tumour follows a particular probability distribution throughout delivery. This was implemented using an averaged dose-influence matrix during optimization, calculated by multiplying the matrix elements  $\tilde{D}_{ij}^{kl}$  by the stationary probability distribution for each tumour position bin  $l$ , and summing over all

$l$ :

$$\tilde{D}_{ij,\text{MKO}}^k = \sum_l \pi_l \tilde{D}_{ij}^{kl}. \quad (4.21)$$

The stationary distribution  $\pi_\alpha$  was obtained from the Markov chain motion model according to Section 3.2.1;  $\pi_l$  then was calculated by accumulating the probabilities from each state  $\alpha$  into their respective position bin using the lookup tables  $l_\alpha$ , defined in Section 3.3.1. The MKO method therefore results in plans for which the apertures encompass the entire range of motion of the tumour throughout delivery. Plans are expected to be very robust to changes in respiratory motion, with an increase in dose to OARs due to the larger irradiated volume.

The 3D-CTV method instead delivered the prescription dose to the CTV only; thus for optimization dose was calculated on the reference geometry only. This method was included to assess the impact of the respiratory motion on a plan without motion compensation. For both the MKO and 3D-CTV methods, the trajectories generated when calculating the PDVHs were sampled under the assumption that the initial position bin followed the stationary distribution  $\pi_l$ . This was chosen to reflect the fact that if no tumour tracking is available, treatment delivery would begin at a random tumour position defined by this distribution.

Finally, the two previously-published methods of 4D-VMAT optimization mentioned in Section 4.1 were considered: single trajectory optimization (STO) [52] and direct aperture deformation (DAD) [51]. The former method

assumes that the patient's breathing motion is unchanging from one fraction to the next, while the latter makes no assumptions or provisions for the patient motion pattern, except to simply deform the aperture to follow the tumour motion. The STO method is not expected to be robust to variations in the respiratory trajectory, while higher doses to OARs may be delivered with DAD.

STO was accomplished by assuming that the trajectory of the tumour did not change from one respiratory cycle; the trajectory was chosen by assuming that at each Markov chain time interval  $\tau$ , the transition with the highest probability was taken. Following this, only the corresponding apertures which would be selected from the entire aperture library for this particular trajectory  $\Omega$  were optimized.

The DAD aperture library was constructed using the 3D-CTV VMAT plan as a basis. The aperture at each gantry angle was then deformed to track the relative motion of the tumour for each patient geometry, as described in Section [4.2.4](#).

## 4.3 Results

### 4.3.1 Validation and verification

#### 4.3.1.1 Expectation value and variance

A histogram of the normalized dose difference for voxel  $i$  as calculated in Equation (4.14),  $\epsilon_i$ , is shown in Figure 4.7. A total of  $N_{\text{vox}} = 1672$  entries are shown, which represent all voxels with a dose at least as high as 50% of the maximum dose. As described in Section 4.2.5.1, assuming that the analytical calculations of both the mean dose  $\bar{d}_i$  and variance  $\sigma_{d_i}^2$  are correct,  $\epsilon_i$  should follow the standard normal distribution (also shown). Visually this seems to be the case, as the histograms agrees well with the probability density function of the standard normal.

Following Equation (4.15),  $\chi^2$  is calculated as 1425 for this particular set of  $\epsilon_i$ . The corresponding  $p$ -value is 1.000, using  $N_{\text{vox}} = 1672$  degrees of freedom ( $1 - p = 4 \cdot 10^{-6}$ ). This demonstrates that the form and implementation of the analytical calculations for both the mean dose  $\bar{d}_i$  and variance  $\sigma_{d_i}^2$  are highly likely to be correct.

#### 4.3.1.2 Dependence on dose calculation angular and number of dose calculation geometries

The results of the dose angle interval study are shown in Table 4.2. The dose distribution calculated using each value of  $\Delta\theta_{\text{dose}}$  was compared to

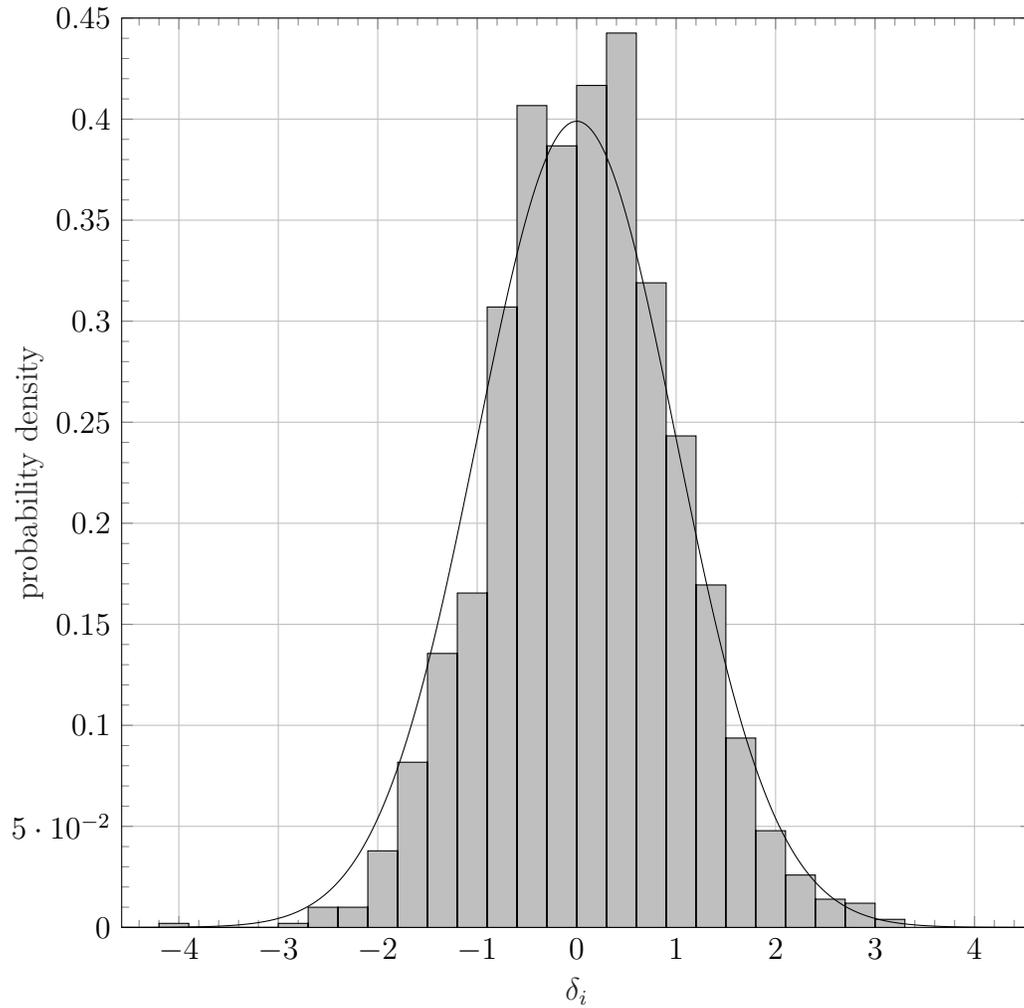


Figure 4.7: Histogram of the normalized dose difference  $\epsilon_i$ , scaled to display the probability density (the total area of the bars is equal to 1). Also shown for comparison is a graph of the probability density function of the standard normal distribution (mean of 0 and variance of 1).

Table 4.2: Results of the dose distribution comparison for the dose angle interval study in lung. Each dose distribution was compared to the reference dose calculated with  $\Delta\theta_{\text{dose}} = 0.5^\circ$ , using the method described in Section 4.2.5.2. Along with the parameters fitted to the distribution of Equation (4.17), the maximum RCSU and the relative systematic dose differences,  $\Delta_1$  and  $\Delta_2$ , are reported ( $\Delta_i = \text{RCSU}_{\text{max}} \cdot |\delta_i|$ ).

$\Delta\theta_{\text{dose}}/^\circ$	$\text{RCSU}_{\text{max}}/\%$	$\alpha_1$	$\delta_1$	$\Delta_1/\%$	$\alpha_2$	$\delta_2$	$\Delta_2/\%$
1	0.25	0.49	0.12	0.03	0.51	0.12	0.03
2	0.33	0.50	-0.09	0.03	0.50	-0.09	0.03
4	0.41	0.96	-0.03	0.01	0.02	0.76	0.31

the reference dose, calculated with  $\Delta\theta_{\text{dose}} = 0.5^\circ$ , using the statistical model described in Section 4.2.5.2. The results of the comparison show that  $\Delta\theta_{\text{dose}} = 4^\circ$  is sufficiently small to give an accurate dose distribution in lung, since the systematic dose difference quantities of  $\Delta_1$  and  $\Delta_2$  are both less than 1%. This dose angle interval was therefore used for the lung phantoms throughout rest of the study.

Table 4.3 shows the results of the study conducted to determine the number of patient geometries to use when calculating dose to a moving patient. The same methods of analysis were used here to compare each resulting dose distribution to the reference dose, which used 10 patient geometries corresponding to independent tumour positions. The results of the comparison show that 4 patient geometries are sufficient to obtain a systematic dose difference of less than 1%. Note that although for 5 patient geometries  $\Delta_1 = 2.00\%$ ,  $\alpha_1$ , which is the proportion of voxels which display a

systematic difference of  $\delta_1$ , is equal to 0. For this number of patient geometries therefore, the only voxels which can be said to have a systematic difference are the  $\alpha_2 = 0.04$  fraction, with  $\Delta_1 = 0.68\%$ .

Although 4 patient geometries are seen to be sufficient for accurate dose calculation, 5 geometries were used throughout the optimization study. This was done for the purposes of similarity with the standard clinical practice of acquiring 10 phase-binned frames in a 4D-CT. Since this type of 4D-CT distinguishes between the inhaling and exhaling portions of the respiratory cycle, while the geometries used for this study do not (see Section 4.2.3), 10 phase-binned frames correspond to 5 patient geometries.

Table 4.3: Results of the dose distribution comparison for the study investigating the number of patient geometries. Each dose distribution was compared to the reference dose calculated with 10 patient geometries, using the method described in Section 4.2.5.2. Along with the parameters fitted to the distribution of Equation (4.17), the maximum RCSU and the relative systematic dose differences,  $\Delta_1$  and  $\Delta_2$ , are reported ( $\Delta_i = \text{RCSU}_{\max} \cdot |\delta_i|$ ).

Number of patient geometries	$\text{RCSU}_{\max}/\%$	$\alpha_1$	$\delta_1$	$\Delta_1/\%$	$\alpha_2$	$\delta_2$	$\Delta_2/\%$
1	0.63	0.14	13.81	8.68	0.05	35.97	22.60
2	0.47	0.03	-15.65	7.30	0.09	89.48	3.95
3	0.40	0.19	1.93	0.77	0.02	4.94	1.98
4	0.38	0.05	1.74	0.65	0.05	1.74	0.65
5	0.34	0.00	-5.80	2.00	0.04	-1.96	0.68
6	0.34	0.04	0.81	0.27	0.04	0.81	0.27
7	0.32	0.46	0.01	0.00	0.46	0.01	0.00
8	0.31	0.00	-1.73	0.54	0.01	-1.73	0.54
9	0.30	0.50	-0.03	0.01	0.50	-0.03	0.01

### 4.3.1.3 Fluence calculation angle interval

The dose distribution calculated with  $\Delta\theta_{\text{flu}} = 1.3^\circ$  was compared to the reference dose, calculated with  $\Delta\theta_{\text{flu}} = 0.1465^\circ$ , using the statistical method described in Section 4.2.5.3. These angles were selected to correspond to the Markov chain time interval  $\tau = 386.9$  ms and the linac update time interval of 40 ms, respectively.

The weighting parameters  $\alpha_1$  and  $\alpha_2$  were found to be equal at 0.45, with equal systematic difference parameters of  $-20.14$ . The maximum RCSU of these two dose distributions is  $2.5 \cdot 10^{-7}$ , leading to relative systematic dose differences of  $\Delta_1, \Delta_2 = 5.1 \cdot 10^{-6}$ . This is far below the target value of a 1% systematic dose difference, indicating that the sources of errors presented in Section 4.2.5.3 have very little effect on the dose when the fluence calculation angle interval  $\Delta\theta_{\text{flu}}$  is related to the Markov chain time interval  $\tau = 386.9$  ms through Equation (4.9). The fluence calculation interval for the SBRT plans was also calculated using this relation. Therefore the results gathered for this conventionally-fractionated plan were deemed to also apply to the SBRT plans.

## 4.3.2 Comparison of 4D methods

### 4.3.2.1 Patient 1: conventional fractionation

The results of the optimization and robustness study for conventionally-fractionated plans on patient 1 are shown in Figures 4.8 to 4.10. In these

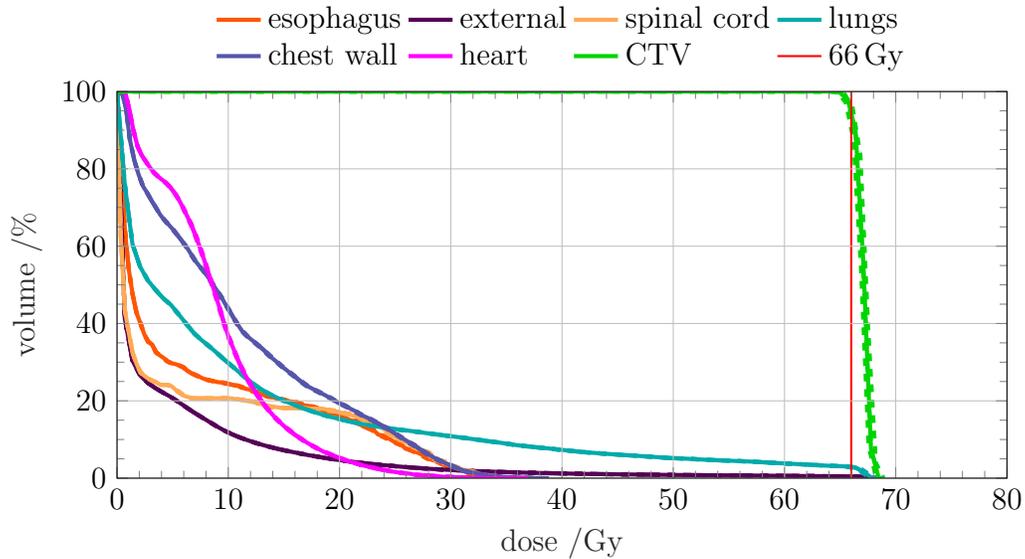
figures, the DVHs show the plan quality of each plan after optimization. The PDVHs, on the other hand, are included to quantify the robustness of each plan with respect to changes in the patient motion trajectory; specifically, they show percentiles of the simulated delivered dose. PDVHs for  $p = 5\%$ ,  $50\%$ , and  $95\%$ , all for the full treatment course of 33 fractions. The CTV DVH objectives used at The Ottawa Hospital are shown in Table 4.4; results for both the originally optimized DVH as well as the PDVHs for  $p = 5\%$ ,  $50\%$ , and  $95\%$  are shown.

Throughout Figures 4.8 to 4.10, the PDVHs for OARs mostly align with the DVHs for the originally optimized plan. Exceptions include a small increase in the dose for the PDVHs for the DAD plan in Figure 4.9b, and a small decrease in dose for the 3D-CTV plan in Figure 4.10b. For this reason, discussion of the results with respect to robustness are focussed on the CTV PDVHs, which show much more variation than the OAR DVHs.

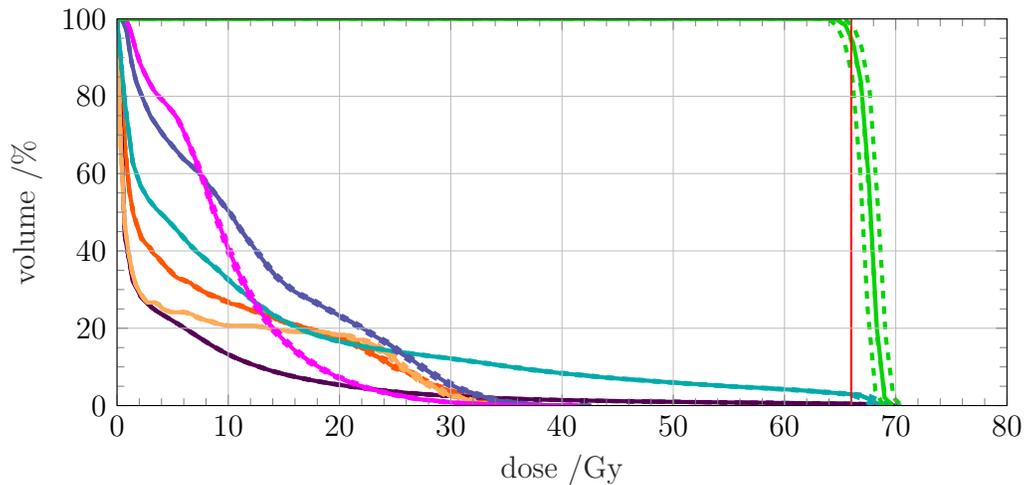
The small width of the CTV PDVHs around the prescription dose line for the CO and MKO plans in Figures 4.8a and 4.10a indicate the high degree of robustness of these plans. For both of these, the median PDVH (i.e., for  $p = 50\%$ ) agrees exceptionally well with the DVH of the optimized dose. However, depending on which specific percentile is used, either plan can be seen as delivering greater target coverage than the other. For example, from Table 4.4, the MKO plan has a higher 5th percentile prescription dose volume,  $94.0\%$  versus  $91.7\%$ , while the latter has a higher 95th percentile prescription dose volume of  $97.1\%$  versus  $95.3\%$ . Ideally this volume would

Table 4.4: CTV DVH objectives and thresholds for patient 1, conventionally fractionated, for each optimization method considered. The first row for each optimization method indicates the DVH results of the optimized plan, while the second row indicates the PDVH results for  $p = 5\%$  to  $95\%$ .

Optimization method	DVH objectives and their thresholds		
	$V_{66\text{Gy}}/\%$ >95	$D_{98\%}/\text{Gy}$ >62.7	$D_{2\%}/\text{Gy}$ <72.6
CO	94.8	65.7	68.1
	91.7 to 97.1	65.4 to 65.8	67.9 to 68.4
PO	95.0	65.6	69.0
	86.8 to 98.9	64.8 to 66.2	68.3 to 69.8
STO	94.9	65.7	68.1
	78.7 to 82.4	64.3 to 64.6	68.9 to 69.1
DAD	95.0	65.6	68.3
	81.4 to 82.6	63.0 to 63.2	71.9 to 72.1
MKO	95.0	65.6	67.8
	94.0 to 95.3	65.5 to 65.6	67.6 to 68.0
3D-CTV	95.0	65.6	68.3
	65.6 to 66.9	38.5 to 39.0	68.3 to 68.9



(a) CO DVHs.



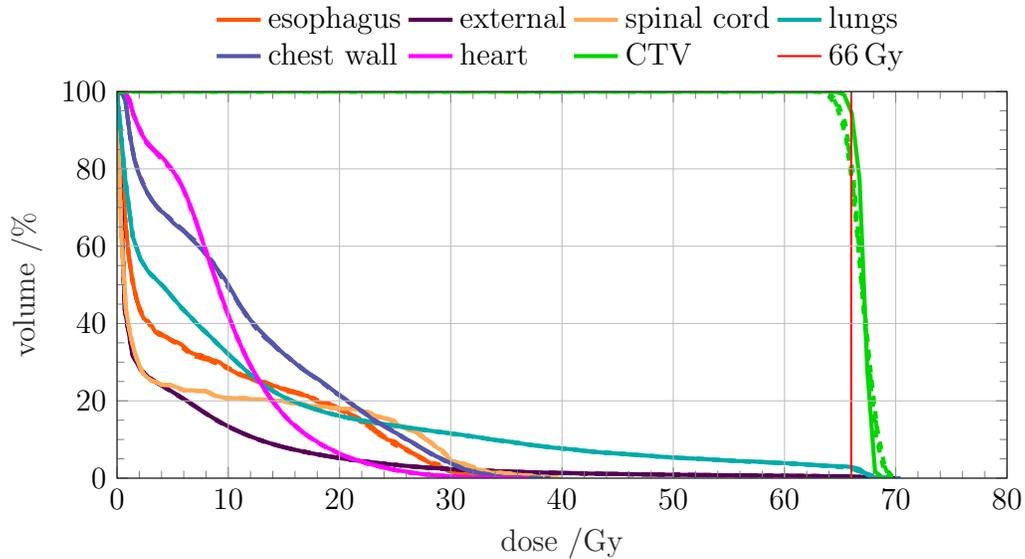
(b) PO DVHs.

Figure 4.8: DVHs for patient 1, conventionally fractionated, optimized with CO and PO. The solid lines indicate the DVHs of the optimized plan, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ . The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH. The prescription dose of 66 Gy is shown in red.

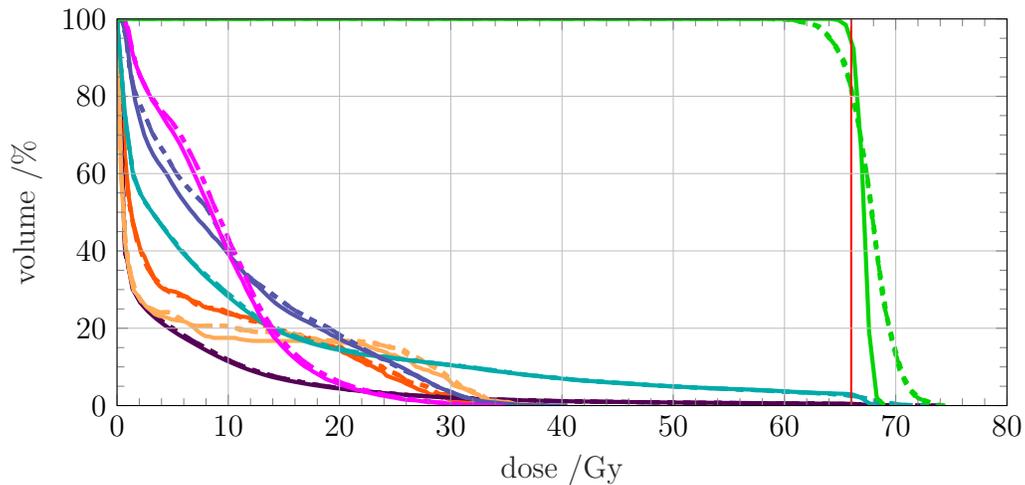
be 95 % or greater with a probability of 100 %; indeed, each optimized plan was rescaled to this DVH point immediately following optimization. The MKO plan therefore delivers a higher minimum prescription dose volume to the CTV, while the CO plan delivers a higher maximum prescription dose volume. Furthermore, the plan optimized with PO, shown in Figure 4.8b, has a smaller 5th percentile for the prescription dose volume of 86.8 % than the CO plan, despite the inclusion of the variance term in the optimization, indicating a less robust plan. Possible reasons for this result are discussed in Section 4.4. This plan yields, however, the largest 95th percentile prescription dose volume, of 98.9 %.

Although the PDVHs of the plan optimized with STO have a narrow width, as seen in Figure 4.9a, the target coverage is not robust, with 5th and 95th percentile prescription dose volumes of only 78.7 % and 82.4 %, respectively. This is due to the fact that the VMAT plan in this method was optimized under the assumption of a single reproducible respiratory trajectory. The DAD method, on the other hand, under-doses part of the tumour and nearly over-doses other parts of the tumour, as seen by the shallow PDVHs in Figure 4.9b. From Table 4.4, 95th percentile for the largest dose given to 2 % of the CTV is 72.1 Gy, whereas the threshold for this objective is 72.6 Gy.

The high robustness of the MKO plan, shown in Figure 4.10a, is due to the use of the probabilistic motion-encompassing dose-influence matrix; this behaviour is repeated with the other studied patient geometries, shown in Section 4.3.2.3. The dose to OARs however is increased as compared to the

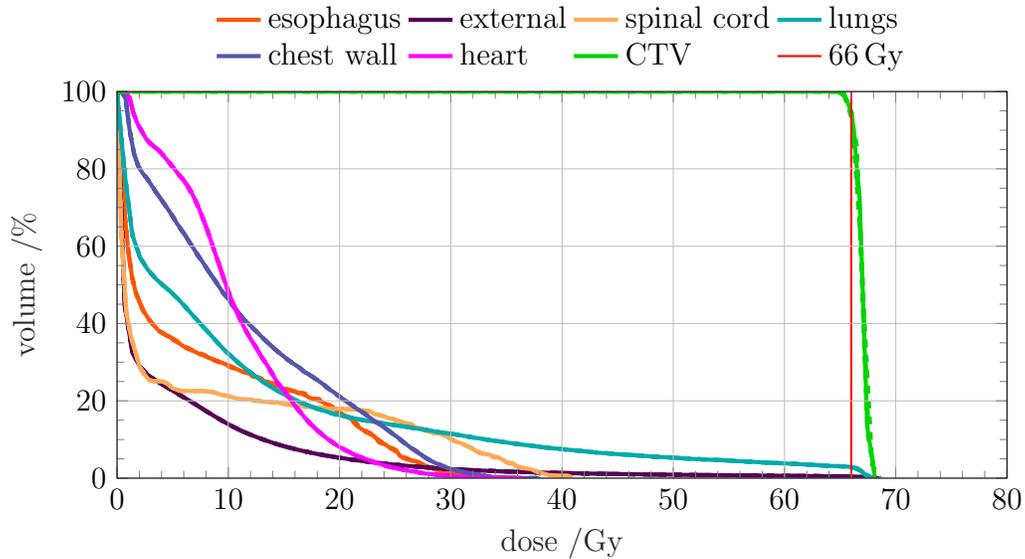


(a) STO DVHs.

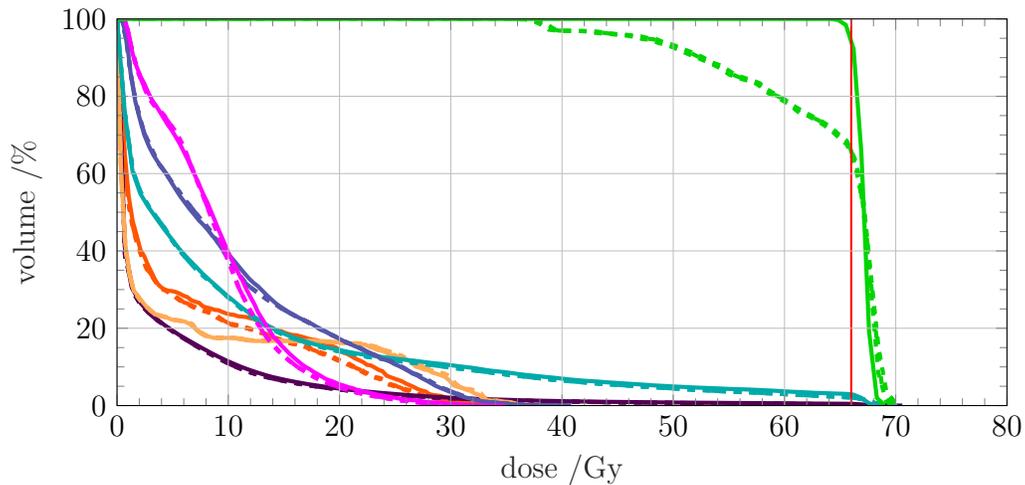


(b) DAD DVHs.

Figure 4.9: DVHs for patient 1, conventionally fractionated, optimized with STO and DAD. The solid lines indicate the DVHs of the optimized plan, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ . The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH. The prescription dose of 66 Gy is shown in red.



(a) MKO DVHs.



(b) 3D-CTV DVHs.

Figure 4.10: DVHs for patient 1, conventionally fractionated, optimized with MKO and 3D-CTV. The solid lines indicate the DVHs of the optimized plan, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ . The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH. The prescription dose of 66 Gy is shown in red.

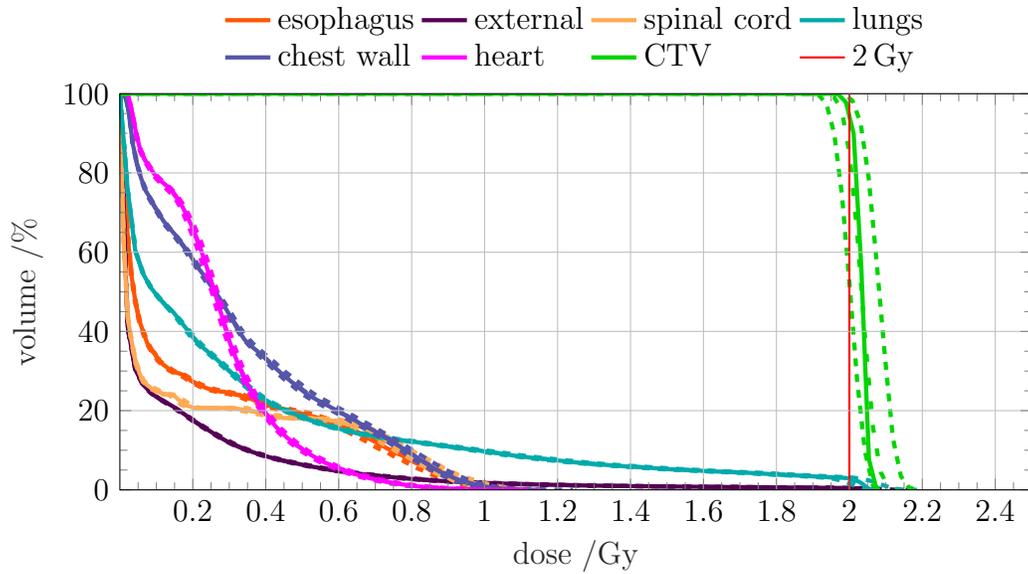
CO plan. For example, the percentage volume of lung receiving 5 Gy increases from  $V_{5\text{Gy}} = 44.2\%$  for CO to  $V_{5\text{Gy}} = 47.6\%$ , whereas the threshold is equal to 50%. The maximum dose to  $0.03\text{ cm}^3$  of the spinal cord is also significantly increased between these two methods from 34.9 Gy to 40.9 Gy, although this is still below the threshold of 48 Gy.

The plans optimized with 3D-CTV offer little robustness to the variation of respiratory trajectory, as the PDVHs in Figure 4.10b show an extremely low target dose. The 5th and 95th percentile prescription dose volumes are 65.4% and 66.8%, reflecting the fact that no compensation or adaptation whatsoever was made for the presence of respiratory motion in this plan.

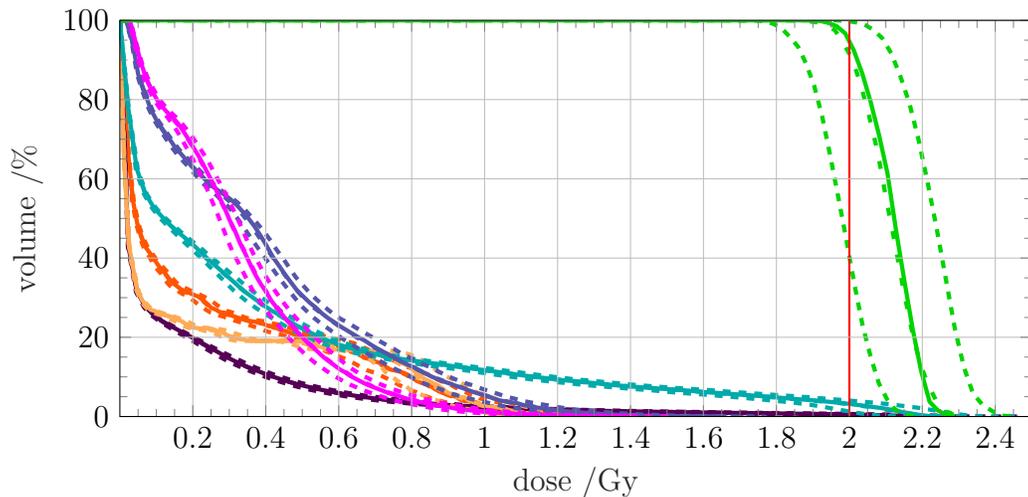
The DVHs and PDVHs for a single fraction only are shown in Figure 4.11, for the two 4D-VMAT methods of CO and PO. For the CO plan, the widening of the PDVHs for all volumes, and in particular the CTV, demonstrates that the dose for a single fraction is much less robust than the dose for the full treatment of 33 fractions, as expected following the discussion in Section 4.2.1.2. The PO plan shown in Figure 4.11b was optimized using a denominator of  $N_{\text{frac}} = 1$  in the second term of Equation (4.10) to account for this effect. As in the conventionally-fractionated plans, the PO plan is significantly less robust than the CO plan.

#### 4.3.2.2 Patient 1: SBRT

Optimization results for the patient SBRT plans for patient 1 are shown in Figures 4.12 to 4.14, with CTV DVH objectives in Table 4.5. PO was not



(a) CO DVHs.



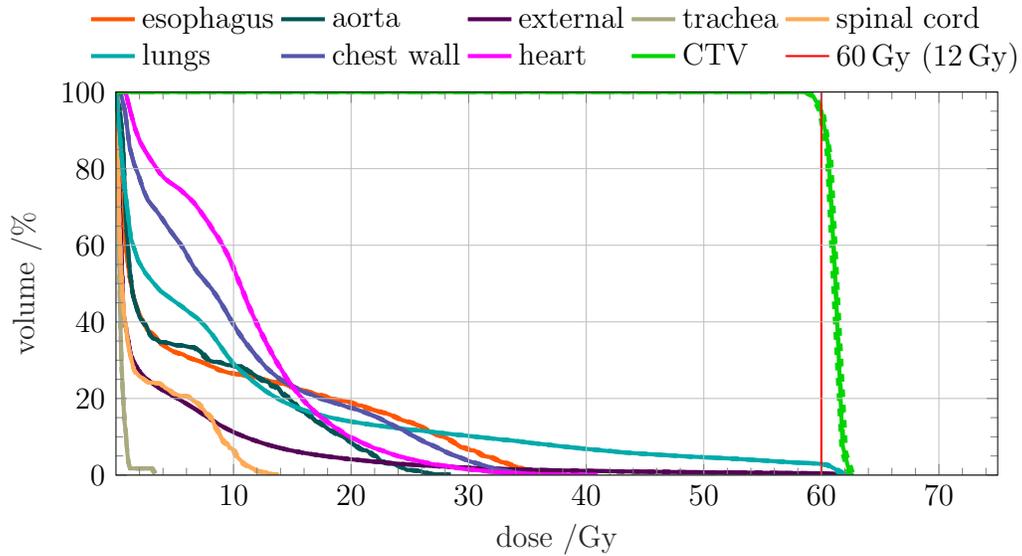
(b) PO DVHs.

Figure 4.11: DVHs for patient 1, conventionally fractionated, optimized with CO and PO for 1 fraction only. The solid lines indicate the DVHs of the optimized plan, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$  and the loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH. The prescription dose of 2 Gy is shown in red.

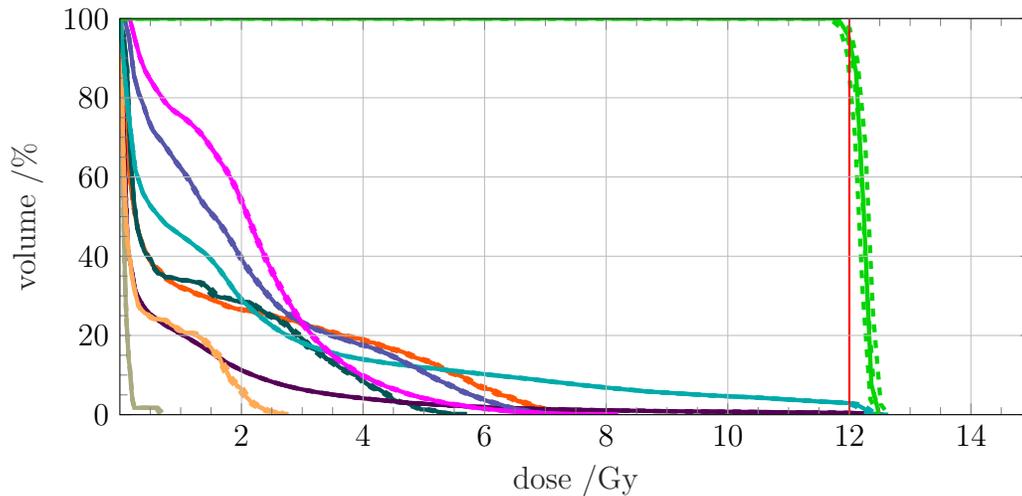
Table 4.5: CTV DVH objectives and thresholds for patient 1, SBRT, for each optimization method considered. The first row for each optimization method indicates the DVH results of the optimized plan, while the second row indicates the PDVH results for  $p = 5\%$  to  $95\%$ .

Optimization method	DVH objectives and their thresholds		
	$V_{60\text{Gy}}/\%$ >95	$D_{98\%}/\text{Gy}$ >57	$D_{99\%}/\text{Gy}$ >54
CO	94.7	59.6	59.3
	92.7 to 96.4	59.4 to 59.8	59.1 to 59.4
STO	94.8	59.6	59.4
	93.2 to 94.8	59.5 to 59.6	59.3 to 59.5
DAD	94.9	59.6	59.5
	88.1 to 88.9	57.9 to 58.0	57.6 to 57.7
MKO	94.8	59.7	59.5
	94.2 to 95.2	59.6 to 59.7	59.5 to 59.6
3D-CTV	94.9	59.6	59.5
	66.8 to 67.4	32.9 to 33.5	32.2 to 32.8

performed for SBRT, since this method yielded a slightly less robust plan than did CO with conventional fractionation, and the small value of  $\Delta\theta_{\text{flu}} = 0.3^\circ$  would lead to a prohibitively long optimization time due to the the variance term in Equation (4.10), estimated at 18 days. For comparison, optimization times for the conventionally-fractionated PO plans are approximately 24 h, compared to 3 h for plans optimized with CO (15 h for SBRT).



(a) CO DVHs for all 5 fractions.



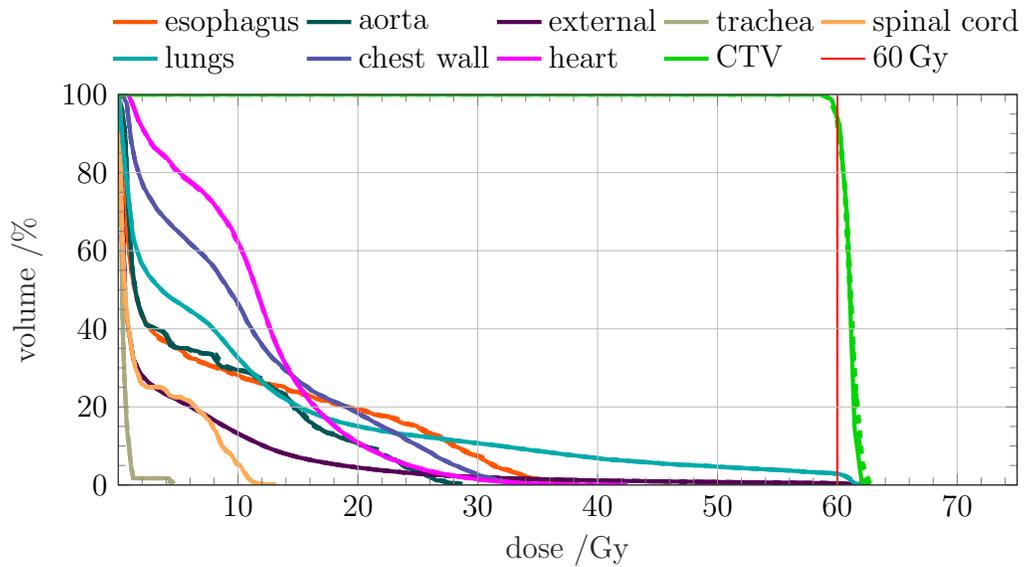
(b) CO DVHs for 1 fraction.

Figure 4.12: DVHs for patient 1, SBRT, optimized with CO, for 5 fractions and 1 fraction only. The solid lines indicate the DVHs of the optimized plan, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ . The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH. The prescription dose of 60 Gy (12 Gy for the single fraction) is shown in red.

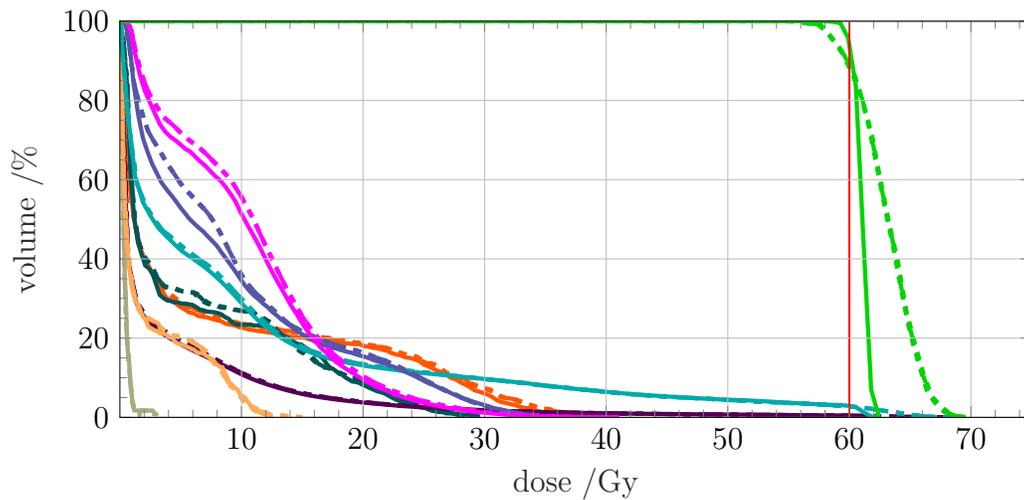
The robustness of the SBRT CO plan is similar to the corresponding conventionally-fractionated plan, with 5th and 95th percentile prescription dose volumes of 92.7% and 96.4%, respectively, as seen in Table 4.5. Both the MKO and STO methods yielded SBRT plans with slightly larger 5th percentile prescription dose volumes of 94.2% and 93.2%, respectively. The 95th percentile prescription dose volume of the MKO and STO plans were 95.2% and 94.8%, respectively, however, both smaller than the CO plan. The DAD and 3D-CTV SBRT plans gave similar patterns of robustness to their conventionally-fractionated counterparts.

Figure 4.12b shows the DVHs and PDVHs optimized with CO for a single SBRT fraction only. Similarly to the single fraction results for the conventionally-fraction plan in Figure 4.11, the PDVHs are wider than the corresponding curves for the full treatment, shown in Figure 4.12a, demonstrating increase in robustness as more fractions are delivered. The difference in patterns of robustness between the SBRT and conventionally-fractionated plans is discussed in Section 4.4.

The doses to the OARs for the SBRT plans for patient 1 optimized with each of the different methods are somewhat varied. The percentage volume of the lung receiving a dose of more than 20 Gy in the STO plan, shown in Figure 4.13a, exceeded the threshold of 15% used at The Ottawa Hospital, with  $V_{20\text{Gy}} = 15.1\%$ ; this threshold was nearly exceeded by the MKO plan, with  $V_{20\text{Gy}} = 14.9\%$ . In contrast, the CO plan yielded a smaller volume of  $V_{20\text{Gy}} = 14.0\%$ . In addition to this, the maximum dose for the heart of 38 Gy



(a) STO DVHs.



(b) DAD DVHs.

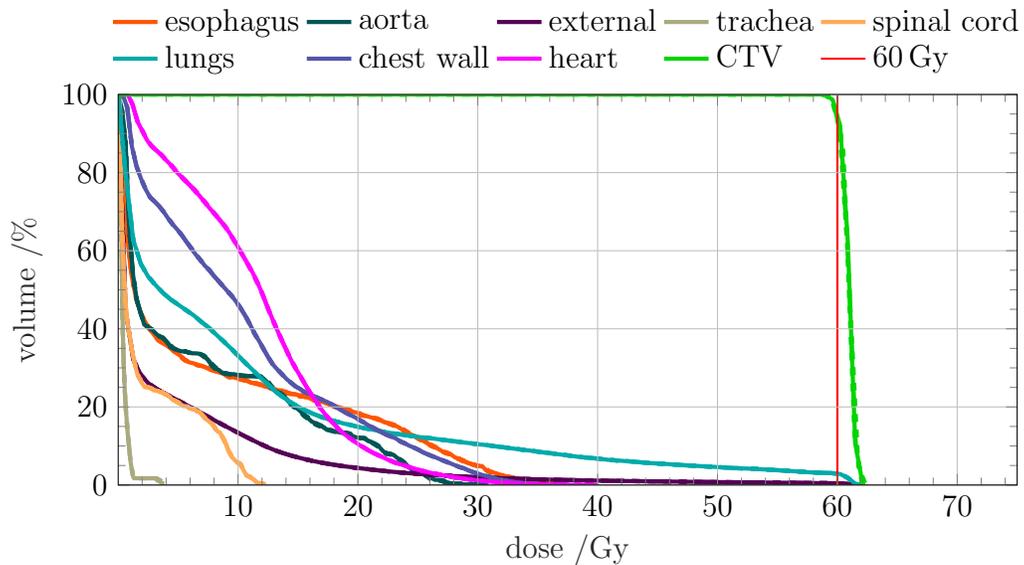
Figure 4.13: DVHs for patient 1, SBRT, optimized with STO and DAD. The solid lines indicate the DVHs of the optimized plan, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ . The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH. The prescription dose of 60 Gy is shown in red.

was exceeded by all plans, however for a very small volume. At most only 7 voxels in the heart, representing a volume of  $0.875 \text{ cm}^3$ , were found to have a dose larger than the maximum, out of a total heart volume of  $873 \text{ cm}^3$ .

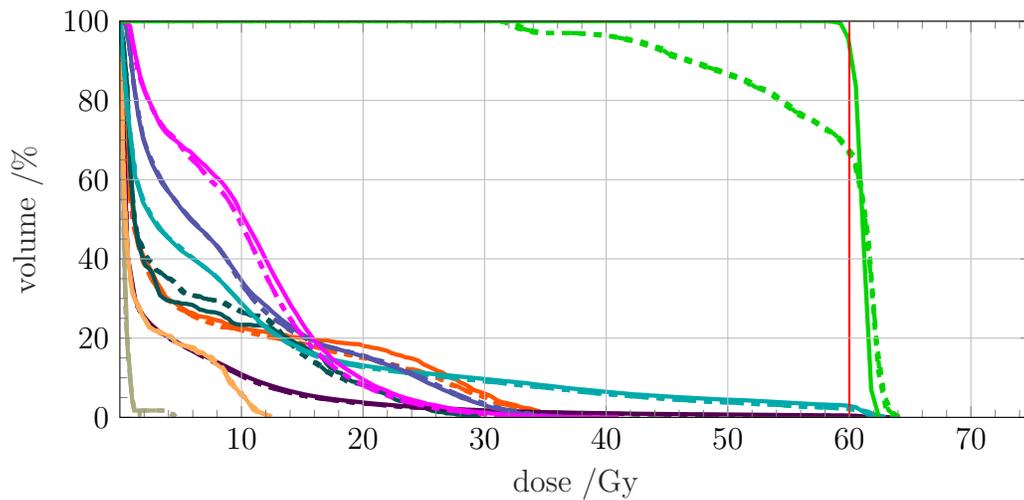
#### 4.3.2.3 Other patients

Similar results for the robustness of the dose to the CTV are observed in patients 2 to 4, with the plans optimized under the CO and MKO methods generally having the greatest robustness. Select optimization methods for each of these are shown in Figures 4.15 to 4.17. The tumour parameters for each patient number studied can be found in Table 4.1.

Depending on which specific percentile is used, the DAD plan for patient 2 has greater or equal target coverage compared to the CO plan, both shown in Figure 4.15. The former plan has a higher 5th percentile prescription dose volume, 97.3% versus 88.3%, while both have equal 95th percentile prescription dose volume of 97.7%. With this patient and with the remaining ones (DAD results not shown), the maximum dose CTV threshold of  $D_{2\%} = 72.6 \text{ Gy}$  was exceeded in all cases by the DAD plan. Even the 5th percentile of  $D_{2\%}$  to the target reaches doses of 75 Gy, signifying significant tumour over-dosage. In comparison, the 95th percentile of this dose objective for the CO plan on this patient is only 68.0 Gy, indicating much greater dose uniformity to the target. The DAD method also delivers larger doses to OARs, as seen with patient 1 in Section 4.3.2.1. 35.1% of the lung receives a dose higher than 5 Gy, compared to 30.3% for the CO plan, and the maximum



(a) MKO DVHs.



(b) 3D-CTV DVHs.

Figure 4.14: DVHs for patient 1, SBRT, optimized with MKO and 3D-CTV. The solid lines indicate the DVHs of the optimized plan, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ . The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH. The prescription dose of 60 Gy is shown in red.

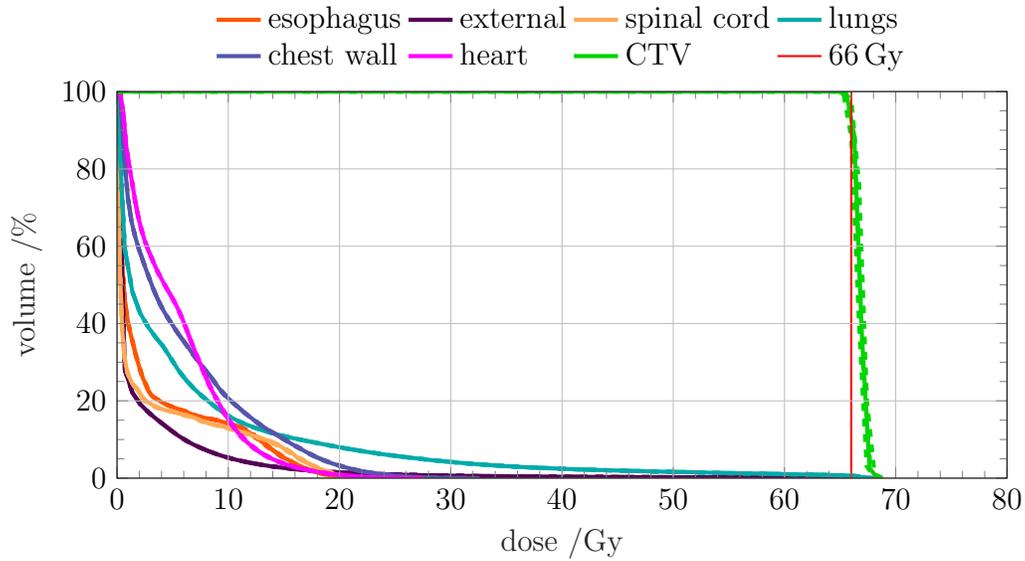
dose to  $0.03 \text{ cm}^3$  has increased from  $D_{0.03 \text{ cm}^3} = 22.5 \text{ Gy}$  to  $23.1 \text{ Gy}$  between these two plans.

Figure 4.16 compares the DVHs for patient 3 plans optimized with CO and MKO. As with the patient 1 plans, the MKO plan has a higher 5th percentile prescription dose volume,  $91.2\%$  versus  $84.5\%$ , while the CO plan has a higher 95th percentile prescription dose volume of  $97.8\%$  versus  $94.8\%$ . The CO plan also gives lower doses to the OARs: the lung  $V_{5 \text{ Gy}}$  was  $33.5\%$  compared to  $35.3\%$  with MKO; while  $D_{0.03 \text{ cm}^3}$  for the spinal cord was  $17.0 \text{ Gy}$  compared to  $20.3 \text{ Gy}$ .

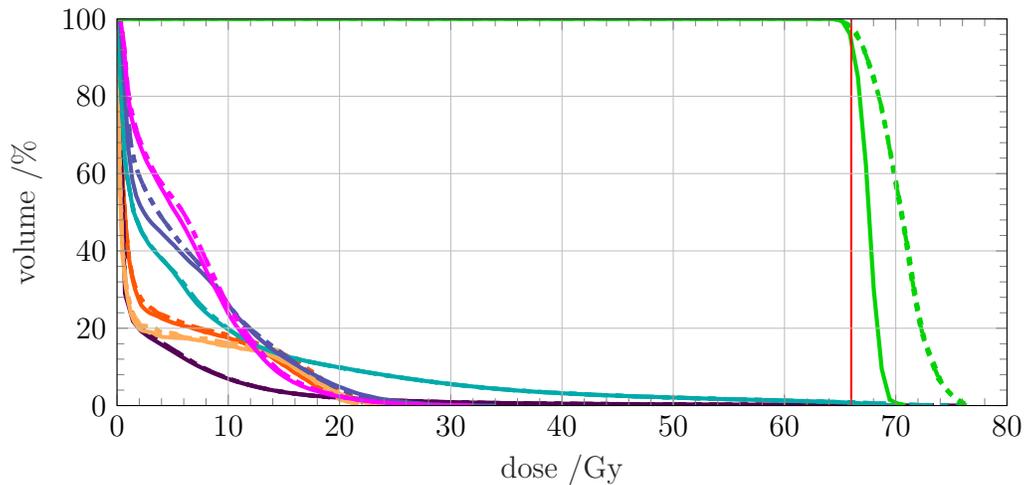
Finally, the DVHs for the CO and MKO plans for patient 4 are shown in Figure 4.17. Robustness is greater with the MKO plan, with a 5th percentile target prescription dose volume of  $91.0\%$  compared to  $82.9\%$  for the CO plan. In case, the CO plan gives a lower dose to the lung, with  $V_{5 \text{ Gy}} = 29.7\%$  compared to  $30.1\%$  for MKO. However, the maximum spinal cord dose is larger for the CO plan than the MKO plan, with  $D_{0.03 \text{ cm}^3} = 39.4 \text{ Gy}$  and  $38.3 \text{ Gy}$ , respectively. Note that both of these doses are significantly elevated compared to the other patients due to the closer proximity of the tumour to the spinal cord in this patient, yet still remain below the threshold of  $48 \text{ Gy}$ .

#### 4.3.2.4 Fraction model intercomparisons

Results of the fraction intercomparison test, done on patient 1, are shown in Figures 4.18 and 4.19. DVHs for the CO, STO, and MKO plans optimized with the fraction A respiratory motion model from Section 3.4.1 are reproduced

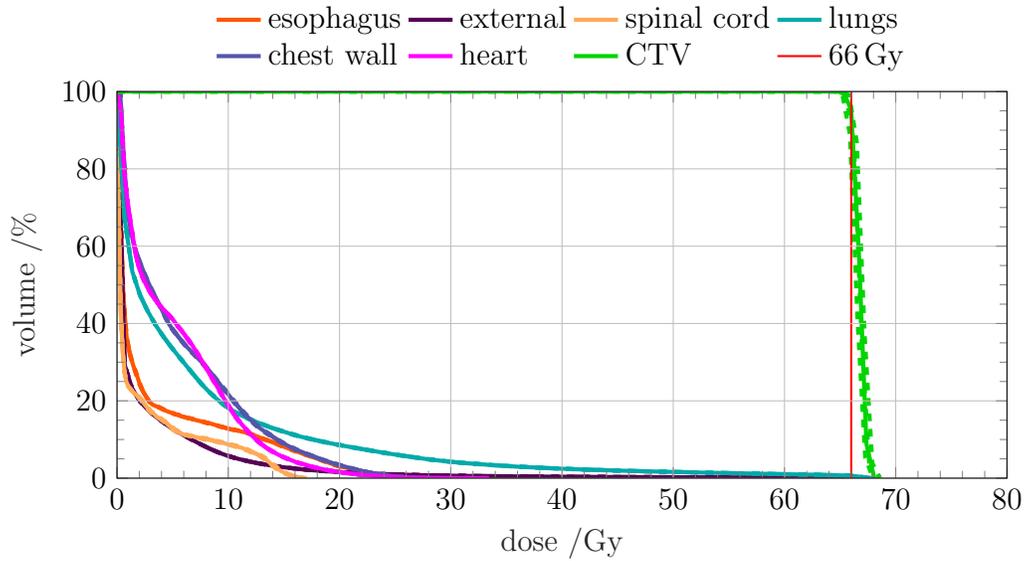


(a) CO DVHs.

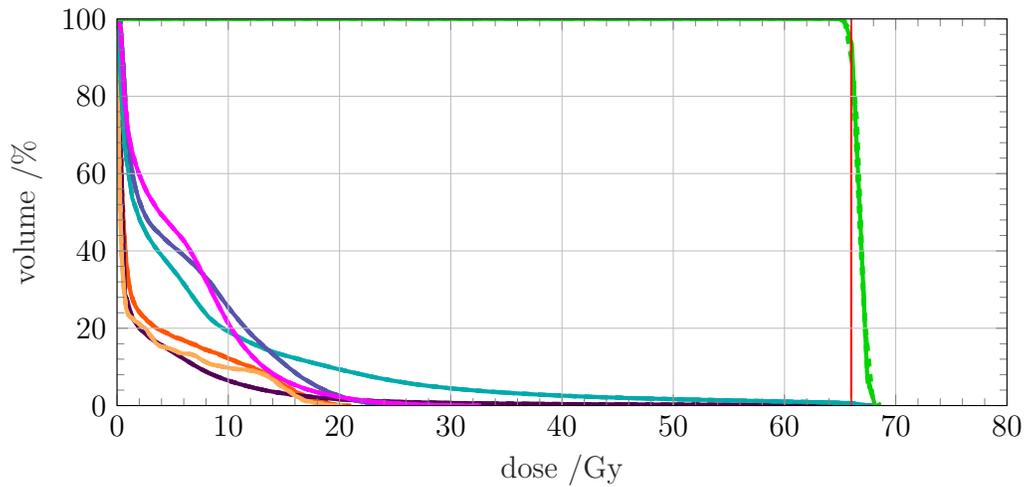


(b) DAD DVHs.

Figure 4.15: DVHs for patient 2, conventionally fractionated, optimized with CO and DAD. The solid lines indicate the DVHs of the optimized plan, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ . The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH. The prescription dose of 66 Gy is shown in red.

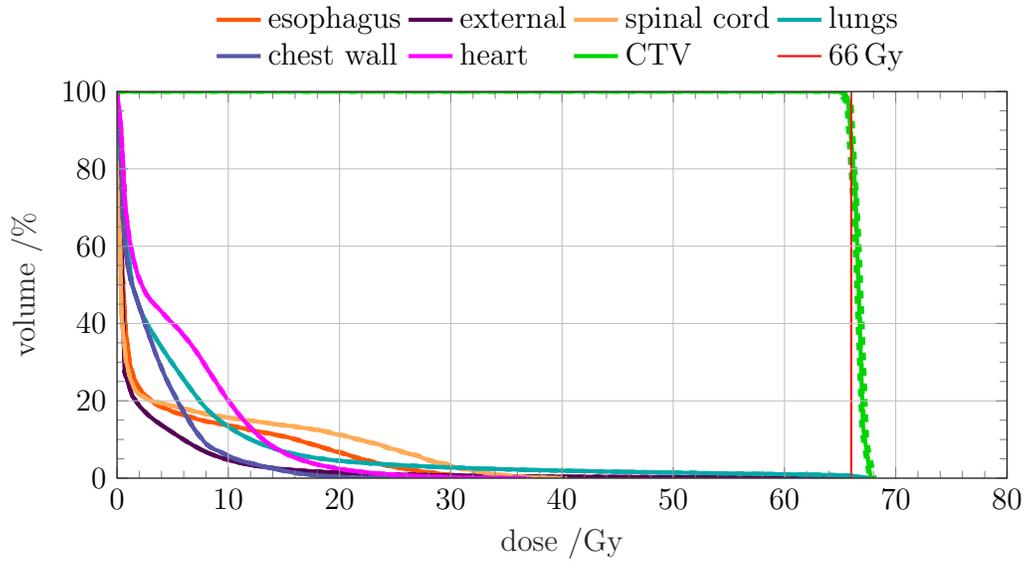


(a) CO DVHs.

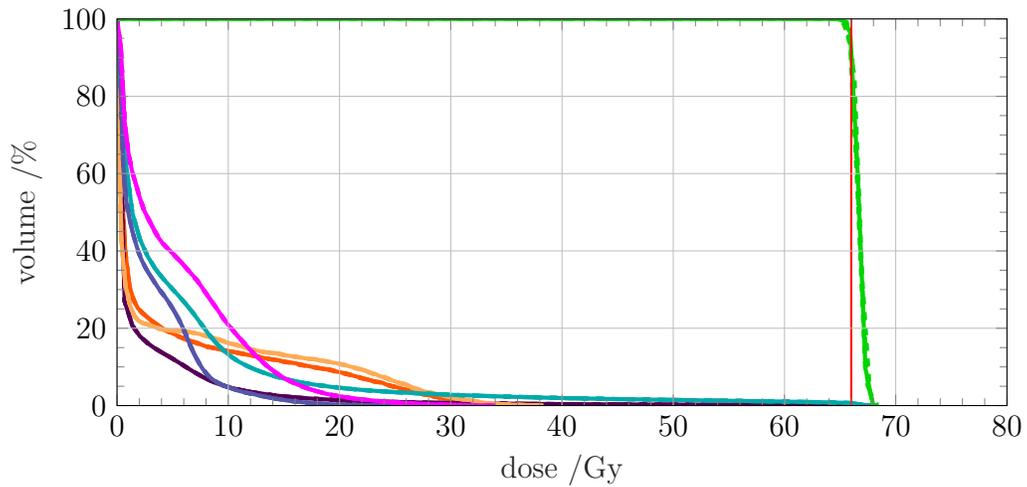


(b) MKO DVHs.

Figure 4.16: DVHs for patient 3, conventionally fractionated, optimized with CO and MKO. The solid lines indicate the DVHs of the optimized plan, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ . The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH. The prescription dose of 66 Gy is shown in red.



(a) CO DVHs.



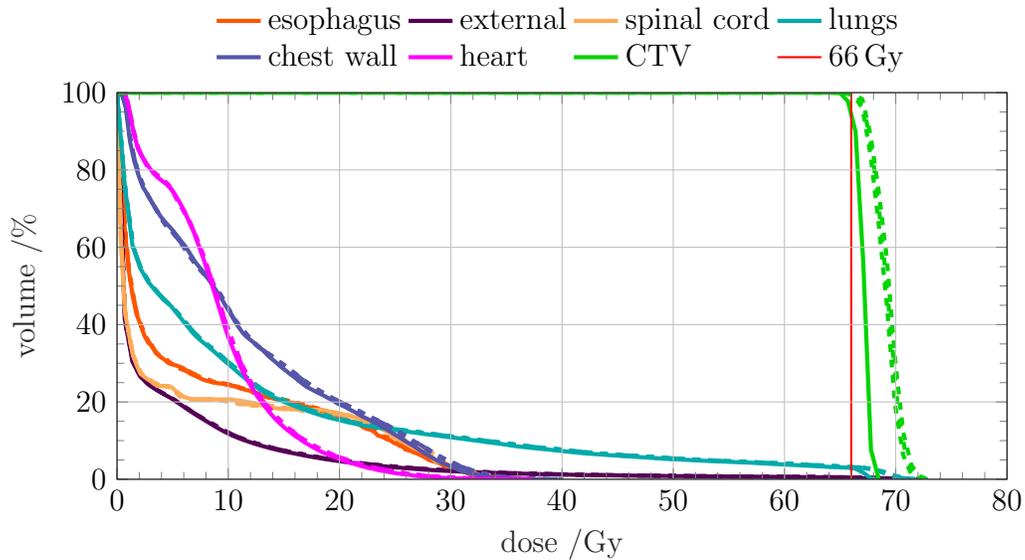
(b) MKO DVHs.

Figure 4.17: DVHs for patient 4, conventionally fractionated, optimized with CO and MKO. The solid lines indicate the DVHs of the optimized plan, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ . The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH. The prescription dose of 66 Gy is shown in red.

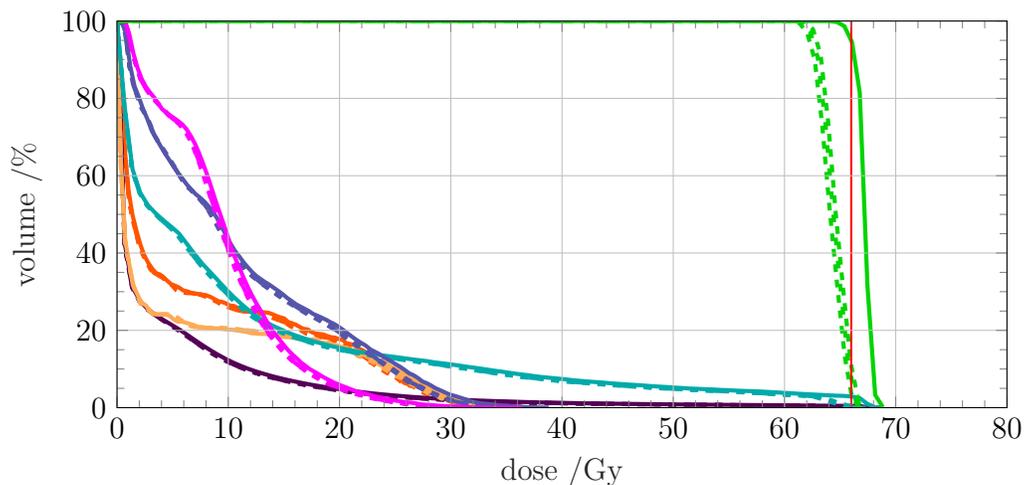
from Figures 4.8a, 4.9a and 4.10a, respectively. PDVHs are also shown, calculated instead with the fraction B respiratory model from Section 3.4.2, with the assumption that the patient motion follows this model for all 33 fractions. Figure 4.18b also shows the DVH for a CO plan optimized with the fraction B respiratory model, and PDVHs calculated with the fraction A model.

In each case investigated, the OAR PDVHs are in excellent agreement with the original DVHs. The dose to the CTV however is not as robust to this change in the patient model. For example, in the CO plan optimized with the fraction A model, the dose is increased by 1 Gy to 4 Gy when the fraction B model is used to simulate delivery. The increase is such that the 5th percentile prescription dose volume is 99.6 %, in excess of the required 95 % coverage. The maximum dose to 2 % of the target volume has also increased, with a 95th percentile of 71.7 Gy, which is yet smaller than the threshold of 72.6 Gy. When a plan is optimized with the fraction B model instead, and the fraction A model used for delivery simulation, the dose is instead decreased by 1 Gy to 4 Gy such that the 5th and 95th percentile prescription dose volumes are 3.1 % and 10.5 %, respectively, severely compromising the plan quality.

The CTV PDVHs for the STO plan shown in Figure 4.19a are nearly the same when calculated using either of the fraction A or B models. Taken from Table 4.4, the 5th and 95 percentile prescription dose volumes with the fraction A model are 78.7 % and 82.4 %, respectively; with the fraction B model, they are 79.6 % and 82.3 %. This plan is therefore not sensitive to



(a) Fraction A optimization, fraction B delivery.



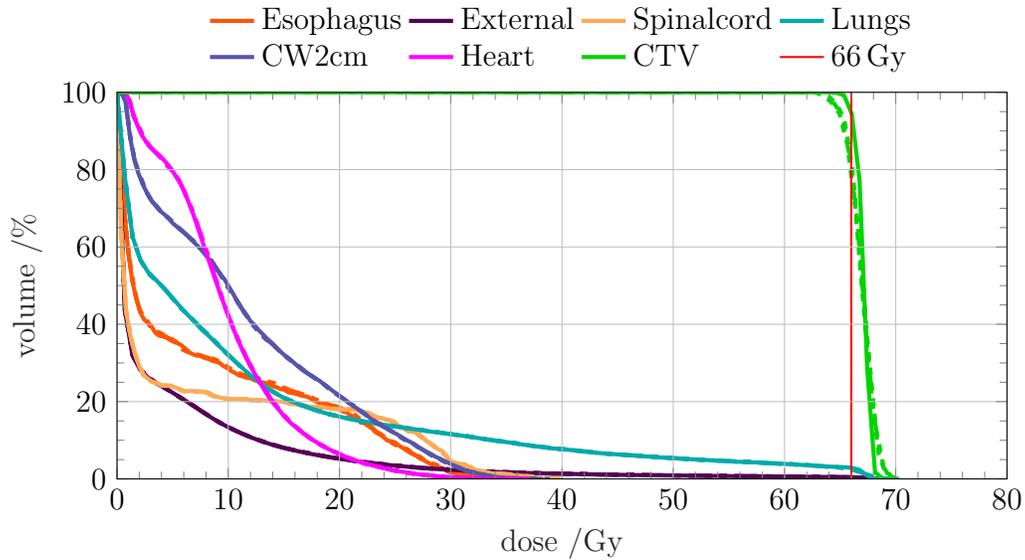
(b) Fraction B optimization, fraction A delivery.

Figure 4.18: DVHs for patient 1, conventionally fractionated, optimized with CO, using different respiratory motion models for optimization and delivery simulation. The solid lines indicate the DVHs of the plan optimized with the fraction A (or B) respiratory model, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ , calculated with the fraction B (or A) respiratory model. The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH, also calculated with the fraction B (or A) respiratory model. The prescription dose of 66 Gy is shown in red.

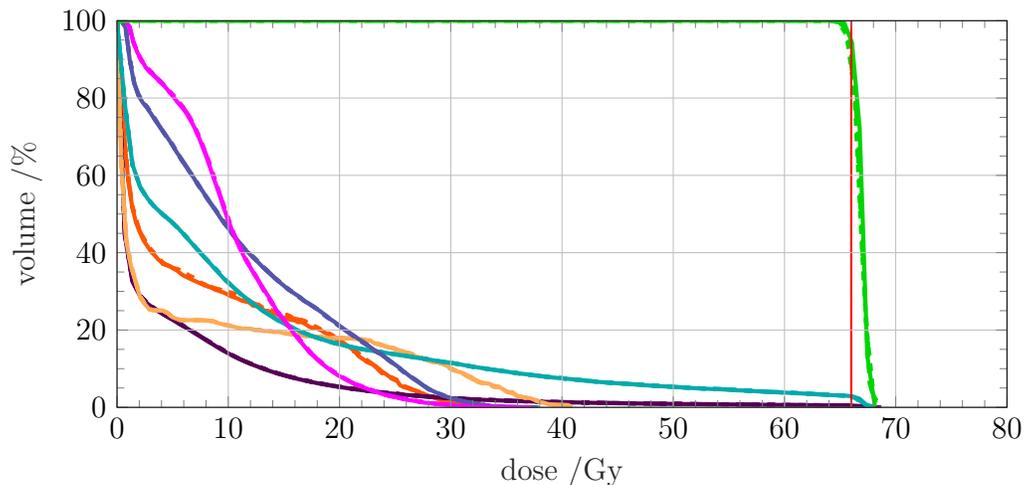
this particular change in the respiratory motion model. Finally, the target coverage for the MKO plan shown in Figure 4.19b dropped slightly. With the fraction A model, 5th and 95 percentile prescription dose volumes are 94.0% and 95.3%; with the fraction B model, 89.7% and 92.7%.

## 4.4 Discussion

The plans optimized with the ALERT-RA CO method deliver doses which are consistently robust with respect to variations in the patient's respiratory motion for all patients investigated, as measured by the minimal departure of the CTV PDVHs from the prescription dose. Although the variance term was included in the objective function according to Equation (4.10), the plans optimized with PO are not as robust, in contradiction with results observed in previous studies [55, 84, 85]. A possible explanation for this may be due to the differing number of iterations of matRad's optimizer, the IPOPT algorithm [67], undergone in each optimization method: an average of 390 iterations were used for CO, while PO used only 100. A smaller number of iterations would lead to a plan which is less optimized; in the case of PO, because of the inclusion of the variance term in the objective function, this would imply a less robust treatment. The number of iterations is not directly controlled in IPOPT, however; instead, optimization concludes when the change in the objective function between successive iterations is below a specified threshold. In this study, the same value of the threshold was used for



(a) STO DVHs.



(b) MKO DVHs.

Figure 4.19: DVHs for patient 1, conventionally fractionated, optimized with STO and MKO, with delivery simulated using the fraction B respiratory model. The solid lines indicate the DVHs of the plan optimized with the fraction A respiratory model, while the tightly dashed lines indicate the PDVHs for  $p = 5\%$  and  $95\%$ , calculated with the fraction B respiratory model. The loosely dashed line indicates the PDVH for  $p = 50\%$ , the median DVH, also calculated with the fraction B respiratory model. The prescription dose of 66 Gy is shown in red.

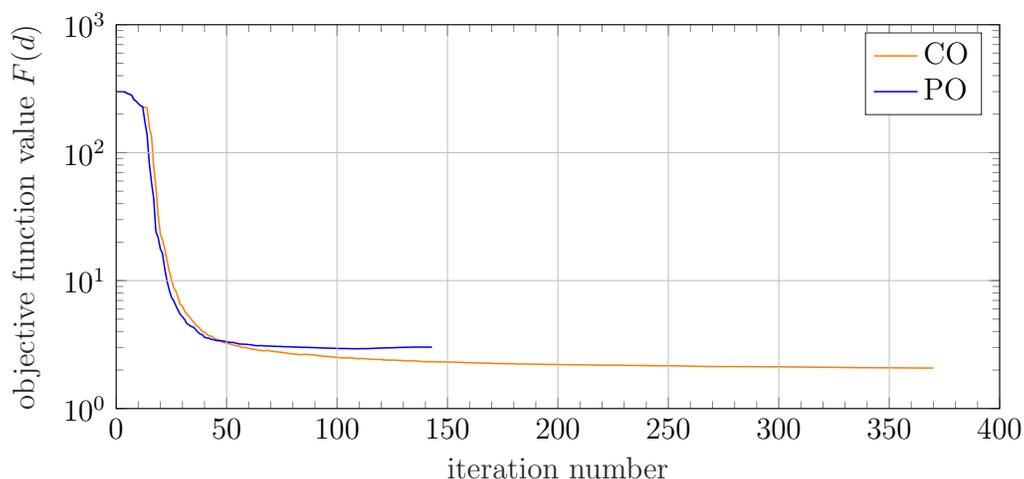


Figure 4.20: Value of the objective function  $F(d)$  as a function of iteration number throughout the DAO step of plan optimization. Conventionally-fractionated CO and PO plans for patient 1 are shown. Note the logarithmic scale of the vertical axis.

all treatment plan optimization methods. A graph of the objective function  $F(d)$  as a function of the optimization iteration number for both CO and PO methods planned on patient 1 is shown in Figure 4.20. Both optimization methods are observed to converge at different values of the objective function before terminating, but it is yet possible that the PO plan could be optimized further with more iterations.

Nevertheless, the CO method yields plans which are consistently robust and deliver a low dose to OARs. In contrast to this, for example, the DAD method, which begins with a 3D-VMAT plan optimized on the reference patient geometry and deforms the aperture to track the tumour motion, can give plans whose robustness varies: for patient 1, this method does not provide adequate coverage of the prescription dose when the dose delivered during

actual trajectories is calculated. This may be due to the varying amount of tissue density and resulting beam attenuation from the reference geometry to the other geometries. Patients 2 to 4, however, which have smaller tumour diameters of 20 mm compared to 40 mm, give DAD plans which have more robust target coverage than the corresponding CO plans. The maximum dose to the target for these DAD plans are however exceeded, in violation of clinical guidelines. This result may also come from the lesser amount of beam attenuation for these smaller tumours, which is not adequately compensated for with strict aperture deformation only. In addition to this, the DAD plans deliver a higher dose to OARs than plans optimized with CO. This demonstrates the importance of considering all respiratory trajectories when optimizing the aperture library, as is done with the CO method.

The 3D-CTV plans, on the other hand, consist of the DAD plans without the benefit of target tracking and aperture adaptation. The extremely low CTV coverage seen in the PDVHs of the 3D-CTV plans demonstrate, then, that without some form of motion compensation treatment outcomes will be severely compromised, as was noted in Section 1.3.

The MKO method, which used a motion-encompassing dose-influence matrix to optimize a 3D-VMAT plan, results in plan which are more robust than the CO method when considering the 5th percentile prescription dose coverage of the tumour. In general, the dose delivered by a plan will vary as a function of the trajectory due to two factors: the sequence of patient

geometries in that trajectory, and the corresponding sequence of MLC apertures. With MKO, the same set of motion-encompassing apertures is always delivered, independent of the patient motion. This is in contrast with the CO plans which use an aperture library, with a different, motion-conformal sequence of apertures for each motion trajectory. This variation in the MLC aperture sequence in the CO plan therefore results in a larger variation in the dose compared to the MKO plans. The constancy of the MKO apertures, however, precludes the possibility of selectively delivering dose to the tumour in patient geometries with greater OAR separation, one of the benefits of 4D optimization [35]. Because of this, MKO plans generally give a larger dose to OARs than do the CO plan, in particular to the lungs and spinal cord. The exception to this was patient 4, in which the MKO plan gave a lower maximum dose to the spinal cord. This particular patient also exhibited the lowest tumour motion amplitude, of only 3.3 mm; the motion-encompassing apertures of the MKO plan therefore did not need to cover as large a range of motion for this particular tumour.

Despite the assumption of only a single, perfectly reproducible respiratory trajectory, the CTV dose for plans optimized with the STO method do not change significantly under the variation of the trajectory. The use of a DAO angle interval of  $\Delta\theta_{\text{DAO}} = 8^\circ$  implied that the time between these angles was 2.1 s, approximately half of the period of the respiratory cycle. Apertures were therefore optimized for the extreme tumour position bins of  $l = 1$  and 5, with the apertures at intermediary angles and position bins interpolated from

the DAO angles. The original STO design instead optimizes apertures at approximately 0.5 s time intervals, representing a finer sampling of the tumour trajectory [52]. These apertures therefore follow the tumour motion more closely throughout the entire respiratory cycle. This change from the original implementation was done to maintain a uniform  $\Delta\theta_{\text{DAO}}$  angle interval for all optimization methods. Because of this, the STO method implemented here is not directly comparable to the original design. Upon closer examination of the STO plans, it was found that the apertures mostly encompassed the entire extent of the tumour motion throughout the delivery, rather than conforming to the current position of the tumour at each angle. There is therefore little dependence of the dose to the target on the realized patient motion. The actual target coverage, however, is not robust, as noted in Section 4.3.2.1.

The SBRT plans were observed to be more robust to variations in the respiratory motion, seemingly in contradiction to the discussion in Section 4.2.1.2 on the effect of the number of fractions. This contradiction can be resolved by examining the results for a single fraction optimized with CO. The CTV PDVHs for the conventionally-fractioned plan in Figure 4.11a are notably wider than those for the SBRT plan, seen in Figure 4.12b. This difference in robustness can be attributed to the longer treatment time, and therefore number of respiratory cycles, for the SBRT plan. With a total delivery time of 410 s compared to 95 s, the patient is expected to undergo approximately 4 times more cycles; the uncertainty in the dose delivered throughout the entire delivery due to any one cycle will therefore be diminished. This effect

leads to the similar degree of robustness observed for both the conventionally-fractionated and SBRT plans.

In addition to this, the SBRT plans optimized using MKO and STO were observed to be slightly more robust than the CO plan. As explained earlier, the dose delivered from a particular trajectory depends on both the sequence of patient geometries in that trajectory, and the corresponding sequence of MLC apertures. Since the MKO method uses the same set of motion-encompassing apertures, the dose to the target is robust to changes in the respiratory motion trajectory. Although STO assumes a single motion trajectory during optimization, the longer treatment times used in SBRT result in longer time intervals between  $\{\theta_{\text{DAO}}\}$  angles of 9.0 s, compared to 2.1 s for the conventionally-fractionated plans. Since this is approximately double the period of the respiratory cycle, the apertures at the DAO control points tend to be optimized for the same tumour position for large portions of the angular arc, while the apertures interpolated at the  $\{\theta_{\text{tu}}\}$  calculation angles deliver dose to the full set of patient geometries at the  $\tau = 386.9$  ms Markov chain time interval. This effectively ensures that each aperture in the STO plans is also motion-encompassing, similar to the MKO apertures, because each optimized aperture must capture the entire range of motion of the tumour. With the longer treatment times used in SBRT, the variation in the MLC aperture sequence in the CO plans therefore result in a larger variation in the dose compared to the MKO and STO plans, which both use a constant set of motion-encompassing apertures. Because of the larger

apertures, and the lack of true 4D optimization, both of these latter plans give a larger dose to OARs, exceeding or nearly exceeding the DVH objective thresholds for the lung volume.

The intercomparison results of Section 4.3.2.4 demonstrate the effect of a patient's respiratory motion model changing in the period from optimization to delivery. In this case, two models were used from Sections 3.4.1 and 3.4.2, both trained on data acquired from the same patient, but on different treatment fractions. For all three optimization methods considered, CO, STO, and MKO, the dose to OARs did not change significantly, nor did the target coverage for the STO plan. When using the fraction A motion model for optimization and the fraction B model for delivery simulation, the target dose was slightly decreased for the MKO plan, while for the CO plan it was increased. When the CO plan was optimized using the fraction B model, it was observed that the dose to the target was instead decreased, resulting in reduced target coverage and compromised treatment outcomes.

As seen from Figure 3.19, there is a higher probability to observe tumour position  $l = 1$  with fraction B model. Apertures in the optimized library with both fraction models at this position also release the largest linac intensity, a total of 766 MU or 775 MU (fraction A or B) when summed over all gantry angles, out of 2988 MU or 2792 MU for the entire library. Note that a particular treatment fraction would only deliver part of this amount of MU, depending on the observed respiratory trajectory: estimated using Monte Carlo simulations, the mean linac output for the fraction A library is

621 MU, with 5th and 95th percentiles of 611 MU and 631 MU, respectively. When using the fraction B motion model to simulate the plan optimized with the fraction A model, the motion pattern therefore shifts to one in which the linac outputs a larger intensity than expected, with a resulting increase in the target dose. When the plan is instead optimized with the fraction B model and model A used for simulation, the motion pattern shifts away from the position bins with high output, resulting in reduced target coverage and compromised treatment outcomes. While the proposed optimization and delivery framework yields plans which are robust to variations in the patient motion governed by a specific motion model, this model may vary between delivered fractions, thereby changing the resulting dose.

It must be noted that the ALERT-RA framework is not specific to any particular choice in the respiratory motion model. While a Markov chain model was used in this study, the only criterion is that the model must calculate the probability to observe any respiratory motion trajectory. A possible solution for the problem of the model variation between fractions then might be to construct a “super”-model, with both fraction A and B models as possible outcomes, among others, with distinct probabilities. This model could then be used to calculate the expectation value of the dose, taking into consideration all possible realized respiratory motion models.

## 4.5 Conclusions

A 4D-VMAT optimization and delivery framework, ALERT-RA, was defined, developed, validated, and tested against other 4D delivery techniques. This method uses the unique combination of: an optimized aperture library, consisting of an MLC aperture for each combination of gantry angle and tumour position; a delivery scheme which tracks the respiratory motion in real time to pick out the corresponding MLC aperture from the library; and the accounting of all possible respiratory motion trajectories, and the resulting tracked apertures, when calculating dose during optimization. This combination results in plans which have a high degree of robustness to variations in the patient motion which is consistent across all patients investigated. The robustness is similar or greater than that seen in other methods when considering the dose to the target, while delivering lower doses to OARs than methods with greater robustness.

## Chapter 5

### Conclusions and future work

In this thesis, a framework for the optimization and delivery of respiratory motion adaptive VMAT was developed and investigated. In Chapter 2 a non-motion adaptive VMAT optimization method was implemented in the open-source treatment planning system matRad. This implementation included a novel dose calculation method which accurately accounts for the continuous MLC leaf motion undergone in VMAT delivery during the optimization itself. This allows for a large reduction, of 4 to 8 depending on the tumour site, in the number of dose calculation angles required for plan optimization compared to the discrete aperture method, while having a similar dose calculation accuracy and final plan quality.

The aperture library-enabled real-time robust adaptation (ALERT-RA) framework presented in Chapter 4 was built from this VMAT optimization method. The core of the ALERT-RA framework lies in the use of a library of MLC apertures and MU rates for each combination of gantry angle and

tumour position. As the gantry rotates around the breathing patient, the position of the tumour would be tracked in real time; this information would then be used to select the corresponding aperture and MU rate from the library, ensuring that the tumour is continuously irradiated with an aperture appropriate to its current position. Because of the additional temporal degree of freedom, the optimization algorithm can selectively increase the MU rate for apertures in the library corresponding to patient geometries in which the target is farther away from a dose-sensitive OAR than in others, potentially decreasing the total dose to OARs for a given tumour dose.

Each individual motion trajectory would then deliver a different dose distribution to the patient, both as a consequence of the sequence of patient geometries varying in time, and the corresponding set of delivered apertures selected from the library. The aperture library was therefore optimized to account for possible variations in the patient motion using robust optimization, under the conventional optimization (CO) and probabilistic optimization (PO) methods. Specifically, with the CO planning method, the expectation value of the dose across all possible trajectories was optimized to ensure the delivery of a dose distribution of ideal plan quality, independent of the realized respiratory motion trajectory. The PO method included an additional term in the minimized objective function representing the variance of the dose to each voxel with respect to fluctuations in the respiratory motion trajectory. Plans optimized with this method were therefore expected to be more robust to these fluctuations.

The calculation of the expectation value and variance requires a probabilistic respiratory motion model which can calculate the probability of arbitrary trajectories. A Markov chain model was constructed and evaluated in Chapter 3 to fulfill this purpose. The model was trained and tested using patient respiratory motion data acquired during the same delivery fraction. Statistical tests were performed to compare the model predictions of the probability distribution of the tumour position to that observed in the testing data. This model was found to be able to adequately estimate the distribution of patient respiratory motion. More importantly, the model can reproduce the evolution of such distribution from fractions of a second to the order of a minute into delivery, which is when the stationary distribution is reached. The model trained on data originating from one fraction was, however, found to be incompatible with testing data originating from a separate fraction, indicating that model retraining may be required when the respiratory pattern changes significantly.

This model was successfully used in Chapter 4 to optimize 4D-VMAT plans under the proposed ALERT-RA framework for a total of 4 simulated patients with varying tumour size and position, for conventionally-fractionated and SBRT plans. Plans were also optimized with 4D methods found in the literature: single trajectory optimization (STO), direct aperture deformation (DAD), and motion kernel optimization (MKO). This comparison was done to compare the plan quality, and, most importantly, the robustness of the plan quality to variations in the respiratory motion.

It was found that for each patient, the plan optimized under the CO method gave similar or less dose to OARs, while delivering a dose to the target which was more robust to changes in the respiratory motion compared to some methods, with less or equal robustness as compared to others. The PO plans were not as robust, despite the inclusion of the variance term in the objective function. Robustness may have improved with a greater number of iterations, but this would have come at the cost of longer optimization times. The MKO method, which optimized a single set of apertures encompassing the range of motion of the tumour, delivered a dose to the target with greater robustness than the CO method. The DAD method also gave plans with more robust target coverage, but only for patients with smaller tumours; for these plans, the maximum dose threshold to the tumour was also exceeded. The increased robustness of both of these methods compared to the CO plan came at a cost, however, of higher dose to OARs.

SBRT plans were observed to be more robust to variations in the patient motion, despite the smaller number of fractions. This was attributed to the longer delivery times used for the SBRT plans, and the corresponding increase in the total number of respiratory cycles; the change in the total dose due to the variation in any one cycle is thus diminished. Both MKO and STO methods yielded SBRT plans which were slightly more robust than the corresponding CO plan, due to the independence of the delivered MLC apertures with respect to the motion trajectory. Because of this, however, neither of these two methods could use the varying spatial relationship between the tumour

and OARs as a function of the respiratory motion to their advantage, unlike the CO plan. This resulted in the MKO and STO plans delivering a higher dose to OARs than the CO plan, exceeding or nearly exceeding the DVH objective thresholds for the lung volume.

All 4 patients were simulated using the XCAT numerical phantom [64], an artificial model of a real patient; the only variation between patients was in the size and location of the implanted lung tumour. All patients exhibited the same respiratory motion throughout the geometry, with the motion of the tumour only varying as its location within the moving lung anatomy changed. Furthermore, the tumour and corresponding CTV were limited to being perfectly spherical in shape. To investigate the effect of different patient sizes, geometries, and respiratory motion patterns on the results observed here, the treatment planning method comparison study should be repeated with actual patient 4D-CT data, exhibiting realistic variations in the OAR and tumour anatomy and respiratory motion.

The effects of varying the patient respiratory motion model between optimization and delivery was investigated. With the two particular models used in this study, this led to a decrease or increase in the dose to the target of 1 Gy to 4 Gy, depending on which model was used for optimization and which for delivery. The increase in dose had the effect of also increasing the 5th percentile of the prescription dose volume to 99.6%; the decrease in dose, however, reduced the 5th and 95th percentile prescription dose volumes to 3.1% and 10.5%. The intercomparison results from both the statistical

evaluation of the model and the dose robustness study call into question the suitability of using a single respiratory motion model from optimization through the delivery of all fractions.

## 5.1 Future work

One of the challenges that needs to be resolved before clinical implementation of the ALERT-RA framework is how to account for changes in the patient respiratory motion model between optimization and delivery. The respiratory motion model used for optimization would be acquired before any planning could commence; a likely time would be to acquire the 5 min of motion data required to train the model just before or after the planning 4D-CT is acquired. A patient whose treatment is planned using the ALERT-RA framework based on this model may then change their respiratory motion model from optimization to delivery, or between individual treatment fractions. In that case, the aperture library that was originally planned would become inappropriate for that particular day of treatment, delivering a higher or lower dose to the tumour than expected.

Adaptive radiation therapy (ART) is a technique which accounts for inter-fractional variations of the patient, usually changing anatomy, by reoptimizing the treatment plan if the originally plan becomes suboptimal [99, 100]. Tumour tracking as a response to intrafractional variations of the anatomy, such as the ALERT-RA method presented in this thesis, could also be considered

methods of real-time ART [101]. ART could be applied to the 4D-VMAT framework presented here by reoptimizing the aperture library if the patient's respiratory motion becomes different from that acquired before optimization. This is similar to the method proposed by Chan and Mišić [102], in which a time-independent tumour position distribution, which is used to perform robust optimization on an IMRT plan, is updated each fraction based on newly-acquired respiratory motion data. Robust optimization is then performed based on this new probability distribution.

The application of ART to ALERT-RA would require acquiring respiratory motion data each fraction on which to train a new model, taking place ideally while the patient is on the treatment couch. Thresholds would then need to be defined to trigger whether or not ART is needed: it may be possible that the patient's model has not changed significantly from before optimization.

Such thresholds should ideally be based on the similarity of the dose delivered from each model. This may be based on the expectation value of the dose, which can be calculated in approximately 20 s. For a more complete picture of the robustness of the plan to variations in the motion under the model, however, PDVHs would need to be calculated, which require approximately 15 min of calculation time per fraction. To decrease the amount of time between data acquisition and delivery, the triggering decision for ART might instead be based on differences in the respiratory motion models themselves which could be evaluated much quicker; for example, differences in the matrix elements themselves, or in the stationary distribution. To do

this, correlations between these model differences and discrepancies in the resulting dose would have to be investigated across an exhaustive range of motion models.

The applicability of ART to the ALERT-RA framework would also need to be tested. The total optimization time for an aperture library planned using CO is currently 3 h; this time does not include calculation of the dose-influence matrices, which could be done once before initial optimization, and then stored on a data drive for fast retrieval. This optimization time is prohibitively long for ART, which can perform reoptimization to adapt to interfractional changes within seconds [103, 100]. This has been done by deforming the original optimized MLC apertures to changes in the tumour geometry, and then by restricting the optimization space to only the MU output by each aperture. Reductions in the reoptimization time of the aperture library might then be realized by initializing the optimization with the original aperture library, and restricting the optimized variables to the MU output by each aperture in the library. This should be tested with different respiratory models used for initial optimization and reoptimization to determine the similarity of the original and reoptimized dose, and also the gains in optimization efficiency.

As indicated in Section 3.3.2, the two respiratory motion models used in this thesis were trained on patient data acquired from two separate delivery fractions on the same patient. The source of this data was a database of clinical patient respiratory data acquired during 102 CyberKnife treatment fractions delivered on 32 unique patients at Georgetown University Hospital [80]. An

exhaustive study of motion models trained with each of these datasets should be conducted, with the aim of discovering patterns in the models and the optimal parameters. This study could also be used to investigate possible models beyond the Markov chain model presented here. This may include the construction and training of a “super”-model, as discussed in Section 4.4, having different Markov chain models as possible outcomes with distinct probabilities.

As mentioned earlier, the treatment planning comparison study should be repeated with a more comprehensive array of patient data. Specifically, real patient 4D-CT images should be used, with varying tumour size and location, respiratory motion, and patient geometry and size. This should be done with the purpose of determining if the results discovered here for one particular dataset hold true when optimizing plans with the various planning methods, or if other results appear. Ideally, 4D-CT patient data would be gathered with the corresponding temporal respiratory motion data required to construct the motion model. This would ensure that both the geometrical data and probabilistic model are specific to the same patient. Furthermore, this framework could also be applied to other radiotherapy treatment sites which are affected by respiratory motion, such as the liver or breast.

Finally, the 4D-VMAT ALERT-RA framework must be tested by performing the delivery with a linac on a phantom, simulating the delivery to an actual patient. A deformable lung phantom with the ability to follow arbitrary motion trajectories, such as the one proposed by Gholampourkashi

et al. [104], should be used. Two-dimensional and point dose measurements can be carried out in or near the tumour in this phantom to verify the accuracy of the predicted dose, dependent on the programmed trajectory. A tracking technique such as PeTrack [39] could then be used to track the motion of the tumour in real time. This tracking system must then be interfaced with the optimized aperture library and the linac, so that the correct aperture for the moving tumour is selected and then delivered by the MLC.

## List of References

- [1] H.-P. Wieser, E. Cisternas, N. Wahl, S. Ulrich, A. Stadler, H. Mescher, L.-R. Müller, T. Klinge, H. Gabrys, L. Burigo, A. Mairani, S. Ecker, B. Ackermann, M. Ellerbrock, K. Parodi, O. Jäkel, and M. Bangert, “Development of the open-source dose calculation and optimization toolkit matRad,” *Medical Physics*, vol. 44, no. 6, pp. 2556–2568, 2017.
- [2] MATLAB, *version 9.5.0.944444 (R2018b)*. Natick: The MathWorks Inc., 2018.
- [3] E. Christiansen, E. Heath, and T. Xu, “Continuous aperture dose calculation and optimization for volumetric modulated arc therapy,” *Physics in Medicine & Biology*, vol. 63, no. 21, p. 21NT01, 2018.
- [4] E. Christiansen, M. Bangert, T. Xu, and E. Heath, “Implementation of VMAT in matRad, an open source treatment planning toolkit,” in *Sixty-Third Annual Scientific Meeting of the Canadian Organization of Medical Physicists*, *Medical Physics*, vol. 44, no. 8, 2017, p. 4394.

- [5] E. Christiansen, E. Heath, and T. Xu, “A Continuous Aperture Approach for VMAT Plan Optimization,” in *Scientific Abstracts and Sessions, Medical Physics*, vol. 45, no. 6, 2018, p. e412.
- [6] E. Christiansen, E. Heath, and T. Xu, “Markov chain modeling of respiratory motion,” in *Sixty-Fifth Annual Scientific Meeting of Canadian Organization of Medical Physicists, Medical Physics*, vol. 46, no. 11, 2019, p. 5395.
- [7] E. Christiansen, E. Heath, and T. Xu, “A framework for robust optimization of 4D-VMAT,” in *Sixty-Fifth Annual Scientific Meeting of Canadian Organization of Medical Physicists, Medical Physics*, vol. 46, no. 11, 2019, p. 5372.
- [8] D. R. Brenner, H. K. Weir, A. A. Demers, L. F. Ellison, C. Louzado, A. Shaw, D. Turner, R. R. Woods, and L. M. Smith, “Projected estimates of cancer in Canada in 2020,” *Canadian Medical Association Journal*, vol. 192, no. 9, pp. E199–E205, 2020.
- [9] Canadian Cancer Statistics Advisory Committee, *Canadian Cancer Statistics 2019*, Toronto, 2019.
- [10] C. de Oliveira, K. E. Bremner, R. Pataky, N. Gunraj, M. Haq, K. Chan, W. Y. Cheung, J. S. Hoch, S. Peacock, and M. D. Krahn, “Trends in use and cost of initial cancer treatment in Ontario: a population-based descriptive study,” *CMAJ Open*, vol. 1, no. 4, pp. E151–E158, 2013.

- [11] M. Cheng, S. Jolly, W. O. Quarshie, N. Kapadia, F. D. Vigneau, and F.-M. Spring Kong, “Modern Radiation Further Improves Survival in Non-Small Cell Lung Cancer: An Analysis of 288,670 Patients,” *Journal of Cancer*, vol. 10, no. 1, pp. 168–177, 2019.
- [12] D. E. Lea, *Actions of Radiations on Living Cells*. Cambridge: Cambridge University Press, 1946.
- [13] M. A. Bender, H. G. Griggs, and J. S. Bedford, “Mechanisms of chromosomal aberration production III. Chemicals and ionizing radiation,” *Mutation Research*, vol. 23, no. 2, pp. 197–212, 1974.
- [14] J. Bergonié and L. Tribondeau, “Interprétation de quelques résultats de la radiothérapie et essai de fixation d’une technique rationnelle,” *Comptes rendus des séances de l’Académie des sciences*, vol. 143, pp. 983–985, 1906.
- [15] B. Emami, J. Lyman, L. Coia, M. Goitein, J. E. Munzenrider, B. Shank, L. J. Solin, and M. Wesson, “Tolerance of normal tissue to therapeutic irradiation,” *International Journal of Radiation Oncology • Biology • Physics*, vol. 21, no. 1, pp. 109–122, 1991.
- [16] B. Emami, “Tolerance of Normal Tissue to Therapeutic Irradiation,” *Reports of Radiotherapy and Oncology*, vol. 1, no. 1, pp. 35–48, 2013.
- [17] M. Patel, S. Maraboyina, and T. Kim, “Brachytherapy vs External

- Beam Radiation in the Management of Non Small Cell Lung Cancer,” *Brachytherapy*, vol. 18, no. 3, p. S39, 2019.
- [18] A. Taylor and M. E. B. Powell, “Intensity-modulated radiotherapy — what is it?” *Cancer Imaging*, vol. 4, no. 2, pp. 68–73, 2004.
- [19] M. M. Elkind, H. Sutton-Gilbert, W. B. Moses, T. Alescio, and R. W. Swain, “Radiation Response of Mammalian Cells Grown in Culture: V. Temperature Dependence of the Repair of X-Ray Damage in Surviving Cells (Aerobic and Hypoxic),” *Radiation Research*, vol. 25, no. 2, pp. 359–376, 1965.
- [20] C. P. Karger, “Biological Models in Treatment Planning,” in *New Technologies in Radiation Oncology*, W. Schlegel, T. Bortfeld, and A.-L. Grosu, Eds. Berlin, Heidelberg: Springer, 2006, ch. 18, pp. 221–235.
- [21] J. M. Brown, D. J. Carlson, and D. J. Brenner, “The Tumor Radiobiology of SRS and SBRT: Are More Than the 5 Rs Involved?” *International Journal of Radiation Oncology • Biology • Physics*, vol. 88, no. 2, pp. 254–262, 2014.
- [22] E. J. Moding, M. B. Kastan, and D. G. Kirsch, “Strategies for optimizing the response of cancer and normal tissues to radiation,” *Nature Reviews Drug Discovery*, vol. 12, no. 7, pp. 526–542, 2013.
- [23] C. X. Yu, “Intensity-modulated arc therapy with dynamic multileaf

- collimation: an alternative to tomotherapy,” *Physics in Medicine & Biology*, vol. 40, no. 9, pp. 1435–1449, 1995.
- [24] K. Otto, “Volumetric modulated arc therapy: IMRT in a single gantry arc,” *Medical Physics*, vol. 35, no. 1, pp. 310–317, 2008.
- [25] J. Unkelbach, T. Bortfeld, D. Craft, M. Alber, M. Bangert, R. Bokrantz, D. Chen, R. Li, L. Xing, C. Men, S. Nill, D. Papp, E. Romeijn, and E. Salari, “Optimization approaches to volumetric modulated arc therapy planning,” *Medical Physics*, vol. 42, no. 3, pp. 1367–1377, 2015.
- [26] ICRU, “Report 62. Prescribing, Recording and Reporting Photon Beam Therapy (Supplement to ICRU Report 50),” International Commission on Radiation Units and Measurements, Bethesda, MD, Tech. Rep., 1999.
- [27] T. Krieger and O. A. Sauer, “Monte Carlo- versus pencil-beam-/collapsed-cone-dose calculation in a heterogeneous multi-layer phantom,” *Physics in Medicine & Biology*, vol. 50, no. 5, pp. 859–868, 2005.
- [28] S. Webb, “The physical basis of IMRT and inverse planning,” *British Journal of Radiology*, vol. 76, no. 910, pp. 678–689, 2003.
- [29] P. J. Keall, G. S. Mageras, J. M. Balter, R. S. Emery, K. M. Forster, S. B. Jiang, J. M. Kapatoes, D. A. Low, M. J. Murphy, B. R. Murray, C. R. Ramsey, M. B. Van Herk, S. S. Vedam, J. W. Wong, and E. Yorke,

- “The management of respiratory motion in radiation oncology report of AAPM Task Group 76,” *Medical Physics*, vol. 33, no. 10, pp. 3874–3900, 2006.
- [30] S. C. Erridge, Y. Seppenwoolde, S. H. Muller, M. van Herk, K. De Jaeger, J. S. A. Belderbos, L. J. Boersma, and J. V. Lebesque, “Portal imaging to assess set-up errors, tumor motion and tumor shrinkage during conformal radiotherapy of non-small cell lung cancer,” *Radiotherapy and Oncology*, vol. 66, pp. 75–85, 2003.
- [31] T. Bortfeld, S. B. Jiang, and E. Rietzel, “Effects of Motion on the Total Dose Distribution,” *Seminars in Radiation Oncology*, vol. 14, no. 1, pp. 41–51, 2004.
- [32] 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 & Biology*, vol. 43, no. 1, pp. 91–104, 1998.
- [33] 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*, vol. 47, no. 4, pp. 1121–1135, 2000.
- [34] S. J. Thomas, B. J. Evans, L. Harihar, H. J. Chantler, A. G. R. Martin, and S. V. Harden, “An evaluation of the mid-ventilation method for

- the planning of stereotactic lung plans,” *Radiotherapy and Oncology*, vol. 137, pp. 110–116, 2019.
- [35] 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 & Biology*, vol. 50, no. 12, pp. 2779–2798, 2005.
- [36] R. I. Berbeco, S. Nishioka, H. Shirato, G. T. Y. Chen, and S. B. Jiang, “Residual motion of lung tumours in gated radiotherapy with external respiratory surrogates,” *Physics in Medicine & Biology*, vol. 50, no. 16, pp. 3655–3667, 2005.
- [37] C. R. Ramsey, D. Scaperoth, and D. Arwood, “Clinical experience with a commercial respiratory gating system,” *International Journal of Radiation Oncology • Biology • Physics*, vol. 48, no. 3, pp. 164–165, 2000.
- [38] H. Shirato, S. Shimizu, T. Shimizu, T. Nishioka, and K. Miyasaka, “Real-time tumour-tracking radiotherapy,” *Lancet*, vol. 353, no. 9161, pp. 1331–1332, 1999.
- [39] T. Xu, J. T. Wong, P. M. Shikhaliev, J. L. Ducote, M. S. Al-Ghazi, and S. Molloy, “Real-time tumor tracking using implanted positron emission markers: Concept and simulation study,” *Medical Physics*, vol. 33, no. 7, pp. 2598–2609, 2006.

- [40] M. Chamberland, R. Wassenaar, B. Spencer, and T. Xu, “Performance evaluation of real-time motion tracking using positron emission fiducial markers,” *Medical Physics*, vol. 38, no. 2, pp. 810–819, 2011.
- [41] O. L. Green, L. J. Rankine, B. Cai, A. Curcuru, R. Kashani, V. Rodriguez, H. H. Li, P. J. Parikh, C. G. Robinson, J. R. Olsen, S. Mutic, S. M. Goddu, and L. Santanam, “First clinical implementation of real-time, real anatomy tracking and radiation beam control,” *Medical Physics*, vol. 45, no. 8, pp. 3728–3740, 2018.
- [42] J. Hanley, M. M. Debois, D. Mah, G. S. Mageras, A. Raben, K. Rosenzweig, B. Mychalczak, L. H. Schwartz, P. J. Gloeggler, W. Lutz, C. C. Ling, S. A. Leibel, Z. Fuks, and G. J. Kutcher, “Deep inspiration breath-hold technique for lung tumors: the potential value of target immobilization and reduced lung density in dose escalation,” *International Journal of Radiation Oncology • Biology • Physics*, vol. 45, no. 3, pp. 603–611, 1999.
- [43] J. W. Wong, M. B. Sharpe, D. A. Jaffray, V. R. Kini, J. M. Robertson, J. S. Stromberg, and A. A. Martinez, “The use of active breathing control (ABC) to reduce margin for breathing motion,” *International Journal of Radiation Oncology • Biology • Physics*, vol. 44, no. 4, pp. 911–919, 1999.
- [44] I. Lax, H. Blomgren, I. Näslund, and R. Svanström, “Stereotactic

- Radiotherapy of Malignancies in the Abdomen: Methodological aspects,” *Acta Oncologica*, vol. 33, no. 6, pp. 677–683, 1994.
- [45] Y. Negoro, Y. Nagata, T. Aoki, T. Mizowaki, N. Araki, K. Takayama, M. Kokubo, S. Yano, S. Koga, K. Sasai, Y. Shibamoto, and M. Hiraoka, “The effectiveness of an immobilization device in conformal radiotherapy for lung tumor: reduction of respiratory tumor movement and evaluation of the daily setup accuracy,” *International Journal of Radiation Oncology • Biology • Physics*, vol. 50, no. 4, pp. 889–898, 2001.
- [46] M. J. Murphy, “Tracking Moving Organs in Real Time,” *Seminars in Radiation Oncology*, vol. 14, no. 1, pp. 91–100, 2004.
- [47] B. W. Raaymakers, J. J. W. Lagendijk, J. Overweg, J. G. M. Kok, A. J. E. Raaijmakers, E. M. Kerkhof, R. W. van der Put, I. Meijding, S. P. M. Crijns, F. Benedosso, M. van Vulpen, C. H. W. de Graaff, J. Allen, and K. J. Brown, “Integrating a 1.5 T MRI scanner with a 6 MV accelerator: proof of concept,” *Physics in Medicine & Biology*, vol. 54, no. 12, pp. N229–N237, 2009.
- [48] A. Schweikard, G. Glosser, M. Bodduluri, M. J. Murphy, and J. R. Adler, “Robotic Motion Compensation for Respiratory Movement during Radiosurgery,” *Computer Aided Surgery*, vol. 5, no. 4, pp. 263–277, 2000.
- [49] M. J. Murphy, M. Isaksson, and J. Jalden, “Adaptive filtering to predict

- lung tumor breathing motion during image-guided radiation therapy,” in *Proceedings of the 16th International Congress on Computer-assisted Radiology and Surgery*, 2002, pp. 539–544.
- [50] P. J. Keall, V. R. Kini, S. S. Vedam, and R. Mohan, “Motion adaptive x-ray therapy: a feasibility study,” *Physics in Medicine & Biology*, vol. 46, no. 1, pp. 1–20, 2001.
- [51] Y. Niu, G. T. Betzel, X. Yang, M. Gui, W. C. Parke, B. Yi, and C. X. Yu, “Planning 4D intensity-modulated arc therapy for tumor tracking with a multileaf collimator,” *Physics in Medicine & Biology*, vol. 62, no. 4, pp. 1480–1500, 2017.
- [52] E. Chin and K. Otto, “Investigation of a novel algorithm for true 4D-VMAT planning with comparison to tracked, gated and static delivery,” *Medical Physics*, vol. 38, no. 5, pp. 2698–2707, 2011.
- [53] J. Unkelbach, M. Alber, M. Bangert, R. Bokrantz, T. C. Y. Chan, J. O. Deasy, A. Fredriksson, B. L. Gorissen, M. van Herk, W. Liu, H. Mahmoudzadeh, O. Nohadani, J. V. Siebers, M. Witte, and H. Xu, “Robust radiotherapy planning,” *Physics in Medicine & Biology*, vol. 63, no. 22, p. 22TR02, 2018.
- [54] A. Fredriksson, A. Forsgren, and B. Hårdemark, “Minimax optimization for handling range and setup uncertainties in proton therapy,” *Medical Physics*, vol. 38, no. 3, pp. 1672–1684, 2011.

- [55] J. Unkelbach and U. Oelfke, “Inclusion of organ movements in IMRT treatment planning via inverse planning based on probability distributions,” *Physics in Medicine & Biology*, vol. 49, no. 17, pp. 4005–4029, 2004.
- [56] M. A. Earl, D. M. Shepard, S. Naqvi, X. A. Li, and C. X. Yu, “Inverse planning for intensity-modulated arc therapy using direct aperture optimization,” *Physics in Medicine & Biology*, vol. 48, no. 8, pp. 1075–1089, 2003.
- [57] J. L. Bedford, “Treatment planning for volumetric modulated arc therapy,” *Medical Physics*, vol. 36, no. 11, pp. 5128–5138, 2009.
- [58] K. Bzdusek, H. Friberger, K. Eriksson, B. Hårdemark, D. Robinson, and M. Kaus, “Development and evaluation of an efficient approach to volumetric arc therapy planning,” *Medical Physics*, vol. 36, no. 6, pp. 2328–2339, 2009.
- [59] L. Masi, R. Doro, V. Favuzza, S. Cipressi, and L. Livi, “Impact of plan parameters on the dosimetric accuracy of volumetric modulated arc therapy,” *Medical Physics*, vol. 40, no. 7, p. 071718, 2013.
- [60] J. Yang, G. Tang, P. Zhang, M. Hunt, S. B. Lim, T. LoSasso, and G. Mageras, “Dose calculation for hypofractionated volumetric-modulated arc therapy: approximating continuous arc delivery and

- tongue-and-groove modeling,” *Journal of Applied Clinical Medical Physics*, vol. 17, no. 2, pp. 3–13, 2016.
- [61] J.-Y. Park, F. Li, J. Li, D. Kahler, J. C. Park, G. Yan, C. Liu, and B. Lu, “Angular under-sampling effect on VMAT dose calculation: An analysis and a solution strategy,” *Medical Physics*, vol. 44, no. 6, pp. 2096–2114, 2017.
- [62] D. Papp and J. Unkelbach, “Direct leaf trajectory optimization for volumetric modulated arc therapy planning with sliding window delivery,” *Medical Physics*, vol. 41, no. 1, p. 011701, 2014.
- [63] T. Bortfeld, W. Schlegel, and B. Rhein, “Decomposition of pencil beam kernels for fast dose calculations in three-dimensional treatment planning,” *Medical Physics*, vol. 20, no. 2, pp. 311–318, 1993.
- [64] W. P. Segars, G. Sturgeon, S. Mendonca, J. Grimes, and B. M. W. Tsui, “4D XCAT phantom for multimodality imaging research,” *Medical Physics*, vol. 37, no. 9, pp. 4902–4915, 2010.
- [65] I. Kawrakow, “VMC++, Electron and Photon Monte Carlo Calculations Optimized for Radiation Treatment Planning,” in *Advanced Monte Carlo for Radiation Physics, Particle Transport Simulation and Applications*, A. Kling, F. J. C. Barão, M. Nakagawa, L. Távora, and P. Vaz, Eds. Springer, Berlin, Heidelberg, 2001, pp. 229–236.

- [66] E. Heath, F. Tessier, and I. Kawrakow, “Investigation of voxel warping and energy mapping approaches for fast 4D Monte Carlo dose calculations in deformed geometries using VMC++,” *Physics in Medicine & Biology*, vol. 56, no. 16, pp. 5187–5202, 2011.
- [67] A. Wächter and L. T. Biegler, “On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming,” *Mathematical Programming*, vol. 106, no. 1, pp. 25—57, 2006.
- [68] R. A. C. Siochi, “Minimizing static intensity modulation delivery time using an intensity solid paradigm,” *International Journal of Radiation Oncology • Biology • Physics*, vol. 43, no. 3, pp. 671–680, 1999.
- [69] E. Wild, M. Bangert, S. Nill, and U. Oelfke, “Noncoplanar VMAT for nasopharyngeal tumor: Plan quality versus treatment time,” *Medical Physics*, vol. 42, no. 5, pp. 2157–2168, 2015.
- [70] D. Craft, M. Bangert, T. Long, D. Papp, and J. Unkelbach, “Shared data for intensity modulated radiation therapy (IMRT) optimization research: the CORT dataset,” *GigaScience*, vol. 3, no. 37, 2014.
- [71] G. A. Ezzell, J. W. Burmeister, N. Dogan, T. J. LoSasso, J. G. Mechalakos, D. Mihailidis, A. Molineu, J. R. Palta, C. R. Ramsey, B. J. Salter, J. Shi, P. Xia, N. J. Yue, and Y. Xiao, “IMRT commissioning: Multiple institution planning and dosimetry comparisons, a report from AAPM Task Group 119,” *Medical Physics*, vol. 36, no. 11, pp. 5359–5373, 2009.

- [72] G. C. Sharp, S. B. Jiang, S. Shimizu, and H. Shirato, “Prediction of respiratory tumour motion for real-time image-guided radiotherapy,” *Physics in Medicine & Biology*, vol. 49, no. 3, pp. 425–440, 2004.
- [73] A. Kalet, G. Sandison, H. Wu, and R. Schmitz, “A state-based probabilistic model for tumor respiratory motion prediction,” *Physics in Medicine & Biology*, vol. 55, no. 24, pp. 7615–7631, 2010.
- [74] T. C. Y. Chan, T. Bortfeld, and J. N. Tsitsiklis, “A robust approach to IMRT optimization,” *Physics in Medicine & Biology*, vol. 51, no. 10, pp. 2567–2583, 2006.
- [75] A. E. Lujan, E. W. Larsen, J. M. Balter, and R. K. T. Haken, “A method for incorporating organ motion due to breathing into 3D dose calculations,” *Medical Physics*, vol. 26, no. 5, pp. 715–720, 1999.
- [76] R. George, S. S. Vedam, T. D. Chung, V. Ramakrishnan, and P. J. Keall, “The application of the sinusoidal model to lung cancer patient respiratory motion,” *Medical Physics*, vol. 32, no. 9, pp. 2850–2861, 2005.
- [77] A. Tolver, “An Introduction to Markov Chains,” 2016, lecture notes, University of Copenhagen, <http://web.math.ku.dk/noter/filer/stoknoter.pdf>.
- [78] F. Noé, B. Keller, and J.-H. Prinz, “Lecture Notes on Stochastic Processes,” 2013, lecture notes, Free University of Berlin, [http:](http://)

[//www.mi.fu-berlin.de/wiki/pub/CompMolBio/MarkovKetten15/stochastic\\_processes\\_2011.pdf](http://www.mi.fu-berlin.de/wiki/pub/CompMolBio/MarkovKetten15/stochastic_processes_2011.pdf).

- [79] K. Pearson, “On lines and planes of closest fit to systems of points in space,” *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, vol. 2, no. 11, pp. 559–572, 1901.
- [80] F. Ernst, “Compensating for Quasi-periodic Motion in Robotic Radio-surgery,” Ph.D. dissertation, University of Lübeck, 2011.
- [81] Digital Signal Processing Committee of the IEEE Acoustics, Speech, and Signal Processing Society, Ed., *Programs for Digital Signal Processing*. New York: IEEE Press, 1979.
- [82] H. Wu, G. C. Sharp, B. Salzberg, D. Kaeli, H. Shirato, and S. B. Jiang, “A finite state model for respiratory motion analysis in image guided radiation therapy,” *Physics in Medicine & Biology*, vol. 49, no. 23, pp. 5357–5372, 2004.
- [83] N. Higham, *Functions of Matrices: Theory and Computation*. Philadelphia: Society for Industrial and Applied Mathematics, 2008.
- [84] J. Unkelbach and U. Oelfke, “Incorporating organ movements in IMRT treatment planning for prostate cancer: Minimizing uncertainties in the inverse planning process,” *Medical Physics*, vol. 32, no. 8, pp. 2471–2483, 2005.

- [85] N. Wahl, P. Henning, H.-P. Wieser, and M. Bangert, “Analytical incorporation of fractionation effects in probabilistic treatment planning for intensity-modulated proton therapy,” *Medical Physics*, vol. 45, no. 4, pp. 1317–1328, 2018.
- [86] L. Xing, J. Siebers, and P. Keall, “Computational Challenges for Image-Guided Radiation Therapy: Framework and Current Research,” *Seminars in Radiation Oncology*, vol. 17, no. 4, pp. 245–257, 2007.
- [87] H. Zhong, E. Weiss, and J. V. Siebers, “Assessment of dose reconstruction errors in image-guided radiation therapy,” *Physics in Medicine & Biology*, vol. 53, no. 3, pp. 719–736, 2008.
- [88] H. Zhong and J. V. Siebers, “Monte Carlo dose mapping on deforming anatomy,” *Physics in Medicine & Biology*, vol. 54, no. 19, pp. 5815–5830, 2009.
- [89] C. L. Williams, P. Mishra, J. Seco, S. S. James, R. H. Mak, R. I. Berbeco, and J. H. Lewis, “A mass-conserving 4D XCAT phantom for dose calculation and accumulation,” *Medical Physics*, vol. 40, no. 7, p. 071728, 2013.
- [90] S. Gholampourkashi, J. E. Cygler, J. Belec, M. Vujicic, and E. Heath, “Monte Carlo and analytic modeling of an Elekta Infinity linac with Agility MLC: Investigating the significance of accurate model parameters

- for small radiation fields,” *Journal of Applied Clinical Medical Physics*, vol. 20, no. 1, pp. 55–67, 2019.
- [91] B. Eiben, J. Bertholet, M. J. Menten, S. Nill, U. Oelfke, and J. R. McClelland, “Consistent and invertible deformation vector fields for a breathing anthropomorphic phantom: a post-processing framework for the XCAT phantom,” *Physics in Medicine & Biology*, vol. 65, no. 16, p. 165005, 2020.
- [92] Y. Feng, C. Castro-Pareja, R. Shekhar, and C. Yu, “Direct aperture deformation: An interfraction image guidance strategy,” *Medical Physics*, vol. 33, no. 12, pp. 4490–4498, 2006.
- [93] G. Cowan, *Statistical Data Analysis*. Oxford: Clarendon Press, 1998.
- [94] I. Kawrakow and M. Fippel, “Investigation of variance reduction techniques for Monte Carlo photon dose calculation using XVMC,” *Physics in Medicine & Biology*, vol. 45, no. 8, pp. 2163–2183, 2000.
- [95] J. Gardner, J. Siebers, and I. Kawrakow, “Dose calculation validation of VMC++ for photon beams,” *Medical Physics*, vol. 34, no. 5, pp. 1809–1818, 2007.
- [96] M. F. Fast, S. Nill, J. L. Bedford, and U. Oelfke, “Dynamic tumor tracking using the Elekta Agility MLC,” *Medical Physics*, vol. 41, no. 11, p. 111719, 2014.

- [97] M. Glitzner, S. P. M. Crijs, B. D. de Senneville, J. J. W. Lagendijk, and B. W. Raaymakers, “On the suitability of Elekta’s Agility 160 MLC for tracked radiation delivery: closed-loop machine performance,” *Physics in Medicine & Biology*, vol. 60, no. 5, pp. 2005–2017, 2015.
- [98] J. J. Gordon, N. Sayah, E. Weiss, and J. V. Siebers, “Coverage optimized planning: Probabilistic treatment planning based on dose coverage histogram criteria,” *Medical Physics*, vol. 37, no. 2, pp. 550–563, 2010.
- [99] D. Yan, F. Vicini, J. Wong, and A. Martinez, “Adaptive radiation therapy,” *Physics in Medicine & Biology*, vol. 42, no. 1, pp. 123–132, 1997.
- [100] S. Lim-Reinders, B. M. Keller, S. Al-Ward, A. Sahgal, and A. Kim, “On-line Adaptive Radiation Therapy,” *International Journal of Radiation Oncology • Biology • Physics*, vol. 99, no. 4, pp. 994–1003, 2017.
- [101] P. Keall, P. Poulsen, and J. T. Booth, “See, Think, and Act: Real-Time Adaptive Radiotherapy,” *Seminars in Radiation Oncology*, vol. 29, no. 3, pp. 228–235, 2019.
- [102] T. C. Y. Chan and V. V. Mišić, “Adaptive and robust radiation therapy optimization for lung cancer,” *European Journal of Operational Research*, vol. 231, no. 3, pp. 745–756, 2013.
- [103] O. Ates, E. E. Ahunbay, M. Moreau, and X. A. Li, “Technical Note:

- A fast online adaptive replanning method for VMAT using flattening filter free beams,” *Medical Physics*, vol. 43, no. 6, pp. 2756–2764, 2016.
- [104] S. Gholampourkashi, J. E. Cygler, B. Lavigne, and E. Heath, “Development of a deformable phantom for experimental verification of 4D Monte Carlo simulations in a deforming anatomy,” *Physica Medica*, vol. 51, pp. 81–90, 2018.
- [105] R. Svensson, P. Källman, and A. Brahme, “An analytical solution for the dynamic control of multileaf collimators,” *Physics in Medicine & Biology*, vol. 39, no. 1, pp. 37–61, 1994.

## Appendix A

# Derivation of bixel-based fluence in continuous aperture VMAT

The following is a derivation of the fluence delivered to a bixel using the continuous aperture fluence calculation method. This derivation was partly inspired by Park et al. [61] and Svensson et al. [105]. Although some of the derivation was performed in Section 2.2.1.1, the entire process is described here for completeness.

According to Equation (2.3), the fluence of a particular bixel  $j$  delivered in an arc section surrounding  $\theta_k$  is equal to (the bixel subscript  $j$  and angle subscript  $\theta_k$  are omitted for simplicity)

$$\phi = \int_{\theta_{\text{init}}}^{\theta_{\text{fin}}} \frac{d\phi}{d\theta}(\theta) d\theta, \quad (\text{A.1})$$

where  $\theta_{\text{init}}$  and  $\theta_{\text{fin}}$  represent the initial and final angles of the arc section. To simplify the calculation, introduce two components of the fluence:  $\phi_L$ ,

defined to be the component of the bixel fluence released by the left leaf in the absence of the right leaf; and  $\phi_R$ , defined to be the component of the bixel fluence blocked by the right leaf in the absence of the left leaf. Then  $\phi = \phi_L - \phi_R$ , with

$$\phi_L = \int_{\theta_{\text{init}}}^{\theta_{\text{fin}}} \frac{d\phi_L}{d\theta}(\theta) d\theta, \quad (\text{A.2a})$$

$$\phi_R = \int_{\theta_{\text{init}}}^{\theta_{\text{fin}}} \frac{d\phi_R}{d\theta}(\theta) d\theta. \quad (\text{A.2b})$$

Assume that both left and right leaves travel from left to right at constant speed (the speeds of the two leaves may be different). Let the initial and final positions of the left and right leaves be  $L_{\text{init}}$ ,  $L_{\text{fin}}$ ,  $R_{\text{init}}$ , and  $R_{\text{fin}}$ , and assume that the linac intensity rate  $\dot{\psi}$  is constant over the arc section surrounding  $\theta$ .

## Component of fluence released by left leaf

Begin with the evaluation of  $\phi_L$ , which will only depend on the motion of the left leaf in the row to which the current bixel belongs.

Starting from Equation (A.2a), change variables to  $x$ , the position of the left leaf:

$$\phi_L = \int_{L_{\text{init}}}^{L_{\text{fin}}} \frac{d\phi_L}{dx}(x) dx. \quad (\text{A.3})$$

$\frac{d\phi_L}{dx}(x)$  can be related to  $\frac{d\phi_L}{dt}(x)$  using the chain rule:

$$\begin{aligned}\frac{d\phi_L}{dx}(x) &= \frac{d\phi_L}{dt}(x) \frac{dt}{dx}(x) \\ &= \frac{1}{v_L} \frac{d\phi_L}{dt}(x) \\ &= \frac{\Delta t}{L_{\text{fin}} - L_{\text{init}}} \frac{d\phi_L}{dt}(x),\end{aligned}\tag{A.4}$$

where  $v_L$  is the constant speed of the left leaf in the arc section, and  $\Delta t$  is the time taken for the gantry to rotate through the arc section.

$\frac{d\phi_L}{dt}(x)$  is the rate of fluence released to the bixel by the left leaf when the left leaf is at position  $x$ . At any instant in time, this is given by the linac intensity rate  $\dot{\psi}$  multiplied by the fraction of the bixel not blocked by the left leaf ( $l$  and  $r$  denote the locations of the left and right boundaries of the bixel, respectively):

$$\frac{d\phi_L}{dt}(x) = \begin{cases} \dot{\psi} & \text{for } x < l \\ \dot{\psi} \frac{r-x}{r-l} & \text{for } l \leq x < r \\ 0 & \text{for } x \geq r. \end{cases}\tag{A.5}$$

Inserting Equation (A.5) into Equation (A.4) and substituting  $\psi = \dot{\psi}\Delta t$ , where  $\psi$  is the total linac intensity delivered in the arc section, yields:

$$\frac{d\phi_L}{dx}(x) = \begin{cases} \psi \frac{1}{L_{\text{fin}} - L_{\text{init}}} & \text{for } x < l \\ \psi \frac{r-x}{(L_{\text{fin}} - L_{\text{init}})(r-l)} & \text{for } l \leq x < r \\ 0 & \text{for } x \geq r. \end{cases}\tag{A.6}$$

This expression for the differential fluence released to the bixel by the left leaf is integrated along the leaf trajectory from  $L_{\text{init}}$  to  $L_{\text{fin}}$ , according to Equation (A.3). There are six cases for this integral, depending on the initial and final left leaf positions  $L_{\text{init}}$  and  $L_{\text{fin}}$  relative to the left and right edges  $l$  and  $r$  of the bixel.

**Case 1:**  $L_{\text{init}} < l$  &  $r < L_{\text{fin}}$

In this case, both the initial and final leaf positions fall outside of the bixel. The integral of Equation (A.3) is evaluated thus:

$$\begin{aligned}
 \phi_L &= \int_{L_{\text{init}}}^{L_{\text{fin}}} \frac{d\phi_L}{dx}(x) dx \\
 &= \int_{L_{\text{init}}}^l \psi \frac{1}{L_{\text{fin}} - L_{\text{init}}} dx + \int_l^r \psi \frac{r - x}{(L_{\text{fin}} - L_{\text{init}})(r - l)} dx + \int_r^{L_{\text{fin}}} 0 dx \\
 &= \psi \frac{1}{L_{\text{fin}} - L_{\text{init}}} \left[ l - L_{\text{init}} + \frac{1}{r - l} \left( r^2 - \frac{1}{2}r^2 - rl + \frac{1}{2}l^2 \right) \right] \\
 &= \psi \frac{(r + l)/2 - L_{\text{init}}}{L_{\text{fin}} - L_{\text{init}}}.
 \end{aligned} \tag{A.7}$$

Define this quantity to be  $\phi_L^1$ , which will be used in Cases 2-4.

**Case 2:**  $l < L_{\text{init}} < r < L_{\text{fin}}$

In this case, the initial leaf position is between the bixel edges. The integral of Equation (A.3) is equal to:

$$\begin{aligned}\phi_L &= \int_{L_{\text{init}}}^{L_{\text{fin}}} \frac{d\phi_L}{dx}(x) dx \\ &= \int_{L_{\text{init}}}^r \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx + \int_r^{L_{\text{fin}}} 0 dx.\end{aligned}\quad (\text{A.8})$$

Adding and subtracting the quantity  $\phi_L^1 = \int_{L_{\text{init}}}^l \psi \frac{1}{L_{\text{fin}}-L_{\text{init}}} dx + \int_l^r \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx$  yields:

$$\begin{aligned}\phi_L &= \phi_L^1 - \int_{L_{\text{init}}}^l \psi \frac{1}{L_{\text{fin}}-L_{\text{init}}} dx - \int_l^r \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx \\ &\quad + \int_{L_{\text{init}}}^r \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx \\ &= \phi_L^1 + \psi \frac{L_{\text{init}}-l}{L_{\text{fin}}-L_{\text{init}}} - \int_l^{L_{\text{init}}} \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx \\ &= \phi_L^1 + \psi \frac{L_{\text{init}}-l}{L_{\text{fin}}-L_{\text{init}}} - \psi \frac{1}{(L_{\text{fin}}-L_{\text{init}})(r-l)} \left( L_{\text{init}}r - \frac{1}{2}L_{\text{init}}^2 - rl + \frac{1}{2}l^2 \right) \\ &= \phi_L^1 + \psi \frac{(L_{\text{init}}-l)^2}{2(L_{\text{fin}}-L_{\text{init}})(r-l)}.\end{aligned}\quad (\text{A.9})$$

Note that the bixel fluence for this case is acquired by adding a positive corrective term to the base value  $\phi_L^1$ .

**Case 3:**  $L_{\text{init}} < l < L_{\text{fin}} < r$

In this case, the final leaf position is between the bixel edges. The integral of Equation (A.3) is equal to:

$$\begin{aligned}\phi_L &= \int_{L_{\text{init}}}^{L_{\text{fin}}} \frac{d\phi_L}{dx}(x) dx \\ &= \int_{L_{\text{init}}}^l \psi \frac{1}{L_{\text{fin}} - L_{\text{init}}} dx + \int_l^{L_{\text{fin}}} \psi \frac{r-x}{(L_{\text{fin}} - L_{\text{init}})(r-l)} dx.\end{aligned}\quad (\text{A.10})$$

Similarly to Case 2, adding and subtracting the quantity  $\phi_L^1 = \int_{L_{\text{init}}}^l \psi \frac{1}{L_{\text{fin}} - L_{\text{init}}} dx + \int_l^r \psi \frac{r-x}{(L_{\text{fin}} - L_{\text{init}})(r-l)} dx$  yields:

$$\begin{aligned}\phi_L &= \phi_L^1 - \int_l^r \psi \frac{r-x}{(L_{\text{fin}} - L_{\text{init}})(r-l)} dx + \int_l^{L_{\text{fin}}} \psi \frac{r-x}{(L_{\text{fin}} - L_{\text{init}})(r-l)} dx \\ &= \phi_L^1 - \int_{L_{\text{fin}}}^r \psi \frac{r-x}{(L_{\text{fin}} - L_{\text{init}})(r-l)} dx \\ &= \phi_L^1 - \psi \frac{1}{(L_{\text{fin}} - L_{\text{init}})(r-l)} \left( \frac{1}{2}r^2 - L_{\text{fin}}r + \frac{1}{2}L_{\text{fin}}^2 \right) \\ &= \phi_L^1 - \psi \frac{(r - L_{\text{fin}})^2}{2(L_{\text{fin}} - L_{\text{init}})(r-l)}.\end{aligned}\quad (\text{A.11})$$

Note again that the bixel fluence for this case is acquired by adding a negative corrective term to the base value  $\phi_L^1$ .

**Case 4:**  $l < L_{\text{init}} < L_{\text{fin}} < r$

In this case, both initial and final leaf positions are between the bixel edges.

The integral of Equation (A.3) is equal to:

$$\begin{aligned}\phi_L &= \int_{L_{\text{init}}}^{L_{\text{fin}}} \frac{d\phi_L}{dx}(x) dx \\ &= \int_{L_{\text{init}}}^{L_{\text{fin}}} \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx.\end{aligned}\quad (\text{A.12})$$

Similarly to Cases 2 and 3, adding and subtracting the quantity  $\phi_L^1 = \int_{L_{\text{init}}}^l \psi \frac{1}{L_{\text{fin}}-L_{\text{init}}} dx + \int_l^r \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx$  yields:

$$\begin{aligned}\phi_L &= \phi_L^1 - \int_{L_{\text{init}}}^l \psi \frac{1}{L_{\text{fin}}-L_{\text{init}}} dx - \int_l^r \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx \\ &\quad + \int_{L_{\text{init}}}^{L_{\text{fin}}} \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx \\ &= \phi_L^1 - \psi \frac{L_{\text{init}}-l}{L_{\text{fin}}-L_{\text{init}}} - \int_l^{L_{\text{init}}} \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx \\ &\quad - \int_{L_{\text{fin}}}^r \psi \frac{r-x}{(L_{\text{fin}}-L_{\text{init}})(r-l)} dx.\end{aligned}\quad (\text{A.13})$$

Comparing this equation to the second lines of Equation (A.9) and Equation (A.11) shows that both the positive and negative corrective terms from Cases 2 and 3, respectively, are added to the base value  $\phi_L^1$  in this case. The final result is:

$$\phi_L = \phi_L^1 + \psi \frac{(L_{\text{init}}-l)^2}{2(L_{\text{fin}}-L_{\text{init}})(r-l)} - \psi \frac{(r-L_{\text{fin}})^2}{2(L_{\text{fin}}-L_{\text{init}})(r-l)}.\quad (\text{A.14})$$

**Case 5:**  $L_{\text{init}} < L_{\text{fin}} < l$ 

In this case, both initial and final leaf positions are to the left of the left bixel edge. The integral of Equation (A.3) is equal to:

$$\begin{aligned}
 \phi_L &= \int_{L_{\text{init}}}^{L_{\text{fin}}} \frac{d\phi_L}{dx}(x) dx \\
 &= \int_{L_{\text{init}}}^{L_{\text{fin}}} \psi \frac{1}{L_{\text{fin}} - L_{\text{init}}} dx \\
 &= \psi.
 \end{aligned} \tag{A.15}$$

This implies that the bixel receives the full linac intensity, which is to be expected, since the left leaf does not block the bixel at any point along its trajectory.

**Case 6:**  $r < L_{\text{init}} < L_{\text{fin}}$ 

In this final case, both initial and final leaf positions are to the right of the right bixel edge. The integral of Equation (A.3) is equal to:

$$\begin{aligned}
 \phi_L &= \int_{L_{\text{init}}}^{L_{\text{fin}}} \frac{d\phi_L}{dx}(x) dx \\
 &= \int_{L_{\text{init}}}^{L_{\text{fin}}} 0 dx \\
 &= 0.
 \end{aligned} \tag{A.16}$$

This implies that the bixel receives 0 fluence, which is to be expected, since the left leaf blocks the bixel during its entire trajectory.

## Component of fluence blocked by right leaf

The derivation of  $\phi_R$  follows a similar procedure as that for  $\phi_L$ , with six separate cases requiring consideration, now depending on the initial and final right leaf positions  $R_{\text{init}}$  and  $R_{\text{fin}}$  relative to the bixel edges. Beginning with Equation (A.2b), change variables to  $x$ , the position of the right leaf:

$$\phi_R = \int_{R_{\text{init}}}^{R_{\text{fin}}} \frac{d\phi_R}{dx}(x) dx. \quad (\text{A.17})$$

$\frac{d\phi_R}{dx}(x)$  can be related to  $\frac{d\phi_R}{dt}(x)$  using the chain rule:

$$\begin{aligned} \frac{d\phi_R}{dx}(x) &= \frac{d\phi_R}{dt}(x) \frac{dt}{dx}(x) \\ &= \frac{1}{v_R} \frac{d\phi_R}{dt}(x) \\ &= \frac{\Delta t}{R_{\text{fin}} - R_{\text{init}}} \frac{d\phi_R}{dt}(x), \end{aligned} \quad (\text{A.18})$$

where  $v_R$  is the constant speed of the right leaf in the arc section, and  $\Delta t$  is the time taken for the gantry to rotate through the arc section.

$\frac{d\phi_R}{dt}(x)$  is the rate of bixel fluence blocked by the right leaf when this leaf is at position  $x$ . At any instant in time, this is given by the linac intensity rate  $\dot{\psi}$  multiplied by the fraction of the bixel blocked by the right leaf ( $l$  and  $r$

denote the locations of the left and right boundaries of the bixel, respectively):

$$\frac{d\phi_R}{dt}(x) = \begin{cases} \dot{\psi} & \text{for } x < l \\ \dot{\psi} \frac{r-x}{r-l} & \text{for } l \leq x < r \\ 0 & \text{for } x \geq r. \end{cases} \quad (\text{A.19})$$

Inserting Equation (A.19) into Equation (A.18) and substituting  $\psi = \dot{\psi}\Delta t$ , where  $\psi$  is the total linac intensity delivered in the arc section, yields:

$$\frac{d\phi_R}{dx}(x) = \begin{cases} \psi \frac{1}{R_{\text{fin}} - R_{\text{init}}} & \text{for } x < l \\ \psi \frac{r-x}{(R_{\text{fin}} - R_{\text{init}})(r-l)} & \text{for } l \leq x < r \\ 0 & \text{for } x \geq r. \end{cases} \quad (\text{A.20})$$

The rest of the derivation is greatly simplified by noting that  $\frac{d\phi_R}{dx}(x)$  in Equation (A.20) can be obtained from the expression of  $\frac{d\phi_L}{dx}(x)$  in Equation (A.6) by performing the following substitutions:  $L_{\text{init}} \rightarrow R_{\text{init}}$ , and  $L_{\text{fin}} \rightarrow R_{\text{fin}}$ . These same substitutions can therefore be made in the expressions for the integrated fluence, yielding six cases for the fluence blocked by the right leaf, as follows:

**Case 1:**  $R_{\text{init}} < l$  &  $r < R_{\text{fin}}$

$$\phi_R = \psi \frac{(r+l)/2 - R_{\text{init}}}{R_{\text{fin}} - R_{\text{init}}}. \quad (\text{A.21})$$

Define this quantity to be  $\phi_R^1$ , which will be used in Cases 2 and 3.

**Case 2:**  $l < R_{\text{init}} < r < R_{\text{fin}}$

$$\phi_R = \phi_R^1 + \psi \frac{(R_{\text{init}} - l)^2}{2(R_{\text{fin}} - R_{\text{init}})(r - l)}. \quad (\text{A.22})$$

**Case 3:**  $R_{\text{init}} < l < R_{\text{fin}} < r$

$$\phi_R = \phi_R^1 - \psi \frac{(r - R_{\text{fin}})^2}{2(R_{\text{fin}} - R_{\text{init}})(r - l)}. \quad (\text{A.23})$$

**Case 4:**  $R_{\text{init}} < l < R_{\text{fin}} < r$

$$\phi_R = \phi_R^1 + \psi \frac{(R_{\text{init}} - l)^2}{2(R_{\text{fin}} - R_{\text{init}})(r - l)} - \psi \frac{(r - R_{\text{fin}})^2}{2(R_{\text{fin}} - R_{\text{init}})(r - l)}. \quad (\text{A.24})$$

**Case 5:**  $R_{\text{init}} < R_{\text{fin}} < l$

$$\phi_R = \psi. \quad (\text{A.25})$$

**Case 6:**  $r < R_{\text{init}} < R_{\text{fin}}$

$$\phi_R = 0 \quad (\text{A.26})$$

## Summary of continuous aperture fluence calculation

The calculation of the total fluence released at a particular bixel as the right and left leaves sweep from one position to the next can be summarized thus:

1. Decompose the fluence  $\phi$  into components  $\phi_L$  and  $\phi_R$ ;
2. Determine which of the six cases each component falls into, and use the corresponding formula;
3. The final fluence is given by the difference  $\phi = \phi_L - \phi_R$ .

## Appendix B

# Derivation of the gradient in continuous aperture VMAT

The following is a derivation of the gradient of the fluence delivered to a bixel using the continuous aperture fluence calculation method.

For the DAO step in continuous aperture VMAT, derivatives of the objective function  $F$  with respect to the following variables are required: left and right leaf positions at the beginning and end of an arc section ( $L_{\text{init}}$ ,  $L_{\text{fin}}$ ,  $R_{\text{init}}$ , and  $R_{\text{fin}}$ ), and the total linac intensity  $\psi$  delivered during the arc section surrounding each  $\theta_k$ . These derivatives can be calculated using the chain rule:

$$\frac{\partial F}{\partial \xi} = \sum_{\theta_k} \sum_j \frac{\partial F}{\partial \phi_{j\theta_k}} \frac{\partial \phi_{j\theta_k}}{\partial \xi}, \quad (\text{B.1})$$

where  $\xi$  is some arbitrary variable. The derivatives of the objective function with respect to the bixel intensities,  $\frac{\partial F}{\partial \phi_{j\theta_k}}$ , are calculated by matRad [1]; the derivatives of the bixel fluences with respect to the optimization variables

of linac intensity and leaf positions must therefore be calculated. As in Appendix A, the bixel subscript  $j$  and angle subscript  $\theta_k$  are omitted for simplicity.

## Gradient of $\phi$ with respect to the total fluence

Begin with the derivative of the total fluence with respect to intensity  $\psi$ . Following Appendix A, the fluence may be separated into two components:  $\phi_L$ , defined to be the component of the bixel fluence released by the left leaf in the absence of the right leaf; and  $\phi_R$ , defined to be the component of the bixel fluence blocked by the right leaf in the absence of the left leaf. As seen in Appendix A, both  $\phi_L$  and  $\phi_R$  are directly proportional to  $\psi$ . Since the total fluence  $\phi$  is the difference of these two, the derivative of the bixel fluence with respect to the linac intensity is:

$$\frac{\partial \phi}{\partial \psi} = \frac{\phi_L - \phi_R}{\psi}. \quad (\text{B.2})$$

## Gradient of $\phi_L$ with respect to the leaf positions

The derivatives with respect to the leaf positions require more care. As before, decompose the fluence  $\phi$  into the components  $\phi_L$  and  $\phi_R$ , and calculate the gradient of each one separately. The gradient of  $\phi$  with respect to a variable

$\xi$  is then equal to:

$$\frac{\partial \phi}{\partial \xi} = \frac{\partial \phi_L}{\partial \xi} - \frac{\partial \phi_R}{\partial \xi}. \quad (\text{B.3})$$

Begin with the evaluation of the derivatives of  $\phi_L$ , which depends only on the initial and final positions of the left leaf,  $L_{\text{init}}$  and  $L_{\text{fin}}$ . The same six cases that were used for the evaluation of the bixel fluence in Appendix A are used for the derivatives, depending on the positions  $L_{\text{init}}$  and  $L_{\text{fin}}$  relative to the bixel edges. In the following, the symbol  $c = (r + l)/2$  is used for the centre of the bixel, and the symbol  $d = r - l$  is used for the bixel width.

**Case 1:**  $L_{\text{init}} < l$  &  $r < L_{\text{fin}}$

The bixel fluence for case 1 was derived to be (from Equation (A.7)):

$$\phi_L = \psi \frac{c/2 - L_{\text{init}}}{L_{\text{fin}} - L_{\text{init}}}. \quad (\text{B.4})$$

The derivatives of this expression with respect to  $L_{\text{init}}$  and  $L_{\text{fin}}$  are given by:

$$\begin{aligned} \frac{\partial \phi_L}{\partial L_{\text{init}}} &= \psi \frac{c - L_{\text{fin}}}{(L_{\text{fin}} - L_{\text{init}})^2}, \\ \frac{\partial \phi_L}{\partial L_{\text{fin}}} &= \psi \frac{L_{\text{init}} - c}{(L_{\text{fin}} - L_{\text{init}})^2}. \end{aligned} \quad (\text{B.5})$$

Define these two quantities to be  $\frac{\partial \phi_L^1}{\partial L_{\text{init}}}$  and  $\frac{\partial \phi_L^1}{\partial L_{\text{fin}}}$ , respectively.

**Case 2:**  $l < L_{\text{init}} < r < L_{\text{fin}}$

The bixel fluence for case 2 was derived to be (from Equation (A.9)):

$$\phi_L = \phi_L^1 + \psi \frac{(L_{\text{init}} - l)^2}{2d(L_{\text{fin}} - L_{\text{init}})}. \quad (\text{B.6})$$

The derivatives of this expression with respect to  $L_{\text{init}}$  and  $L_{\text{fin}}$  are given by:

$$\begin{aligned} \frac{\partial \phi_L}{\partial L_{\text{init}}} &= \frac{\partial \phi_L^1}{\partial L_{\text{init}}} + \psi \frac{(L_{\text{init}} - l)(2L_{\text{fin}} - L_{\text{init}} - l)}{2d(L_{\text{fin}} - L_{\text{init}})^2}, \\ \frac{\partial \phi_L}{\partial L_{\text{fin}}} &= \frac{\partial \phi_L^1}{\partial L_{\text{fin}}} - \psi \frac{(L_{\text{init}} - l)^2}{2d(L_{\text{fin}} - L_{\text{init}})^2}. \end{aligned} \quad (\text{B.7})$$

**Case 3:**  $L_{\text{init}} < l < L_{\text{fin}} < r$

The bixel fluence for case 3 was derived to be (from Equation (A.11)):

$$\phi_L = \phi_L^1 - \psi \frac{(r - L_{\text{fin}})^2}{2d(L_{\text{fin}} - L_{\text{init}})}. \quad (\text{B.8})$$

The derivatives of this expression with respect to  $L_{\text{init}}$  and  $L_{\text{fin}}$  are given by:

$$\begin{aligned} \frac{\partial \phi_L}{\partial L_{\text{init}}} &= \frac{\partial \phi_L^1}{\partial L_{\text{init}}} - \psi \frac{(r - L_{\text{fin}})^2}{2d(L_{\text{fin}} - L_{\text{init}})^2}, \\ \frac{\partial \phi_L}{\partial L_{\text{fin}}} &= \frac{\partial \phi_L^1}{\partial L_{\text{fin}}} + \psi \frac{(r - L_{\text{fin}})(L_{\text{fin}} + r - 2L_{\text{init}})}{2d(L_{\text{fin}} - L_{\text{init}})^2}. \end{aligned} \quad (\text{B.9})$$

**Case 4:**  $l < L_{\text{init}} < L_{\text{fin}} < r$

The bixel fluence for case 4 was derived to be (from Equation (A.13)):

$$\phi_L = \phi_L^1 + \psi \frac{(L_{\text{init}} - l)^2}{2d(L_{\text{fin}} - L_{\text{init}})} - \psi \frac{(r - L_{\text{fin}})^2}{2d(L_{\text{fin}} - L_{\text{init}})}. \quad (\text{B.10})$$

The derivative of this expression with respect to  $L_{\text{init}}$  is given by (inserting the definitions of  $c$  and  $d$  where required):

$$\begin{aligned} \frac{\partial \phi_L}{\partial L_{\text{init}}} &= \frac{\partial \phi_L^1}{\partial L_{\text{init}}} + \psi \frac{(L_{\text{init}} - l)(2L_{\text{fin}} - L_{\text{init}} - l)}{2d(L_{\text{fin}} - L_{\text{init}})^2} - \psi \frac{(r - L_{\text{fin}})^2}{2d(L_{\text{fin}} - L_{\text{init}})^2} \\ &= \psi \frac{c - L_{\text{fin}}}{(L_{\text{fin}} - L_{\text{init}})^2} + \psi \frac{(L_{\text{init}} - l)(2L_{\text{fin}} - L_{\text{init}} - l)}{2d(L_{\text{fin}} - L_{\text{init}})^2} - \psi \frac{(r - L_{\text{fin}})^2}{2d(L_{\text{fin}} - L_{\text{init}})^2} \\ &= \psi \frac{1}{2d(L_{\text{fin}} - L_{\text{init}})^2} [2d(c - L_{\text{fin}}) + 2L_{\text{init}}L_{\text{fin}} - L_{\text{init}}^2 - L_{\text{init}}l - 2L_{\text{fin}}l \\ &\quad + L_{\text{init}}l + l^2 - r^2 + 2L_{\text{fin}}r - L_{\text{fin}}^2] \\ &= \psi \frac{1}{2d(L_{\text{fin}} - L_{\text{init}})^2} [r^2 - l^2 - 2L_{\text{fin}}r + 2L_{\text{fin}}l + 2L_{\text{init}}L_{\text{fin}} - L_{\text{init}}^2 \\ &\quad - 2L_{\text{fin}}l + l^2 - r^2 + 2L_{\text{fin}}r - L_{\text{fin}}^2] \\ &= -\psi \frac{L_{\text{fin}}^2 - 2L_{\text{init}}L_{\text{fin}} + L_{\text{init}}^2}{2d(L_{\text{fin}} - L_{\text{init}})^2} \\ \frac{\partial \phi_L}{\partial L_{\text{init}}} &= -\frac{\psi}{2d}. \end{aligned} \quad (\text{B.11})$$

Similarly, the derivative of  $\phi_L$  with respect to  $L_{\text{fin}}$  is given by (inserting the definitions of  $c$  and  $d$  where required):

$$\begin{aligned}
\frac{\partial \phi_L}{\partial L_{\text{fin}}} &= \frac{\partial \phi_L^1}{\partial L_{\text{fin}}} - \psi \frac{(L_{\text{init}} - l)^2}{2d(L_{\text{fin}} - L_{\text{init}})^2} + \psi \frac{(r - L_{\text{fin}})(L_{\text{fin}} + r - 2L_{\text{init}})}{2d(L_{\text{fin}} - L_{\text{init}})^2} \\
&= \psi \frac{L_{\text{init}} - c}{(L_{\text{fin}} - L_{\text{init}})^2} - \psi \frac{(L_{\text{init}} - l)^2}{2d(L_{\text{fin}} - L_{\text{init}})^2} + \psi \frac{(r - L_{\text{fin}})(L_{\text{fin}} + r - 2L_{\text{init}})}{2d(L_{\text{fin}} - L_{\text{init}})^2} \\
&= \psi \frac{1}{2d(L_{\text{fin}} - L_{\text{init}})^2} \left[ 2d(L_{\text{init}} - c) - L_{\text{init}}^2 + 2L_{\text{init}}l - l^2 + L_{\text{fin}}r + r^2 \right. \\
&\quad \left. - 2L_{\text{init}}r - L_{\text{fin}}^2 - L_{\text{fin}}r + 2L_{\text{init}}L_{\text{fin}} \right] \\
&= \psi \frac{1}{2d(L_{\text{fin}} - L_{\text{init}})^2} \left[ 2L_{\text{init}}r - 2L_{\text{init}}l - r^2 + l^2 - L_{\text{init}}^2 + 2L_{\text{init}}l - l^2 \right. \\
&\quad \left. + r^2 - 2L_{\text{init}}r - L_{\text{fin}}^2 + 2L_{\text{init}}L_{\text{fin}} \right] \\
&= -\psi \frac{L_{\text{fin}}^2 - 2L_{\text{init}}L_{\text{fin}} + L_{\text{init}}^2}{2d(L_{\text{fin}} - L_{\text{init}})^2} \\
\frac{\partial \phi_L}{\partial L_{\text{fin}}} &= -\frac{\psi}{2d}. \tag{B.12}
\end{aligned}$$

Equations (B.11) and (B.12) give the derivatives of  $\phi_L$  when both initial and final left leaf positions are inside the same bixel. In particular, these expressions are true in the limiting case  $L_{\text{init}} \rightarrow L_{\text{fin}}$ , representing the discrete aperture VMAT approximation. Define  $L$  to be the initial and final leaf position in this limit:  $L = L_{\text{init}} = L_{\text{fin}}$ . Then the derivative with respect to  $L$

is given by the chain rule:

$$\begin{aligned}
 \frac{\partial \phi_L}{\partial L} &= \frac{\partial \phi_L}{\partial L_{\text{init}}} \frac{\partial L_{\text{init}}}{\partial L} + \frac{\partial \phi_L}{\partial L_{\text{fin}}} \frac{\partial L_{\text{fin}}}{\partial L} \\
 &= -\frac{\psi}{2d} - \frac{\psi}{2d} \\
 &= -\frac{\psi}{d}.
 \end{aligned} \tag{B.13}$$

This corresponds precisely with the derivative found, for example, by Wild et al. [69], proving that the discrete aperture VMAT approximation can be considered to be a limiting case of the more general continuous aperture VMAT calculation method.

**Case 5:**  $L_{\text{init}} < L_{\text{fin}} < l$

The bixel fluence for case 4 was derived to be (from Equation (A.15)):

$$\phi_L = \psi. \tag{B.14}$$

In this case, the bixel fluence does not depend on either the initial or final left leaf positions, therefore both derivatives are 0:

$$\begin{aligned}
 \frac{\partial \phi_L}{\partial L_{\text{init}}} &= 0, \\
 \frac{\partial \phi_L}{\partial L_{\text{fin}}} &= 0.
 \end{aligned} \tag{B.15}$$

**Case 6:**  $r < L_{\text{init}} < L_{\text{fin}}$ 

The bixel fluence for case 5 was derived to be (from Equation (A.16)):

$$\phi_L = 0 \tag{B.16}$$

In this case, the bixel fluence does not depend on either the initial or final left leaf positions, therefore both derivatives are 0:

$$\begin{aligned} \frac{\partial \phi_L}{\partial L_{\text{init}}} &= 0, \\ \frac{\partial \phi_L}{\partial L_{\text{fin}}} &= 0. \end{aligned} \tag{B.17}$$

## Gradient of $\phi_R$ with respect to the leaf positions

As in the derivation of the fluence in Appendix A, the derivatives of  $\phi_R$  with respect to  $R_{\text{init}}$  and  $R_{\text{fin}}$  can be obtained by performing the following substitutions in the  $\phi_L$  derivatives:  $L_{\text{init}} \rightarrow R_{\text{init}}$ , and  $L_{\text{fin}} \rightarrow R_{\text{fin}}$ .

## Appendix C

### Discrete expression for $\bar{d}_i$ , the mean dose in 4D-VMAT

The continuous expression for the expectation value of the dose derived in Equation (4.6) is re-written here:

$$\bar{d}_i = \iiint \Pr[\Omega] \delta(\omega - \Omega(\Theta^{-1}(\theta))) \tilde{D}_{ij}(\theta, \omega) \frac{d\phi_j}{d\theta}(L_J, R_J) d\theta d\omega \mathcal{D}\Omega. \quad (\text{C.1})$$

The integrals over angle  $\theta$  and tumour position  $\omega$  must be converted into discrete sums for the purposes of dose calculation using traditional dose-influence matrices.

Since  $\tilde{D}_{ij}(\theta, \omega)$  does not depend on  $\Omega(t)$ , the dose-influence matrix can be removed from the integral over  $\Omega$ . Equation (C.1) is then converted into a discrete equation, where  $\tilde{D}_{ij}(\theta, \omega)$  is assumed to be relatively constant over

intervals of length  $\Delta\theta$  and  $\Delta\omega$ :

$$\begin{aligned}\bar{d}_i &= \sum_l \sum_k \tilde{D}_{ij}^{kl} \int_{\omega^l - \Delta\omega/2}^{\omega^l + \Delta\omega/2} \int_{\theta^k - \Delta\theta/2}^{\theta^k + \Delta\theta/2} \int \text{Pr}[\Omega] \delta(\omega - \Omega(\Theta^{-1}(\theta))) \\ &\quad \cdot \frac{d\phi_j}{d\theta}(L_J, R_J) \mathcal{D}\Omega d\theta d\omega \\ &= \sum_l \sum_k \tilde{D}_{ij}^{kl} \phi_j^{kl}.\end{aligned}\tag{C.2}$$

$\phi_j^{kl}$  is the effective discrete fluence for bixel  $j$  emitted in the rectangle in  $(\theta, \omega)$  space surrounding the point  $(\theta^k, \omega^l)$  with area  $\Delta\omega \cdot \Delta\theta$ , while  $\tilde{D}_{ij}^{kl}$  is the dose-influence matrix at the central point:  $\tilde{D}_{ij}^{kl} = \tilde{D}_{ij}(\theta^k, \omega^l)$ . In order to obtain the expectation value of the dose, this effective fluence must be multiplied by the dose-influence matrices  $\tilde{D}_{ij}^{kl}$  and summed over all indices. Calculating the effective fluence will be the focus of the remainder of this appendix.

From the implicit definition in Equation (C.2), the effective fluence is given by:

$$\phi_j^{kl} = \int_{\omega^l} \int_{\theta^k} \int \text{Pr}[\Omega] \delta(\omega - \Omega(\Theta^{-1}(\theta))) \frac{d\phi_j}{d\theta}(L_J, R_J) \mathcal{D}\Omega d\theta d\omega.\tag{C.3}$$

Here, the integral over tumour position (angle) is understood to be centred at  $\omega^l$  ( $\theta^k$ ) with a width of  $\Delta\omega$  ( $\Delta\theta$ ). The two first integrals indicate an integration over the  $(\theta^k, \omega^l)$  rectangle. The third integral is over all trajectories  $\Omega(t)$ , while the Dirac delta function picks out only those respiratory trajectories which cross the rectangle in some way. In the integrand there is the probability

of each trajectory,  $\Pr[\Omega]$ , as well as the differential fluence emitted as the leaves sweep along that trajectory.

Before carrying out the integrations in Equation (C.3), three simplifying assumptions are made to the trajectories  $\omega = \Omega(\Theta^{-1}(\theta))$  in the integral over  $\Omega$ : the codomain of  $\Omega(\Theta^{-1}(\theta))$  is limited to the set of discrete tumour positions  $\{\omega^l\}$ ; in each interval  $[\theta^k - \Delta\theta/2, \theta^k + \Delta\theta/2]$ , the trajectories can only take on two values, an initial and final tumour position; and the transition between these positions must occur at the discrete angle  $\theta^k$  (or at the endpoints of the interval  $[\theta^k - \Delta\theta/2, \theta^k + \Delta\theta/2]$ ). A full trajectory is built piecewise out of trajectory segments  $\Omega^{k,l \rightarrow l'}(\Theta^{-1}(\theta))$  which transition from an initial tumour position  $\omega^l$  to a final position  $\omega^{l'}$ . An example of such a trajectory segment is shown graphically in Figure C.1, and can be defined mathematically as:

$$\Omega^{k,l \rightarrow l'}(\Theta^{-1}(\theta)) = \begin{cases} \omega^l & \text{for } \theta \in [\theta^k - \Delta\theta/2, \theta^k] \\ \omega^{l'} & \text{for } \theta \in [\theta^k, \theta^k + \Delta\theta/2]. \end{cases} \quad (\text{C.4})$$

Now the integration over the  $(\theta^k, \omega^l)$  rectangle in the  $\omega$  direction is performed. Because of the discretization of the integral over  $\Omega$ , only trajectories which have an initial or final tumour position of  $\omega^l$  survive this integration; that is, the trajectory must contain the trajectory segment  $\Omega^{k,l \rightarrow l'}$  or  $\Omega^{k,l' \rightarrow l}$  for any  $l'$ . For trajectories which contain a trajectory segment of the type  $\Omega^{k,l \rightarrow l'}$ , the delta function  $\delta(\omega - \Omega(\Theta^{-1}(\theta)))$  ensures that only the part of the  $\theta$  integral from  $\theta^k - \Delta\theta/2$  to  $\theta^k$  survives. Conversely, for trajectories

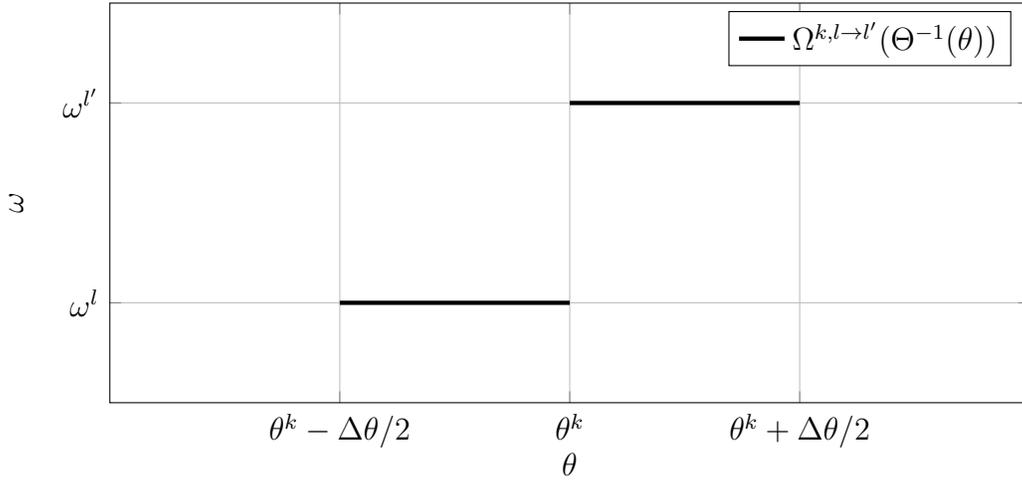


Figure C.1: Allowed trajectory segments  $\omega = \Omega^{k,l \rightarrow l'}(\Theta^{-1}(\theta))$  in the 4D-VMAT dose calculation model.

which feature a trajectory segment of the type  $\Omega^{k,l' \rightarrow l}$ , only the part of the  $\theta$  integral from  $\theta^k$  to  $\theta^k + \Delta\theta/2$  survives. Equation (C.3) then becomes:

$$\phi_j^{kl} = \sum_{l'} \left[ \int_{\overleftarrow{\theta^k}} \int_{\Omega^{k,l \rightarrow l'}} \Pr[\Omega] \frac{d\phi_j}{d\theta}(L_J, R_J) \mathcal{D}\Omega d\theta + \int_{\overrightarrow{\theta^k}} \int_{\Omega^{k,l' \rightarrow l}} \Pr[\Omega] \frac{d\phi_j}{d\theta}(L_J, R_J) \mathcal{D}\Omega d\theta \right]. \quad (\text{C.5})$$

Here two new symbols have been introduced:  $\overleftarrow{\theta^k}$  and  $\overrightarrow{\theta^k}$  represent the intervals  $[\theta^k - \Delta\theta/2, \theta^k]$  and  $[\theta^k, \theta^k + \Delta\theta/2]$ , respectively. The first term concerns trajectory segments which transition from tumour position  $\omega^l$  to  $\omega^{l'}$ , while the second term concerns trajectory segments which transition from position  $\omega^{l'}$  to  $\omega^l$ . Each term is summed over  $l'$ , the index of the final (initial) tumour position. The subscript  $\Omega^{k,l \rightarrow l'}$  in the integral over  $\Omega$  of each term indicates

that the integration is performed only over trajectories which include that particular trajectory segment.

Next the integration over  $\Omega$  is performed:

$$\begin{aligned}
& \int_{\Omega^{k,l \rightarrow l'}} \Pr[\Omega] \frac{d\phi_j}{d\theta}(L_J, R_J) \mathcal{D}\Omega \\
&= \int_{\Omega^{k,l \rightarrow l'}} \Pr[\Omega] \mathcal{D}\Omega \frac{d\phi_j}{d\theta}(L_J^{k,l \rightarrow l'}, R_J^{k,l \rightarrow l'}) \\
&= \Pr(\Omega^{k,l \rightarrow l'}) \frac{d\phi_j}{d\theta}(L_J^{k,l \rightarrow l'}, R_J^{k,l \rightarrow l'}). \tag{C.6}
\end{aligned}$$

The definition of the aperture library in Section 4.2.1.1 indicates that the right and left leaf positions are defined on a set of discrete points in  $(\theta, \omega)$  space which are now identified as  $\{(\theta^k, \omega^l)\}$ . The leaf positions take on interpolated values between these control points depending on the trajectories  $\Theta$  and  $\Omega$  which are followed. Specifically, the leaf positions for any trajectory which contains a trajectory segment of the type  $\Omega^{k,l \rightarrow l'}$  all follow the same function of  $\theta$  in the interval  $[\theta^k - \Delta\theta/2, \theta^k + \Delta\theta/2]$ , which is the exact domain of integration in Equation (C.5). These functions are denoted by the symbols  $L_J^{k,l \rightarrow l'}(\theta)$  and  $R_J^{k,l \rightarrow l'}(\theta)$  for the left and right leaves, respectively. Then the differential fluence may be taken out of the integral, as in the second line of Equation (C.6), so that only the probability of a particular trajectory  $\Omega$  is left in the integral over  $\Omega$ . The result of this integral is given in the third line as the probability of observing the trajectory segment  $\Omega^{k,l \rightarrow l'}$ .

In order to be consistent with the restrictions placed on the possible trajectories (Figure C.1), in the calculation of  $\Pr(\Omega^{k,l \rightarrow l'})$  the trajectory

segment  $\Omega^{k,l \rightarrow l'}$  is considered to be achieved whenever two conditions are fulfilled:  $\omega(\Theta^{-1}(\theta^k - \Delta\theta/2)) = \omega^l$  and  $\omega(\Theta^{-1}(\theta^k + \Delta\theta/2)) = \omega^{l'}$ , thus:

$$\begin{aligned} \Pr(\Omega^{k,l \rightarrow l'}) &= \Pr(\omega(\Theta^{-1}(\theta^k - \Delta\theta/2)) = \omega^l \cap \omega(\Theta^{-1}(\theta^k + \Delta\theta/2)) = \omega^{l'}) \\ &= \Pr(\omega(\Theta^{-1}(\theta^k - \Delta\theta/2)) = \omega^l) \\ &\quad \cdot \Pr(\omega(\Theta^{-1}(\theta^k + \Delta\theta/2)) = \omega^{l'} | \omega(\Theta^{-1}(\theta^k - \Delta\theta/2)) = \omega^l). \end{aligned} \tag{C.7}$$

The second line is due to the definition of conditional probability: the probability of observing the trajectory segment  $\Omega^{k,l \rightarrow l'}$  is given by the probability of being observing tumour position  $\omega^l$  at angle  $\theta^k - \Delta\theta/2$ , multiplied by the conditional or transition probability of observing  $\omega^{l'}$  at angle  $\theta^k + \Delta\theta/2$  given that the tumour position was  $\omega^l$  at angle  $\theta^k - \Delta\theta/2$ . These probabilities are obtained from a model of the tumour motion, such as the one introduced in Chapter 3. It is important to note that the derivation of the dose up until this point has not depended on the particulars of the motion model chosen, only on the existence of such a model that can calculate the required probabilities.

The final expression for the effective fluence for bixel  $j$  emitted in the  $(\theta^k, \omega^l)$  rectangle is given by inserting Equation (C.6) into Equation (C.5),

with the prescription for probabilities given in Equation (C.7):

$$\phi_j^{kl} = \sum_{l'} \left[ \Pr(\Omega^{k,l \rightarrow l'}) \int_{\overleftarrow{\theta}^k} \frac{d\phi_j}{d\theta} (L_J^{k,l \rightarrow l'}, R_J^{k,l \rightarrow l'}) d\theta \right. \\ \left. + \Pr(\Omega^{k,l' \rightarrow l}) \int_{\overrightarrow{\theta}^k} \frac{d\phi_j}{d\theta} (L_J^{k,l' \rightarrow l}, R_J^{k,l' \rightarrow l}) d\theta \right]. \quad (\text{C.8})$$

The remaining integrals over  $\theta$  can be calculated using the continuous aperture fluence method [3], described in Section 2.2.1.1 and Appendix A.

## Appendix D

### Discrete expression for $\overline{d_i^2}$ , the mean squared dose in 4D-VMAT

The continuous expression for the expectation value of the square of the dose derived in Equation (4.12) is re-written here:

$$\begin{aligned} \overline{d_i^2} = & \iiint \iiint \text{Pr}[\Omega] \delta(\omega_1 - \Omega(\Theta^{-1}(\theta_1))) \tilde{D}_{ij_1}(\theta_1, \omega_1) \frac{d\phi_{j_1}}{d\theta_1} \\ & \cdot \delta(\omega_2 - \Omega(\Theta^{-1}(\theta_2))) \tilde{D}_{ij_2}(\theta_2, \omega_2) \frac{d\phi_{j_2}}{d\theta_2} d\theta_1 d\omega_1 d\theta_2 d\omega_2 \mathcal{D}\Omega. \end{aligned} \quad (\text{D.1})$$

Equation (D.1) is converted to a discrete equation in a similar way going from Equation (C.1) to Equation (C.2), thus:

$$\overline{d_i^2} = \sum_{l_1} \sum_{k_1} \sum_{l_2} \sum_{k_2} \tilde{D}_{ij_1}^{k_1 l_1} \tilde{D}_{ij_2}^{k_2 l_2} \Phi_{j_1 j_2}^{k_1 k_2 l_1 l_2}. \quad (\text{D.2})$$

Here,  $\Phi_{j_1 j_2}^{k_1 k_2 l_1 l_2}$  is an effective fluence tensor. To obtain the expectation value of the dose squared, two multiplications of this tensor with the dose-influence matrix are necessary.

This fluence tensor is given by:

$$\begin{aligned} \Phi_{j_1 j_2}^{k_1 k_2 l_1 l_2} = & \int_{\omega^{l_1}} \int_{\theta^{k_1}} \int_{\omega^{l_2}} \int_{\theta^{k_2}} \int \text{Pr} [\Omega] \delta (\omega_1 - \Omega (\Theta^{-1} (\theta_1))) \frac{d\phi_{j_1}}{d\theta_1} \\ & \cdot \delta (\omega_2 - \Omega (\Theta^{-1} (\theta_2))) \frac{d\phi_{j_2}}{d\theta_2} \mathcal{D}\Omega d\theta_2 d\omega_2 d\theta_1 d\omega_1. \quad (\text{D.3}) \end{aligned}$$

The notation from Appendix C is used here: each integral over tumour position (angle) is understood to be centred at  $\omega^{l_{1,2}}$  ( $\theta^{k_{1,2}}$ ) with a width of  $\Delta\omega$  ( $\Delta\theta$ ). The first four integrals indicate an integration over the two rectangles  $(\theta^{k_1}, \omega^{l_1})$  and  $(\theta^{k_2}, \omega^{l_2})$ . As in Equation (C.3), the fifth integral is over all trajectories  $\Omega$ , while the Dirac delta functions pick out only those trajectories which cross both rectangles in some way.

These delta functions can be dealt with one at a time, similarly to the progression from Equation (C.3) to Equation (C.5). First, the delta function

in  $\omega_1$  is eliminated by performing the integration over  $\omega_1$ :

$$\begin{aligned} \Phi_{j_1 j_2}^{k_1 k_2 l_1 l_2} = & \sum_{l'_1} \left[ \int_{\overleftarrow{\theta}^{k_1}} \int_{\omega^{l_2}} \int_{\theta^{k_2}} \int_{\Omega^{k_1, l_1 \rightarrow l'_1}} \Pr [\Omega] \frac{d\phi_{j_1}}{d\theta_1} \delta (\omega_2 - \Omega (\Theta^{-1} (\theta_2))) \frac{d\phi_{j_2}}{d\theta_2} \mathcal{D}\Omega d\theta_2 d\omega_2 d\theta_1 \right. \\ & \left. + \int_{\overrightarrow{\theta}^{k_1}} \int_{\omega^{l_2}} \int_{\theta^{k_2}} \int_{\Omega^{k_1, l'_1 \rightarrow l_1}} \Pr [\Omega] \frac{d\phi_{j_1}}{d\theta_1} \delta (\omega_2 - \Omega (\Theta^{-1} (\theta_2))) \frac{d\phi_{j_2}}{d\theta_2} \mathcal{D}\Omega d\theta_2 d\omega_2 d\theta_1 \right]. \end{aligned} \quad (\text{D.4})$$

The symbols  $\overleftarrow{\theta}^k$  and  $\overrightarrow{\theta}^k$  introduced in Appendix C (just below Equation (C.5)) are used here.

Next the integrations over  $\omega_2$  may be performed, yielding a sum of four terms:

$$\begin{aligned} \Phi_{j_1 j_2}^{k_1 k_2 l_1 l_2} = & \sum_{l'_1} \sum_{l'_2} \left[ \int_{\overleftarrow{\theta}^{k_1}} \int_{\overleftarrow{\theta}^{k_2}} \int_{\Omega^{k_1, l_1 \rightarrow l'_1} \cap \Omega^{k_2, l_2 \rightarrow l'_2}} \Pr [\Omega] \frac{d\phi_{j_1}}{d\theta_1} \frac{d\phi_{j_2}}{d\theta_2} \mathcal{D}\Omega d\theta_2 d\theta_1 \right. \\ & + \int_{\overleftarrow{\theta}^{k_1}} \int_{\overrightarrow{\theta}^{k_2}} \int_{\Omega^{k_1, l_1 \rightarrow l'_1} \cap \Omega^{k_2, l'_2 \rightarrow l_2}} + \int_{\overrightarrow{\theta}^{k_1}} \int_{\overleftarrow{\theta}^{k_2}} \int_{\Omega^{k_1, l'_1 \rightarrow l_1} \cap \Omega^{k_2, l_2 \rightarrow l'_2}} \\ & \left. + \int_{\overrightarrow{\theta}^{k_1}} \int_{\overrightarrow{\theta}^{k_2}} \int_{\Omega^{k_1, l'_1 \rightarrow l_1} \cap \Omega^{k_2, l'_2 \rightarrow l_2}} \right]. \end{aligned} \quad (\text{D.5})$$

The four terms have the same integrands but different domains of integration, therefore the integrands have been suppressed in the latter three terms. The domain of integration in the first  $\Omega$ -integral is the set of all trajectories which contain both trajectory segments  $\Omega^{k_1, l_1 \rightarrow l'_1}$  and  $\Omega^{k_2, l_2 \rightarrow l'_2}$ , and similarly for the other three.

The integration over  $\Omega$  is performed as in Equation (C.6):

$$\begin{aligned}
& \int_{\Omega^{k_1, l_1 \rightarrow l'_1} \cap \Omega^{k_2, l_2 \rightarrow l'_2}} \Pr[\Omega] \frac{d\phi_{j_1}}{d\theta_1}(L_{J_1}, R_{J_1}) \frac{d\phi_{j_2}}{d\theta_2}(L_{J_2}, R_{J_2}) \mathcal{D}\Omega \\
&= \int_{\Omega^{k_1, l_1 \rightarrow l'_1} \cap \Omega^{k_2, l_2 \rightarrow l'_2}} \Pr[\Omega] \mathcal{D}\Omega \frac{d\phi_{j_1}}{d\theta}(L_{J_1}, R_{J_1}) \frac{d\phi_{j_2}}{d\theta}(L_{J_2}, R_{J_2}) \\
&= \Pr\left(\Omega^{k_1, l_1 \rightarrow l'_1} \cap \Omega^{k_2, l_2 \rightarrow l'_2}\right) \frac{d\phi_{j_1}}{d\theta}\left(L_{J_1}^{k_1, l_1 \rightarrow l'_1}, R_{J_1}^{k_1, l_1 \rightarrow l'_1}\right) \\
&\quad \cdot \frac{d\phi_{j_2}}{d\theta}\left(L_{J_2}^{k_2, l_2 \rightarrow l'_2}, R_{J_2}^{k_2, l_2 \rightarrow l'_2}\right). \tag{D.6}
\end{aligned}$$

Here, the dependence of the differential fluence on the left and right leaf positions has been made explicit.

Next the probability to observe both trajectory segments  $\Omega^{k_1, l_1 \rightarrow l'_1}$  and  $\Omega^{k_2, l_2 \rightarrow l'_2}$  must be calculated. Two cases are considered:  $\theta^{k_1}$  is equal to  $\theta^{k_2}$ ; and  $\theta^{k_1}$  is achieved by  $\theta = \Theta(t)$  before  $\theta^{k_2}$  (the reverse can be calculated by swapping the indices). The first case can be calculated with Equation (C.7) by using the simplifying relation  $\Omega^{k_1, l_1 \rightarrow l'_1} \cap \Omega^{k_2, l_2 \rightarrow l'_2} = \Omega^{k_1, l_1 \rightarrow l'_1}$ . In the second case, successive applications of the definition of conditional probability are

performed:

$$\begin{aligned}
\Pr\left(\Omega^{k_1, l_1 \rightarrow l'_1} \cap \Omega^{k_2, l_2 \rightarrow l'_2}\right) &= \Pr\left(\omega\left(\Theta^{-1}\left(\theta^{k_1} - \Delta\theta/2\right)\right) = \omega^{l_1}\right. \\
&\quad \cap \omega\left(\Theta^{-1}\left(\theta^{k_1} + \Delta\theta/2\right)\right) = \omega^{l'_1} \cap \omega\left(\Theta^{-1}\left(\theta^{k_2} - \Delta\theta/2\right)\right) = \omega^{l_2} \\
&\quad \left. \cap \omega\left(\Theta^{-1}\left(\theta^{k_2} + \Delta\theta/2\right)\right) = \omega^{l'_2}\right) \\
&= \Pr\left(\omega\left(\Theta^{-1}\left(\theta^{k_1} - \Delta\theta/2\right)\right) = \omega^{l_1}\right) \\
&\quad \cdot \Pr\left(\omega\left(\Theta^{-1}\left(\theta^{k_1} + \Delta\theta/2\right)\right) = \omega^{l'_1} \mid \omega\left(\Theta^{-1}\left(\theta^{k_1} - \Delta\theta/2\right)\right) = \omega^{l_1}\right) \\
&\quad \cdot \Pr\left(\omega\left(\Theta^{-1}\left(\theta^{k_2} - \Delta\theta/2\right)\right) = \omega^{l_2} \mid \omega\left(\Theta^{-1}\left(\theta^{k_1} + \Delta\theta/2\right)\right) = \omega^{l'_1}\right) \\
&\quad \cdot \Pr\left(\omega\left(\Theta^{-1}\left(\theta^{k_2} + \Delta\theta/2\right)\right) = \omega^{l'_2} \mid \omega\left(\Theta^{-1}\left(\theta^{k_2} - \Delta\theta/2\right)\right) = \omega^{l_2}\right).
\end{aligned} \tag{D.7}$$

The transition probabilities shown above can be calculated using a suitable model of the tumour motion.

The final expression for the effective tensor is given by inserting Equation (D.6) into Equation (D.5), with the formula for probabilities given in Equations (C.7) and (D.7):

$$\begin{aligned}
\Phi_{j_1 j_2}^{k_1 k_2 l_1 l_2} &= \sum_{l'_1} \sum_{l'_2} \left[ \Pr\left(\Omega^{k_1, l_1 \rightarrow l'_1} \cap \Omega^{k_2, l_2 \rightarrow l'_2}\right) \cdot \int_{\theta^{k_1} \leftarrow} \frac{d\phi_{j_1}}{d\theta_1} d\theta_1 \int_{\theta^{k_2} \leftarrow} \frac{d\phi_{j_2}}{d\theta_2} d\theta_2 \right. \\
&\quad + \Pr\left(\Omega^{k_1, l_1 \rightarrow l'_1} \cap \Omega^{k_2, l_2 \rightarrow l_2}\right) \cdot \int_{\theta^{k_1} \leftarrow} \int_{\theta^{k_2} \rightarrow} + \Pr\left(\Omega^{k_1, l'_1 \rightarrow l_1} \cap \Omega^{k_2, l_2 \rightarrow l'_2}\right) \cdot \int_{\theta^{k_1} \rightarrow} \int_{\theta^{k_2} \leftarrow} \\
&\quad \left. + \Pr\left(\Omega^{k_1, l'_1 \rightarrow l_1} \cap \Omega^{k_2, l'_2 \rightarrow l_2}\right) \cdot \int_{\theta^{k_1} \rightarrow} \int_{\theta^{k_2} \rightarrow} \right]. \tag{D.8}
\end{aligned}$$

The integrals over  $\theta_1$  and  $\theta_2$  can be separated at this point, since the two

factors in the integrand  $\frac{d\phi_{j1}}{d\theta_1}$  and  $\frac{d\phi_{j1}}{d\theta_1}$  depend only on their respective angular variable. These integrals can be calculated using the continuous aperture fluence method [3], described in Section 2.2.1.1 and Appendix A.