Mixed electron-photon radiation therapy treatment planning and delivery

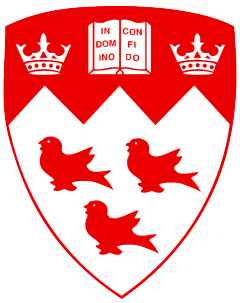

Marc-André Renaud

Department of Physics
McGill University

A thesis submitted to McGill University in partial fulfillment of the requirements of the degree of
Doctor of Philosophy

 January 2019

Acknowledgements

Firstly, I would like to thank my supervisor, Dr. Jan Seuntjens for directing and supporting my work while allowing me considerable freedom to pursue the research opportunities that presented themselves to me over the course of my PhD. Without your support, I would not have found the motivation to perform the amount of work necessary for the success of this ambitious project.

I would also like to acknowledge the essential clinical insights provided by Monica Serban, which served to ensure that the work in this thesis is relevant to the broader Medical Physics community. Your presence as part of this project helped me to never lose sight of the ultimate goal: improving the lives of cancer patients.

I am fortunate to have crossed paths with Stavros Korokithakis, who taught me to take pride in my work and to design it for reuse beyond its immediate purpose in order to benefit as many people as possible.

I benefited greatly from doing my PhD as part of the Medical Physics Unit at McGill University, which allowed me to be surrounded by so many smart and motivated staff members and fellow students. I would like to extend my sincere gratitude to Margery Knewstubb & Tatjana Nisic for their administrative support. I am grateful to Joe Larkin and Bhavan Siva for contributing tools needed to perform the experiments in this work. Dr. François DeBlois, Dr. Andrew Alexander, Dr. Pavlos Papaconstadopoulos, Russell Ruo, Dr. Tanner Connell, Dr. Piotr Pater, Dr. James Renaud and Dr. Susannah Hickling all contributed valuable knowledge relating to measurement procedures on linear accelerators, for which I am infinitely grateful.

The work of translating research into clinical applications is a multidisciplinary task which could not be achieved without the support of physicians. I would like to acknowledge the contributions that Drs. Carolyn Freeman and Christine Lambert have made to this project by helping identify the patient cases and cancer sites which could benefit the most from our research.

I would like to express my gratitude to the National Research Council Canada (NRC) for the effort they have put over the years into developing and maintaining the EGSnrc codebase, one of the widely acknowledged *gold standard* Monte Carlo particle transport libraries.

My years spent as a PhD student were made much more enjoyable due to ongoing friendships and enlightening discussions with fellow colleagues Andrew, Pavlos, Kyu, Eunah, Tanner, James, Pete, Ali, Martin, Joel, Yana, Michael, Jessica, Susannah, Andre, Logan, Simon, Mikaël, Gabriel, Véronique, Veng.

Final thanks to my parents Nicole and Paul, and my sister Julie for their unwavering love and support, and last but not least, thanks to Judy for sharing this period of my life with me and lending your support at every step of the way.

Abstract

In linac-based radiation therapy, the bulk of recent innovations have involved delivering photon radiation over an ever-growing range of beam delivery angles around the patient, rather than incorporating more particle types into the treatment planning process.

Mixed beam radiation therapy (MBRT) is the concept of creating and delivering treatment plans involving multiple modalities, such as multiple photon and electron energies. Combined electron-photon treatment plans can reduce the dose to normal tissue compared to photon-only treatments for patients with a superficial tumour due to the limited range of electrons in matter. However, MBRT has yet to be implemented widely in a clinical environment due to primitive optimisation models and the cumbersome collimation devices typically used in electron delivery. In this work, the components necessary for the creation and validation of mixed electron-photon radiation therapy plans delivered with a photon multi-leaf collimator (pMLC) as the sole collimation device were developed, characterised and validated.

As neither MBRT nor pMLC-collimated electron deliveries are part of routine clinical practice, it was necessary to first create an academic treatment planning system (TPS) capable of visualising patient images and dose distributions and coordinating the creation of the beamlet dose distributions necessary for treatment plan optimisation. A collapsed cone convolution superposition dose calculation engine was implemented to generate photon beamlet dose distributions. Electron beamlet dose distributions were created using the DOSXYZnrc Monte Carlo dose calculation engine coupled to a validated electron beam model. The TPS was programmed to interface with the two dose calculation engines in order to facilitate the beamlet generation process.

Secondly, a column generation-based (CG) direct aperture optimisation model for MBRT planning was derived and subsequently implemented. Using beamlets created for each modality, the algorithm iteratively creates a treatment plan by producing candidate electron and photon apertures at each iteration, and selecting the apertures most likely to improve the cost function.

A study of the algorithm showed that it could successfully produce MBRT plans with a better cost function compared to photon-only and electron-only plans.

Next, the robustness of MBRT dose distributions to patient positioning uncertainties was studied. The CG-based MBRT model was extended to explicitly ensure that clinical target volume (CTV) coverage is preserved under a series of setup error *scenarios*. The study showed that traditional PTV-based optimisation is inadequate for MBRT planning, whereas robust CTV-based optimisation was able to produce plans which maintained adequate coverage under all setup error scenarios considered.

Finally, two MBRT plans were delivered on a solid water slab phantom and an acrylic cylinder, respectively. The accuracy of the deliveries was assessed using film dosimetry and by measuring the dose on the central axis of the acrylic cylinder using an ionisation chamber. The solid water slab film agreed with the simulated dose distribution with a 96.4 % pass rate for a 2 %/2 mm gamma criteria. On the cylindrical acrylic phantom, the measured dose for the combined MBRT delivery agreed within 0.2 %. Broken down by modality, the agreement between measured and simulated dose in the ionisation chamber varied between -0.8 % and +2.7 %. We conclude that our beam models can accurately simulate MBRT dose distributions, however, further investigation is needed to confirm that the deliveries are sufficiently robust to positioning uncertainties.

This body of work demonstrates that a mixed electron-photon treatment planning algorithm based on the column generation method can produce realistic, deliverable and verifiable treatment plans which offer superior normal tissue sparing compared to current state-of-the-art single-modality photon or electron plans, without sacrificing target coverage. At the outset of this project, no published algorithm existed which could manage the added computational burden of simultaneously optimising both electrons and photons when the ever-increasing degrees of freedom involved in modern deliveries were considered. This work paves the way towards the implementation of MBRT techniques as part of routine clinical practice for any patient with a superficial tumour.

Résumé

En radiothérapie à base d'accélérateur linéaire (linac), la majorité des innovations récentes portent sur l'incorporation de plus en plus d'angles d'émission de rayonnement photonique, au lieu de l'ajout de différents types de particules dans le processus de planification de traitement.

La radiothérapie à faisceau mixte (MBRT) constitue la création et l'administration de plans de traitement impliquant plusieurs modalités, telles que les multiples énergies de photons et d'électrons produites par un linac. La combinaison d'électrons et de photons dans un seul plan de traitement peut réduire la dose administrée aux tissus normaux pour tous les patients ayant une tumeur superficielle grâce au parcours limité des faisceaux d'électrons dans la matière. Par contre, l'adoption générale de la technique du MBRT n'a pas encore été realisée en clinique et demeure largement un sujet de recherche en raison des collimateurs emcombrants utilisés lors de l'administration de rayonnement à base d'électrons ainsi que d'un manque de sophistication des modèles d'optimisation. Dans ces travaux de thèse, les composantes nécessaires à la création et la validation de plans de traitement à faisceau mixte administrés uniquement avec un collimateur multilame pour photons (pMLC) ont été construites et characterisées.

Étant donné que ni la technique MBRT ni l'administration d'électrons à l'aide d'un pMLC ne fait partie de la pratique routinière clinique, un système de planification de traitement (TPS) a dû être créé pour la visualisation de données de patients et de distributions de doses ainsi que pour la coordination du processus de création de mini-faisceaux nécessaires pour l'optimisation de plans de traitement. Un moteur de calcul de dose *collapsed cone* a été implémenté pour la création de mini-faisceaux de photons. Pour la création de mini-faisceaux d'électrons, un modèle de faisceau Monte Carlo a été couplé au moteur de calcul de dose DOSXYZnrc. Le TPS a été programmé pour communiquer avec ces deux moteurs de calcul de dose pour faciliter le processus de génération de mini-faisceaux.

En second lieu, un modèle d'optimisation *direct aperture* (DAO) basé sur la méthode de génération de colonne a été dérivé pour résoudre le problème de planification de traitement

MBRT. En utilisant les mini-faisceaux créés pour chaque modalité, l'algorithme obtient un plan de traitement de façon itérative en proposant des ouvertures de champs pour chaque modalité et en selectionnant les champs les plus suspectibles d'améliorer la fonction de coût pour ajouter au plan de traitement à chaque itération. Une étude du modèle d'optimisation a montré que cette méthode pouvait produire des plans MBRT avec une fonction coût améliorée comparativement aux plans à modalité unique.

Ensuite, la robustesse des distributions de doses MBRT aux incertitudes de positionnement des patients a été étudié. Le modèle d'optimisation MBRT a été étendu pour s'assurer que le volume-cible clinique (CTV) demeure couvert de façon adéquate dans une série de *scénarios* d'erreurs de positionnement. Cette étude a démontré que l'optimisation traditionnelle basé sur un volume-cible de planification (PTV) est inadéquate pour la technique MBRT. Par contre, l'optimisation robuste basé sur le CTV a été en mesure de produire des plans de traitement qui maintiennent une couverture adéquate du CTV pour tous les scénarios d'erreurs considérés.

Finalement, deux plans MBRT ont été administrés sur un fantôme à base de *Solid Water* et un cylindre d'acrylique, respectivement. L'exactitude de l'administration du rayonnement a été évaluée à l'aide de la dosimétrie à base de film ainsi qu'en mesurant la dose sur l'axe centrale du cylindre avec une chambre d'ionisation. La dose mesurée par le film était en accord avec la dose prédite par simulations, avec un taux de passage de 96.4 % pour un critère γ de 2 %/2 mm. Pour le fantôme cylindrique, la différence entre la dose mesurée par la chambre d'ionisation pour tous les modalités combinées et la dose prédite était de 0.2 %. Lorsque chaque modalité est évaluée individuellement, la différence entre dose mesurée et simulée varie entre -0.8 % et 2.7 %. Nous concluons que nos modèles de calcul de dose peuvent prédire la distribution de dose des plans MBRT de façon précise. Par contre, d'autres études seront nécessaires pour confirmer que l'administration de plans MBRT est suffisamment robuste aux incertitudes de positionnement.

Ce travail de thèse démontre que notre modèle d'optimisation MBRT peut produire des plans de traitement réalistes, administrables et vérifiables, offrant la possibilité d'épargner davantage les tissus normaux comparativement aux plans à modalité unique sans compromettre la couverture de la cible. À l'origine de ce travail, il n'existait pas d'algorithmes publiés pouvant gérer la charge de calcul associée à l'optimisation simultanée des électrons et des photons, lorsque le nombre élevé de degrés de liberté associés aux traitements moderne sont pris en compte. En conclusion, ce travail ouvre la voie à l'intégration de la technique MBRT dans la routine clinique pour tout patient ayant une tumeur superficielle.

Contents

List of Figures

List of Tables

List of Abbreviations

APBI	Accelerated Partial Breast Irradiation
BT	Brachytherapy
CCCS	Collapsed Cone Convolution Superposition
CG	Column Generation
CH	Condensed History
CSDA	Continuous Slowing Down Approximation
CT	Computed Tomography
CTV	Clinical Target Volume
DAO	Direct Aperture Optimisation
DNA	Deoxyribonucleic acid
DVH	Dose Volume Histogram
EAT	Electron Arc Therapy
EBRT	External Beam Radiation Therapy
EDK	Energy Deposition Kernel
EUD	Equivalent Uniform Dose
FLEC	Few Leaf Electron Collimator
FMO	Fluence Map Optimisation

FWHM	Full Width at Half Maximum
GPU	Graphical Processing Unit
GTV	Gross Tumour Volume
GUI	Graphical User Interface
HU	Hounsfield Unit
IMPT	Intensity Modulated Proton Therapy
IMRT	Intensity Modulated Radiation Therapy
LP	Linear Program / Linear Programming
MBRT	Mixed Beam Radiation Therapy
MC	Monte Carlo
MERT	Modulated Electron Radiation Therapy
MLC	Multi-Leaf Collimator
MP	Master Problem
MRI	Magnetic Resonance Imaging
MU	Monitor Unit
NLP	Nonlinear Program / Nonlinear Programming
NRC	National Research Council
NTCP	Normal Tissue Complication Probability
OAR	Organ At Risk
PDD	Percent Depth Dose
PET	Positron Emission Tomography
PMMA	Polymethyl methacrylate

PTV	Planning Target Volume
PP	Pricing Problem
RMP	Restricted Master Problem
RT	Radiation Therapy
SA	Simulated Annealing
SAD	Source-Axis Distance
SBRT	Stereotactic Body Radiation Therapy
SPR	Stopping Power Ratio
SSD	Source to Surface Distance
TCP	Tumour Control Probability
TPS	Treatment Planning System
VMAT	Volumetric Modulated Arc Therapy

Preface

Contribution to original knowledge

This thesis contains three manuscripts: two published and one in preparation. All manuscripts represent original contributions to research. More specifically, the work presented in this thesis is the first application of the column generation method to the problem of mixed electron-photon radiation therapy (MBRT) treatment planning. Prior to this work, the literature on MBRT planning focused on simplified deliveries which produced inferior plans to state-of-the-art photon-only planning techniques. In addition, the second manuscript constitutes the first study of the robustness of mixed electron-photon deliveries to patient positioning errors, and the first application of robust optimisation to MBRT. Finally, the third manuscript establishes a quality assurance procedure for the delivery of MBRT plans which will assist in introducing MBRT as part of routine clinical practice.

Contribution of Authors

Chapters 1-5 provide introductory and contextual information for the work performed in this thesis, and Chapter 9 summarises the findings. I wrote all the chapters, and they were proofread by Judy Wang and Jan Seuntjens. The three manuscripts featured in this thesis are presented in Chapters 6-8. The contribution of authors was as follows:

1. *Marc-André Renaud, Monica Serban and Jan Seuntjens, "On mixed electron-photon radiation therapy optimization using the column generation approach", Medical Physics, vol. 44, no.8, pp. 4287-4298, 2017.* (Chapter 6)
   I designed the study, researched and implemented the optimisation model, developed the beamlet creation and optimisation workflow, analysed the optimisation results, and wrote the manuscript. Monica Serban selected the patient cases for the study, produced contoured structures for each patient case, analysed and fine-tuned the treatment planning

objectives for each patient, provided expert analysis on the quality of the resulting treatment plans compared to clinical plans and reviewed the manuscript. Jan Seuntjens contributed to study design and framing of the work within the broader field of MERT and MBRT planning, assisted with the analysis of the aperture addition scheme data and interpretation of the individual modality contributions to the total dose distribution of MBRT plans, and reviewed the manuscript.

2. *Marc-André Renaud, Monica Serban and Jan Seuntjens, "Robust mixed electron-photon radiation therapy optimisation", Medical Physics, vol. 46, no.3, pp. 1384-1396, 2019.* (Chapter 7)
   I researched and implemented the robust extension of the MBRT planning algorithm. I also established and implemented the workflow for robust optimisation beamlet creation, generated and analysed the results, and wrote the manuscript. Monica Serban proposed to study robust optimisation, formulated the planning objectives for each patient in the study, produced contoured structures for each patient, and iteratively arrived at the final optimisation constraints. She also provided expert analysis and interpretation of the results and assisted in writing the manuscript. Jan Seuntjens contributed to the study design, interpretation of the results and reviewed the manuscript.

3. *Marc-André Renaud, Monica Serban and Jan Seuntjens, "Delivery and quality assurance of mixed electron-photon radiation therapy plans", in preparation for submission to Medical Physics, 2019.* (Chapter 8)
   I wrote the software to recalculate aperture dose distributions using a full MC beam model and obtain monitor units for delivery. I designed and assisted in the creation of the cylindrical phantom. I modified the SPRRZnrc MC code to create SPRXYZnrc, the tool used to calculate beam quality conversion factors. I took part in all measurements. Finally, I wrote the manuscript. Monica Serban assisted in the design and simulation of the cylindrical phantom, produced treatment plans for delivery, designed the positioning protocol for measurements and was present during most measurements. Jan Seuntjens oversaw the study, assisted with the measurements, the interpretation of ion chamber results, provided theoretical insights on the beam conversion factors and reviewed the manuscript.

Chapter 1

Introduction

1.1 Cancer treatments

Cancer is a broadly defined and complex disease. All cancers develop through a similar process of abnormal cell growth which eventually impedes the normal functioning of bodily functions. However, they vary greatly in cause, risk factors and outcomes. According to Statistics Canada, half of all Canadians will develop cancer in their lifetime, and it is the leading cause of death at approximately 30% of all deaths per year [1].

According to data obtained in 2008, 60% of all people diagnosed with cancer will survive beyond 5 years post-diagnosis, compared to 53% in 1992 [1]. Improvements in all methods utilised in the treatment of cancer contributed to the increase in survival rate. In addition, more frequent screening procedures resulted in earlier detection of cancer which is typically associated with higher survival. The primary options for the treatment of cancer are: chemotherapy, hormonal therapy, surgery and radiation therapy. Cancer patients receive a subset of those options based on factors such as cancer site, affected organs, overall health of the patient, etc.

In broad terms, chemotherapy is the use of drugs to kill rapidly dividing cancer cells while inflicting less damage to healthy cells. Hormone therapy aims to slow the growth of cancer by preventing certain types of cancer cells from binding to hormones such as oestrogen. Surgery is used to remove tumours and surrounding cancerous tissue during an operation. Finally, radiation therapy (RT) is the delivery of ionising radiation targeted towards cancer cells to damage their DNA and prevent them from dividing.

1.2 Radiation therapy

Radiation therapy is beneficial to an estimated 50% of cancer patients, making it one of the major cancer treatment techniques available [2]. It can be administered alone or in combination with other treatment techniques such as surgery and chemotherapy. The purpose of RT is to destroy cancer cells to prevent them from replicating uncontrollably. To first order, both healthy and cancer cells are susceptible to the damage caused by RT, meaning that all radiation delivery systems must grapple with the issue of maximising radiation dose to cancerous tissue while minimising dose to healthy tissue.

Radiation can be classified as either directly or indirectly energy depositing. Directly energy depositing radiation is made up of charged particles such as electrons and protons which deposit energy in matter through Coulombic interactions. This energy deposition is responsible for causing DNA damage such as single and double strand breaks which ultimately can lead to cell death. Uncharged particles such as photons and neutrons make up the indirectly energy depositing radiation category as these particles must first interact and release secondary charged particles which will then deposit energy in matter through Coulomb interactions. Radiation dose is defined as the energy deposited per unit mass. The SI unit is called the *Gray* (abbreviated: Gy), and is equivalent to 1 J / kg (i.e., 1 J kg^{-1} = 1 Gy).

The two broadest types of RT are brachytherapy (BT) and external beam radiation therapy (EBRT). In BT, radioactive sources are introduced invasively in close proximity to cancer tissue to deliver radiation. EBRT involves the production of a radiation beam externally to the patient. The beam is then aimed towards the tumour inside the patient, irradiating a portion of normal tissue in the process. For that reason, EBRT is typically delivered from multiple angles to prevent a specific region of healthy tissue from receiving excessive radiation. Despite delivering more radiation dose to healthy tissue than BT, EBRT can be delivered without invasive procedures and with less personnel. Therefore, the majority of RT treatments are delivered using EBRT rather than BT.

EBRT beams have been produced with photons, neutrons, electrons, protons, and heavier ions. The majority of EBRT treatments are delivered using a medical linear accelerator (linac) which can produce both photon and electron beams. Proton and heavy ion beams must be produced by a cyclotron or synchrotron in order to reach the high energies required to reach tumours inside the body, making them more costly and logistically difficult to implement clinically.

1.3 Treatment planning in external beam radiation therapy

Due to the large variations in tumour sizes and characteristics among all cancer sites, every patient receiving radiation therapy requires a treatment plan customised to their anatomy.

As mentioned above, the primary goal of radiation therapy is to deliver a maximum amount of radiation to cancerous tissue while limiting the dose to critical organs. The response of tumours and critical organs to radiation can be viewed in terms of the tumour control probability (TCP) and normal tissue complication probability (NTCP) as a function of absorbed dose. The aim of treatment planning can be expressed as a maximisation problem of the *therapeutic ratio*, defined as the ratio of TCP / NTCP for a specific TCP (e.g 90%) or NTCP (e.g 5%).

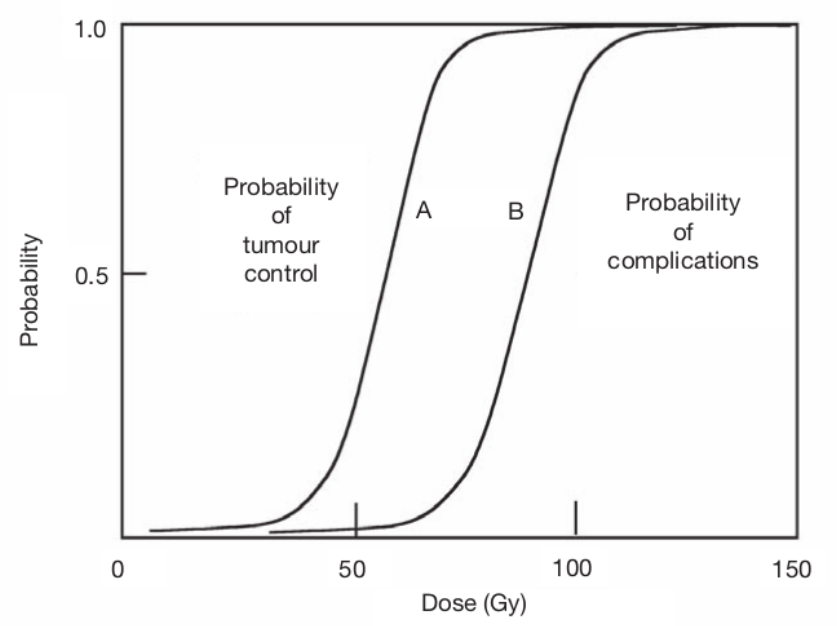

Figure 1.1 Ideal TCP (A) and NTCP (B) curves demonstrating a situation where both high probability of tumour control and low probability of normal tissue complication can be achieved. Reproduced from Podgorsak [3].

Figure 1.1 shows ideal TCP and NTCP dose response curves where both a high probability of tumour control and low probability of normal tissue complication are simultaneously achievable. In practice the two goals are not always achievable. To improve the therapeutic ratio, multiple beams or arcs are used to optimise radiation dose to the planning target volume while keeping the dose to healthy organs as low as possible. In addition, in conventional EBRT, radiation is often delivered in a fractionated schedule, where a fraction of the total prescription dose is

delivered every day for multiple days or weeks primarily to allow time for sublethal damage repair in healthy tissue between fractions [3].

Before a treatment plan can be created, both the tumour and organs at risk must be identified. The workflow for treatment planning is dictated by the inputs required for plan creation, which will briefly be explained.

1.3.1 Computed tomography simulation scan

Prior to treatment, a computed tomography (CT) simulation scan of the patient is typically acquired to provide a full 3D anatomical representation of the patient. The patient is set up in the exact position required for treatment, and all accessories such as patient immobilising devices and bolus are present. Radio-opaque fiducial markers are also placed on the patient to define a coordinate system which will be used for positioning during treatment delivery. The CT dataset not only provides anatomical information, but also electron density and mass density information which can be crucial in the dose calculation process described below.

1.3.2 Volume definition

Volume definition is the contouring of regions of interest on the CT images. At this step, images from other imaging modalities such as magnetic resonance imaging (MRI) or positron emission tomography (PET) may be registered to the CT dataset to assist in properly visualising certain structures or subvolumes of interest or suspected disease sites in the patient. Critical structures (organs or substructures) are identified by trained medical professionals and contoured on every CT slice in which they are present.

The principal volumes related to 3D treatment planning are the gross tumour volume (GTV), the clinical target volume (CTV), and the planning target volume (PTV). [4–7] The GTV represents all visible or clinically demonstrable tumour regions on the imaging dataset. To form the CTV, a margin is then added to the GTV based on clinical knowledge of the presence of microscopic disease extending beyond the visible parts of the tumour. Finally, the PTV is determined by adding an additional margin to the CTV representing the extent of set-up uncertainties, machine tolerances and expected tumour motion during treatment. [3–7] Typical values for the PTV margin are in the 3-10 mm range. In addition to the tumour volumes, sensitive organs at risk (OAR) are also contoured to allow planners to accurately determine

how much radiation dose is being delivered to those organs and limit the dose below prescribed thresholds.

1.3.3 Treatment planning

In EBRT, a treatment plan is defined by a series of *control points* which represent a snapshot of the state of the linac during treatment. In practice, a control point defines the gantry angle, treatment couch angle, the position of all the collimating devices involved in treatment, and the amount of radiation delivered in a given configuration. There are two broad approaches to treatment planning: forward and inverse planning.

Forward planning

Under the forward planning approach, a planner manually selects beam angles, energies and their collimation to meet the prescribed dose constraints in the patient. This approach is typically reserved for simple cases where few beams are necessary as it rapidly becomes unwieldy for a planner to manually tune each beam.

Inverse planning

In the inverse planning approach, the role of the planner is to specify dose volume constraints on contoured structures. Dose volume constraints take the form "at most Y% of the volume of a contoured structure can receive more/less than X Gy". Depending on the specific algorithm used, the planner may also be required to pre-select beam delivery angles. The dose-volume constraints are then combined into a cost function which mathematically quantifies the quality of a dose distribution against the planning constraints. An optimisation algorithm produces MLC apertures for each beam angle and finds their respective intensities in order to minimise the cost function.

The computational burden of inverse planning is orders of magnitude greater than forward planning due to the fact that the algorithm often must explore thousands-to-millions of combinations of the degrees of freedom offered by medical linacs to identify the optimal treatment plan. Modern treatment planning algorithms strive to include as many degrees of freedom as possible, creating an ever-increasing demand for computing power. For example, finding delivery trajectories which combine gantry and couch motion is an active area of research

[8–10]. The inverse planning approach is at the core of the work presented in this thesis. As such, the methodology will be elaborated in greater detail in chapter 3.

1.4 Linac-based radiation therapy

The process of producing a radiation beam from a linac begins with an electron gun producing a low energy beam of electrons through thermionic emission. These electrons are then accelerated through a waveguide using microwave radiation produced by a radiofrequency power generator. The nominal accelerating potential of the electron beam is determined by the exact frequency of the microwave radiation used in the waveguide, with a typical range of 4-25 MeV. Once accelerated, the beam transport system guides the electron beam to the head of the gantry using a bending magnet.

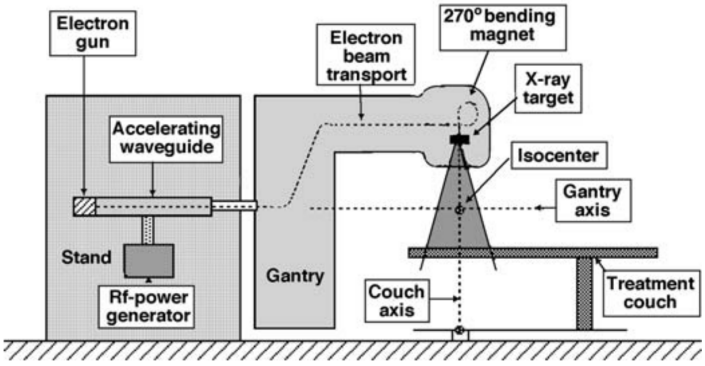

Figure 1.2 Schematic of a modern medical linear accelerator [3].

For clinical electron treatments, scattering foils are put in the path of the electron beam to broaden its intensity profile. To create a photon beam, a high atomic number target (e.g., tungsten) is put in the path of the electron beam instead of scattering foils. The target is responsible for producing x-ray photons primarily through bremsstrahlung interactions between the electrons and the target material. Both electron and photon beams are typically collimated further by one or many accessories such as tungsten jaws, a multi-leaf collimator (MLC), applicators and wedges to delivery conformal radiation to the tumour. The gantry is mounted on a ring which allows it to rotate a full 360° around the patient.

1.4.1 Photon and electron radiation therapy

Typically, electron radiation therapy is used to treat superficial tumours while megavoltage photons are reserved for deep seated lesions. Figure 1.3 highlights the differences between photon and electron beams by showing percent depth dose (PDD) curves in water for typical beam energies used clinically. The PDD curve is defined as the ratio of the dose as a function of depth normalised to the dose at the depth of maximum dose d_{max}. Electron beams have a

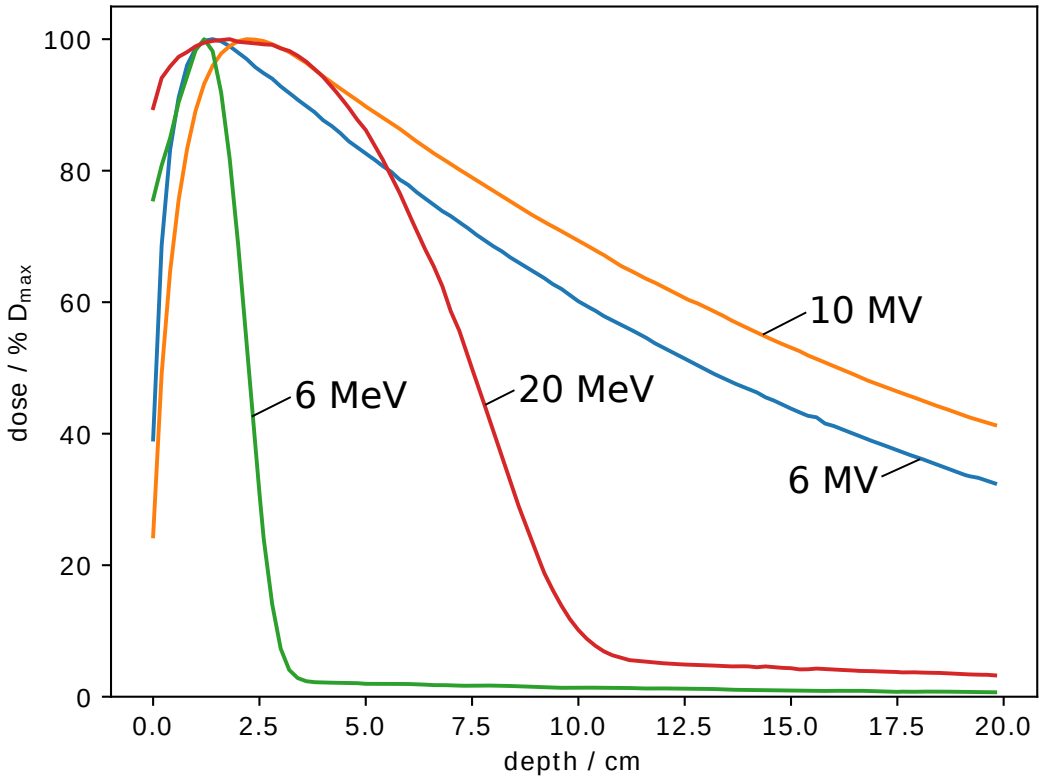

Figure 1.3 Percent depth dose curves in water for 6 and 10 MV photon beams compared to 6 and 20 MeV electron beams.

well-defined range, determined by the nominal energy of the beam, beyond which a negligible amount of dose is deposited. This feature is helpful to deliver high doses directly to superficial tumours while sparing normal tissue downstream from the tumour. However, electron beam profiles degrade significantly faster than photon beams due to scattering, requiring special accessories to collimate the beam as closely to the patient surface as possible. Photon beams, on

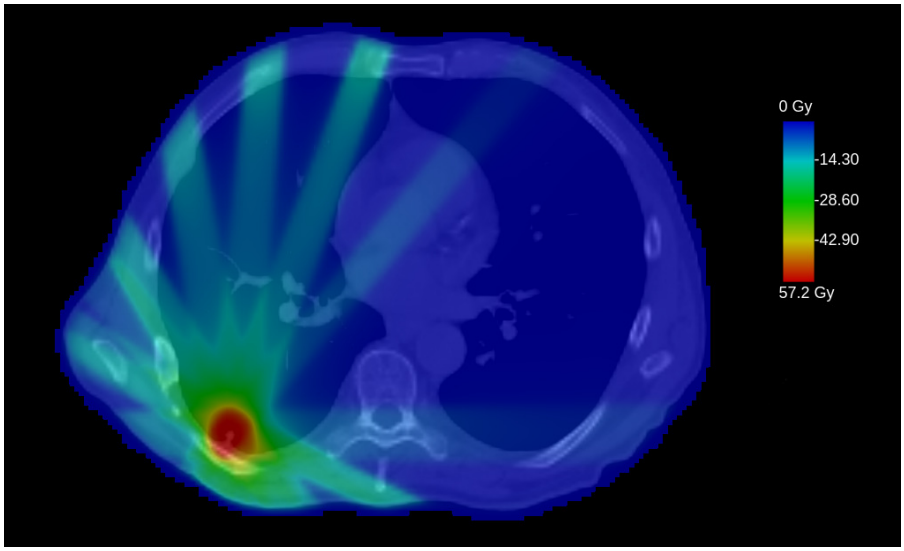

Figure 1.4 Dose colourwash showing multiple beam directions being focused on a common point to achieve high dose in the target while maintaining clinically acceptable levels of normal tissue dose.

the other hand, can be used to deliver high doses at any depth by focusing beams to a common point from multiple angles around the patient, as shown in Figure 1.4.

The increased logistical burden associated with electron beam delivery, in addition to the relatively high quality of photon-only treatment plans even for superficial tumours, has meant that electrons are sparingly used in a clinical environment despite their potential for superior normal tissue sparing.

1.4.2 The multi-leaf collimator

Practically all modern medical linacs are equipped with a MLC similar to the one shown in figure 1.5. It is a dynamic beam modifying accessory capable of collimating the beam into complex shapes to create arbitrary fluence profiles.

MLC leaves are able to move continuously during treatment, offering a large number of degrees of freedom to customise the delivery of radiation specifically to each patient's anatomy. These degrees of freedom are essential to deliver modern treatment techniques such as intensity modulated radiation therapy (IMRT) [11] and volumetric modulated arc therapy (VMAT) [12]. One of the key computational problems in modern EBRT treatment planning is the task of

Figure 1.5 A multi-leaf collimator. Individual leaves are controlled by stepper motors and can move independently to form arbitrary shapes and collimate the beam of radiation [11].

finding high quality MLC-defined shapes (apertures), and the amount of radiation to be delivered from each aperture to achieve a desired dose distribution.

1.5 Mixed beam radiation therapy

Linear accelerators offer a large array of degrees of freedom which can be programmed to produce a treatment plan. The treatment couch angle, gantry rotation, collimator head rotation, MLC leaf positions and jaws positions can all be controlled individually and adjusted dynamically during treatment.

In addition to the motional degrees of freedom, the particle type and energy are also parameters that can be exploited to improve the quality of treatment plans, as medical linacs can produce both electron and photon radiation beams. While both photons and electron treatment planning has been widely studied and implemented clinically, the concept of combining modalities–electrons and photons–has received less attention. This is in part due to the fact that electrons historically required additional accessories to produce acceptable dose distributions, making the delivery of mixed beam (MBRT) electron-photon plans cumbersome. However,

studies have shown that modulated electron radiation therapy (MERT) can be delivered with a standard photon MLC on a modern linac provided the patient is brought closer to the treatment head [13, 14]. By using the photon MLC to collimate electron beams, a MBRT treatment plan could be delivered without staff intervention between modality changes. With MBRT plans, the tissue sparing capabilities of electrons could be combined with the sharp penumbra of photon beams to lower the dose to normal tissue without compromising target coverage for tumours with a superficial component.

The final challenge, then, is to develop algorithms that are able to produce MBRT plans. In order to produce competitive plans compared to high quality photon-only IMRT and VMAT plans, MBRT treatment planning algorithms must incorporate the recent advances in photon planning which greatly increased the computational complexity of the planning process.

1.6 Thesis hypotheses

The commonly used algorithms for photon-only treatment planning cannot be trivially extended to produce MBRT treatment plans due to the large increase in computational complexity from the addition of the electron components. The accuracy of combined intensity modulated photon and electron delivery has also never been demonstrated clinically on modern linear accelerators. Therefore, We therefore make the following hypotheses:

1. It is feasible to combine simultaneously optimised intensity-modulated photon radiation (IMRT) and modulated electron radiation therapy (MERT) beams into a single mixed beam plan using a photon MLC as the principal beam collimation device. Feasibility is considered demonstrated when MBRT plans can be created that provide superior normal tissue sparing compared to photon-only treatment plans without compromising tumour coverage.

2. It is feasible to include setup uncertainties in simultaneously optimised MBRT plans to explicitly ensure that target coverage is preserved under setup errors. Feasibility is considered demonstrated when robust MBRT plans are produced that are significantly more robust to setup errors than standard photon plans.

3. Modern linear accelerators can accurately deliver MBRT plans using a photon MLC. Accuracy in this context is assessed in comparison with clinically accepted tolerances in conventional photon-only modulated deliveries.

1.7 Thesis objectives

To address the issue of computational complexity associated with MBRT planning, we will employ the column generation method, an iterative optimisation technique especially well-suited for problems where a large number of possible variables—MLC apertures in our case—exist but only a few are expected to be needed to approach an optimal solution. The objectives of the thesis are related to the creation and validation of MBRT plans created using the column generation method, and are as follows:

1. *Develop a framework for the creation of MBRT treatment plans.* The complete framework consists of multiple components which must first be developed and validated:

   (a) A graphical user interface to visualise CT images, contoured structures and simulated radiation doses.

   (b) Fast dose calculation engines to generate beamlet dose distributions for both photons and electrons.

   (c) A column generation-based optimisation algorithm to produce MBRT treatment plans.

   (d) A module to convert the output of the optimisation algorithm into a format readable by the linear accelerator to deliver the treatment plan.

2. *Assess the performance and behaviour of the column generation method applied to MBRT planning.* Prior to performing large scale planning studies, we will confirm that the algorithm behaves as expected with respect to the modalities chosen in the final treatment plan for a variety of dose planning constraints. We will also seek to understand the effect of heuristic choices associated with the column generation method on the plans produced. Factors such as the cost function value as a function of the number of apertures in the treatment plan, the total running time of the algorithm and the average aperture size per modality will be studied for various choices of heuristics.

3. *Assess the feasibility of making MBRT treatment plans robust to positioning errors.* Electrons are losing energy at a much faster rate than photons as they travel through a medium. For example, the mean free path between interactions for a 10 MeV electron in water is $\approx 10^{-5}$ cm, compared to 20 cm for a photon of the same energy. Unlike photon beams, electron fluence is therefore significantly altered by minute changes in

medium density, type and thickness. For that reason, small (5 mm) patient positioning errors could compromise tumour coverage in a manner that cannot be prevented simply by treating a larger volume. In addition, there is already a consensus in the intensity modulated proton therapy (IMPT) literature that IMPT plan dose distributions are more sensitive to positioning errors than photon-only plans because protons are also charged particles [15–19]. Therefore, there are good reasons to believe that MBRT plans are also more sensitive to positioning errors than photon-only plans. We will assess the robustness of PTV-based MBRT plans and extend the MBRT column generation model to produce CTV-based robust plans in a similar manner to what is done in the proton therapy literature.

4. *Assess the deliverability of MBRT plans.* Before patients can be treated using the MBRT technique, we must assess the accuracy of MBRT plan delivery and identify the steps of a patient-specific QA protocol to build confidence in the delivery technique. As a surrogate for patient treatment delivery, we will deliver MBRT plans on both Solid Water[1] slabs and a cylindrical phantom meant to represent a simple, inherently robust delivery and a complex delivery, respectively. With film and ionisation chamber dosimetry, we will assess the accuracy of our delivered dose distributions by comparing the measured dose distributions to Monte Carlo calculated doses. We will also develop a method to determine beam quality conversion factors for ion chambers that correct for changes in calibration coefficient between reference calibration conditions and measurement conditions in MBRT fields.

1.8 Thesis outline

An overview of radiation dose simulation methods and the underlying physics is provided in Chapter 2. In Chapter 3, optimisation methods and their application to radiation therapy treatment planning are presented. A review of the modulated electron radiation therapy literature is presented in chapter 4.

Implementation details for the MBRT treatment planning framework, the dose calculation engines and the optimisation library are given in chapter 5. Three original manuscripts are

[1] Solid Water is a plastic with radiological properties and density similar to water and is a plastic trademarked by Sun Nuclear Corporation (Melbourne, Florida).

presented in chapters 6-8. Chapter 6 describes the application of the column generation method to MBRT treatment planning. Chapter 7 features the extension of the MBRT optimisation model to robust optimisation in order to produce treatment plans which are not degraded by small positioning errors. The deliverability of MBRT plans and a potential patient-specific QA procedure is presented in Chapter 8.

Finally, Chapter 9 presents a summary of the work featured in this thesis and discusses the future of MBRT research.

References

[1] M. J. Thun, J. O. DeLancey, M. M. Center, A. Jemal, and E. M. Ward, "The global burden of cancer: Priorities for prevention," Tech. Rep. 1, 2009.

[2] R. Atun, D. A. Jaffray, M. B. Barton, F. Bray, M. Baumann, B. Vikram, T. P. Hanna, F. M. Knaul, Y. Lievens, T. Y. Lui, M. Milosevic, B. O'Sullivan, D. L. Rodin, E. Rosenblatt, J. Van Dyk, M. L. Yap, E. Zubizarreta, and M. Gospodarowicz, "Expanding global access to radiotherapy," *The Lancet Oncology*, vol. 16, no. 10, pp. 1153–1186, 2015.

[3] E. Podgorsak, *Radiation Oncology Physics: A Handbook for Teachers and Students*. Vienna: International Atomic Energy Agency, 2005.

[4] ICRU, "Report 50: Prescribing, Recording, and Reporting Photon Beam Therapy," *Journal of the ICRU*, vol. os26, no. 1, 1993.

[5] ICRU, "Report 62 - Prescribing, Recording and Reporting Photon Beam Therapy (Supplement to ICRU Report 50)," *Journal of ICRU*, vol. os32, no. 1, 1999.

[6] ICRU, "Report 83: Prescribing, Recording, and Reporting Photon-Beam Intensity-Modulated Radiation Therapy (IMRT)," *Journal of the ICRU*, vol. 19, no. 1, 2010.

[7] ICRU, "Report 91: Prescribing, recording, and reporting of stereotactic treatments with small photon beams," *Journal of the International Commission on Radiation Units and Measurements*, vol. 14, no. 2, 2014.

[8] C. B. Locke and K. K. Bush, "Trajectory optimization in radiotherapy using sectioning (TORUS):," *Medical Physics*, vol. 44, no. 7, pp. 3375–3392, 2017.

[9] A. Tran, J. Zhang, K. Woods, V. Yu, D. Nguyen, G. Gustafson, L. Rosen, and K. Sheng, "Treatment planning comparison of IMPT, VMAT and 4Π radiotherapy for prostate cases," *Radiation Oncology*, vol. 12, no. 1, 2017.

[10] P. Dong, H. Liu, and L. Xing, "Monte Carlo tree search -based non-coplanar trajectory design for station parameter optimized radiation therapy (SPORT)," *Physics in Medicine and Biology*, vol. 63, no. 13, 2018.

[11] T. Bortfeld, "IMRT: A review and preview," *Physics in Medicine and Biology*, vol. 51, no. 13, 2006.

[12] K. Otto, "Volumetric modulated arc therapy: IMRT in a single gantry arc," *Medical Physics*, vol. 35, no. 1, pp. 310–317, 2008.

[13] F. C. Du Plessis, A. Leal, S. Stathakis, W. Xiong, and C. M. Ma, "Characterization of megavoltage electron beams delivered through a photon multi-leaf collimator (pMLC)," *Physics in Medicine and Biology*, vol. 51, no. 8, pp. 2113–2129, 2006.

[14] D. Henzen, P. Manser, D. Frei, W. Volken, H. Neuenschwander, E. J. Born, a. Joosten, K. Lössl, D. M. Aebersold, C. Chatelain, M. F. M. Stampanoni, and M. K. Fix, "Beamlet based direct aperture optimization for MERT using a photon MLC.," *Medical physics*, vol. 41, p. 121711, dec 2014.

[15] D. Pflugfelder, J. J. Wilkens, and U. Oelfke, "Worst case optimization: A method to account for uncertainties in the optimization of intensity modulated proton therapy," *Physics in Medicine and Biology*, vol. 53, no. 6, pp. 1689–1700, 2008.

[16] 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.

[17] W. Liu, X. Zhang, Y. Li, and R. Mohan, "Robust optimization of intensity modulated proton therapy," *Medical Physics*, vol. 39, no. 2, pp. 1079–1091, 2012.

[18] M. Zaghian, W. Cao, W. Liu, L. Kardar, S. Randeniya, R. Mohan, and G. Lim, "Comparison of linear and nonlinear programming approaches for "worst case dose" and "minmax" robust optimization of intensity-modulated proton therapy dose distributions," *Journal of Applied Clinical Medical Physics*, vol. 18, no. 2, pp. 15–25, 2017.

[19] J. Unkelbach and H. Paganetti, "Robust Proton Treatment Planning: Physical and Biological Optimization," *Seminars in Radiation Oncology*, vol. 28, no. 2, pp. 88–96, 2018.

Chapter 2

Dose calculation methods in radiation therapy and underlying physics

As shown in chapter 1, TCP and NTCP curves have a sigmoidal shape with regions where the probability of tumour control or normal tissue toxicity varies rapidly as a function of absorbed dose. In the high gradient region of the TCP curve, a small (e.g 5%) dose inaccuracy can result in large (e.g 10-20%) changes in TCP. [1]. The accuracy requirements vary with the delivery techniques used, but in this work, we will focus on modern delivery methods which attempt to deliver dose to the target as conformally as possible.

The IAEA report 31 on accuracy requirements and uncertainties in radiotherapy recommends that treatment planning dose calculation uncertainties be below 3% in high dose, low gradient regions and below 5% in low dose regions. The end-to-end accuracy of dose delivery to patients should be within 5% or 4 mm when all sources of uncertainty are considered, based on tests made with anthropomorphic phantoms. However, the IAEA notes that 5% uncertainty may be difficult to achieve on real patients due to additional uncertainties such as respiratory motion [1].

2.1 General concepts and nomenclature

There are recurring concepts and nomenclature that appear in the majority of dose calculation methods. This section provides brief descriptions of each concept.

- Fluence (ϕ), dN/dA: the number of particles incident upon surface area dA.

- Energy fluence (ψ), dE/dA: the radiant energy incident on surface area dA, can be expressed as $\psi(E) = \phi$E, where E is the energy of the particles.
- Mass attenuation coefficient, μ/ρ: Linear attenuation coefficient for a given material representing the removal of primary photons from a monoenergetic narrow photon beam (also known as narrow beam attenuation coefficient).
- Total energy released per unit mass (TERMA): Total energy released by primary photons in matter following an interaction by primary photons. For example, in a Compton scattering event, TERMA includes both the kinetic energy of the electron set in motion, and the remaining scattered photon energy. TERMA can be expressed as TERMA $= \psi\mu/\rho$.
- Mass energy transfer coefficient, μ_{tr}/ρ: represents the fraction of released photon energy transferred into kinetic energy of charged particles.
- Kinetic energy released per unit mass (KERMA): KERMA represents the fraction of released photon energy that is transferred to the kinetic energy of charged particles. For example, in a pair production interaction, the rest mass energy of the electron positron pair is omitted from KERMA, but is included in TERMA as it originated from the photon. KERMA can be expressed as KERMA $= \psi\mu_{tr}/\rho$.
- Absorbed dose, D: Energy deposited in matter per unit mass. For photon beams, absorbed dose differs from KERMA because the charged particles set in motion by photons can lose part of their energy through radiative processes which will carry energy away from the medium. In addition, for both photon and electron beams, charged particles can have enough kinetic energy to deposit the bulk of their energy away from the point of energy release.

2.2 Monte Carlo dose calculation

Monte Carlo dose calculation methods represent the gold standard in dose calculation accuracy. Macroscopic quantities such as dose distributions in patients are obtained by averaging the *histories* of a large number of individual particles. One by one, particles are transported according to physical models describing their interactions with matter. These models appear

in the form of cross sections—interaction probability distributions—that are sampled using a pseudo-random number generator on a computer in order to produce a particle track. The name Monte Carlo, associated with a casino in Monaco bearing the name, derives from the fact that the sampling processes involve generating large quantities of random numbers to simulate the stochastic nature of particle interactions in nature [2].

For patient dose calculations, the main concern is recording the position and amount of energy deposition in matter. Because quantities of interest are averaged over a number of particle histories, there is an associated statistical uncertainty with any quantity calculated with the MC method. In general, the statistical uncertainties are proportional to the inverse square root of the total number of particles simulated, and achieving relative uncertainties below 1% in patient dose calculations typically requires millions of particles to be transported.

The basic principles of particle transport are the same for most types of particles. First, a particle is initialised with a position, energy and direction. The distance to the next interaction is then sampled from the relevant physical interaction cross sections. Interaction type, scattering angles, production of secondary particles and energy loss are finally sampled to form a particle *step*. The process is repeated until the energy of the particle falls below a threshold, or the particle leaves the geometry of interest.

2.2.1 Cross sections

The majority of physical interaction between two particles have an associated cross section σ. Cross sections have units of area, and are proportional to the probability of interaction. The concept is best illustrated for scattering processes, for which the differential cross section, $d\sigma/d\Omega$, is defined as:

$$\frac{d\sigma}{d\Omega} = \frac{\text{number of particles scattered into d}\Omega}{\text{number of incident particles per area}} \tag{2.1}$$

Figure 2.1 illustrates an experimental setup for measuring $d\sigma/d\Omega$ for a scattering process. The total cross section σ is obtained by integrating over the 4π solid angle. In MC simulations, the probability of a particle undergoing a given interaction is $\sigma/\sigma_{\text{tot}}$ where σ_{tot} is simply the sum of the cross sections for all possible interactions for the particle.

Differential cross sections are also used to sample scattering angles since the probability of a particle scattering at a given angle (θ, ϕ) is simply $\frac{d\sigma}{d\Omega}(\theta, \phi)/\sigma$. For non-scattering processes

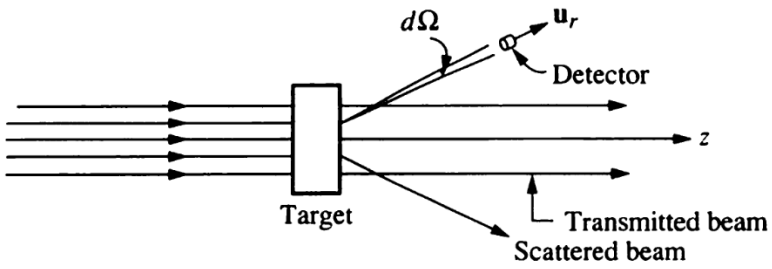

Figure 2.1 Experimental setup to measure $d\sigma/d\Omega$ for a scattering experiment [3].

such as characteristic x-ray emission, the cross section is interpreted differently than in eq. 2.1 but it remains proportional to the probability of the process to occur.

2.2.2 Photon interactions in matter

Photons must interact with matter to produce charged particles before any energy can be deposited. The physical models needed to accurately simulate photon transport must therefore represent the processes by which photons create these charged particles. In the energy range used in linac-based radiation therapy, there are three processes of interest: photoelectric effect, Compton scattering and pair production [4]. Photo-nuclear interactions also occur above the threshold energy for materials present in RT beams, but the dose deposited resulting from this interaction type is usually neglected in external beam RT dose calculation, except in the case where activation or neutron shielding is of interest (personnel dosimetry).

- The *photoelectric effect* is an interaction between a photon and a tightly bound orbital electron of an absorbing medium atom. When the photoelectric effect occurs, the photon is entirely absorbed and the orbital electron is ejected with a kinetic energy equal to the incoming photon energy, minus the binding energy of the electron that had to be overcome for the electron to be ejected. After the electron is ejected, the atom is left in an excited state and the remaining orbital electrons rearrange their energy levels to return to the ground state for the atom, releasing characteristic x-rays and Auger, Coster-Kronig and super Coster-Kronig electrons as a result (atomic relaxation).

- *Compton scattering* is an interaction between a photon and a loosely bound orbital electron. As a result of the interaction, the incoming photon is scattered with an energy

lower than its initial energy, and an electron is ejected from the atom with a kinetic energy depending on the scattering angle. This process is also followed by atomic relaxation.

- *Pair production* is an interaction resulting in the complete absorption of a photon and the creation of an electron-positron pair. This interaction can only occur if the photon has an energy greater than $2m_ec^2(1+m_e/M)$—twice the rest mass of an electron corrected for electron mass to nuclear mass ratio —for energy and momentum conservation to be satisfied. Photons cannot undergo pair production in a vacuum, as excess momentum from the photon must be absorbed by the medium. If the photon has energy greater than $4m_ec^2$, there is a probability that the excess momentum is absorbed by an orbital electron, causing a third charged particle to be released. This interaction is sometimes referred to as "triplet production" or pair production in the orbital electron field.

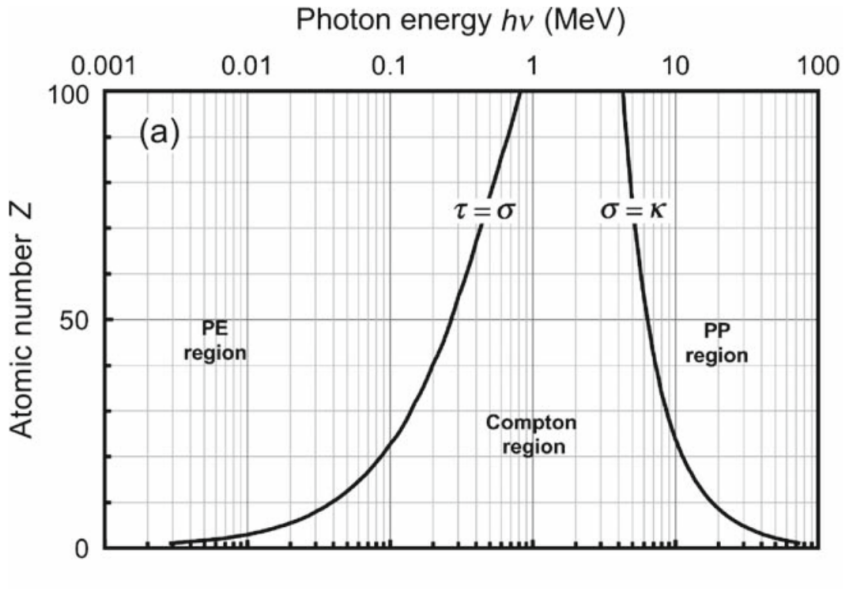

Figure 2.2 Dominant photon interaction process as a function of incoming photon energy and absorber material atomic number [4].

Figure 2.2 shows that in the energy range of interest in external beam radiation therapy, 1-20 MeV, the dominant photon interaction type is Compton scattering, with pair production featuring more prominently at high photon energies and high atomic number materials. Compared to a typical length scale of $\approx$ 30 cm in patient dose calculations, the mean free path of photons

is on the order of 10 cm in human tissue, meaning that photons will undergo relatively few interactions before leaving the geometry of interest or being absorbed [2].

The fact that Compton scattering is dominant at clinical treatment energies has important consequences for photon dose distributions and historical simplicity of dose calculations. The Compton scattering contribution to the mass energy transfer coefficient μ_{tr}/ρ has a relatively negligible dependence on the atomic number Z, meaning that heterogeneities such as bone and lung have a relatively minor effect on the overall dose distribution as the energy transferred $dE \sim \mu_{tr}^{\text{Compton}} \propto \rho$ and the dose distribution is $D \sim dE/\rho$.

In comparison, the photoelectric effect contribution to μ/ρ has a $\approx Z^3$ dependence, leading to a larger increase in energy deposited in higher-Z materials such as bone for low energy kilovoltage (kV) photon beams where the photoelectric effect is more prominent [4]. Figure 2.3 compares the dose colour wash for lung tumours treated with kV and MV photons, showing that the presence of bone leads to an increase in local dose deposition for the kV treatment plan. When photon interactions are independent of atomic number, dose distributions can be simulated by only keeping track of the mass density in the patient without a major loss of accuracy, i.e., by treating every voxel in the patient as having the elemental composition of water with varying density obtained from CT data.

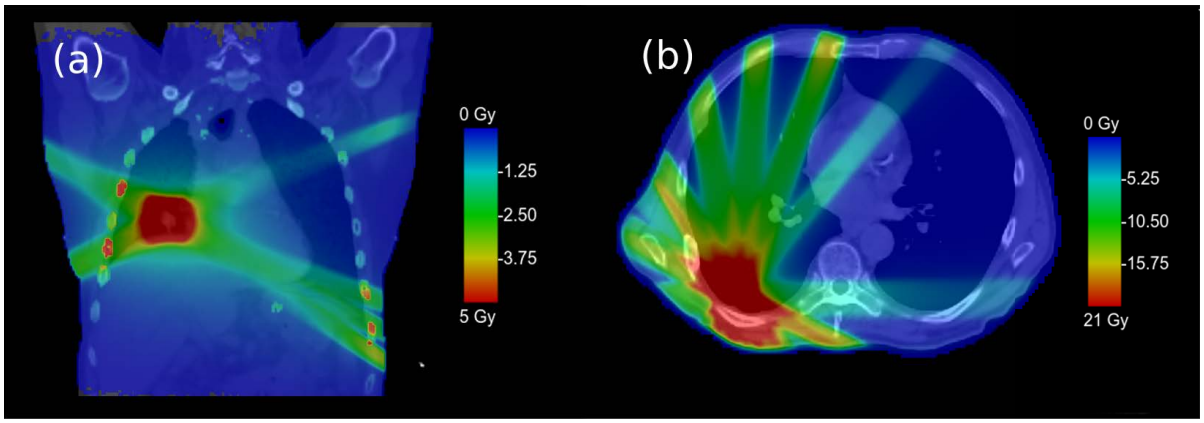

Figure 2.3 Simulated lung treatment plans on two different patients with (a) 200 kV and (b) 6 MV photons. The colourwash upper bound is set to 40% of the prescription dose for both plans to emphasise the effect of bone on the dose distribution. The dose deposited in bone is larger for kV photons due to the higher probability of photoelectric effect.

2.2.3 Electron interactions in matter

Unlike photons, charged particles undergo a large number of individual interactions per unit length. A 1 MeV electron, for example, will undergo on the order of 10^5 interactions in water before losing all of its kinetic energy [5]. Coulombic interactions are typically separated into soft (elastic) and hard (inelastic) collisions with the absorber media, in addition to radiative interactions known as *bremsstrahlung* resulting in energy loss in the form of photon radiation. The types of hard interactions possible are electron-electron (Møller scattering), positron-electron (Bhabha scattering), and electron impact ionisation which creates inner shell vacancies in atomic nuclei. Positrons can be treated identically to electrons, with the exception that they can also undergo electron-positron annihilation.

Condensed history

Explicitly sampling every individual charged particle interaction renders Monte Carlo methods too computationally intensive for practical dose calculation in patient geometries. However, an analysis of the scattering properties of electrons in matter shows that single collisions cause only small changes in direction and energy. As a result, Berger et al. proposed two condensed history (CH) schemes, dubbed class I and class II, which combine multiple interactions into larger bulk "steps" using a multiple scattering theory [6]. The two schemes differ by how the creation of secondary particles such as knock-on electrons and bremsstrahlung emission is treated.

In a class I scheme, the effects of secondary particle creation on the primary particle are decoupled from the primary particle stepping process. For primary particles, the step size and angular deflection are determined by the multiple scattering theory, and the creation of a secondary knock-on electron does not explicitly affect the direction of the primary electron. Class II schemes incorporate both the deflection due to multiple scattering theory and the creation of secondary particles above an energy threshold. Figure 2.4 compares the logic flows of class I and class II schemes for electron particle steps.

In EGSnrc, the class II MC code used in this work, electron interactions resulting in either bremsstrahlung emission or "hard" collisions generating secondary electrons above an energy threshold are considered "catastropic" interactions, and the secondary particles are explicitly transported. Individual interactions resulting in energy losses below the catastrophic interaction threshold are combined and their cumulative energy loss is treated under the continuous slowing

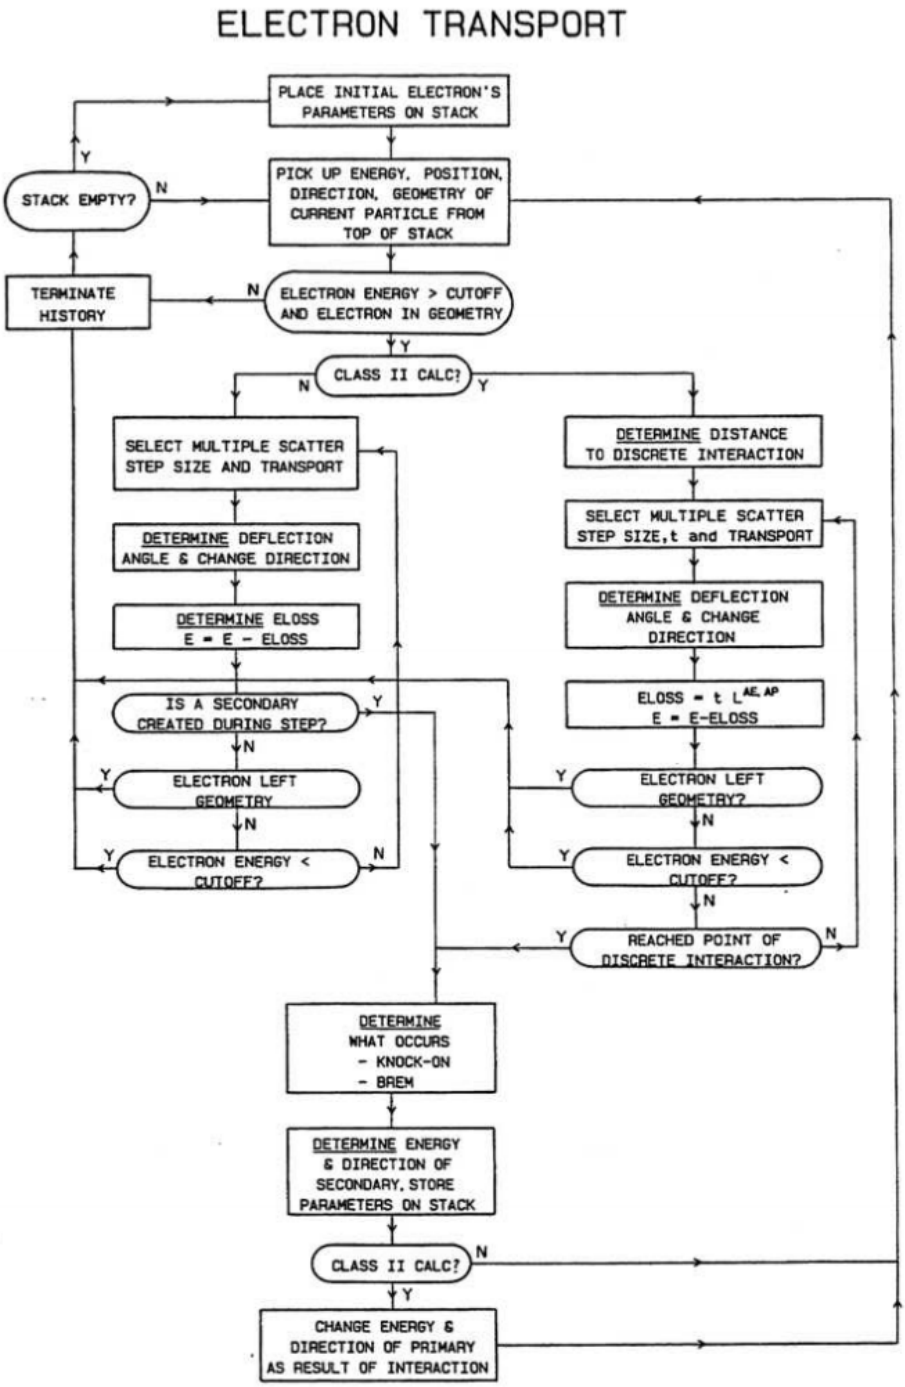

Figure 2.4 Simplified logic flow comparing class I to class II electron transport schemes. Reproduced from Rogers et al. [7]

down approximation (CSDA) where the energy is deposited uniformly between two steps, taking into account the curved step length. As a result of this scheme, one of the major parameters affecting the accuracy of CH schemes is step size selection. In addition, a carefully tuned boundary crossing algorithm must be used when transporting particles close to the interface between two regions [8].

2.2.4 Monte Carlo codes

There are many general-purpose MC codes released and maintained by large organisations. This work makes use of EGSnrc, a MC code used to transport electrons, positrons and photons maintained by the National Research Council (NRC) of Canada [8]. EGSnrc is based on EGS4, an older MC code released by the Stanford Linear Accelerator Center (SLAC) [9]. Other Major MC codes include MCNP [10], distributed by the Oak Ridge National Laboratory, GEANT4, maintained by CERN [11], FLUKA [12], maintained by the Italian National Institute for Nuclear Physics and PENELOPE [13], maintained by the University of Barcelona.

These MC codes serve as a foundation upon which *user codes* can be built. For example, BEAMnrc [14] is a library of predefined geometric component modules to produce an accurate physical model of a linear accelerator. These modules define boundary conditions to restrict steps taken by the EGSnrc engine. DOSXYZnrc [15], another EGSnrc user code, is used to calculate patient dose distributions by transporting particles in voxelised geometries such as phantoms created from patient CTs. Other examples of user-codes include GATE [16] and TOPAS [17], both using GEANT4 as the underlying particle transport engine.

BEAMnrc

As mentioned above, BEAMnrc provides predefined *component modules* (CM) whose parameters can be defined to model the most important components of medical linear accelerators. Figure 2.5 shows an example BEAMnrc accelerator with standard linac components.

Using BEAMnrc, a radiation beam can be fully modelled by the incoming electron source with a minimum number of tunable parameters, provided that an representation of the linac geometry is provided by the manufacturer. When the geometry is well known, the remaining parameters are the energy and full width half maximum (FWHM) of the electron beam prior to impinging on the target, and the angular divergence of the beam.

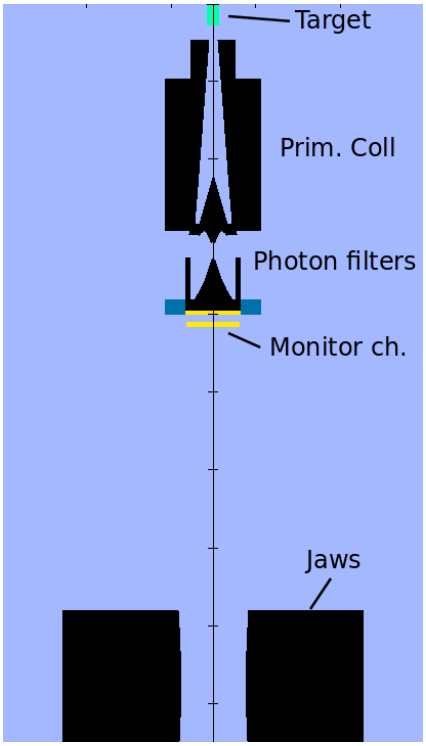

Figure 2.5 Example medical linear accelerator as modelled using BEAMnrc. The photon target, primary collimator, photon flattening and hardening filters, a monitor chamber and collimating jaws are shown.

The output of BEAMnrc is a phase space file accumulated at a user-defined plane. The file contains positions, directions and energies of all particles that were not absorbed by the linac components or left the geometry prior to reaching the phase space plane.

DOSXYZnrc

DOSXYZnrc is an EGSnrc user code with predefined routines to transport particles and score dose in voxelised, rectangular geometries. It can be used to reproduce realistic clinical reference calibration conditions, and also accurate patient representations typically obtained from a CT dataset. The major inputs are the particle source, and the *phantom* defining the material and density of every voxel in the simulation grid. One of the core features of DOSXYZnrc is that a BEAMnrc accelerator can be used directly as a particle source, allowing the dose distribution from full radiation therapy treatment plans to be calculated. The position of linac jaws and MLC leaves defined in BEAMnrc can be coordinated with the particle source angle and phantom

position, allowing for accurate simulations of dynamic deliveries. For example, the dose distributions shown in figure 2.3 which feature multiple beam angles and linac configurations were calculated using DOSXYZnrc.

2.3 Analytical dose calculation methods

Unlike MC methods, analytical dose calculation methods produce dose distributions that are free of statistical fluctuations. The convolution-superposition methods presented in this section rely on macroscopic characterisations of various components of radiotherapy beams such as energy fluence and spectrum, along with precomputing energy deposition "point-spread" or "line-spread" kernels.

2.3.1 Convolution-superposition

The convolution-superposition (CS) method is an analytical photon dose calculation method first introduced in 1985 by Mackie et al. [18]. CS methods are primarily applied to patient dose calculations on a voxelised grid, where the density of each voxel is obtained from the planning CT acquired for that patient. The CS method typically calculates dose distributions in patients using a two-step process: TERMA calculation followed by convolution of with a "point spread" kernel [19]. This section is aimed at providing a brief theoretical overview of convolution-based methods. Specific implementation details for the CS-based dose calculation engine developed as part of the work performed in this thesis are instead presented in chapter 5.

TERMA calculation

First, the total energy released in matter (TERMA) is raytraced from the source to every voxel in the dose calculation grid. TERMA simply represents the distribution of the energy released in the patient through attenuation of photons coming from the source. For a monoenergetic beam, the TERMA can be calculated as:

$$T(r) = \frac{\mu}{\rho}(E)\Phi(E,r)E \tag{2.2}$$

In eq. 2.2, $\Phi(E_i, r)$ is the primary photon fluence at point $\boldsymbol{r}$:

$$\Phi(E,r) = \Phi_0(E,r_0)\left(\frac{r_0}{r}\right)^2 e^{-\mu(E)d(r)} \tag{2.3}$$

where Φ_0 is the primary photon fluence at r_0 on reference plane upstream from the calculation volume, $\left(\frac{r_0}{r}\right)^2$ accounts for the fact that clinical photon beams are divergent and $d(r)$ is the radiation path length traced by the TERMA ray at point $\boldsymbol{r}$:

$$d(\boldsymbol{r}) = \int_0^r \rho(t)dt \tag{2.4}$$

Table 2.1 defines the major quantities featured in CS calculations.

Symbol	Description
T	Total energy released in matter by primary photons
ρ	Mass density
η	Electron density relative to water
μ/ρ	Mass attenuation coefficient
Φ	Particle fluence
r_0	Distance between photon source and reference plane
Φ_0	Particle fluence at reference plane
d	Radiation path length
h	Differential energy deposition kernel
D	Dose

Table 2.1 Quantities of interest in CS calculations

Kernel convolution

The total dose distribution is obtained by convolving the TERMA distribution with a point kernel describing the spread of energy from the interaction point. Energy deposition kernels (EDK) are obtained from MC simulations by forcing photons to interact at a single point and recording the energy deposition of primary and secondary scatter particles as a function of distance from the interaction point, and angle from the incident particle direction. Mathematically, the EDK $h(r)$ is defined as:

$$h(\boldsymbol{r};\boldsymbol{s}) = \frac{d\varepsilon}{RdV} \tag{2.5}$$

where $d\varepsilon$ is the mean energy imparted in volume element dV out of the total radiant energy R originating from the source point s. EDKs are commonly normalised such that they integrate to unity over all space,

$$\int h(\boldsymbol{r};\boldsymbol{s})dV \equiv 1 \tag{2.6}$$

in order for the kernel to represent the fraction of energy deposited at point $\boldsymbol{r}$ given energy release at point $\boldsymbol{s}$. Kernels cannot be obtained or validated experimentally as it is impossible in practice to force photon interactions at a specific point, making MC calculation codes and models essential even for analytical CS dose calculation.

Total dose distribution

The total dose distribution in a homogeneous phantom for a monoenergetic photon beam is obtained by numerically performing the convolution:

$$D(\boldsymbol{r}) = \int_V T(\boldsymbol{s})h(\boldsymbol{r}-\boldsymbol{s})d^3s \tag{2.7}$$

where $D(\boldsymbol{r})$ is the dose at point $\boldsymbol{r}$, T and h have the definitions given above, and the integration is performed over the patient volume. However, realistic linac-produced photon beams are polyenergetic, requiring a generalisation of eq. 2.7 for polyenergetic beams:

$$D(\boldsymbol{r}) = \int_E \int_V T_E(\boldsymbol{s})h(E,\boldsymbol{r}-\boldsymbol{s})\mathrm{d}^3s\ \mathrm{d}E \tag{2.8}$$

In the name of computational efficiency, many approximate methods have been investigated to avoid performing the repeated CS steps (TERMA and convolution) for each energy in a polyenergetic spectrum. The main sources of inaccuracies introduced by these approximations come from depth hardening, where the average energy of the photon beam spectrum increases as a function of depth due to the fact that the linear attenuation coefficient, μ, is higher at lower photon energies. For that reason, lower energy primary photons are removed from the beam at a faster rate than higher energy photons. A single EDK calculated for the entrance spectrum will be inaccurate when convolving energy released at larger depths inside the patient. Similarly, the photon spectrum varies laterally as a function of off-axis distance, resulting in a softer beam due to the flattening filter providing less beam hardening for off-axis photons [19].

Dose calculations in heterogeneous media are typically handled by using radiological distance rather than physical distance between between energy release and deposition points,

$d = \int_{\boldsymbol{s}}^{\boldsymbol{r}} \rho(\boldsymbol{s'})ds'$. In the heterogeneous case, equation 2.7 is no longer strictly a convolution integral in the mathematical sense, as the kernel is not translationally invariant.

Computational complexity

The computational complexity of the CS method can be viewed in terms of how the number of operations required scale as a function of the number of voxels. Assuming a cubic grid of voxels of N^3 voxels, the complexity of TERMA calculations is $O(N^3)$ for homogeneous phantoms. However, the convolution step involves assessing the dose contribution from every voxel (N^3) to every other voxel (N^3), resulting in $O(N^6)$ operations. Heterogeneous phantoms result in $O(N^7)$ complexity due to the raytracing required to calculating radiological distance.

Collapsed cone convolution superposition

The collapsed cone convolution superposition (CCCS) algorithm greatly improves the computational efficiency of the CS algorithm by introducing an angular discretisation and analytical parameterisation of the kernel [20].

The kernel is first discretised into angular sectors (cones) Ω_i,

$$k_{\Omega_i}(r) = \int_{\Omega_i} h(r,\Omega) r^2 \mathrm{d}^2\Omega \tag{2.9}$$

and each cone is then parameterised by two exponentials:

$$k_{\Omega_i}(r) = A_{\Omega_i} e^{-a_{\Omega_i} r} + B_{\Omega_i} e^{-b_{\Omega_i} r} \tag{2.10}$$

where the parameters A_{Ω_i}, a_{Ω_i}, B_{Ω_i}, b_{Ω_i} are obtained through a fitting process. The first component is said to model primary energy deposition, that is, energy deposited from particles created by the primary photons. The second component captures the energy deposition from all secondary sources such as electrons created from photons which have already interacted, and knock-on electrons. In the CCCS method, energy is only transported along the discretised kernel directions, and all energy transported within a solid angle cone is "collapsed" by being deposited along the axis of the solid angle. As shown in Fig. 2.6, the CCCS method will result in misplaced energy deposition far from the interaction point where the solid angle spans multiple voxels. However, the fraction of energy deposited far from the interaction point is small enough for the method to remain accurate in patient dose calculations.

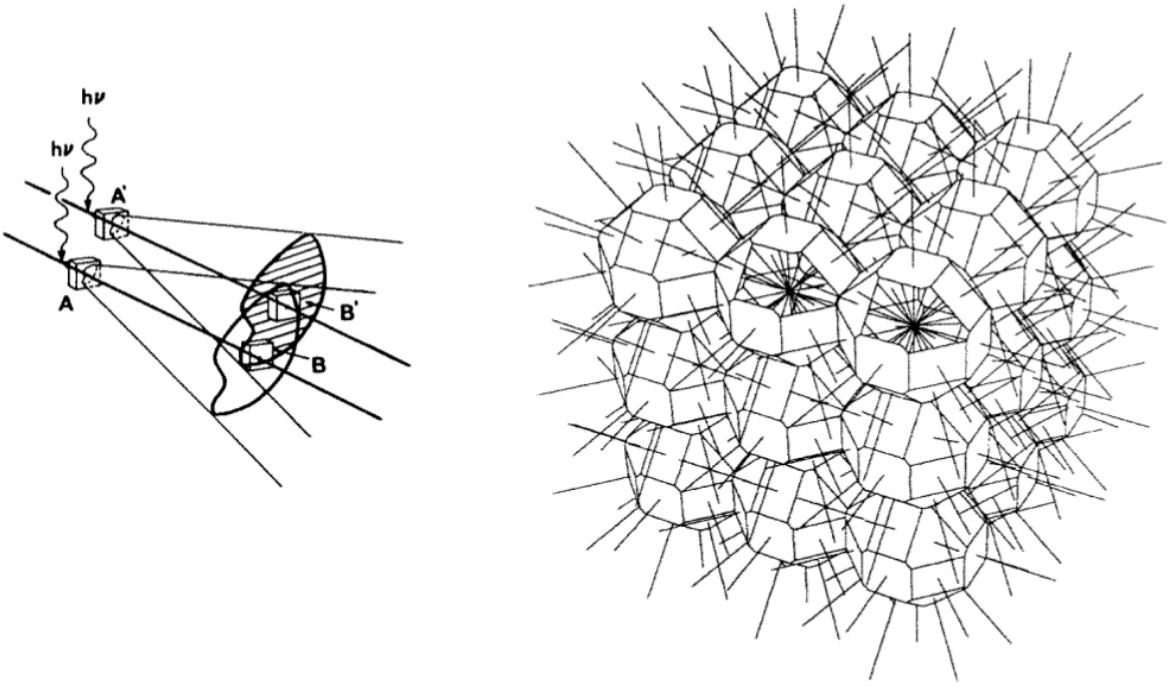

Figure 2.6 Left: Collapsed cone energy deposition from two adjacent TERMA voxels. Energy that should have been deposited from A to B' is deposited in B, and vice versa. Right: Example discretisation of the kernel and their transport directions superimposed on a grid of TERMA voxels. Reproduced from Ahnesjö et al. [20]

The total dose distribution is obtained in a similar fashion as the CS method. However, the convolution step is only performed over the discrete cones, hence the complexity of the CCCS method is $O(MN^4)$, where M is the number of angle segments. Given the form of the analytical kernel in eq. 2.10, recursive formulas can be derived to obtain the dose in a voxel by reusing quantities calculated in previous voxels, reducing the complexity to $O(MN^3)$ [20, 21]. It should be noted that the CCCS algorithm can be implemented with tabulated (non-analytical) kernels as well, with the understanding that the lack of recursive formulas will increase the computational complexity back to $O(MN^4)$.

2.3.2 Pencil beam convolution

Alternatively to the point-spread kernel approach, a line-spread *pencil beam* kernel approach can be used which greatly simplifies the dose computation process at the expense of accuracy.

Rather than performing a point kernel convolution with each TERMA voxel, dose is obtained by convolving line kernels with the incoming fluence distribution at a reference plane [22, 23]. The difference between the two kernels is shown in figure 2.7.

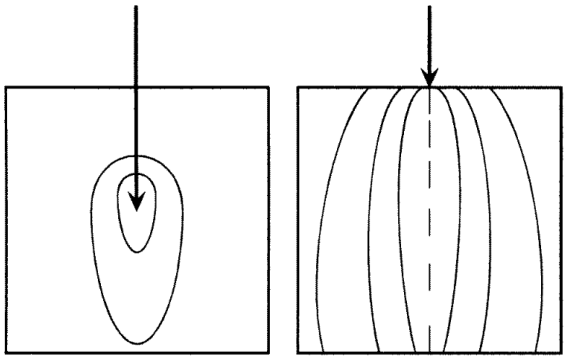

Figure 2.7 Graphical representation of a point kernel (left) and a line kernel (right). In pencil beam calculations, multiple pencil beams are superimposed to obtain dose distributions for macroscopic fields. Reproduced from Ahnesjö et al. [19]

Dose distributions in macroscopic fields are obtained by summing pencil *beamlets*, as shown in equation 2.11.

$$D(r,z) = \sum_p D_p(r,z) = \sum_p I_p(z) H_p(r,z) \tag{2.11}$$

where $D_p(r,z)$ is the dose from pencil beam p, which is expressed as the contribution $H_p(r,z)$ to point (r,z) from a line source $I_p(z)$. In general, pencil beam methods are considered inaccurate in heterogeneous situations, as the simple summation of pencil beams cannot account for perturbations on lateral scattering induced by heterogeneities.

2.4 Differences between Monte Carlo and analytical methods

MC dose calculations performed with an accurate geometric representation of a linear accelerator are considered the gold standard in dose calculation accuracy due to the accurate modelling of all relevant physical phenomenon involved in the creation of the radiation beam and subsequent transport in the patient. However, historically, MC dose calculations have been considered significantly slower than analytical algorithms such as CCCS. Analytical algorithms, despite

having faster running time, must approximate the effects of important physical phenomenon such as beam hardening, or the perturbation of collimating devices on the radiation beam.

For broad fields where the exact construction of beam modifiers is not expected to contribute large differences in calculated dose, the major sources of discrepancies between MC and CCCS doses occur in areas of lateral charged particle disequilibrium in low density regions, as shown in figure 2.8. Dose differences between MC and CCCS in the low density region reach up to 10% [24]. In patient geometries, these large differences in densities between adjacent voxels happen at the junction between the ribs and the lungs.

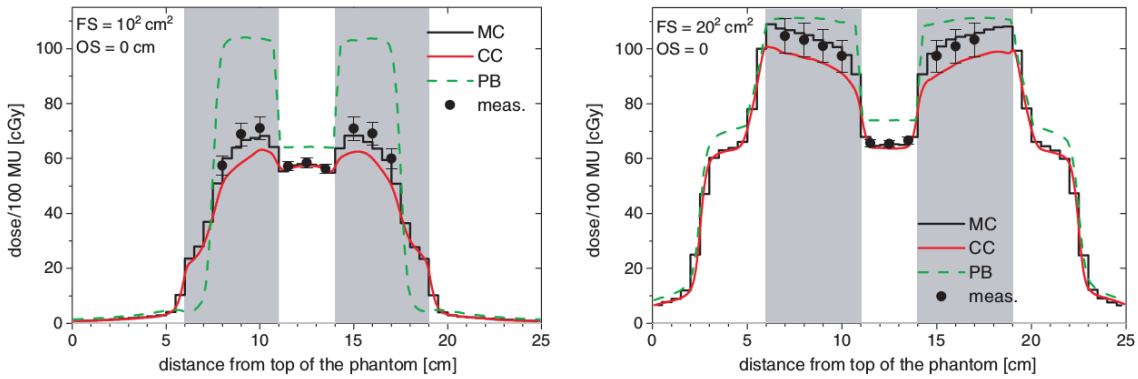

Figure 2.8 Dose profile comparisons between measurements, MC, CCCS and the pencil beam algorithm on a heterogeneous phantom for a (left) 10x10 cm^2 field and (right) 20x20 cm^2 field. The grey areas represent slabs of Styrofoam ($\rho = 0.035\ g/cm^3$) and white areas are RW3 white polystyrene ($\rho = 1.045\ g/cm^3$). Reproduced from Krieger et al. 2005 [24].

The pencil beam results show the largest discrepancies as expected. However, it should be noted that pencil beam methods can be made accurate enough for clinical use by incorporating lateral scattering correction factors. The Eclipse clinical treatment planning system (Varian Medical Systems, Palo Alto California) currently uses a pencil beam-based algorithm as the main dose calculation engine in the platform, with accurate results in all but the most challenging cases such as small radiation fields combined with bone-lung interfaces [23, 1]. Nonetheless, a study comparing seven different photon dose calculation algorithms distributed by various radiation therapy vendors showed that, in general, CS/CCCS methods were closer to MC-calculated dose distributions in challenging heterogeneous phantoms. In general, the agreement between analytical methods and MC was worse at higher energies such as 10, 15 and 18 MV photon beams [25].

References

[1] INTERNATIONAL ATOMIC ENERGY AGENCY, *Accuracy Requirements and Uncertainties in Radiotherapy*. No. 31 in IAEA Human Health Series, Vienna: INTERNATIONAL ATOMIC ENERGY AGENCY, 2016.

[2] J. Seco and F. Verhaegen, *Monte Carlo Techniques in Radiation Therapy*. 2013.

[3] J. S. Townsend, *A modern approach to quantum mechanics*. Sausalito: University Science Books, 2000.

[4] E. B. Podgoršak, *Radiation physics for medical physicists*. Berlin: Springer, second ed., 2010.

[5] P. Andreo, D. T. Burns, A. E. Nahum, J. Seuntjens, and F. H. Attix, *Fundamentals of ionizing radiation dosimetry*. John Wiley & Sons, 2017.

[6] M. Berger, *Monte Carlo calculation of the penetration and diffusion of fast charged particles*, vol. 1. New york: Academic, 1 ed., 1963.

[7] D. Rogers and A. Bielajew, "Monte Carlo Techniques of Electron and Photon Transport for Radiation Dosimetry," in *The Dosimetry of Ionizing Radiation* (K. R. Kase, B. E. Bjarngard, and F. H. Attix, eds.), p. Chapter 5, Academic Press, 1990.

[8] I. Kawrakow, "Accurate condensed history Monte Carlo simulation of electron transport. I. EGSnrc, the new EGS4 version," *Medical Physics*, vol. 27, no. 3, pp. 485–498, 2000.

[9] W. R. Nelson, H. Hirayama, and D. W. Rogers, "The EGS4 Code System," tech. rep., SLAC, 1985.

[10] C. J. Werner, "MCNP Users Manual-Code Version 6.2," tech. rep., Los Alamos National Laboratory, 2017.

[11] S. Agostinelli, J. Allison, K. Amako, J. Apostolakis, H. Araujo, P. Arce, M. Asai, D. Axen, S. Banerjee, G. Barrand, F. Behner, L. Bellagamba, J. Boudreau, L. Broglia, A. Brunengo, H. Burkhardt, S. Chauvie, J. Chuma, R. Chytracek, G. Cooperman, G. Cosmo, P. Degtyarenko, A. Dell'Acqua, G. Depaola, D. Dietrich, R. Enami, A. Feliciello, C. Ferguson, H. Fesefeldt, G. Folger, F. Foppiano, A. Forti, S. Garelli, S. Giani, R. Giannitrapani,

D. Gibin, J. J. Gomez Cadenas, I. Gonzalez, G. Gracia Abril, G. Greeniaus, W. Greiner, V. Grichine, A. Grossheim, S. Guatelli, P. Gumplinger, R. Hamatsu, K. Hashimoto, H. Hasui, A. Heikkinen, A. Howard, V. Ivanchenko, A. Johnson, F. W. Jones, J. Kallenbach, N. Kanaya, M. Kawabata, Y. Kawabata, M. Kawaguti, S. Kelner, P. Kent, A. Kimura, T. Kodama, R. Kokoulin, M. Kossov, H. Kurashige, E. Lamanna, T. Lampen, V. Lara, V. Lefebure, F. Lei, M. Liendl, W. Lockman, F. Longo, S. Magni, M. Maire, E. Medernach, K. Minamimoto, P. Mora de Freitas, Y. Morita, K. Murakami, M. Nagamatu, R. Nartallo, P. Nieminen, T. Nishimura, K. Ohtsubo, M. Okamura, S. O'Neale, Y. Oohata, K. Paech, J. Perl, A. Pfeiffer, M. G. Pia, F. Ranjard, A. Rybin, S. Sadilov, E. di Salvo, G. Santin, T. Sasaki, N. Savvas, Y. Sawada, S. Scherer, S. Sei, V. Sirotenko, D. Smith, N. Starkov, H. Stoecker, J. Sulkimo, M. Takahata, S. Tanaka, E. Tcherniaev, E. Safai Tehrani, M. Tropeano, P. Truscott, H. Uno, L. Urban, P. Urban, M. Verderi, A. Walkden, W. Wander, H. Weber, J. P. Wellisch, T. Wenaus, D. C. Williams, D. Wright, T. Yamada, H. Yoshida, and D. Zschiesche, "GEANT4 - A simulation toolkit," *Nuclear Instruments and Methods in Physics Research, Section A: Accelerators, Spectrometers, Detectors and Associated Equipment*, vol. 506, no. 3, pp. 250–303, 2003.

[12] A. Ferrari and Others, "FLUKA: A multi-particle transport code (Program version 2005)," tech. rep., CERN, 2005.

[13] J. Baró, J. Sempau, J. M. Fernández-Varea, and F. Salvat, "PENELOPE: An algorithm for Monte Carlo simulation of the penetration and energy loss of electrons and positrons in matter," *Nuclear Inst. and Methods in Physics Research, B*, vol. 100, no. 1, pp. 31–46, 1995.

[14] D. W. O. Rogers, B. Walters, and I. Kawrakow, "BEAMnrc Users Manual," tech. rep., National Research Council of Canada, 2011.

[15] B. Walters, I. Kawrakow, and D. W. O. Rogers, "DOSXYZnrc Users Manual," tech. rep., National Research Council of Canada, 2016.

[16] S. Jan, G. Santin, D. Strul, S. Staelens, K. Assié, D. Autret, S. Avner, R. Barbier, M. Bardiès, P. M. Bloomfield, D. Brasse, V. Breton, P. Bruyndonckx, I. Buvat, A. F. Chatziioannou, Y. Choi, Y. H. Chung, C. Comtat, D. Donnarieix, L. Ferrer, S. J. Glick, C. J. Groiselle, D. Guez, P. F. Honore, S. Kerhoas-Cavata, A. S. Kirov, V. Kohli, M. Koole,

M. Krieguer, D. J. van der Laan, F. Lamare, G. Largeron, C. Lartizien, D. Lazaro, M. C. Maas, L. Maigne, F. Mayet, F. Melot, C. Merheb, E. Pennacchio, J. Perez, U. Pietrzyk, F. R. Rannou, M. Rey, D. R. Schaart, C. R. Schmidtlein, L. Simon, T. Y. Song, J. M. Vieira, D. Visvikis, R. Van de Walle, E. Wieërs, and C. Morel, "GATE: A simulation toolkit for PET and SPECT," *Physics in Medicine and Biology*, vol. 49, no. 19, pp. 4543–4561, 2004.

[17] J. Perl, J. Shin, J. Schümann, B. Faddegon, and H. Paganetti, "TOPAS: An innovative proton Monte Carlo platform for research and clinical applications," *Medical Physics*, vol. 39, no. 11, pp. 6818–6837, 2012.

[18] T. R. Mackie, J. W. Scrimger, and J. J. Battista, "A convolution method of calculating dose for 15 MV x rays," *Medical Physics*, vol. 12, no. 2, pp. 188–196, 1985.

[19] A. Ahnesjö and M. M. Aspradakis, "Dose calculations for external photon beams in radiotherapy," *Physics in Medicine and Biology*, vol. 44, no. 11, 1999.

[20] A. Ahnesjö, "Collapsed cone convolution of radiant energy for photon dose calculation in heterogeneous media," *Medical Physics*, vol. 16, no. 4, pp. 577–592, 1989.

[21] Q. Chen, M. Chen, and W. Lu, "Ultrafast convolution/superposition using tabulated and exponential kernels on GPU," *Medical Physics*, vol. 38, no. 3, pp. 1150–1161, 2011.

[22] A. Ahnesjö, M. Saxner, and A. Trepp, "A pencil beam model for photon dose calculation," *Medical physics*, vol. 19, no. 2, pp. 263–273, 1992.

[23] L. Tillikainen, H. Helminen, T. Torsti, S. Siljamäki, J. Alakuijala, J. Pyyry, and W. Ulmer, "A 3D pencil-beam-based superposition algorithm for photon dose calculation in heterogeneous media," *Physics in Medicine and Biology*, vol. 53, no. 14, pp. 3821–3839, 2008.

[24] T. Krieger and O. A. Sauer, "Monte Carlo- versus pencil-beam-/collapsed-cone-dose calculation in a heterogeneous multi-layer phantom," in *Physics in Medicine and Biology*, vol. 50, pp. 859–868, 2005.

[25] A. Fogliata, E. Vanetti, D. Albers, C. Brink, A. Clivio, T. Knöös, G. Nicolini, and L. Cozzi, "On the dosimetric behaviour of photon dose calculation algorithms in the presence of simple geometric heterogeneities: Comparison with Monte Carlo calculations," *Physics in Medicine and Biology*, vol. 52, no. 5, pp. 1363–1385, 2007.

Chapter 3

Optimisation in radiation therapy treatment planning

A radiation therapy treatment plan is described by a set of control points fully defining the state (gantry angle, couch angle, jaw and MLC positions, etc) of the linear accelerator and its accessories at a point in time. During delivery, the linac proceeds through each control point and delivers a preset amount of radiation, measured in monitor units (MU), from each control point. The aim of a treatment planning algorithm is to identify the optimal (or near-optimal) set of control points and number of MUs for each control point. To achieve this goal, optimisation techniques are used to minimise a cost function representing the deviation of the treatment plan from its prescribed dose distribution. An overview of optimisation techniques will be presented in this chapter, along with their applications to radiation therapy treatment planning.

In their general form, optimisation problems can be written as

$$\begin{aligned} \underset{x}{\text{minimise}} \quad & f(\boldsymbol{x}) \\ \text{subject to} \quad & g_i(\boldsymbol{x}) \leq 0, i = 1, \ldots, m \\ & h_j(\boldsymbol{x}) = 0,\ j = 1, \ldots, p \\ & \boldsymbol{x} \in S \end{aligned} \tag{3.1}$$

where the vector $\boldsymbol{x} \in \mathbb{R}^N$ represents the *decision variables* that must be determined to minimise the *cost function* $f(\boldsymbol{x})$. $g_i(\boldsymbol{x})$ are inequality constraints, $h_j(\boldsymbol{x})$ are equality constraints, and S is the set of allowable values that $\boldsymbol{x}$ can take. Note that maximisation and minimisation

are interchangeable without loss of generality as maximising $f(\boldsymbol{x})$ is identical to minimising $-f(\boldsymbol{x})$, and vice versa.

3.1 Linear optimisation

A linear optimisation problem (or linear program (LP)) is a problem in which both the cost function and the set of constraints in the problem are linear. All LPs can be written in the standard form [1],

$$\begin{aligned} &\underset{x}{\text{minimise}} \quad \boldsymbol{c}^T\boldsymbol{x} \\ &\text{subject to} \quad \boldsymbol{A}\boldsymbol{x} = \boldsymbol{b},\ \boldsymbol{x} \geq 0 \end{aligned} \tag{3.2}$$

where $\boldsymbol{x}$ and $\boldsymbol{c}$ are n-dimensional vectors, $\boldsymbol{A}$ is a $m \times n$ matrix and $\boldsymbol{b}$ is a m-dimensional vector. Inequalities are interpreted component-wise, e.g. the inequality $\boldsymbol{x} \geq 0$ is to be interpreted as meaning that all components of $\boldsymbol{x}$ must be non-negative.

All linear optimisation problems can be transformed into the standard form. Equality and inequality constraints can easily be transformed into one another. For example, equality constraints such as $h(\boldsymbol{x}) = 0$ can simply be replaced by two inequality constraints, $g_1(\boldsymbol{x}) \leq 0$ and $g_2(\boldsymbol{x}) \geq 0$ where $g_1(\boldsymbol{x}) = g_2(\boldsymbol{x}) = h(\boldsymbol{x})$. If a decision variable x_f is free to take negative values, then a substitution can be made $x_f = u_f - v_f$ where u_f and v_f satisfy the non-negativity constraint. Inequality constraints can be cast as equality constraints by the addition of a "slack variable". For example, the inequality

$$x_1 + 2x_2 - x_3 \leq 5$$

becomes the equality constraint:

$$x_1 + 2x_2 - x_3 + s_1 = 5$$

where s_1 is an added slack variable. For m inequality constraints, the resulting LP in standard form will have $n+m$ decision variables, $\boldsymbol{x} = (x_1, ..., x_n, s_1, ..., s_m)$.

Of particular interest in radiation therapy optimisation, a cost function of the form $f(\boldsymbol{x}) = \max(a(\boldsymbol{x}), b(\boldsymbol{x}))$ can be included in a LP by introducing a new decision variable, t, and two inequality constraints, $t \geq a(\boldsymbol{x}), t \geq b(\boldsymbol{x})$. The LP then becomes *min* $f(t) = t$ subject to the

two new constraints, which can be cast into standard form provided $a(\boldsymbol{x})$ and $b(\boldsymbol{x})$ are linear functions.

Interpreted geometrically, the system of inequalities $\boldsymbol{Ax} = \boldsymbol{b},\ \boldsymbol{x} \geq 0$ defines a *feasible region* known as a convex polytope which represents the set of all points satisfying the constraints of the problem. For example, figure 3.1 shows the feasible region defined by the constraints $x_1 + x_2 + x_3 = 1, x_1 \geq 0, x_2 \geq 0, x_3 \geq 0$.

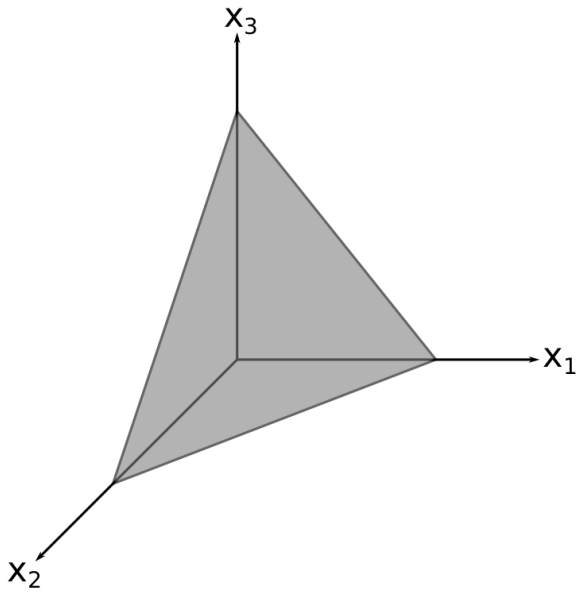

Figure 3.1 Feasible region defined by the constraints $x_1 + x_2 + x_3 = 1, x_1 \geq 0, x_2 \geq 0, x_3 \geq 0$. The solution of any bounded LP is found on one of the extreme points of its feasible region [2].

It can be shown that the solution to any bounded LP is found on the extreme points of the feasible region [1]. One of the well known methods for solving LPs, the simplex method, exploits this property by intelligently traversing the extreme points of the feasible region until the optimal feasible solution is found [2].

3.2 Nonlinear optimisation

Nonlinear optimisation (or nonlinear programming (NLP)) is concerned with solving problems of the more general form,

$$\begin{aligned} \underset{x}{\text{minimise}} \quad & f(\boldsymbol{x}) \\ \text{subject to} \quad & g_i(\boldsymbol{x}) \leq 0, i = 1, \ldots, m \\ & h_j(\boldsymbol{x}) = 0,\ j = 1, \ldots, p \\ & \boldsymbol{x} \in S \end{aligned} \tag{3.3}$$

where $f(\boldsymbol{x})$, $g_i(\boldsymbol{x})$, $h_j(\boldsymbol{x})$, can be nonlinear and S is a subset of $\boldsymbol{R}^n$. In this work, we only consider continuous functions.

3.2.1 Optimality conditions

Optimal solutions to a NLP must satisfy the necessary optimality conditions known as the Karush-Kuhn-Tucker (KKT) conditions [3]:

Let $\boldsymbol{x}^*$ be a minimum of the general NLP shown in 3.3, then there exists KKT multipliers $\boldsymbol{\mu} = (\mu_1, \ldots, \mu_m)$ and $\boldsymbol{\lambda} = (\lambda_1, \ldots, \lambda_p)$ such that

Stationarity

$$\nabla f(\boldsymbol{x}^*) + \sum_{i=1}^{m} \mu_i \nabla g_i(\boldsymbol{x}^*) + \sum_{j=1}^{p} \lambda_j \nabla h_j(\boldsymbol{x}^*) = 0$$

Primal feasibility

$$\begin{aligned} & g_i(\boldsymbol{x}^*) \leq 0,\ i = 1, \ldots, m \\ & h_j(\boldsymbol{x}^*) = 0,\ j = 1, \ldots, p \end{aligned}$$

Dual feasibility

$$\mu_i \geq 0, i = 1, \ldots, m$$

Complementary slackness

$$\mu_i g_i(\boldsymbol{x}^*) = 0, i = 1, \ldots, m$$

(3.4)

Note that in the case where a NLP is unconstrained, the stationarity condition becomes the familiar condition that a function is at an extremum when $\nabla f(\boldsymbol{x}^*) = 0$. The KKT conditions are a generalisation of this notion to constrained problems. The primal feasibility condition re-states the original constraints of the problem, and for equality-constrained problems, the KKT multipliers $\boldsymbol{\lambda}$ are known as Lagrange multipliers.

The complementary slackness condition is a compact way to express the condition that only binding inequalities ($g_i(\boldsymbol{x}*) = 0$) contribute to the optimality conditions. If, at an optimal value $\boldsymbol{x}^*$, an inequality constraint is non-binding ($g_i\boldsymbol{x}^* < 0$), then the associated KKT multiplier μ_i must be 0, and vice versa. The KKT multipliers are sometimes known as "prices" as they are related to the amount of change in the cost function of the optimal solution from changing the inequality constraint associated with the KKT multiplier.

In general, the NLP described in 3.3 can have multiple *local minima* which satisfy the KKT conditions. For example, in Figure 3.2, points A, B, C are local minima which satisfy the stationarity condition, but point A is the *global* minimum of the function in the range given.

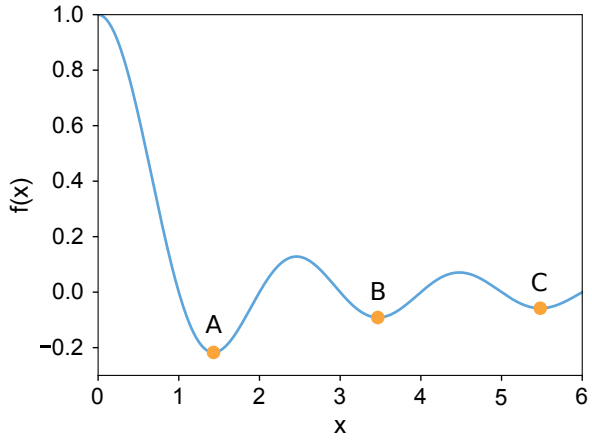

Figure 3.2 Simple one-dimensional non-convex function. The local minima are shown as dots.

There is a special case of NLPs, known as convex problems, where a point satisfying the KKT conditions is guaranteed to be a global minimum for the problem. A problem is convex if the cost function and constraint functions are convex. For functions, the mathematical definition of convexity is given by [3]:

$$f(\alpha\boldsymbol{x_1} + (1-\alpha)\boldsymbol{x_2}) \leq \alpha f(\boldsymbol{x_1}) + (1-\alpha)f(\boldsymbol{x_2}),\ \forall\, \boldsymbol{x_1}, \boldsymbol{x_2} \in \boldsymbol{R}^n, 0 \leq \alpha \leq 1 \tag{3.5}$$

In words, for any interval, the value of the function at the midpoint cannot exceed the mean of the function values evaluated at the ends of the interval. In more practical terms, a function is convex on an interval if its second derivative is non-negative throughout the interval. For non-convex problems, the vast majority of NLP solving algorithms will only provide a local minimum which depends on the initial conditions of the algorithm.

3.2.2 Gradient descent

To solve unconstrained NLP problems, one of the major ideas is to start from a candidate point $\boldsymbol{x}$ and identify a *descent* direction which will lead to a reduction of the cost function. The candidate point is then moved by a step t in that direction, and the process is repeated iteratively until no descent direction can be found. Mathematically, at the kth iteration, we seek a descent direction $\boldsymbol{d}^{(k)}$ where

$$\boldsymbol{x}^{(k+1)} = \boldsymbol{x}^{(k)} + t^{(k)}\boldsymbol{d}^{(k)} \implies f(\boldsymbol{x}^{(k+1)}) < f(\boldsymbol{x}^{(k)})$$

The idea can be concretised by inspection of the Taylor expansion of the cost function f(x) around a point $\boldsymbol{x_0}$:

$$f(\boldsymbol{x}) = f(\boldsymbol{x_0}) + (\boldsymbol{\nabla} f(\boldsymbol{x})|_{\boldsymbol{x_0}})^T \cdot (\boldsymbol{x}-\boldsymbol{x_0}) + \frac{1}{2}(\boldsymbol{x}-\boldsymbol{x_0})^T \boldsymbol{H}(\boldsymbol{x}-\boldsymbol{x_0}) + \mathscr{O}(\boldsymbol{x}^3) \tag{3.6}$$

where $\boldsymbol{H}$ is the matrix of mixed second order partial derivatives known as the Hessian. Hence, for a small change about $\boldsymbol{x}^{(k)}$,

$$f(\boldsymbol{x}^{(k+1)}) = f(\boldsymbol{x}^{(k)}) + t\boldsymbol{\nabla} f(\boldsymbol{x})|_{\boldsymbol{x}^{(k)}} \cdot \boldsymbol{d} + \ldots \tag{3.7}$$

By inspection, if $\boldsymbol{\nabla} f(\boldsymbol{x})|_{\boldsymbol{x}^{(k)}} \cdot \boldsymbol{d} < 0$, then there exists t such that $f(\boldsymbol{x}^{(k+1)}) < f(\boldsymbol{x}^{(k)})$. One of the natural choices is to choose $\boldsymbol{d} = -\boldsymbol{\nabla} f(\boldsymbol{x})$, leading to the method of *gradient descent*:

Algorithm 1 Gradient descent algorithm

Require: Starting point $\boldsymbol{x_0}$, tolerance ε

1: **while** $||\boldsymbol{\nabla} f(\boldsymbol{x})|_{\boldsymbol{x}^{(k)}}|| > \varepsilon$ **do**
2: $\quad \boldsymbol{d} := -\boldsymbol{\nabla} f(\boldsymbol{x})|_{\boldsymbol{x}^{(k)}}$
3: $\quad$ Find appropriate step size t
4: $\quad \boldsymbol{x}^{(k+1)} = \boldsymbol{x}^{(k)} + t\boldsymbol{d}$
5: **end while**
6: **return** $(\boldsymbol{x}^*)$

The step size t can be found by explicitly solving the *exact line search* problem for the value of t yielding the largest decrease in $f(\boldsymbol{x})$:

$$t = argmin_{s \geq 0}\ f(\boldsymbol{x} + s\boldsymbol{d})$$

or, alternatively, by using a practical algorithm such as the *backtracking line search* which simply finds a step size yielding "good enough" improvement in $f(\boldsymbol{x})$ [3]:

Algorithm 2 Backtracking line search

Require: $0 < \alpha < 0.5, 0 < \beta < 1$

1: $t := 1$
2: **while** $f(\boldsymbol{x} + t\boldsymbol{d}) > f(\boldsymbol{x}) + \alpha t \nabla f(\boldsymbol{x})^T \boldsymbol{d}$ **do**
3: $\quad t := \beta t$
4: **end while**
5: **return** t

The candidate step size t is reduced by a factor β until $f(\boldsymbol{x} + t\boldsymbol{d})$ is lower than a fraction of the improvement (determined by α) predicted by linear extrapolation.

However, the method of gradient descent with $d = \boldsymbol{\nabla} f(\boldsymbol{x})|_{\boldsymbol{x}^{(k)}}$ is known to exhibit poor convergence properties, meaning it can take many iterations to arrive at the solution [3].

3.2.3 Newton's method

An alternative choice of $\boldsymbol{d}$ is to select $\boldsymbol{d} = -\boldsymbol{H}^{-1}\boldsymbol{\nabla} f(\boldsymbol{x})$. To motivate the choice, we expand $f(\boldsymbol{x})$ around a point x to second order,

$$f(\boldsymbol{x} + \Delta\boldsymbol{x}) \approx f(\boldsymbol{x}) + \boldsymbol{\nabla} f(\boldsymbol{x})^T \cdot \boldsymbol{\Delta x} + \frac{1}{2}\Delta\boldsymbol{x}^T \boldsymbol{H} \Delta\boldsymbol{x} \tag{3.8}$$

The optimality condition for the second order expansion of $f(\boldsymbol{x})$ is

$$\nabla f(\boldsymbol{x} + \Delta\boldsymbol{x}) \approx \nabla f(\boldsymbol{x}) + H\Delta\boldsymbol{x} = 0 \tag{3.9}$$

which is solved by $\Delta\boldsymbol{x} = -\boldsymbol{H}^{-1}\boldsymbol{\nabla} f(\boldsymbol{x})$. In other words, this choice of $\boldsymbol{d}$ is the direction along which the second order Taylor expansion of $f(\boldsymbol{x})$ is minimised. The broad class of algorithms using this choice of direction is known as *Newton's method*. The Newton direction can be

difficult to compute in practice due to the computational cost associated with a straightforward calculation and inversion of the Hessian. It also requires that the cost function be twice differentiable. However, Newton's method is known to lead to faster convergence than gradient descent, and it can also be adapted to solve equality-constrained problems [3].

3.2.4 Interior point method

Interior point methods are a class of algorithms able to solve the general NLP problem with inequality constraints as defined in 3.3. The intuition behind interior point methods is to incorporate the inequality constraints into the cost function and heavily penalise violations of the inequality constraints with the aim of ensuring that the solution of the transformed NLP is also a solution to the original NLP. Mathematically, problem 3.3 becomes:

$$\begin{aligned} \underset{x}{\text{minimise}} \quad & f(\boldsymbol{x}) + \sum_{i=1}^{m} I\_(g_i(\boldsymbol{x})) \\ \text{subject to} \quad & h_j(\boldsymbol{x}) = 0,\ j = 1, \ldots, p \\ & \boldsymbol{x} \in S \end{aligned} \tag{3.10}$$

where, ideally,

$$I\_(u) = \begin{cases} 0 & \text{if } u \leq 0 \\ \infty & \text{if } u > 0 \end{cases}$$

However, this definition of $I\_(u)$ is not differentiable, hence a common substitute is the logarithmic barrier function,

$$\hat{I}\_(u) = -(1/t)log(-u) \tag{3.11}$$

where $t > 0$ is a parameter controlling how closely $\hat{I}$ approximates I, as shown in Fig. 3.3. The strategy is then to apply the adapted Newton's method to the transformed NLP with progressively larger values of t to refine the approximation. The reason for increasing t with each iteration rather than starting with a large value of t is a practical one, related to the fact that the Hessian begins to vary rapidly when the solution is near a boundary when t is large, which renders applications of Newton's method difficult.

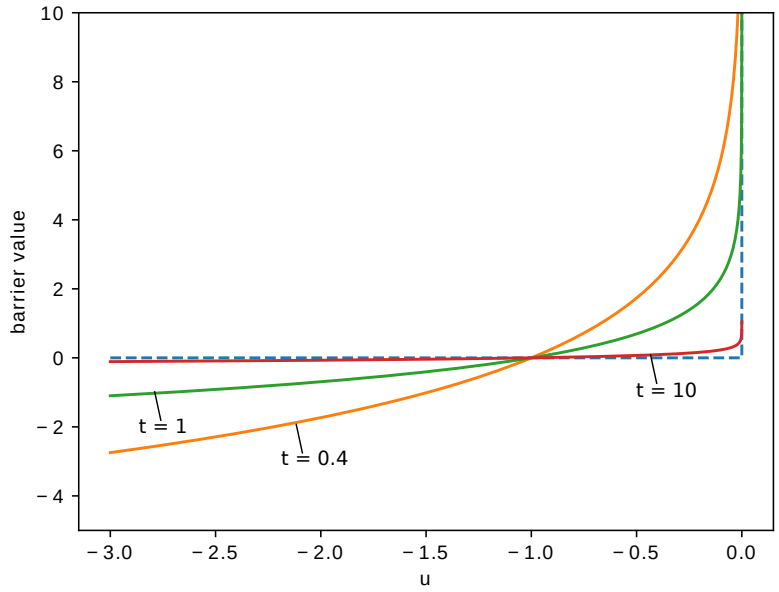

Figure 3.3 Approximate form of the barrier function $\hat{I}_{-}(u)$ compared to the ideal form as a function of t.

3.3 Column generation

Column generation is a method for solving large optimisation problems iteratively without having to consider all possible decision variables at once [4]. The method is particularly useful when heuristic knowledge about the specific optimisation problem leads one to believe that only a small subset of the decision variables are expected to be non-zero. For example, in radiation therapy, the number of possible aperture shapes can quickly exceed 10^{17} even for simple cases [5]. In practice, IMRT photon plans satisfying clinical constraints can be produced with fewer than 200 apertures. Column generation is therefore a good candidate for treatment planning.

When using the column generation method, the original problem that we wish to solve is labelled the *master problem* (MP). Instead of solving the MP, which is assumed to be too computationally demanding, a *restricted master problem* (RMP) is solved multiple times. The RMP is identical to the MP with the exception that only a subset of the decision variables is considered. Each time the RMP is solved, the solution is used to construct a *pricing problem* (PP) which produces a candidate decision variable that can be added to the RMP to further improve the cost function. The PP serves dual purposes: it produces a new decision variable

and a *price*, which is an estimate of the improvement on the cost function from adding the new variable to the RMP. If the price is positive, then the algorithm terminates as no more further decision variables can be added to the RMP to improve the cost function.

Algorithm 3 Outline of the column generation method.

Require: Initial set of optimal decision variables in the RMP, $\hat{\boldsymbol{x}}^*$

1: **loop**
2: Solve the PP
3: **if** $PP \geq 0$ **then return** $\hat{\boldsymbol{x}}^*$
4: **end if**
5: Add a new decision variable to the RMP
6: Solve the RMP
7: **end loop**

Column generation is typically associated with linear optimisation, where each decision variable corresponds to a column in the matrix representation of the standard form of the LP 3.2. However, the method can also be applied to non-linear programs, and a mathematical derivation of column generation applied to radiation therapy treatment planning is provided in chapter 6.

3.4 Simulated annealing

Simulated annealing (SA) is a general optimisation method that is, in principle, able to solve non-convex problems to global optimality. SA also does not require a continuous or differentiable cost function, unlike gradient descent methods. The outline of the algorithm is as follows:

Algorithm 4 Outline of the simulated annealing.

Require: Initial guess for the solution, $\boldsymbol{x}$, cost function $f(\boldsymbol{x})$, max no. of iterations N_{max}

for $I = 1 \rightarrow N_{max}$ **do**
    Sample a new candidate for the solution, $\boldsymbol{x'}$
    if $f(\boldsymbol{x'}) < f(\boldsymbol{x})$ **then**
        Set $\boldsymbol{x} := \boldsymbol{x'}$
    else
        Accept worse solution (set $\boldsymbol{x} := \boldsymbol{x'}$) with probability $P(I)$
    end if
    Lower probability $P(I)$ of accepting a worse solution
end for
return $\boldsymbol{x}$

Accepting a worse solution allows the algorithm to escape local minima and explore the solution space more thoroughly. The major parameters of the SA algorithm are the sampling method, and the *cooling schedule* responsible for lowering the probability of accepting worse solutions. Typically the sampling method will attempt to find a candidate solution in the neighbourhood of the previous one. For example, in radiation therapy, the new candidate solution could be a change in the amount of radiation delivered from an aperture, sampled from a Gaussian distribution centered around the previous value, or a change in the radiation field by shifting the position of an MLC leaf pair [6].

3.5 Optimisation strategies in radiation therapy

In this section, the optimisation theory described in the previous sections is applied to the problem of IMRT planning. IMRT plans have pre-selected number of beam angles, and multiple apertures are delivered from each beam angle. Inverse planning strategies in radiation therapy fall within two broad categories: fluence map optimisation (FMO) and direct aperture optimisation (DAO).

3.5.1 Fluence map optimisation

For FMO techniques, the radiation field at each beam angle is decomposed into small *beamlets*, representing the smallest unit of field size appropriate for the problem. In modern TPS, 5 x 5 mm^2 is a common beamlet size.

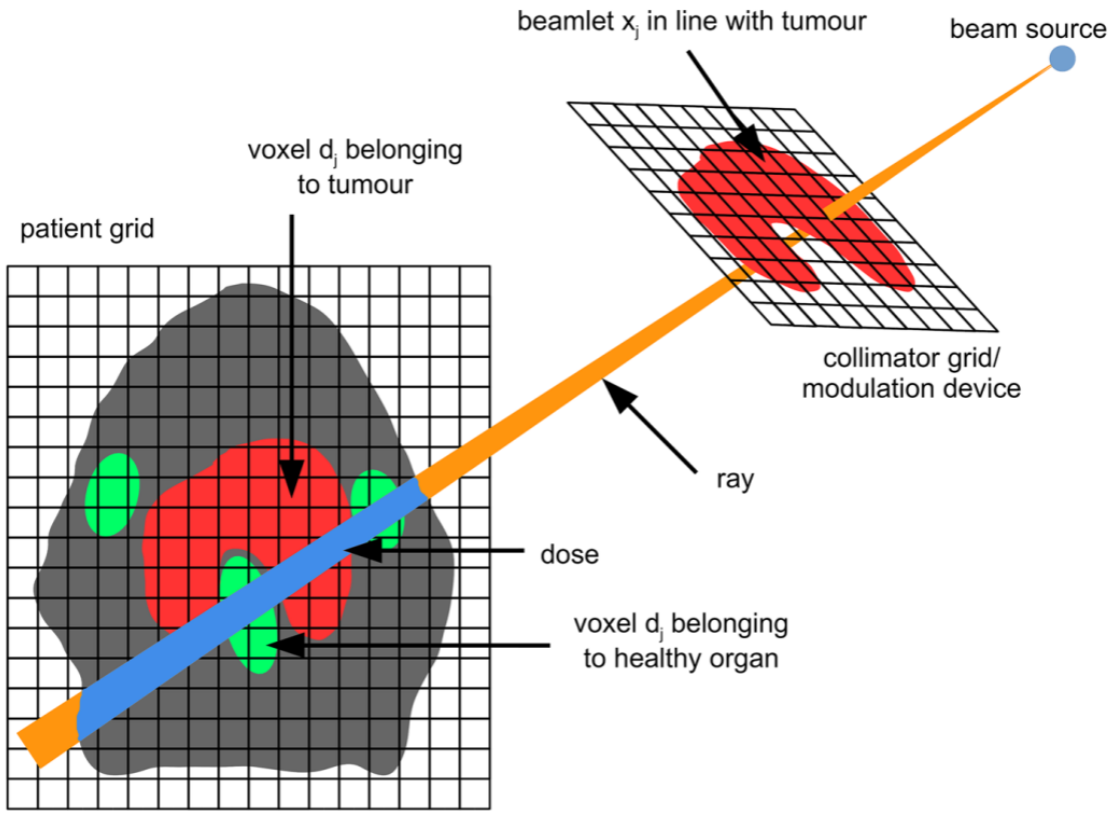

Figure 3.4 Illustration of the fluence map optimisation model. The radiation field is decomposed into beamlets at the collimator plane, and the intensity of each beamlet is optimised independently. Reproduced from Breedveld et al., 2016 [7].

The cost function is then minimised as a function of individual beamlet intensities. The resulting fluence map indicates how much radiation must be delivered from each beamlet, but additional steps must be performed to identify the sequence of MLC apertures needed to deliver the fluence map. To limit plan complexity and treatment time, the number of MLC apertures is often restricted, leading to a degradation between the deliverable plan and the fluence map.

Fluence map optimisation in radiation therapy can be cast as a linear program. Typically, the dose distributions are normalised to represent the dose delivered per MU. Let x_i denote the

weight (or number of MUs) of beamlet i, and D_{ij} the dose contribution per MU from beamlet i to voxel $j \in V$, where V is the voxel representation of the patient. The total dose distribution in the patient is given by:

$$D_j(\boldsymbol{x}) = \sum_i^N D_{ij} x_i, \ j \in V$$

where N is the total number of beamlets over all beam angles.

The quality of a treatment plan is assessed by comparing the total dose distribution to the prescribed doses in the target and organs at risk. We seek to penalise deviations from the prescribed dose limits on a voxel-per-voxel basis. The penalty function for a voxel will have the form

$$F_j(D_j) = \beta_j^U \max\left(0, T_j^U - D_j\right) + \beta_j^O \max\left(0, D_j - T_j^O\right)$$

where β_j^U and β_j^O are the slopes of the linear penalties for under-dosing and over-dosing the voxel, respectively, and T_j^U and T_j^O are the prescribed dose thresholds for the voxel. Often, all voxels belonging to the same structure will have identical penalties applied to them, i.e. $\beta_j^U = \beta_s^U \ \forall j \in V_s$ where V_s is the set of voxels belonging to structure s. For organs at risk, any voxel with dose greater than the prescribed limit is penalised, but doses below the threshold are not ($\beta_s^U = 0$). Targets will typically have both minimum and maximum dose limits to ensure dose coverage but prevent hot spots. Figure 3.5 illustrates the form of the cost function for targets and organs at risk.

It is also possible to specify dose thresholds L_j, U_j that are not allowed to be violated for a treatment plan to be valid. The total cost function is then given by the sum over all voxels of individual voxel-based cost functions, hence before being put into standard form, the optimisation model will have the following form:

$$\begin{aligned}
\underset{\boldsymbol{x}}{\text{minimise}} \quad & \sum_j F_j(D_j) \\
\text{subject to} \quad & D_j(\boldsymbol{x}) = \sum_i^N D_{ij} x_i, j = 1, \ldots, |V|, \\
& D_j \geq L_j, j = 1, \ldots, |V| \\
& D_j \leq U_j, j = 1, \ldots, |V| \\
& x_i \geq 0, i = 1, \ldots, N, \\
& D_j \geq 0, j = 1, \ldots, |V|
\end{aligned}$$

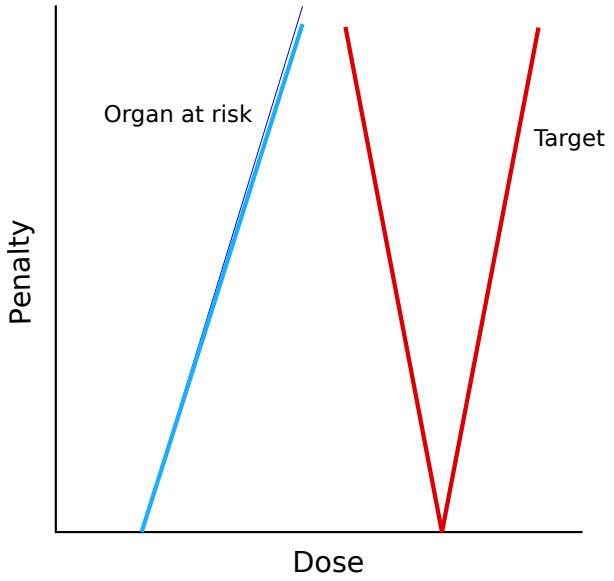

Figure 3.5 Typical linear cost functions for targets and organs at risk.

where $F_j(D_j)$ is a linear function or can be transformed into a linear function.

In principle, the number of decision variables in the problem is equal to the number of beamlets, as the intensity of each beamlet fully determines the dose distribution. However, in standard form, a typical radiation therapy linear problem as shown here will have on the order of 10^5 decision variables and roughly the same order of magnitude of constraints, meaning that the dimensionality of matrix $\boldsymbol{A}$ in eq. 3.2 will be $10^5 \times 10^5$ [8].

In practice, it is desirable to penalise large deviations from the prescribed dose limits more than small deviations, which is not possible with the simple linear programming model presented here. Non-linear cost functions can be approximated by linear segments to produce a LP model, but the number of constraints and decision variables can become too large to be solved efficiently [9]. With a NLP model, cost functions of the form $\max\left(0, T_j^U - D_j\right)^m$ where $m > 1$ can be used, in addition to more complex functions of dose such as tumour control probability (TCP) and equivalent uniform dose (EUD) [5].

3.5.2 Direct aperture optimisation

DAO was first introduced as a problem solved using the simulated annealing method [6]. A preselected number of apertures from all beam angles were initialised such that the apertures corresponded to the projection of the target on the MLC plane. The SA algorithm then adjusted

the intensities and shapes of each aperture according to the algorithm outlined in section 3.4. The mathematical representation of the optimisation problem is similar to the beamlet-based model,

$$\begin{aligned} \underset{\boldsymbol{x}}{\text{minimise}} \quad & \sum_j F_j(D_j) \\ \text{subject to} \quad & D_j(\boldsymbol{x}) = \sum_i^{|K|} D_{kj} x_k, j = 1, \ldots, |V|, \\ & x_i \geq 0, i = 1, \ldots, |K|, \\ & \vdots \end{aligned}$$

Where K is the set of apertures and D_{kj} is the dose contribution to voxel j from aperture k. All the cost function components described in LP and NLP approaches can be used with the SA method, including the non-convex components like TCP. The aperture dose distribution $\boldsymbol{D_k}$ can be calculated using a full dose calculation algorithm, however, the number of aperture changes required to arrive at a solution can be large, and calculating a full aperture dose distribution each time can dominate the running time of the algorithm. Alternatively, the aperture dose can also be constructed from the sum of individual precalculated beamlet dose distributions. However, beamlet-defined aperture dose distributions can differ from the "true" aperture dose distribution as the perturbation of the MLC leaves on the aperture is not modelled.

DAO methods are advantageous in that the output of the optimisation does not need to go through a post-processing step which could degrade the plan, as with FMO. In the case of beamlet-based DAO, a final recalculation step is often performed to properly account for the presence of the collimator for each aperture, but the treatment plans are typically not degraded [10]. One of the drawbacks of the SA method requires a finely tuned sampling method and cooling schedule to ensure good performance, and the running time of the algorithm does not scale favourably with a large number of degrees of freedom [11]. Column generation methods can be used to perform DAO optimisation by forming apertures as part of the pricing problem. In fact, column generation-based DAO is at the core of the optimisation models used in this work, and the algorithm is explained in detail in chapter 6.

References

[1] D. J. Bertsimas and J. J. N. J. Tsitsiklis, *Introduction to Linear Optimization*. Belmont: Athena Scientific, 1997.

[2] D. G. Luenberger and Y. Ye, *Linear and Nonlinear Programming*. New York: Springer, fourth ed., 2016.

[3] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge: Cambridge University Press, 2004.

[4] G. Desaulniers, J. Desrosiers, and M. M. Solomon, *Column generation*. New York: Springer US, 2005.

[5] H. E. Romeijn, J. F. Dempsey, and J. G. Li, "A unifying framework for multi-criteria fluence map optimization models," *Physics in Medicine and Biology*, vol. 49, no. 10, pp. 1991–2013, 2004.

[6] D. M. Shepard, M. A. Earl, X. A. Li, S. Naqvi, and C. Yu, "Direct aperture optimization: A turnkey solution for step-and-shoot IMRT," *Medical Physics*, vol. 29, no. 6, pp. 1007–1018, 2002.

[7] S. Breedveld, B. van den Berg, and B. Heijmen, "An interior-point implementation developed and tuned for radiation therapy treatment planning," *Computational Optimization and Applications*, vol. 68, no. 2, pp. 209–242, 2017.

[8] H. E. Romeijn, R. K. Ahuja, J. F. Dempsey, and A. Kumar, "A New Linear Programming Approach to Radiation Therapy Treatment Planning Problems," *Operations Research*, vol. 54, no. 2, pp. 201–216, 2006.

[9] H. E. Romeijn, R. K. Ahuja, J. F. Dempsey, A. Kumar, and J. G. Li, "A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning," *Physics in Medicine and Biology*, vol. 48, no. 21, pp. 3521–3542, 2003.

[10] D. Henzen, P. Manser, D. Frei, W. Volken, H. Neuenschwander, E. J. Born, a. Joosten, K. Lössl, D. M. Aebersold, C. Chatelain, M. F. M. Stampanoni, and M. K. Fix, "Beamlet

based direct aperture optimization for MERT using a photon MLC.," *Medical physics*, vol. 41, p. 121711, dec 2014.

[11] A. Alexander, E. Soisson, M.-A. Renaud, and J. Seuntjens, "Direct aperture optimization for FLEC-based MERT and its application in mixed beam radiotherapy.," *Medical physics*, vol. 39, pp. 4820–31, aug 2012.

Chapter 4

Review of modulated electron radiation therapy

While IMRT has evolved from a research concept to a major mode of radiation delivery in the past two decades, the electron counterpart, MERT, has lagged behind in clinical application due to the lack of convenient collimation methods and treatment planning algorithms. Obtaining accurate dose distributions in patients for electron beams requires more advanced and time consuming dose calculation methods compared to photon beams [1], and the high plan quality achievable by IMRT has also reduced the clinical pressure to implement MERT techniques.

4.1 Bolus electron conformal therapy

The use of bolus was the first suggested means of modulating electron beams. An electron bolus is a tissue-equivalent material shaped specifically for each patient and placed either in contact with the patient or in close proximity. The electron beam can be made to conform to an irregular target by varying the thickness of the bolus [2]. In addition, partial bolusing was introduced as a way to homogenise the dose near the skin by including a bolus only for a fraction of each daily treatment in order to degrade the energy of the beam and cause more dose to be deposited near the skin [3]. Support for custom electron bolus in treatment planning systems and automated bolus design algorithms were introduced [4–6], but they are not part of routine clinical use due to the added time and logistics burden associated with the technique.

In recent years, interest in Bolus-based MERT has been rekindled due to the advent of 3D printing. Su et al. presented an algorithm using the dose obtained from a commercial MC electron dose calculation engine as an input to automatically design patient-specific bolus that are 3D-printed using solid polylactic acid (PLA). The authors found that 3D-printed bolus improved dose conformity in the target compared to uniform bolus, and concluded that 3D-printed bolus are a practical and low-cost approach to performing MERT [7].

Zou et al. investigated the performance of two 3D-printing fabrication methods, fused deposition modeling (FDM) and selective laser sintering (SLS) and two bolus materials, PLA and polyamide, though they only verified FDM with PLA for electron purposes. The authors found that material inhomogeneities can appear in the bolus due to differences in heat dissipation during the FDM 3D-printing process, however, they estimated that the variations in density result in at most a 1 mm change in the depth of the 90% isodose line of an electron beam [8].

4.2 Modulated electron arc therapy

Electron arc therapy (EAT) was first proposed as an alternative to static photon and electron beams for post-mastectomy chest wall irradiation. At the time, static photon techniques resulted in an unacceptable amount of dose delivered to normal tissue, whereas electron techniques were not able to achieve the required target dose homogeneity for clinical acceptability [9]. Initially, modulation was achieved by changing the beam energy and varying the dose rate during the arc delivery. Field shaping was limited to open jaws and aluminium blocks as secondary collimation [10]. Each patient required a 9-11 kg *breastplate* casted specifically to their anatomy.

An electron-specific MLC was subsequently introduced to further improve dose uniformity by adding more modulation capabilities [11]. The EAT technique was implemented clinically with favourable patient outcomes compared to contemporary photon techniques [12]. Recent additions to the electron arc therapy literature investigated the use of simultaneous gantry and couch motion to allow for the electron applicators to remain as close as possible to the patient without collisions [13].

4.3 MLC-based electron therapy

Research into MLC-based electron therapy has followed two broad categories: incorporation of an electron-specific MLC (eMLC) in the beam path, or collimating electron beams using the

photon MLC (pMLC) typically present in all modern linear accelerators. One of the central topics of MERT research is how the degradation of the dose distribution due to electron scatter from rounded MLC leaf ends and from travel in air between collimation devices and the patient is handled.

4.3.1 Electron-specific MLC

The first eMLC model was proposed by Lee et al. as a tertiary collimation device to be inserted with a 25 x 25 cm^2 applicator to limit the distance between the MLC and the patient skin. The dose profiles produced by eMLC-collimated beams were of the same quality as cutout-defined fields. However, the leaves were controlled manually, preventing dynamic delivery [14]. The feasibility of MERT on this prototype eMLC was also demonstrated by manually delivering various field sizes and energies on film, and comparing the results to MC simulations [15]. The same work also demonstrated that analytical pencil beam algorithms are inadequate for MERT treatment planning dose calculations and recommended MC dose calculation methods.

Hogstrom et al. developed a retractable eMLC with the aim of completely eliminating the need for electron-specific accessories. The eMLC was designed with the goal of allowing three possible vertical positions leaving an increasing air gap between the machine and the patient. The fully retracted position had a source-to-collimator distance (SCD) of 67 cm, which left enough space between the collimator and the patient to perform rotational arc delivery without collision risks. The other two positions were 80 cm SCD for isocentric treatments and 90 cm SCD for traditional electron treatments. For preliminary dosimetric analysis, the authors' initial prototype was fixed at 90 cm SCD. The quality of dose distributions collimated with the retractable eMLC was found to match the quality of dose delivered by cutout-defined fields [16]. However, the MLC leaves were operated manually which did not allow for dynamic deliveries.

Gauer et al. set out to design and develop an eMLC specifically meeting the technical requirements for automated MERT, notably, computer-controlled beam collimation, optimised SCD and minimal total weight. They selected brass for the leaf material to limit the weight of the eMLC and designed the eMLC to allow both 72 and 84 cm SCD [17]. The same authors then characterised a manufactured prototype of their eMLC design with fully motorised MLC leaves. The total weight of their prototype was 30 kg, which was above their stated design goal of 20 kg. The authors measured the field displacement at different gantry angles and attributed 0.3 mm of difference to the weight of their prototype [18].

A third type of MLC was proposed by Al-Yahya et al. as an alternative to the full-featured eMLCs previously introduced. The few-leaf electron collimator (FLEC) consisted of two pairs of trimmer bars to allow arbitrary rectangular fields. It was fitted on a 15 x 15 cm^2 applicator and featured motorised leaves for dynamic deliveries. The prototype weighted less than 2.5 kg, compared to over 8 kg for other prototypes [19]. The high weight of full featured eMLCs was considered a significant downside due to the effects of *gantry sag*, *applicator sag* and the added burden on radiation therapy technicians associated with inserting and removing the eMLC. Connell et al. performed the fully automated delivery of a MERT plan on a solid water phantom using the FLEC, and measured the delivered absorbed dose using film. Comparisons between measurements and simulations showed that an accurate delivery of MERT plans was feasible, with gamma pass rates above 97% for a 3%/3 mm criteria for all modalities except 9 MeV [20].

4.3.2 Photon MLC

The use of a pMLC for electron beam collimation has been investigated in detail due to the significant logistical advantages of having a single collimating device for all modalities. In modern linear accelerators, the photon MLC is installed at an SCD of roughly 50 cm, leaving 30-50 cm of air between the collimation device and the patient in standard isocentric deliveries. However, unlike photon beams, electron beam profiles are degraded significantly by scatter in air, which compromises the quality of pMLC-collimated electron plans. Early studies when pMLCs were introduced showed that treatment SSDs of 70 cm were necessary to maintain an acceptable penumbra for electron beams [21]. Lee et al., in the same work introducing their prototype eMLC, assessed the performance of pMLC-collimated beams and concluded that they are inferior to eMLC beams even at reduced SSDs of 80 cm. They also concluded that replacing the atmosphere in the beam path with helium to reduce air scatter did not improve the profiles enough to justify the modifications needed to realise the setup in practice [14]. The same conclusions about pMLC-collimated beams requiring the shortest SSD possible were reproduced in later studies [22, 23]. For reference, figure 4.1 demonstrates the effect of SSD on the penumbra of a 9 x 9 cm^2 profile.

Salguero et al. investigated pMLC-collimated electron treatments for four post-mastectomized chest wall cases using a Siemens PRIMUS linear accelerator which allows for shortened SSDs as low as 60 cm. The authors compared static block-collimated electrons to static pMLC-collimated electrons and found that the target homogeneity was poorer for the pMLC-collimated

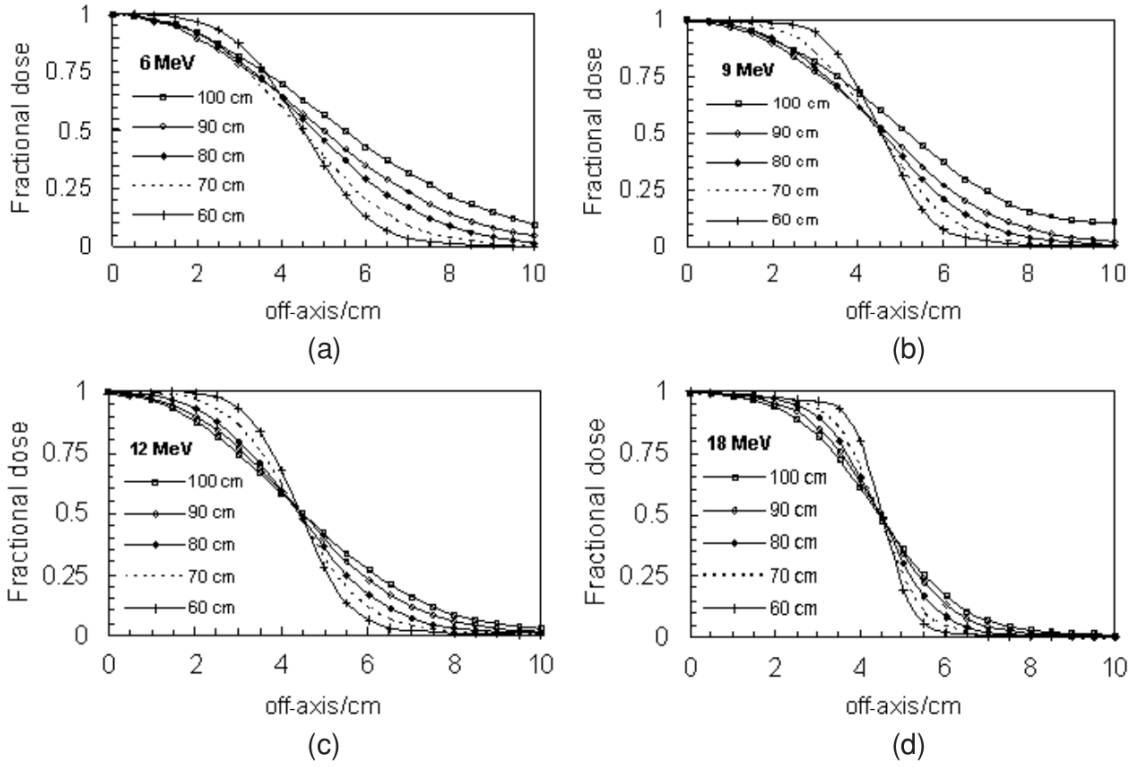

Figure 4.1 Effect of varying the SSD on a 9 x 9 cm^2 pMLC-collimated electron field for 6, 9, 12, 18 MeV beams. Reproduced from Duplessis et al. [22]

cases, though the coverage is similar. In the same work, the authors also investigated MERT deliveries from a single beam angle using 6, 9 and 15 MeV electrons and concluded that pMLC-collimated MERT yields superior target coverage and lower doses to normal tissue compared to static blocks [24]. Extending their work other cancer sites, the authors then investigated the quality of pMLC-collimated MERT plans for head & neck cases. Once again, they concluded that MERT was able to achieve clinically acceptable target homogeneity. The MERT plans also achieved superior normal tissue sparing compared to IMRT plans for the same cases [25].

More recently, Mueller et al. argued that simulated pMLC-collimated electron plans on realistic patient cases can remain clinically acceptable at SSDs above 70 cm when using DVH parameters to judge acceptability. However, the same degradation of the penumbra was observed as with previous investigations, leading to higher normal tissue doses laterally from the target compared to cutout-collimated electron beams. The authors also show that the addition of a

photon component can mitigate the degradation of the penumbra at higher SSDs as electrons make up a smaller portion of the total dose delivered [26].

4.4 Mixed beam radiation therapy

MLC-collimated mixed electron-photon treatments have been proposed shortly after the introduction of the photon MLC. Early experiments were performed using the MM50 racetrack microtron which featured a treatment head filled with helium, leading to less electron beam degradation [27]. Jansson et al. investigated MBRT for loco-regional radiotherapy after breast cancer surgery based on a planning study where 12 out of 30 patients were found to receive significantly lower doses to normal tissue from MBRT compared to the standard parallel-opposed photon treatment. The authors used static pMLC-defined fields for both photons and electrons [28]. Li et al. also concluded that MBRT could provide superior OAR sparing compared to 9-field IMRT in a planning study involving two patients with early stage breast cancer [29].

In more recent years, Miguez et al. implemented an MBRT technique combining MERT and IMRT for accelerated partial breast irradiation on seven patients and reported no incidence of moderate to severe grade of toxicity. A single electron energy was used in five out of the seven patients, and two energies were used for the remaining two patients. The photon component for the majority of patients was a simple parallel-opposed arrangement with approximately seven fields [30–32]. Finally, Mueller et al. showed a planning study where MBRT outperformed MERT, IMRT and VMAT on chest wall irradiation and squamous cell carcinoma clinical cases by delivering lower OAR doses without sacrificing target coverage [33].

4.5 Treatment planning

Compared to photon treatment planning, where a single energy is sufficient for the vast majority of treatments, MLC-based MERT planning typically involves identifying high quality apertures for multiple electron energies per beam angle. For a standard modern linear accelerator with, for example, 6 electron energies, even a simple 3 beam angle delivery can mean simultaneously performing 18 fluence map optimisations. Nonetheless, advances in computing power have made MERT and MBRT planning achievable within reasonable computing times.

4.5.1 Modulated electron radiation therapy inverse planning

Early treatment planning efforts for MLC-based MERT involved forward planning of preselected MLC aperture shapes [34, 35]. However, early in the development of MERT, Hyodynmaa et al. introduced a gradient descent-based optimisation algorithm to optimise the fluence of scanned electron pencil beams delivered by a racetrack microtron. While the dose calculation algorithm employed did not take heterogeneities or lateral scatter into account, the authors showed that MERT could replace bolus in controlling the depth distribution of the overall electron treatment [36].

Lee et al. introduced a MC beamlet-based gradient descent optimisation algorithm involving fluence map optimisation of small beamlets followed by leaf sequencing for an eMLC. The final apertures were then simulated to account for leaf end effects on the aperture dose distributions, and the aperture intensities were re-optimised. Comparisons of the final treatment plans between beamlet-defined apertures and fully simulated apertures confirmed the need for full aperture recalculation and weight re-optimisation [37]. Salguero et al. also performed beamlet-based gradient descent optimisation for MERT delivered using the pMLC on a Siemens Primus linear accelerator. The authors performed radiochromic film measurements for 4 plans delivered on a solid water phantom and obtained excellent gamma pass rates with a criteria of 3%/3 mm [24].

Al-Yahya et al. developed a simulated annealing-based inverse planning algorithm for the FLEC which optimised the intensities of preselected MC-calculated FLEC *fieldlet* dose distributions [35]. Unlike beamlet-based optimisation, the fieldlet doses did not require a recalculation step as they already included the effects of the collimator. Alexander et al. expanded the optimisation model to remove the need for aperture preselection by precalculating a large number of FLEC-defined fieldlets and leaving the simulated annealing algorithm free to determine which fieldlets to include in the final treatment plan [38].

Henzen et al. studied MC beamlet-based MERT inverse planning delivered with a photon MLC. The authors used a DAO simulated annealing algorithm to optimise a user-specified number of apertures. At each iteration, the algorithm randomly selected either an aperture weight adjustment or an aperture shape adjustment. After optimisation, the apertures were recalculated to take into account the MLC leaf perturbation on the dose, and the aperture weights reoptimised in a similar fashion as Lee et al. The authors found that the final *deliverable* dose distribution obtained from recalculated apertures did not degrade significantly compared the *optimised* dose distribution obtained from beamlet-defined apertures, meaning that beamlet-based DAO is adequate for MERT [39].

4.5.2 Mixed beam radiation therapy inverse planning

In many cases, MBRT studies did not simultaneously optimise both the electron and photon components together, instead having elements of forward planning or producing a plan for one modality and then using it as an input in the creation of the other plan components [40, 41, 35]. For example, Alexander et al. first performed simulated annealing DAO MERT optimisation to obtain a set of electron apertures for the FLEC, and then incorporated a photon fluence map optimisation step [42].

To produce plans for their accelerated partial breast irradiation MBRT technique described above, Palma et al. extended their MERT beamlet-based gradient descent optimisation model to also optimise photon fluence maps and thus perform truly simultaneous MBRT optimisation. The resulting plans contained at most 17 apertures spread among all modalities [30].

Finally, during the course of performing the work presented in this thesis related to MBRT optimisation, Mueller et al. independently introduced a simultaneous MBRT optimisation algorithm based on the beamlet-based DAO algorithm from Henzen et al. [39] to produce plans with as many as 60 apertures spread across all modalities [33]. The authors applied their algorithm to an academic situation, a clinical chest wall irradiation case and a squamous cell carcinoma case. They also created IMRT and MERT plans for each case using the same algorithm. For all cases, the MBRT plans had the lowest cost function, and as expected, the MBRT plans were able to significantly reduce the overall dose to normal tissue without sacrificing target coverage.

4.6 Conclusions

At the beginning of this work, there were no published algorithms that could simultaneously optimise electron and photon components in such a way that the results are no worse than state of the art plans in any single modality. Attempts at simultaneous MBRT planning typically involved simplified deliveries, for example considering only two photon beam delivery angles, or a limited number of electron apertures.

In addition, much of the work already published had been performed on linear accelerators which are no longer available and had physical characteristics that are not achievable on modern linacs, such as SSDs of 60 cm, or were using tertiary electron-specific collimators that would be seen as too cumbersome to implement in routine clinical use.

Finally, none of the published work addressed the issue of dose distribution robustness to positioning uncertainties. The intensity modulated proton literature showed that robust optimisation was essential to ensure that target coverage is preserved under realistic patient setup uncertainties and range uncertainties [43]. Unlike photons, charged particles are constantly scattering and losing energy as they travel within a medium. Given that electrons are also charged particles, electron dose distributions are more perturbed by changes in effective depth induced by setup errors, compared to photons. The robustness of MBRT plans should therefore be investigated, and an MBRT optimisation algorithm should be capable of producing robust plans. Most published methods for robust treatment planning require a dramatic increase in computational burden compared to non-robust approaches [44], meaning that simple extensions of previously published MBRT methods would most likely prove inadequate as they are already suffering from excessive computation times for simple, non-robust plans.

References

[1] J. E. Cygler, G. M. Daskalov, G. H. Chan, and G. X. Ding, "Evaluation of the first commercial Monte Carlo dose calculation engine for electron beam treatment planning," *Medical Physics*, vol. 31, no. 1, pp. 142–153, 2004.

[2] D. A. Low, G. Starkschall, S. W. Buinowski, L. L. Wang, and K. R. Hogstrom, "Electron Bolus Design for Radiotherapy Treatment Planning: Bolus Design Algorithms," *Medical Physics*, vol. 19, no. 1, pp. 115–124, 1992.

[3] D. M. Galbraith and J. A. Rawlinson, "Partial bolussing to improve the depth doses in the surface region of low energy electron beams," *International Journal of Radiation Oncology, Biology, Physics*, vol. 10, no. 2, pp. 313–317, 1984.

[4] D. A. Low, G. Starkschall, N. E. Sherman, S. W. Bujnowski, J. R. Ewton, and K. R. Hogstrom, "Computer-aided design and fabrication of an electron bolus for treatment of the paraspinal muscles," *International Journal of Radiation Oncology, Biology, Physics*, vol. 33, no. 5, pp. 1127–1138, 1995.

[5] G. H. Perkins, M. D. McNeese, J. A. Antolak, T. A. Buchholz, E. A. Strom, and K. R. Hogstrom, "A custom three-dimensional electron bolus technique for optimization of post-

mastectomy irradiation," *International Journal of Radiation Oncology Biology Physics*, vol. 51, no. 4, pp. 1142–1151, 2001.

[6] R. J. Kudchadker, K. R. Hogstrom, A. S. Garden, M. D. McNeese, R. A. Boyd, and J. A. Antolak, "Electron conformal radiotherapy using bolus and intensity modulation," *International Journal of Radiation Oncology Biology Physics*, vol. 53, no. 4, pp. 1023–1037, 2002.

[7] S. Su, K. Moran, and J. L. Robar, "Design and production of 3D printed bolus for electron radiation therapy," *Journal of Applied Clinical Medical Physics*, vol. 15, no. 4, pp. 194–211, 2014.

[8] W. Zou, T. Fisher, M. Zhang, L. Kim, T. Chen, V. Narra, B. Swann, R. Singh, R. Siderit, L. Yin, B. K. K. Teo, M. Mckenna, J. McDonough, and Y. J. Ning, "Potential of 3D printing technologies for fabrication of electron bolus and proton compensators," *Journal of Applied Clinical Medical Physics*, vol. 16, no. 3, pp. 90–98, 2015.

[9] L. M. Peacock, D. D. Leavitt, F. A. Gibbs, and J. R. Stewart, "Electron arc therapy: Clinical experience with chest and irradiation," *International Journal of Radiation Oncology, Biology, Physics*, vol. 10, no. 11, pp. 2149–2153, 1984.

[10] D. D. Leavitt, L. M. Peacock, F. A. Gibbs, and J. R. Stewart, "Electron arc therapy: Physical measurement and treatment planning techniques," *International Journal of Radiation Oncology, Biology, Physics*, vol. 11, no. 5, pp. 987–999, 1985.

[11] D. D. Leavitt, J. R. Stewart, J. H. Moeller, W. L. Lee, and G. A. Takach, "Electron arc therapy: design, implementation and evaluation of a dynamic multi-vane collimator system," *International Journal of Radiation Oncology, Biology, Physics*, vol. 17, no. 5, pp. 1089–1094, 1989.

[12] D. K. Gaffney, J. Prows, D. D. Leavitt, M. J. Egger, J. G. Morgan, and J. R. Stewart, "Electron arc irradiation of the postmastectomy chest wall: Clinical results," *Radiotherapy and Oncology*, vol. 42, no. 1, pp. 17–24, 1997.

[13] A. Rodrigues, F.-f. Yin, and Q. Wu, "Dynamic electron arc radiotherapy (DEAR): a feasibility study," *Physics in medicine and biology*, vol. 327, no. 2, p. 327, 2013.

[14] M. C. Lee, S. B. Jiang, and C. M. Ma, "Monte Carlo and experimental investigations of multileaf collimated electron beams for modulated electron radiation therapy," *Medical Physics*, vol. 27, no. 12, pp. 2708–2718, 2000.

[15] C. M. Ma, T. Pawlicki, M. C. Lee, S. B. Jiang, J. S. Li, J. Deng, B. Yi, E. Mok, and A. L. Boyer, "Energy- and intensity-modulated electron beams for radiotherapy," *Physics in Medicine and Biology*, vol. 45, no. 8, pp. 2293–2311, 2000.

[16] K. R. Hogstrom, R. A. Boyd, J. A. Antolak, M. M. Svatos, B. A. Faddegon, and J. G. Rosenman, "Dosimetry of a prototype retractable eMLC for fixed-beam electron therapy," *Medical Physics*, vol. 31, no. 3, pp. 443–462, 2004.

[17] T. Gauer, D. Albers, F. Cremers, R. Harmansa, R. Pellegrini, and R. Schmidt, "Design of a computer-controlled multileaf collimator for advanced electron radiotherapy," *Physics in Medicine and Biology*, vol. 51, no. 23, pp. 5987–6003, 2006.

[18] T. Gauer, J. Sokoll, F. Cremers, R. Harmansa, M. Luzzara, and R. Schmidt, "Characterization of an add-on multileaf collimator for electron beam therapy," *Physics in Medicine and Biology*, vol. 53, no. 4, pp. 1071–1085, 2008.

[19] K. Al-Yahya, F. Verhaegen, and J. Seuntjens, "Design and dosimetry of a few leaf electron collimator for energy modulated electron therapy," *Medical Physics*, vol. 34, no. 12, pp. 4782–4791, 2007.

[20] T. Connell, A. Alexander, P. Papaconstadopoulos, M. Serban, S. Devic, and J. Seuntjens, "Delivery validation of an automated modulated electron radiotherapy plan," *Medical Physics*, vol. 41, no. 6, 2014.

[21] E. E. Klein, Z. Li, and D. A. Low, "Feasibility study of multileaf collimated electrons with a scattering foil based accelerator," *Radiotherapy and Oncology*, vol. 41, no. 2, pp. 189–196, 1996.

[22] F. C. Du Plessis, A. Leal, S. Stathakis, W. Xiong, and C. M. Ma, "Characterization of megavoltage electron beams delivered through a photon multi-leaf collimator (pMLC)," *Physics in Medicine and Biology*, vol. 51, no. 8, pp. 2113–2129, 2006.

[23] E. E. Klein, M. Vicic, C. M. Ma, D. A. Low, and R. E. Drzymala, "Validation of calculations for electrons modulated with conventional photon multileaf collimators," *Physics in Medicine and Biology*, vol. 53, no. 5, pp. 1183–1208, 2008.

[24] F. J. Salguero, B. Palma, R. Arrans, J. Rosello, and A. Leal, "Modulated electron radiotherapy treatment planning using a photon multileaf collimator for post-mastectomized chest walls," *Radiotherapy and Oncology*, vol. 93, no. 3, pp. 625–632, 2009.

[25] F. J. Salguero, R. Arrns, B. A. Palma, and A. Leal, "Intensity- and energy-modulated electron radiotherapy by means of an xMLC for head and neck shallow tumors," *Physics in Medicine and Biology*, vol. 55, no. 5, pp. 1413–1427, 2010.

[26] S. Mueller, M. K. Fix, D. Henzen, D. Frei, D. Frauchiger, K. Loessl, M. F. M. Stampanoni, and P. Manser, "Electron beam collimation with a photon MLC for standard electron treatments," *Physics in Medicine & Biology*, vol. 63, no. 2, p. 025017, 2018.

[27] M. Karlsson and B. Zackrisson, "Matching of electron and photon beams with a multi-leaf collimator," *Radiotherapy and Oncology*, vol. 29, no. 3, pp. 317–326, 1993.

[28] T. Jansson, H. Lindman, K. Nygård, C. V. Dahlgren, A. Montelius, C. Öberg-Kreuger, S. Asplund, and J. Bergh, "Radiotherapy of breast cancer after breast-conserving surgery: An improved technique using mixed electron-photon beams with a multileaf collimator," *Radiotherapy and Oncology*, vol. 46, no. 1, pp. 83–89, 1998.

[29] J. G. Li, S. S. Williams, D. R. Goffinet, A. L. Boyer, and L. Xing, "Breast-conserving radiation therapy using combined electron and intensity-modulated radiotherapy technique," *Radiotherapy and Oncology*, vol. 56, no. 1, pp. 65–71, 2000.

[30] B. A. Palma, A. U. Sánchez, F. J. Salguero, R. Arráns, C. M. Sánchez, A. W. Zurita, M. I. R. Hermida, and A. Leal, "Combined modulated electron and photon beams planned by a Monte-Carlo-based optimization procedure for accelerated partial breast irradiation," *Physics in Medicine and Biology*, vol. 57, no. 5, pp. 1191–1202, 2012.

[31] E. Jimenez-Ortega, C. Miguez-Sanchez, B. Palma, A. Ureba, H. Miras, R. Arrans, A. Barbeiro, J. Baeza, F. Carrasco, and A. L. Plaza, "SU-E-T-593: Outcomes and Toxicities From a Clinical Trial of APBI Using MERT+IMRT with the Same XMLC," *Medical Physics*, vol. 42, no. 6, pp. 3472–3472, 2015.

[32] C. Míguez, E. Jiménez-Ortega, B. A. Palma, H. Miras, A. Ureba, R. Arráns, F. Carrasco-Peña, A. Illescas-Vacas, and A. Leal, "Clinical implementation of combined modulated electron and photon beams with conventional MLC for accelerated partial breast irradiation," *Radiotherapy and Oncology*, vol. 124, no. 1, pp. 124–129, 2017.

[33] S. Mueller, M. K. Fix, A. Joosten, D. Henzen, D. Frei, W. Volken, R. Kueng, D. M. Aebersold, M. F. Stampanoni, and P. Manser, "Simultaneous optimization of photons and electrons for mixed beam radiotherapy," *Physics in Medicine and Biology*, vol. 62, no. 14, pp. 5840–5860, 2017.

[34] E. E. Klein, M. Mamalui-Hunter, and D. A. Low, "Delivery of modulated electron beams with conventional photon multi-leaf collimators," *Physics in Medicine and Biology*, vol. 54, no. 2, pp. 327–339, 2009.

[35] K. Al-Yahya, D. Hristov, F. Verhaegen, and J. Seuntjens, "Monte Carlo based modulated electron beam treatment planning using a few-leaf electron collimator - Feasibility study," in *Physics in Medicine and Biology*, vol. 50, pp. 847–857, 2005.

[36] S. Hyödynmaa, A. Gustafsson, and A. Brahme, "Optimization of conformal electron beam therapy using energy- and fluence- modulated beams," *Medical Physics*, vol. 23, no. 5, pp. 659–666, 1996.

[37] M. C. Lee, J. Deng, J. Li, S. B. Jiang, and C. M. Ma, "Monte Carlo based treatment planning for modulated electron beam radiation therapy," *Physics in Medicine and Biology*, vol. 46, no. 8, pp. 2177–2199, 2001.

[38] A. Alexander, F. Deblois, and J. Seuntjens, "Toward automatic field selection and planning using Monte Carlo-based direct aperture optimization in modulated electron radiotherapy," *Physics in Medicine and Biology*, vol. 55, no. 16, pp. 4563–4576, 2010.

[39] D. Henzen, P. Manser, D. Frei, W. Volken, H. Neuenschwander, E. J. Born, a. Joosten, K. Lössl, D. M. Aebersold, C. Chatelain, M. F. M. Stampanoni, and M. K. Fix, "Beamlet based direct aperture optimization for MERT using a photon MLC.," *Medical physics*, vol. 41, p. 121711, dec 2014.

[40] E. W. Korevaar, B. J. M. Heijmen, E. Woudstra, H. Huizenga, and A. Brahme, "Mixing intensity modulated electron and photon beams: Combining a steep dose fall-off at depth

with sharp and depth-independent penumbras and flat beam profiles," *Physics in Medicine and Biology*, vol. 44, no. 9, pp. 2171–2181, 1999.

[41] X. Mu, L. Olofsson, M. Karlsson, R. Sjögren, and B. Zackrisson, "Can photon IMRT be improved by combination with mixed electron and photon techniques?," *Acta Oncologica*, vol. 43, no. 8, pp. 727–735, 2004.

[42] A. Alexander, E. Soisson, M.-A. Renaud, and J. Seuntjens, "Direct aperture optimization for FLEC-based MERT and its application in mixed beam radiotherapy.," *Medical physics*, vol. 39, pp. 4820–31, aug 2012.

[43] W. Liu, X. Zhang, Y. Li, and R. Mohan, "Robust optimization of intensity modulated proton therapy," *Medical Physics*, vol. 39, no. 2, pp. 1079–1091, 2012.

[44] J. Unkelbach and H. Paganetti, "Robust Proton Treatment Planning: Physical and Biological Optimization," *Seminars in Radiation Oncology*, vol. 28, no. 2, pp. 88–96, 2018.

Chapter 5

Development of a framework for the creation of MBRT treatment plans

As described in chapter 1, treatment planning requires a large number of non-trivial inputs in order to produce a customised treatment plan for every patient. Photon-MLC based electron deliveries are not performed as part of mainstream clinical practice, therefore clinical treatment planning systems do not offer the ability to calculate dose distributions needed to perform MBRT planning. This chapter describes implementation details for the treatment planning system (TPS) components that had to be created prior to studying MBRT optimisation algorithms.

5.1 Radify

The first component required for treatment planning is a graphical user interface (GUI) to visualise patient anatomical images and analyse dose distributions. Instead of working with a clinical treatment planning system using an API, for reasons of flexibility, we decided to create a web-based academic treatment planning system called *Radify*.

Radify is made up of two separate parts, the front-end and the back-end. The front-end is the graphical interface that is accessed from the user's web browser, as shown in figure 5.1. It contains all the logic for visualising radiation therapy data such as CT images, contoured structures, dose distributions and statistics. The front-end is coded using the AngularJS 1.7 JavaScript framework, and data is shown using the WebGL browser extension which allows the user's graphics hardware to be used to render images on the web application.

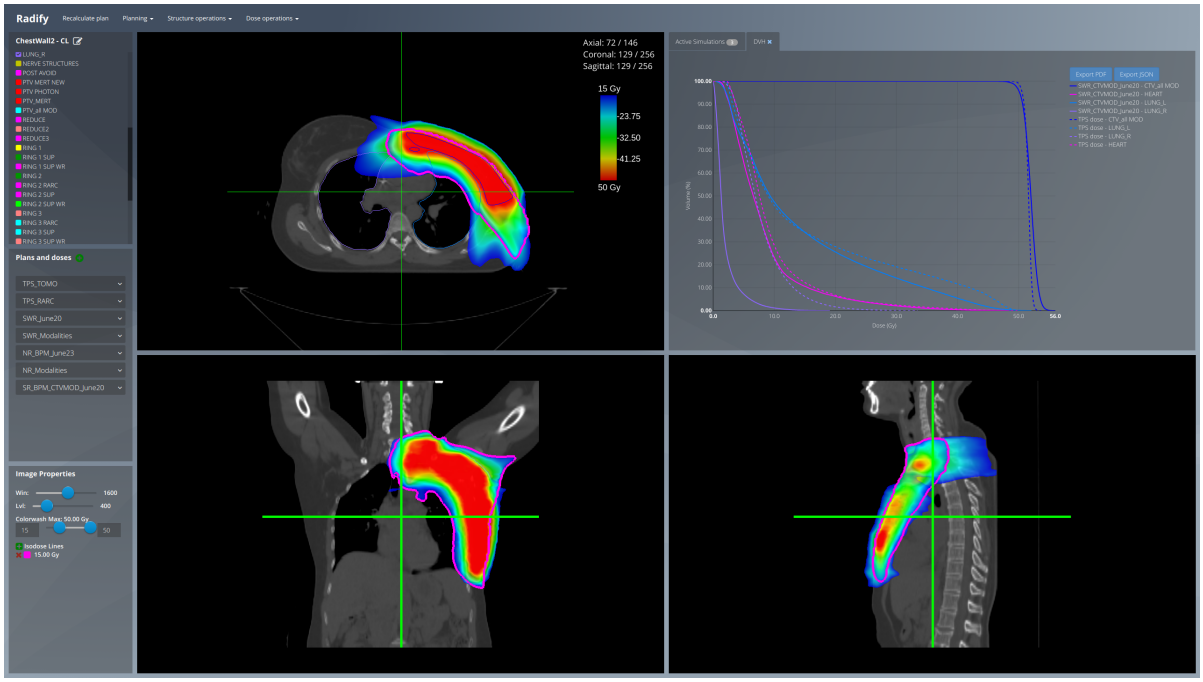

Figure 5.1 Example of patient data being visualised using Radify. Dose colourwash and isodose lines are shown as an example of the visualisation options, along with DVH curves for selected contoured structures comparing two different plans.

The back-end includes the web server, database and data processing libraries. It is based on the Django python framework, in which the patient database is defined along with user access rights. An application program interface (API) is programmed to perform data processing tasks and return the results in JSON format, which can then be interpreted by the front-end for display. The data processing library relies heavily on the pydicom module to interpret and process DICOM data [1]. The API is able to perform the following data processing tasks on patient data, in addition to submitting and managing dose calculation simulations:

- Dose arithmetic (summing dose distributions, percent difference) supporting mismatched grid sizes through trilinear interpolation
- Dose statistics such as dose-volume histograms, profiles, min-max-mean doses inside structures
- Structure boolean, growth, shrinking and translation operations

- Plan transposition (transferring a plan from one patient to another, typically for quality assurance purposes, also known as hybrid plan)
- DICOM exports of plans and doses for reimportation into commercial planning systems
- XML plan export to deliver plans using the TrueBeam developer mode

While these actions are typically performed from the GUI, the back-end operates independently from the front-end, and the API also responds to any properly formatted request. This allows data processing tasks to be performed on a large number of patients simultaneously using user-made scripts.

5.1.1 Workflow for plan creation

Patient CT images and contoured structures are first exported from a clinical treatment planning system and imported into Radify. Within Radify, dose distributions can be compared both visually and via common data analysis metrics such as dose volume histograms (DVH).

Before creating beamlets for optimisation, the planner must specify beam angles as shown in figure 5.2. With beam angles defined, beamlets can be created by choosing from a list of dose calculation engines. Photon beamlet generation is performed using a collapsed cone convolution superposition (CCCS) engine, described in section 5.2.1. Electron beamlets are created using EGSnrc with in-house commissioned electron beam models for 6, 9, 12, 16 and 20 MeV electron beams. Rather than create a specialised platform for MBRT, Radify was programmed to accommodate multiple dose calculation engines and optimisation algorithms. Currently, the TPS can be used to calculate not only beamlet dose distributions for photon and electron beams, but also brachytherapy dwell position dose distributions.

In addition to the choice of dose calculation engine, the major parameters in the beamlet creation process are related to phantom creation. To simulate beamlet dose distributions, the patient CT data must be processed into a simulation *phantom*. Firstly, a voxel size appropriate for the cancer site must be chosen. The standard size is 2.5 cm^3, but for small tumours such as brain metastases, a voxel size of 1 cm^3 or less is sometimes necessary. The CT Houndsfield units are then resampled onto this grid and transformed into mass densities using a piecewise linear calibration curve. The curve is obtained as part of the CT machine commissioning process.

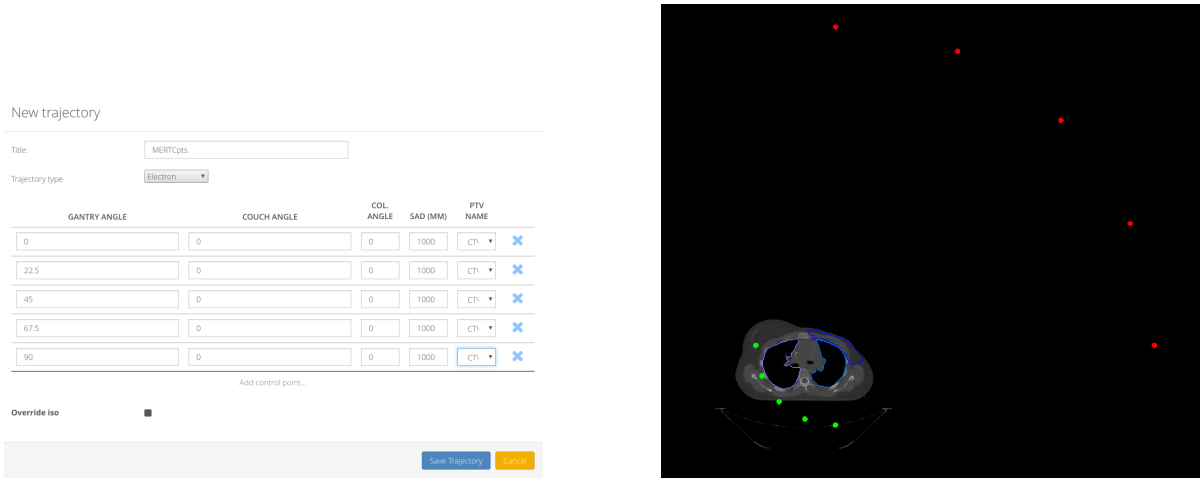

(a) (b)

Figure 5.2 (a) Beam angle selection for MBRT. Gantry, couch and collimator angles are specified. The contoured structure corresponding to the target is also required during beamlet creation. (b) Visualisation of the beam source (red dots) and machine isocenter (green dots) for a set of beam angles.

MC simulations require both mass density and atomic composition information for each voxel. Radify therefore allows planners to assign atomic compositions to ranges of Hounsfield units, or to assign a bulk material to all voxels contained within a contoured structure. Figure 5.3 shows the interface where this information is entered, along with the visualisation options to confirm that the appropriate materials and densities are assigned to the voxels.

Radify is programmed to submit beamlet dose calculation *jobs* to any number of user-defined computation servers. Hundreds of thousands of individual beamlet dose distributions can be generated for a single patient, therefore Radify also continuously tracks the progress of the overall beamlet creation process and notifies the planner when optimisation can be performed.

Finally, once beamlets are generated for all modalities in the desired treatment plan, optimisation can be performed to produce treatment plans. From within Radify, the planner must first specify dose constraints to any contoured structure, as shown in figure 5.4 and then submit the optimisation task to one of the configured computation servers. Planners can assign both *dose* objectives, where any voxel above or below the threshold dose is penalised, and also *dose volume* objectives which allow a fraction of voxels to violate the objective before penalties are applied in the cost function. For voxels belonging to overlapping contoured structures, an

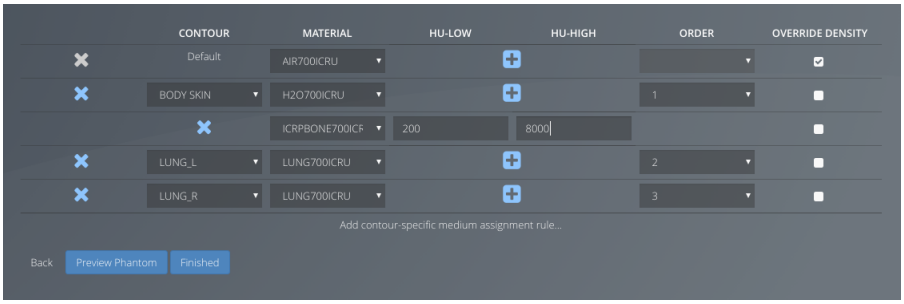

(a)

(b)

Figure 5.3 (a) Atomic composition inputs for phantom creation. (b) Visualisation of the mass densities and materials assigned to voxels within a simulation phantom. The colours indicate the material assigned to the voxels. Clicking on a voxel provides density and material information.

importance value must be specified to define which structures they will be assigned to, as voxels can only belong to a single structure.

Planners can monitor the progress of the optimisation process by observing how the cost function and dose volume histograms evolve with every iteration. This information can then be used to further refine the planning objectives in subsequent optimisation submissions or decide on a maximum number of apertures to include in the final treatment plan to limit complexity.

Following the creation of a treatment plan from the optimisation engine, each aperture dose distribution is recalculated using a full BEAMnrc MC beam model with the dose values normalised to Gy/MU. The shape of each aperture is then held fixed while the weights are reoptimised to obtain the final dose distribution along with monitor units for delivery.

In addition to being used for treatment planning purposes, Radify can also be used to recalculate treatment plans exported from clinical planning systems using in-house commissioned Monte Carlo beam models. The latter feature was used to perform large scale MC

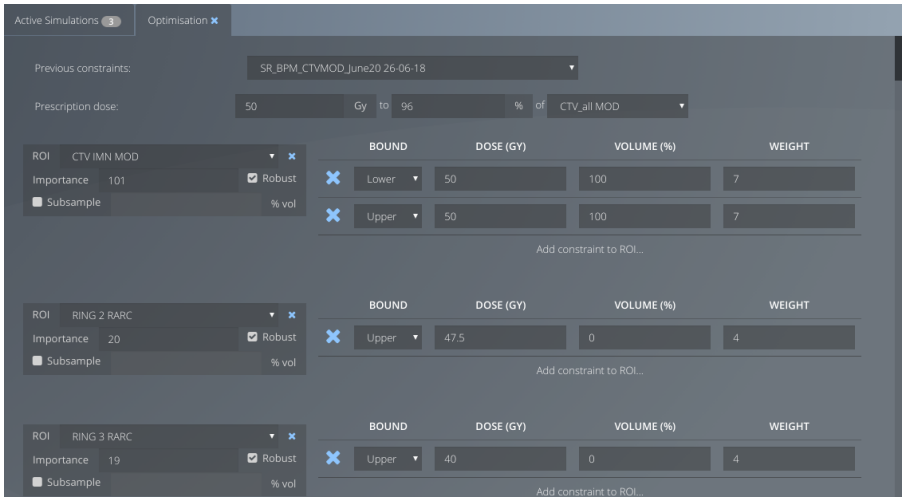

Figure 5.4 Definition of treatment planning objectives within Radify. For each contoured structure (ROI), upper or lower bounds on dose values for voxels inside the structure are specified as dose volume constraints. For example, a dose value of 20 Gy coupled with a volume of 30% will instruct the optimiser to attempt to ensure that at most 30% of the voxels inside the structure can have doses above 20 Gy.

recalculations of lung stereotactic body radiation therapy (SBRT) treatment plans to study the incidence of distant metastases as a function of the dose outside the PTV. The MC treatment plan recalculations served to confirm that the dosimetric characteristics observed in the dose distributions from the clinical treatment planning system were accurate [2]. Similar studies are ongoing using Radify in conjunction with a CyberKnife M6 beam model for noncoplanar treatment plans (Veng Jean Heng, personal communication).

5.2 Dose calculation engines for MBRT

5.2.1 In-house collapsed cone convolution superposition

A CCCS algorithm was developed for the purposes of generating photon beamlets needed as part of the mixed-beam optimization process. Fundamentally, the algorithm calculates dose distributions in voxelised patient geometries by transporting the energy released by photons in each voxel using energy deposition kernels (EDKs) which indicate how much energy is deposited at a point $\boldsymbol{r}$ in the phantom given energy released at another point $\boldsymbol{s}$, as described in

chapter 2. This section describes implementation choices for the purposes of photon beamlet dose calculations.

The algorithm was implemented based on the work of Chen et al. and programmed to run on graphical processing units (GPU) for added speed [3, 4]. As the engine was specifically meant to produce beamlet dose distributions, it used a simplified single-source beam model, as opposed to the more accurate dual source models which are able to accurately reproduce the penumbra of dose profiles [5].

To create the beam models, the energy and fluence spectra for a 6 MV photon beam were binned from a phase space file accumulated above the jaws of an in-house commissioned MC beam model for the Varian TrueBeam linear accelerator. To simplify the TERMA calculation, a lookup table for energy fluence attenuation corrected for beam hardening was pre-calculated,

$$A(\hat{d}) = \sum_n \varphi_0(E_n) \frac{\mu(E_n)}{\rho} e^{-(\mu/\rho)\hat{d}} \delta E \tag{5.1}$$

where $\varphi_0(E_n)$ is the fluence of component for the E_n bin at the reference plane, and $\hat{d}$ is the radiological distance. This allows the TERMA values to be calculated by simply determining the radiological distance in the patient and looking up the value of $A(\hat{d})$,

$$T(u,v,d) = f(u,v)A(\hat{d})a(d) \tag{5.2}$$

where $T(u,v,d)$ is the TERMA value at point (u,v,d) in the beam's eye view (BEV) coordinate system defined by coordinates (u,v) along the reference plane perpendicular to the beam direction and depth d in the beam direction. The fluence at point (u,v) on the reference plane is given by $f(u,v)$ and $a(d)$ corrects for the beam divergence [6].

Primary (H_p) and scatter (H_s) monoenergetic energy deposition kernels (EDK) were calculated using for each 0.2 MeV spectral bins using EDKnrc [7, 8] for the spectrum accumulated at the reference plane. The average spectrum of all photons within a 1 cm radius of the beam central axis was used. The EDKs were then combined into polyenergetic kernels using a TERMA-weighted sum [9, 10], as shown in equation 5.3:

$$H(\boldsymbol{r}) = \frac{\sum_{N=1}^{N} T_n(z_0) H_n(\boldsymbol{r})}{T(z_0)} \tag{5.3}$$

where z_0 is a reference depth (chosen to be at the surface of the phantom, e.g $z_0 = 0$, in this work), the sum is over the spectrum components, T_n is the TERMA value for a specific spectrum component and H_n is either the primary or scatter monoenergetic EDK for that component.

An angular binning scheme featuring fewer cones in the backscattered direction, similar to the one presented in Ahnesjö et al. [11], was used to sample the collapsed cone directions from the EDKnrc kernel. Following the methodology described with Ahnesjö et al. [12], we postulated that EDKs obeyed the analytical form

$$h(r,\theta) = (A_\theta e^{-a_\theta r} + B_\theta e^{-b_\theta r})/r^2 \tag{5.4}$$

which, when integrated over the area of a discrete cone, becomes a *collapsed cone* EDK:

$$k_{m,n}(r) = \int\int_{\Omega_{m,n}} h(r,\theta) r^2 sin\theta d\theta d\phi = \Omega_{m,n} A_m e^{-a_m r} + \Omega_{m,n} B_m e^{-b_m r} \tag{5.5}$$

where the indices (m, n) typically indicate discrete (θ, ϕ) angles identifying the center of each cone. The polyenergetic EDKs calculated from eq. 5.3 were resampled over each cone and the parameters A_m, a_m, B_m, b_m were determined using a nonlinear least squares fitting routine available in the *SciPy* python package, with (A_m, a_m) and (B_m, b_m) corresponding to the primary and scatter EDKs, respectively.

Figure 5.5 shows the quality of the exponential fit compared to the numerical EDKs for forward and backward facing cones. The analytical fit is poorer for the backward facing cones, but the relative amount of energy deposited compared to forward cones is less by an order of magnitude. The fit could be improved by tuning the fitting range for each angle, however, the aim was not to match measured data exactly but rather to generate beamlet dose distributions that are realistic enough that the final treatment plan is not significantly degraded when final aperture dose distributions were recalculated with a full MC beam model.

To generate beamlet dose distributions, an input file defines the beam angle, isocenter coordinates in addition to the size and position of the fluence to be transported. Density values for each voxel are read from a phantom file created using Radify.

The TERMA distribution is first calculated using a *non-voxel-based* approach [6, 3] with the coordinate system aligned along the BEV of the beam. The TERMA value in each voxel of the phantom in Cartesian coordinates is then obtained by interpolating from the BEV TERMA distribution.

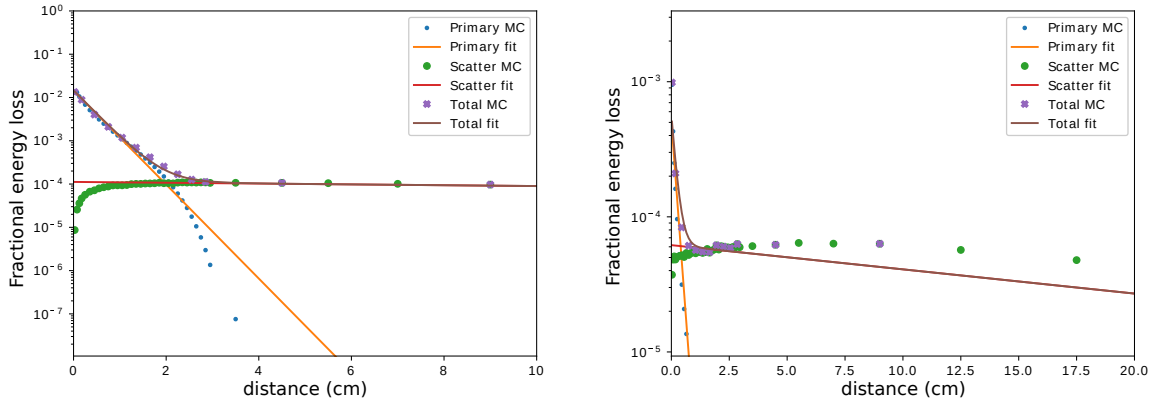

(a) (b)

Figure 5.5 Collapsed cone EDK analytical fit compared to numerical EDK values for (a) a forward facing cone ($\theta = 5°$) and (b) a backward facing cone ($\theta = 135°$).

For each cone direction, an optimal number of rays are cast from the outside of the phantom to transport the energy from the TERMA distribution. A ray spacing of 0.2 cm is used, and the EDKs are sampled every 0.1 cm. The number of collapsed cone rays that are necessary to transport the energy from every voxel with non-zero TERMA varies depending on the cone angle, as shown in figure 5.6, with cone directions aligned with the beam requiring fewer rays. The effects of kernel tilting [13] are not taken into account, hence the EDK angle θ is calculated with respect to the beam axis.

By fitting the EDKs to exponentials, we can obtain the dose contribution at each point from a collapsed cone ray by following a recursive formula [3, 12]. For each ray, the dose contribution to a point is calculated as follows,

$$D_i^{(m,n)} = \frac{A_m}{a_m}\left[(1-g_i)T_i + g_i X_i\right] \tag{5.6}$$

$$X_i = e^{-a_m p_{i-1}} X_{i-1} + (1 - e^{-a_m p_{i-1}}) T_{i-1} \tag{5.7}$$

$$g_i = \frac{1 - e^{-a_m p_i}}{a_m p_i} \tag{5.8}$$

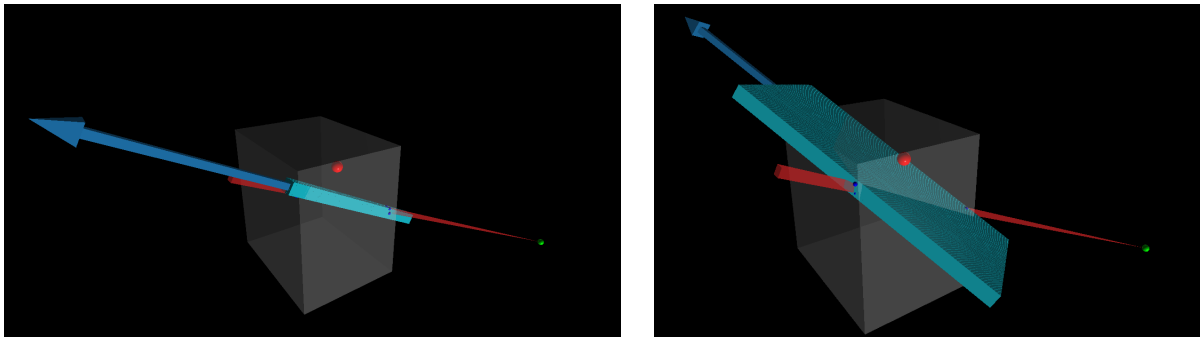

(a) (b)

Figure 5.6 Distribution of collapsed cone rays necessary to transport the energy of every non-zero TERMA voxel for (a) a cone direction aligned with the beam and (b) a cone direction more perpendicular to the beam . The green sphere represents the radiation source, and the red sphere is the machine isocenter. The region with non-zero TERMA values is represented as a red beam, and the collapsed cone rays are shown in blue, with the arrow indicating the cone direction.

where $D_i^{(m,n)}$ is the contribution from cone (m,n) to the dose at point i, p_i is the density-weighted path length along the sampled point, T_i is the TERMA value at point i, and (A,a) are the exponential EDK fitting parameters for a given cone. The dose contribution from the scatter kernel is calculated in the same fashion, with (A_m,a_m) replaced with (B_m,b_m). After the dose contribution from each cone is calculated, the dose contribution from the (m,n) cone is converted back into Cartesian coordinates and accumulated in the coordinate system defined by the phantom.

Compared to a CCCS dose calculation engine designed to accurately model as many of the macroscopic physical qualities of a photon beam as possible, such as the one presented in Ahnesjö et al. [11], the following simplifications are made:

- A single particle source raytraced back to the photon target is modelled. The fraction of photons originating from the flattening filter are therefore modelled as being more forward peaked than reality, resulting in an artificially sharper penumbra.
- The effects of kernel tilting are ignored.
- A single energy independent polyenergetic EDK is calculated for the photon beam spectrum at the entrance of the phantom. The EDKs therefore do not reflect the spectral variations of the beam inside the patient geometry.

- As the monitor units are calculated using a final MC aperture recalculation step, the dose normalisation does not take into account beam calibration or monitor back-scatter correction factors. However, the relative intensity distribution of the beam for off-axis beamlets is modelled in the TERMA component.
- Electron contamination is not taken into account, leading to artificially lower dose in the buildup region of the beam.
- The perturbations of beam collimation devices on the fluence is not taken into account.

Given these simplifications, the CCCS dose calculation engine is able to produce ≈ 2400 beamlet dose distributions per minute on an Nvidia GTX 1080 Ti graphics card. For reference, a target with a 100 cm^2 cross section would require 400 beamlet dose distributions per beam angle with a commonly used 0.5 x 0.5 cm^2 beamlet size.

5.2.2 EGSnrc

Radify can be used to coordinate the generation of EGSnrc input files for electron beamlet creation. For all modalities, a phase space file was collected at the mid-plane of the MLC for a 40×40 cm^2 jaw-defined field by using an in-house commissioned BEAMnrc [14] beam model for the Varian TrueBeam linear accelerator. This phase space file was then split into smaller phase space files representing 1×1 cm^2 beamlets at the machine isocenter based on the position of particles at the mid-plane.

Each beamlet phase space file is used as an input to DOSXYZnrc to generate a patient-specific beamlet dose distribution by transporting particles from the phase space file into the simulation phantom previously created using Radify [15].

5.3 Optimisation software

The optimisation algorithm is implemented as a C++ stand-alone software. The master problem is solved using IPOPT, an open-source interior-point nonlinear optimisation solver [16]. As with Radify, we programmed the optimisation algorithm to be applicable for more general situations than MBRT. For example, the same implementation of the optimisation model described in chapter 6 can also be used to produce photon-only or electron-only plans by simply leaving the set of electron or photon control points empty, respectively. Arc-type VMAT plans can also be

produced by limiting the pricing problem to only consider candidate MLC apertures that are physically achievable based on the position of the MLC of the nearest control points already included in the restricted master problem [17].

The available optimisation types implemented into the software are:

- Fluence map optimisation (FMO)
- Mixed photon IMRT + step-and-shoot MERT (ch. 6)
- Mixed photon arc + step-and-shoot MERT (ch. 6)
- Stochastic robust mixed photon IMRT or arc + step-and-shoot MERT (ch. 7)
- Minimax robust mixed photon IMRT or arc + step-and-shoot MERT (ch. 7)
- High dose rate brachytherapy (Renaud et al. [18])
- Experimental rotating shield high dose rate brachytherapy (Renaud et al. [18], Famulari et al. [19])
- Experimental kV arc therapy (Breitkreutz et al. [20])

By reinterpreting the meaning of a control point, the basic principles of column generation can also be applied to brachytherapy treatment planning, where the pricing problem identifies the most valuable dwell position to add to the restricted master problem, rather than the most valuable MLC aperture(s). The general applicability of CG as implemented in the optimisation software has allowed Famulari et al. to investigate the dosimetric benefits of prototype intensity modulated brachytherapy sources in part by using Radify and the optimisation software produced in this thesis [19]. Breitkreutz et al. have likewise used Radify and the optimisation software to produce and analyse treatment plans for a novel low-cost kV arc therapy machine [20].

References

[1] D. Mason, Scaramallion, Rhaxton, Mrbean-bremen, J. Suever, V. Sochat, G. Lemaitre, D. P. Orfanos, A. Panchal, J. Massich, A. Rothberg, K. van Golen, J. Kerns, T. Robitaille, M. Shun-Shin, Moloney, Pawelzajdel, M. Mattes, F. C. Morency, Huicpc0207, M. D.

Herrmann, K. S. Hahn, H. Meine, J. Stutters, I. de Bruijn, E. Stevens, D. Barreto, C. Bryant, A. Fedorov, and A. Klimont, "pydicom/pydicom: 1.2.1," nov 2018.

[2] A. Diamant, M.-A. Renaud, and J. Seuntjens, "Large scale comparison between AAA and MC-recalculated lung SBRT cases," *Medical Physics*, vol. 44, no. 8, p. 4384, 2017.

[3] Q. Chen, M. Chen, and W. Lu, "Ultrafast convolution/superposition using tabulated and exponential kernels on GPU," *Medical Physics*, vol. 38, no. 3, pp. 1150–1161, 2011.

[4] J. Neylon, K. Sheng, V. Yu, Q. Chen, D. A. Low, P. Kupelian, and A. Santhanam, "A nonvoxel-based dose convolution/superposition algorithm optimized for scalable GPU architectures," *Medical Physics*, vol. 41, no. 10, p. 101711, 2014.

[5] H. H. Liu, T. R. Mackie, and E. C. McCullough, "A dual source photon beam model used in convolution/superposition dose calculations for clinical megavoltage x-ray beams," *Medical Physics*, vol. 24, no. 12, pp. 1960–1974, 1997.

[6] W. Lu, "A non-voxel-based broad-beam (NVBB) framework for IMRT treatment planning," *Physics in Medicine and Biology*, vol. 55, no. 23, pp. 7175–7210, 2010.

[7] T. R. Mackie, A. F. Bielajew, D. W. Rogers, and J. J. Battista, "Generation of photon energy deposition kernels using the EGS Monte Carlo code," *Physics in Medicine and Biology*, vol. 33, no. 1, pp. 1–20, 1988.

[8] D. W. O. Rogers, I. Kawrakow, J. P. Seuntjens, B. R. B. Walters, and E. Mainegra-Hing, "NRC User Codes for EGSnrc," tech. rep., 2003.

[9] P. W. Hoban, D. C. Murray, and W. H. Round, "Photon beam convolution using polyenergetic energy deposition kernels," *Physics in Medicine and Biology*, vol. 39, no. 4, pp. 669–685, 1994.

[10] P. W. Hoban, "Accounting for the Variation in Collision Kerma-To-Terma Ratio in Polyenergetic Photon Beam Convolution," *Medical Physics*, vol. 22, no. 12, pp. 2035–2044, 1995.

[11] A. Ahnesjö, L. Weber, A. Murman, M. Saxner, I. Thorslund, and E. Traneus, "Beam modeling and verification of a photon beam multisource model," in *Medical Physics*, vol. 32, pp. 1722–1737, 2005.

[12] A. Ahnesjö, "Collapsed cone convolution of radiant energy for photon dose calculation in heterogeneous media," *Medical Physics*, vol. 16, no. 4, pp. 577–592, 1989.

[13] H. H. Liu, T. R. Mackie, and E. C. McCullough, "Correcting kernel tilting and hardening in convolution/superposition dose calculations for clinical divergent and polychromatic photon beams," *Medical Physics*, vol. 24, no. 11, pp. 1729–1741, 1997.

[14] D. W. O. Rogers, B. Walters, and I. Kawrakow, "BEAMnrc Users Manual," tech. rep., National Research Council of Canada, 2011.

[15] B. Walters, I. Kawrakow, and D. W. O. Rogers, "DOSXYZnrc Users Manual," tech. rep., National Research Council of Canada, 2016.

[16] 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.

[17] F. Peng, X. Jia, X. Gu, M. A. Epelman, H. E. Romeijn, and S. B. Jiang, "A new column-generation-based algorithm for VMAT treatment plan optimization," *Physics in Medicine and Biology*, vol. 57, no. 14, pp. 4569–4588, 2012.

[18] M.-A. Renaud, G. Famulari, J. Seuntjens, and S. Enger, "OC-0256: Column generation-based Monte Carlo treatment planning for rotating shield brachytherapy," *Radiotherapy and Oncology*, vol. 119, p. S118, 2016.

[19] G. Famulari, M.-A. Renaud, and S. Enger, "An intensity modulated delivery system for prostate brachytherapy using intermediate energy sources," *Medical Physics*, vol. 44, no. 8, p. 4370, 2017.

[20] D. Y. Breitkreutz, M. A. Renaud, J. Seuntjens, M. D. Weil, S. Zavgorodni, and M. Bazalova-Carter, "Inverse optimization of low-cost kilovoltage x-ray arc therapy plans," *Medical Physics*, vol. 45, no. 11, pp. 5161–5171, 2018.

Chapter 6

On mixed electron-photon radiation therapy optimisation using the column generation approach

Marc-andré Renaud, Monica Serban and Jan Seuntjens

Article published in: *Medical Physics, vol. 44, no. 8, pp. 4287-4298, 2017.*
Awarded the 2018 Sylvia Fedoruk prize in Medical Physics for best paper relating to work carried within a Canadian institution.

6.1 Abstract

Purpose: Despite considerable increase in the number of degrees of freedom handled by recent radiotherapy optimisation algorithms, treatments are still typically delivered using a single modality. Column generation is an iterative method for solving large optimisation problems. It is well suited for mixed-modality (e.g., photon-electron) optimisation as the aperture shaping and modality selection problem can be solved rapidly, and the performance of the algorithm scales favourably with increasing degrees of freedom. We demonstrate that the column generation method applied to mixed photon-electron planning can efficiently generate treatment plans and investigate its behaviour under different aperture addition schemes.

Materials and methods: Column generation was applied to the problem of mixed-modality treatment planning for a chest wall case and a leg sarcoma case. 6 MV beamlets (100 cm SAD) were generated for the photon components along with 5 energies for electron beamlets (6, 9, 12, 16 and 20 MeV), simulated as shortened-SAD (80 cm) beams collimated with a photon MLC.

For the chest wall case, IMRT-only, modulated electron radiation therapy (MERT)-only, and mixed electron-photon (MBRT) treatment plans were created using the same planning criteria. For the sarcoma case, MBRT and MERT plans were created to study the behaviour of the algorithm under two different sets of planning criteria designed to favour specific modalities.

Finally, the efficiency and plan quality of four different aperture addition schemes was analysed by creating chest wall MBRT treatment plans which incorporate more than a single aperture per iteration of the column generation loop based on a heuristic aperture ranking scheme.

Results: MBRT plans produced superior target coverage and homogeneity relative to IMRT and MERT plans created using the same optimisation criteria, all the while preserving the normal tissue-sparing advantages of electron therapy. Adjusting the planning criteria to favour a specific modality in the sarcoma case resulted in the algorithm correctly emphasizing the appropriate modality.

As expected, adding a single aperture per iteration yielded the lowest (best) cost function value per aperture included in the treatment plan. However, a greedier scheme was able to converge to approximately the same cost function after 125 apertures in one third of the running time. Electron apertures were on average 50-100% larger than photon apertures for all aperture addition schemes. The distribution of intensities among the available modalities followed a similar trend for all schemes, with the dominant modalities being 6 MV photons along with 6, 9 and 20 MeV electrons.

Conclusion: The column generation method applied to mixed modality treatment planning was able to produce clinically realistic treatment plans and combined the advantages of photon and electron radiotherapy. The running time of the algorithm depended heavily on the choice of mixing scheme. Adding the highest ranked aperture for each modality provided the best trade-off between running time and plan quality for a fixed number of apertures. This work contributes an efficient methodology for the planning of mixed electron-photon treatments.

6.2 Introduction

The introduction of new generation, robotic linear accelerators has been typically accompanied by an increase in the degrees of freedom of motion employed dynamically during treatment. To make use of these additional degrees of freedom, the sophistication of treatment planning software algorithms has increased over time. These advances have enabled the community to meet the stricter clinical dosimetric demands of modern radiotherapy while respecting logistical constraints such as treatment time.

Two major treatment planning innovations that have had a practice-changing impact on modern linac-based radiotherapy are intensity modulated radiotherapy (IMRT) [1] and volumetric modulated arc therapy (VMAT) [2, 3]. In IMRT, the goal of a treatment planning algorithm is to produce a series of multi leaf collimator (MLC)-defined apertures for a fixed set of beam delivery angles, and to find the amount of radiation to deliver from each aperture to satisfy clinical constraints. [4, 5] In order to speed up IMRT treatments, delivering modulated radiation in a continuous coplanar arc around the patient was found to produce dose distributions of comparable quality to IMRT, giving rise to VMAT. For this method, optimisation algorithms have to find the optimal aperture shape at every point in the trajectory while taking machine limitations such as gantry rotation speed, dose rate limits and MLC leaf travel speed into account. [6, 7]

Recent additions to mainstream treatment planning literature have further increased the complexity of the treatment planning process. For example, multi-criteria optimisation (MCO) [8, 9] seeks to evaluate the trade-offs between different, often conflicting clinical constraints and offers the users a choice between these trade-offs. Another avenue involves finding more complex "4π" non-coplanar trajectories around the patient [10, 11], in which case collision detection becomes a part of the optimisation process, on top of adding dimensionality to the set of beam delivery angles. Finally, combined IMRT and VMAT optimisation was also advanced as a way to add flexibility to VMAT plans when more modulation would be beneficial at fixed points within the trajectory. [12, 13]

In most cases, photon radiation is the only modality being considered in the planning process, even though linear accelerators are also capable of delivering electrons. All available photon and electron energies constitute modalities that could potentially be combined to produce a superior treatment plan compared to their single-energy (or single particle) counterpart. A large body of literature has been devoted to intensity and energy modulation with electrons (MERT),

typically involving electron-specific collimators. The clinical adoption of MERT has so far been limited due to the inability of electrons to reach deep-seated tumours, the heterogeneity of the resulting dose distributions and the time-consuming tasks related to their delivery and commissioning. [14–19] However, recent work has shown that high-quality treatment plans could be created by delivering modulated electron beams with a photon MLC, eliminating the accessory requirements (applicators, cut-outs) of conventional electron radiotherapy. [20–23]

In principle, any tumour with a superficial component could potentially benefit from being treated partly with electrons. While the topic of combining photon and electron components has been studied for some time [24–26], the majority of the literature featured an element of manual combination as opposed to simultaneous optimisation. Alexander et al. [18] and Palma et al. [27] have both performed simultaneous fluence-map-based mixed-beam radiotherapy (MBRT) optimisation studies with promising results [28]. However, these authors used optimization methods which would require prohibitively large computing resources in order to take advantage of the large number of degrees of freedom available to modern linear accelerators.

One of the most common methods for obtaining high quality aperture shapes and intensities is a two-part process. First, a "fluence map optimisation" (FMO) step is performed where the radiation field from each beam is discretised into small beamlets, and the optimal intensity for each beamlet is found by minimising a cost function representing clinical dosimetric requirements. The result is a 2D fluence map for each beam. After fluence maps are obtained, a leaf sequencing algorithm must then find a series of MLC-defined apertures and associated intensities to reproduce the fluence maps as closely as possible. A large number of apertures is often needed to precisely reproduce fluence maps, which can lead to exceedingly long treatment times or degraded treatment plans if the number of apertures is limited [29, 30]. This method is typically associated with large amounts of monitor units (MUs) delivered, thereby increasing the concern about out-of-field doses. [31, 32]

Another well-known method, direct aperture optimisation (DAO), eliminates the need for an MLC leaf sequencing step by including aperture shapes in the optimisation problem. DAO was first implemented using simulated annealing to randomly vary aperture shapes and weights during optimisation, with a user-specified number of apertures for each beam. [33] Alternatively, a DAO treatment plan can also be created iteratively, one aperture at a time, by using the column generation method [34, 35]. When applied to the problem of radiotherapy treatment planning, column generation uses the optimal dose distribution from the current iteration to determine the shape of the new aperture to be added in the following iteration. The intensities of each

aperture are optimised at every iteration using established techniques such as the simplex method, gradient descent or variations thereof depending on the choice of cost function. DAO methods can take into account machine restrictions during the whole optimisation process, leading directly to deliverable plans.

Column generation has successfully been applied to contemporary photon optimisation models that incorporate complex trajectories [10] or many more degrees of freedom [13] compared to traditional VMAT deliveries. In this work, the aim is to develop a simultaneous mixed-beam direct aperture optimisation model by using the column generation method to seek a "high-potential" aperture from every delivery angle and from each modality at every iteration. The modalities considered in this work will be 6 MV photons and 6, 9, 12, 16, 20 MeV electrons. The candidate aperture with the greatest instantaneous potential to minimise the cost function will be added to the treatment plan. Alternative aperture selection schemes will be investigated to determine whether the iterative process can be accelerated. In addition, the effects of the aperture selection scheme on the distribution of modalities within a plan will be studied.

6.3 Theory

6.3.1 The direct aperture optimisation model

The notation and formalism in this section closely follows the model presented by Romeijn et al. [34] in their work on direct aperture optimisation (DAO) for photon IMRT. The extension of the approach to multi-modality treatment planning is given in section 6.3.2.

Notation and conventions

Let B be the set of pre-selected beam directions from which radiation will be delivered. The radiation field from each beam is discretised into beamlets (or bixels) represented by the set N_b. For each beam direction $b \in B$, the aim is to generate a set of apertures K_b and find the optimal intensity of each aperture to minimise a cost function. The set of beamlets included in a given aperture $k \in K_b$ will be denoted by A_k.

For every beamlet $n \in N_b$, a dose distribution will be simulated and represented on a voxelised phantom generated from a computed tomography (CT) scan. The set of all voxels in the phantom will be denoted by V and the unit dose deposition coefficient (DDC) in each voxel $j \in V$ from aperture $k \in K_b$ will be denoted by D_{kj}. In this step, it is assumed that the DDC for

aperture k in voxel j can be approximated by summing the contributions to voxel j from each beamlet included in the aperture:

$$D_{kj} = \sum_{n}^{|A_k|} D_{nj}, \;\; n \in A_k \tag{6.1}$$

where D_{nj} is the unit DDC from beamlet n alone. This assumption is typically accurate for photon beams as each beamlet has a sharp penumbra. On the other hand, the penumbra from electron beamlets is less sharp, especially for low energy electrons. However, Henzen et al. verified that the approximation made in equation 6.1 is also adequate for electrons by performing a full MC recalculation step after beamlet-based DAO optimisation and observing no significant degradation of the treatment plan. [23]

A decision variable y_k will be associated with each aperture to control the amount of radiation delivered that aperture. The total dose in voxel j, z_j, is given by the sum of all aperture contributions to voxel j weighted by each aperture's intensity:

$$z_j = \sum_{k}^{|K|} D_{kj} y_k, \;\; K = \cup K_b \tag{6.2}$$

The optimisation problem requires a cost function representing various treatment plan optimisation criteria. Typical examples include voxel-based penalty functions, tumour control probability (TCP), equivalent uniform dose (EUD), mean dose to organs, dose-volume histogram (DVH) criteria and conditional value at risk (CVaR). [5, 29]

The nature of the cost function can fundamentally change the viable optimisation methods based on their properties, such as convexity, or their ability to be represented in a linear optimisation problem. To derive the formalism presented in the next section, the cost function is assumed to be the sum of voxel-based convex functions, $F_j(\boldsymbol{z_j})$. When non-convex functions are used, any optimal value reached by the optimisation algorithm is only guaranteed to be a local optimum, making the problem solution dependent on its initial conditions. Due to their relevance in clinical plan evaluation, dose volume criteria implemented as suggested by Bortfeld et al. [36] are often present in optimisation algorithms despite the cost function not being convex.

Master problem formulation

Although the set of all possible apertures is too large to fully explore ($> 10^{17}$ for typical IMRT cases), it is expected that relatively few apertures will be needed to produce a high quality plan. The master problem (MP) involving all possible apertures can be formally stated as:

$$\min \sum_{j}^{|V|} F_j(z_j) \tag{6.3}$$

subject to the following constraints:

$$\begin{aligned} z_j = \sum_{k}^{|K|} D_{kj} y_k, \quad j = 1, ..., |V| \\ y_k \geq 0, \quad k \in K \end{aligned} \tag{6.4}$$

The aim is to solve a restriction of the MP for a manageable number of apertures, i.e., find a subset of K, $\hat{K}$, which will yield a treatment plan of comparable quality to the one that would be obtained by solving the MP.

Column generation

The column generation method derives its name from its application to linear programming, though the concept can be applied to a more general nonlinear programming problem. In the context of radiotherapy, the method describes a process for iteratively creating a set $\hat{K}$ of "good" apertures, starting from an arbitrary set of initial apertures (which can be the empty set, $\hat{K} = \varnothing$). At every iteration, a restricted master problem (RMP) is solved to optimality for the current aperture set $\hat{K}$, meaning that the aperture shapes in $\hat{K}$ are fixed and the weight of each aperture is found by minimising the cost function.

Evidently, an optimal solution to the RMP is in general not optimal for the unrestricted master problem. The Karush-Kuhn-Tucker (KKT) conditions for optimality [37], when applied to the unrestricted master problem with the restricted solution, provide an approximation for how much improvement to the cost function can be obtained by adding more apertures to $\hat{K}$.

Since $\hat{K}$ is a subset of K, applying the KKT conditions to the master problem implies implicitly setting the weight of apertures not included in $\hat{K}$ to 0. The KKT conditions for the

master problem defined in section 6.3.1 are given by:

$$\pi_j = \frac{\partial F}{\partial z_j}, \qquad j = 1, ..., |V| \tag{6.5a}$$

$$\rho_k = \sum_{i=1}^{|V|} D_{kj}\pi_j, \qquad \forall k \in K \tag{6.5b}$$

$$y_k \geq 0, \qquad \forall k \in K \tag{6.5c}$$

$$z_j = \sum_{k}^{|K|} D_{kj}y_k, \qquad j = 1, ..., |V| \tag{6.5d}$$

$$\rho_k y_k = 0, \qquad \forall k \in K \tag{6.5e}$$

$$\rho_k \geq 0, \qquad \forall k \in K \tag{6.5f}$$

Combining equations 6.5b and 6.5f yields the condition:

$$\sum_{i=1}^{|V|} D_{kj}\pi_j \geq 0, \;\; \forall k \in K \tag{6.6}$$

which serves both as an optimality test and a method for finding "good" apertures. The π_j are obtained after optimising the current set of aperture weights, and the D_{kj} are the DDCs for any valid aperture, including those not currently included in the treatment plan. By definition, equation 6.6 is satisfied for all $k \in \hat{K}$, so the optimality test involves verifying that equation 6.6 is satisfied for apertures **not** included in the treatment plan.

$$\sum_{i=1}^{|V|} D_{kj}\pi_j \geq 0, \;\; \forall k \notin \hat{K} \tag{6.7}$$

If all possible apertures satisfy equation 6.6, then the treatment plan is optimal in the sense that the cost function cannot be improved by adding more apertures. On the other hand, any aperture violating equation 6.6 is a potential candidate for addition to the treatment plan. Recall that the D_{kj} are given by summing individual beamlet dose distributions as shown in equation 6.1. The goal is therefore to find deliverable apertures which violate equation 6.7. In other

words, the aperture which minimises equation 6.8 is sought.

$$\min_{k \notin \hat{K}} \sum_{i=1}^{|V|} D_{kj} \pi_j \tag{6.8}$$

Equation 6.7 can be interpreted heuristically as the potential improvement on the cost function from adding a given aperture, estimated by the linearisation of the cost function centered around the current optimal dose distribution. Using equation 6.8 to obtain aperture shapes is known as solving the "pricing problem" (PP). Strategies for finding apertures are explored in detail in Romeijn et. al. [34], taking into account various types of restrictions on allowable apertures based on the type of multi-leaf collimator in use. The advantage of the column generation method applied to radiotherapy treatment planning lies in the fact that finding potential apertures is a computationally simple problem to solve.

Typically, the aim is to create the highest quality plan with the fewest number of apertures, therefore the aperture which violates equation 6.7 the most (by having the most negative ρ_k) will typically be added to the treatment plan, though multiple apertures could be added at any iteration at the cost of potentially adding redundant apertures.

6.3.2 Application of the model to mixed electron-photon treatment planning

The formalism is presented for electron and photon external beam radiotherapy components with an IMRT-style step-and-shoot delivery. In the original "single-modality" model, solving the pricing problem either yields one deliverable aperture per beam angle or none if the algorithm determines that no further apertures from this beam will improve the cost function. Transforming the single-modality model into a multi-modality model simply means extending the set of all apertures to include apertures from all modalities from every beam angle. The set of all possible apertures, K, becomes

$$K = \left(\bigcup_{b^\gamma} K_{b^\gamma} \right) \cup \left(\bigcup_{b^e} K_{b^e} \right), b^\gamma \in B^\gamma, b^e \in B^e \tag{6.9}$$

where B^γ and B^e are the sets of beam angles from which photons and electrons will be delivered, respectively. K_{b^γ} and K_{b^e} are the sets of all photon and electron apertures from beam angles b^γ

and b^e, respectively. In this work, five electron energies are considered at every electron beam angle, therefore

$$K_{b^e} = \bigcup_{E_i} K_{b^e}^{E_i} \tag{6.10}$$

where $K_{b^e}^{E_i}$ is the set of all electron apertures from beam angle b^e and energy $E_i \in \{6, 9, 12, 16, 20\,\text{MeV}\}$. The same decomposition could be applied to multiple photon energies, however only one energy was considered in this work. The procedure outlined in section 6.3.1 can then be followed to obtain an aperture for each modality's beam angles. For example, if 20 photon beam angles and 5 electron beam angles are specified, a total of 45 apertures are generated (20 from photons, 25 from electrons) from solving the pricing problem. A heuristic decision can then be made to choose which apertures to include at every iteration.

While the pricing problem is useful to construct an aperture for a given beam, the absolute value of the price obtained from equation 6.8 cannot directly be used to compare the relative quality of apertures from different modalities as it depends on the normalization of the aperture dose values, D_{kj}. For photons, summing unit intensity beamlet dose distributions yields an approximately unit intensity aperture dose distribution. The same is not true for electron beamlets due to the large scatter component. Adding all generated apertures to the RMP at every iteration would avoid having to compare the relative quality of each aperture. However, this may lead to "inefficient" plans where many apertures are required to obtain an acceptable plan quality due to the large number of apertures created.

To obtain a "scale-invariant" measure of aperture quality, equation 6.8 is reinterpreted geometrically as

$$\sum_{i=1}^{|V|} D_{kj}\pi_j = \vec{D}_k \cdot \vec{\pi} = ||\vec{D}_k||\ ||\vec{\pi}||\ \cos\theta \tag{6.11}$$

where θ is the angle between a vector of all dose voxels for the aperture and the vector of Lagrange multipliers. Viewing $-\vec{\pi}$ as the direction of steepest descent for the cost function, the aperture with $\cos\theta$ closest to -1 is sought and considered the highest quality aperture. As a heuristic, only the voxels where $\vec{\pi}$ is non-zero are considered when computing $\cos\theta$ from equation 6.11. The pricing problem then becomes:

$$\min_{k \notin \hat{K}} \frac{\vec{D}_k \cdot \vec{\pi}}{||\vec{D}_k||\ ||\vec{\pi}||} \tag{6.12}$$

This new pricing problem is equivalent to the previous scheme in equation 6.8 with each D_{kj} renormalized to $\frac{D_{kj}}{||\vec{D}_k||}$ since $||\vec{\pi}||$ is a constant for all apertures in a given iteration. Using this pricing strategy, four *aperture addition schemes* will be investigated:

- The "all apertures" scheme will include all generated apertures every iteration,
- the "best per modality" scheme will include the highest quality aperture from each modality (6 apertures per iteration),
- the "best per particle type" scheme will include both the best photon and best electron apertures (2 apertures per iteration),
- the "alternating particle type" will alternate between the single best photon aperture and the single best electron aperture (1 aperture per iteration).

Many more schemes could be investigated, such as schemes adding a subset of apertures from different beam angles for a single modality. However, schemes that mix modalities are prioritised as the goal is to capture the beneficial components of each modality in as few apertures as possible.

Finally, in this work, after an aperture dose distribution is constructed from equation 6.12, the dose values in each voxel are renormalized such that the maximum dose is 1 Gy for ease of comparison between modalities. As a result, the decision variables (y_k) are not directly proportional to a physically controllable quantity such as monitor units. Instead, they represent the maximum absolute dose delivered by each aperture and serve as a measure of the relative contribution of each aperture to the total dose. In the remainder of this work, the y_k variables will be referred to as "weights" instead of intensities.

6.3.3 Aperture Pruning

The algorithm utilises the current iteration's optimal dose distribution to determine which aperture will be added next. Initially, when few apertures are included in the treatment plan, the optimal dose distribution changes rapidly between iterations. An aperture which was constructed based on the dose distribution in an early iteration can become obsolete in later iterations.

For example, the initial loop of the algorithm constructs an aperture from the Lagrange multipliers of an "optimal" dose distribution with 0 Gy in every voxel. In that scenario, the

only active component of the cost function is the PTV under-dosing penalty, meaning that the aperture selected will maximise dose in the PTV without taking OARs into account. Such an aperture can be quickly rendered obsolete as more appropriate apertures are added, and its weight will often be optimised down to 0.

To address this situation, an aperture pruning step is performed at the beginning of every iteration to eliminate obsolete apertures from the treatment plan. Any aperture with a weight below 0.1% of the average aperture weight in the treatment plan is removed as it is assumed that the aperture does not contribute significantly to the final dose distribution. Since the restricted master problem optimisation time depends on the number of apertures included in the problem, this step improves the speed of the algorithm. To be fully rigorous, the RMP could be re-solved after this step. However, the cut-off weight for pruning apertures in this work is set such that any pruned apertures have negligible contributions to the total dose distribution and the dual variables, hence the RMP is not re-solved.

The workflow for the algorithm can therefore be described as follows:

1: Initialize $\hat{K} = \varnothing$, $z_j = 0 \; \forall j \in V$
2: **while** $|\hat{K}|$ is below the user-specified number of apertures desired **do**
3: Solve the PP for each beam angle and each modality
4: **if** PP $\geq 0 \; \forall \; K \notin \hat{K}$ **then return** $(\hat{K}, y_k)$
5: **end if**
6: Add one or many candidates apertures to $\hat{K}$ depending on heuristic scheme
7: Solve the RMP to obtain a new set of y_k weights
8: Prune apertures with weight $y_k \leq \overline{y_k}/1000$
9: (Optional) Re-solve the RMP
10: **end while**
11: **return** $(\hat{K}, y_k)$

6.4 Methods

The optimisation model was implemented in C++ using Ipopt, an interior-point optimizer library for large-scale nonlinear optimization, as the solver for the RMP. [38] A chest wall case and a leg sarcoma case were selected as model clinical examples to investigate the performance and behaviour of the mixed modality algorithm. Both sites are typically treated using photons only,

but both feature a shallow target well suited for an electron component. For chest wall cases, the large size of the target along with the proximity of many OARs presents difficulties in generating high quality treatment plans. Since electrons have limited range compared to photons, a mixed photon-electron plan could combine the advantages of both modalities to improve the overall quality of chest wall treatment plans. For leg sarcomas with targets extending to the surface of the skin, the lower surface dose of photon therapy is problematic as it compromises dose homogeneity in the target unless bolus material is used. A combination of electrons and photons could allow homogeneous dose distributions to be delivered without bolus. The emphasis was

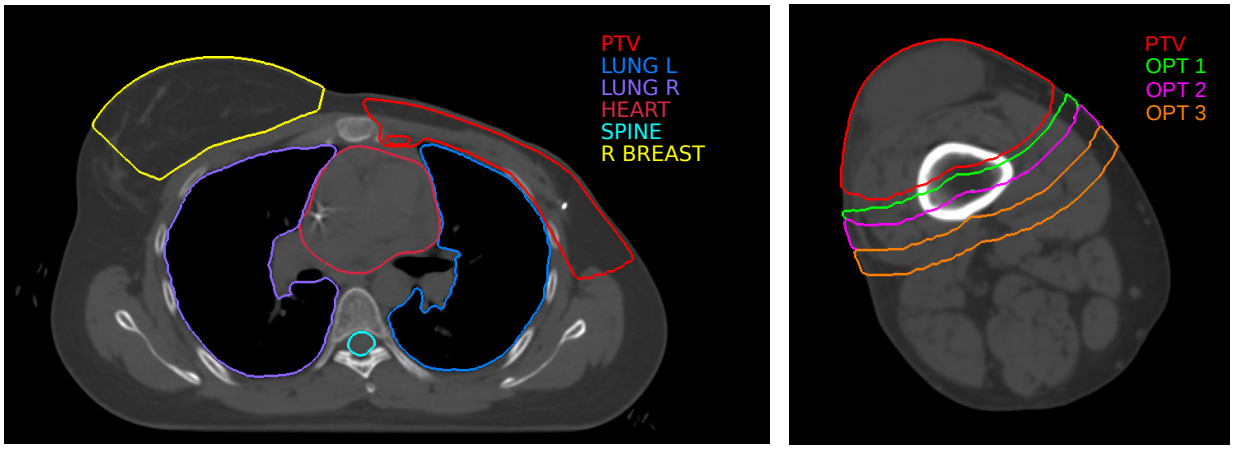

(a) (b)

Figure 6.1 Representative axial CT slice of (a) the chest wall case and (b) the leg sarcoma case with structures used in the optimisation process. Additional chest wall structures were created for optimisation purposes, but they were too numerous to be explicitly labelled.

placed on studying the evolution of the treatment plan as apertures are added to it, rather than rigorously evaluating the plan based on clinical criteria. A representative axial CT slice for each case along with relevant optimisation structures is shown in figure 6.1.

6.4.1 Cost function

Quadratic one-sided voxel-based penalty functions and DVH-based penalties were considered in this work. The same penalties were prescribed to all voxels within contoured structures. For overlapping structures, a priority variable was specified for each structure during a preprocessing step to ensure that a voxel could only belong to one structure. The quadratic one-sided function

for one voxel belonging to a specific structure s had the form:

$$F_{jn}^{s-}(z_j) = \frac{\alpha_n^s}{|V_s|}\max(0, T_n^{s-} - z_j)^2 \tag{6.13}$$

$$F_{jn}^{s+}(z_j) = \frac{\alpha_n^s}{|V_s|}\max(0, z_j - T_n^{s+})^2 \tag{6.14}$$

where V_s is the set of all voxels within structure s, j refers to a specific voxel inside V_s, α_n^s is a user-specified relative weighing factor for that penalty and T_n^{s-} and T_n^{s+} are a user-specified threshold dose controlling the start of the penalty. The over-dosing penalty is expressed by a '+' sign and the under-dosing penalty by a '-' sign. The subscript n is included to indicate that multiple penalties can be applied to a given structure. For example, a target structure would typically have both under and over-dosing penalties with the same threshold dose for homogeneity, whereas only over-dosing penalties would be applied to organs at risk. Voxels which were not part of any structure had no penalties assigned to them.

Dose volume penalties were implemented as described by Bortfeld [36] and Wu [39]. The formulation $V^s(>T) < V_n$, which is interpreted as "the volume of structure s receiving dose greater than T must be less than V_n", can be represented in the optimisation model as

$$F_{jm}^{sDV}(z_j) = \frac{\alpha_m^s}{|V_s|}\max(0, D_{V_m} - z_j) \cdot \max(0, z_j - T)^2 \tag{6.15}$$

where D_{V_m} is the current point on the dose volume histogram such that $V(D_{V_m}) = V_m$. This formulation ensures that only voxels with dose values between D_{V_m} and T are penalized. The subscript m is added to emphasize that multiple dose volume constraints can be specified for the same structure. All other symbols preserve the same definition as in equations 6.13 and 6.14.

The total cost function for one voxel was a weighted sum of all planning criteria for the structure to which the voxel was assigned:

$$F_j^s(z_j) = \sum_n F_{jn}^s(z_j) + \sum_m F_{jm}^{sDV}(z_j) \tag{6.16}$$

6.4.2 Beamlet generation

The treatment planning process was performed entirely within Radify, a web-based academic treatment planning system [40]. Initially, a patient CT scan was imported along with contoured

structures in DICOM format. Beam angles were then pre-specified for each modality as shown in figure 6.2.

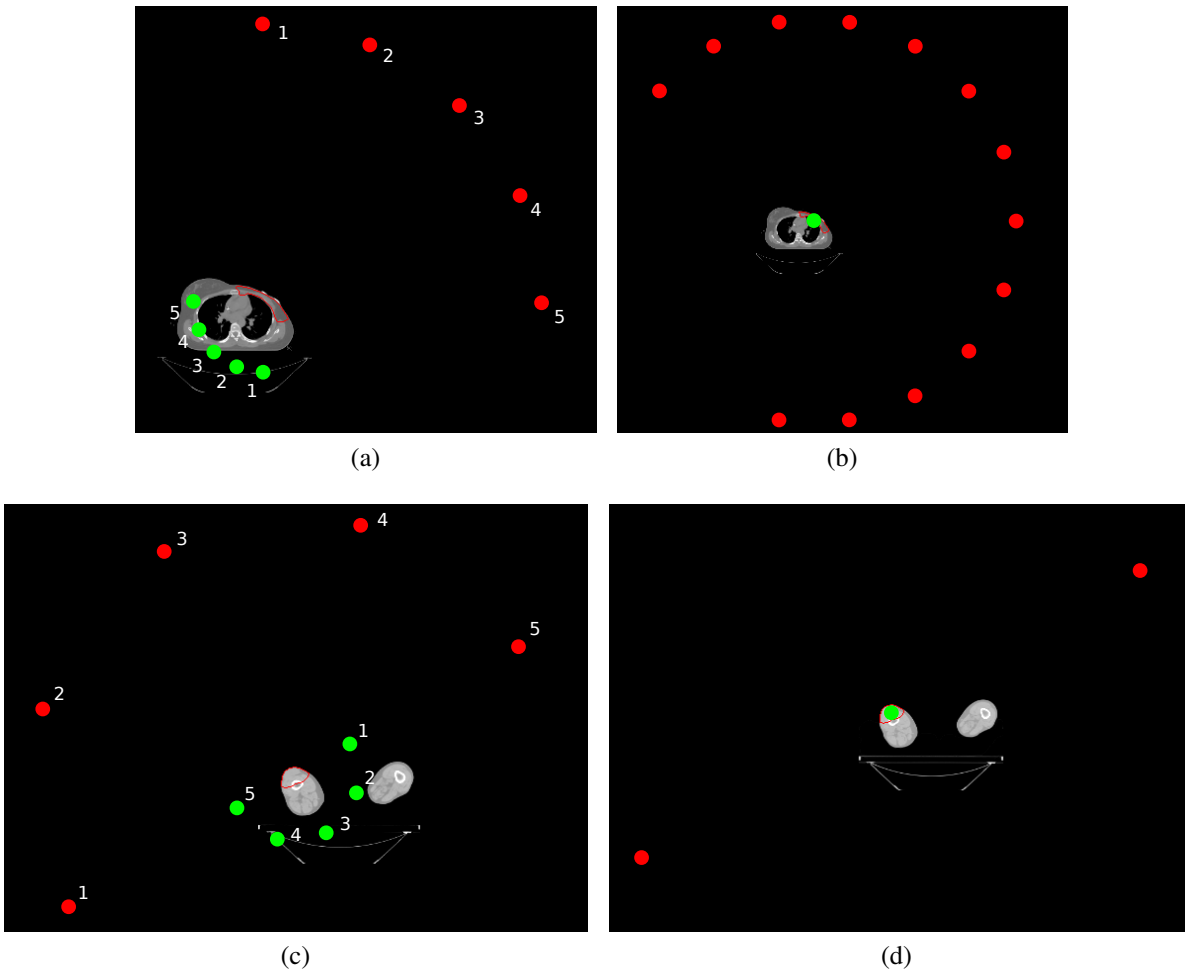

Figure 6.2 Treatment head positions (red dots) and their associated machine isocenter positions (green dots) for: (a) the electron component of the chest wall case, (b) the photon component of the chest wall case, (c) the electron component of the sarcoma case and (d) the photon component of the sarcoma case.

Phantom creation

Phantoms were created from the CT dataset to define the voxels where dose is recorded. In Radify, different material compositions can be explicitly assigned to voxels within any contoured structure, or a set of material assignment rules can be constructed based on Houndsfield units

(HU). In addition, a piecewise linear HU-to-density conversion curve was applied to obtain density values for each voxel. The voxel size was 3 mm^3 for all beamlets. For the electron component, the ICRP cortical bone composition was used for any bones present in the CT and the rest of the body was modelled as water. The specific choice of material compositions can affect the dose distributions [41] but is not expected to alter the conclusions of this work.

Photon beamlets

For photon beamlets, the radiation field of each control point was decomposed into 1 x 1 cm^2 beamlets (when projected to the machine isocenter). The maximum allowable field size was set to 40 x 40 cm^2. To accelerate the beamlet generation process, the convex hull of the PTV was used to create a mask such that only beamlets within the beam's eye view (BEV) projection of the convex hull onto the MLC plane were generated and allowed to participate in the optimisation problem. A one-beamlet-width margin was added to the convex hull to ensure complete beamlet coverage of the PTV from each beam angle. For each beamlet, a dose distribution was created using an in-house convolution-superposition algorithm with a water-based energy deposition kernel and a realistic 6 MV energy spectrum from an accelerator with a flattening filter. The kernel was generated using EDKnrc [42]. A simplified beam model was employed by assuming a uniform particle fluence across the entire beamlet width at the MLC plane. The dose distributions were outputted as sparse binary files where only voxels with non-zero dose values are recorded.

Electron beamlets

The electron component of the plans was modelled as a shortened source-axis distance (80 cm SAD) delivery using a photon MLC as the collimating device, based on the work of Henzen et al. [23] on DAO for modulated electrons. Phase space files for beam energies of 6, 9, 12, 16 and 20 MeV were collected immediately below the (fully open) MLC component module using BeamNRC [42] with a previously commissioned CL21-EX linear accelerator electron beam model [43, 44].

For each energy, the phase space file was split into 1600 individual files representing a 1 x 1 cm^2 projection of a 40 x 40 cm^2 field when projected to the machine isocenter (100 cm SAD), which is the largest available field size for the beam model.

Using these phase space files as particle sources, beamlet dose distributions were generated using DOSXYZnrc for each beam angle and each available electron beam energy. The same BEV projection masking method used to limit the number of photon beamlets generated was applied to electron beamlets as described above. The average dose uncertainty in voxels with over 50% of D_{max} was below 1% for all beamlets and the electron kinetic energy cutoff was 189 keV.

6.5 Results

6.5.1 Clinical model cases

MBRT, IMRT-only and MERT-only plans were created for the chest wall case with an identical set of planning penalties. The penalties were initially selected to match the quality of the clinically delivered plan for this case and then manually refined. The algorithm was terminated after 250 apertures were included in each plan. Figures 6.3 (a, b) show a selection of relevant dose-volume histogram (DVH) curves for all three modalities, along with the evolution of the cost function as the treatment plan is constructed iteratively. The MBRT case was able to achieve the best PTV coverage while preserving the OAR-sparing properties of electrons. A comparison of isodose curves between the three plans is shown in figure 6.4.

The weight distribution among photons and electron energies is shown in figure 6.5. The absolute contribution of the dominant electron energies (6, 9 and 20 MeV) to the treatment plan was sustained as apertures were added, whereas the photon component continued to increase as the same amount of dose was delivered from a larger number of smaller apertures. The intermediate energies (12 and 16 MeV) did not contribute significantly to the final treatment plan, most likely due to the fact that the shallow part of the chest wall could be covered adequately with photons and a mix of 6 and 9 MeV electrons, whereas the deeper section of the PTV (the top half in figure 6.4) benefits from photons and 20 MeV electrons.

For the sarcoma case, two different plans were created using the optimisation structures shown in figure 6.1 (b), a “slack” plan and a “strict” plan. In both plans, the PTV coverage criteria were the same, but the slack plan allowed for more dose spillage outside the PTV by making the threshold dose on the optimisation structures (OPT1, OPT2 and OPT3 in figure 6.1 (b)) higher compared to the strict plan. The planning criteria used in creating the two plans are shown as triangles in figure 6.6 (a).

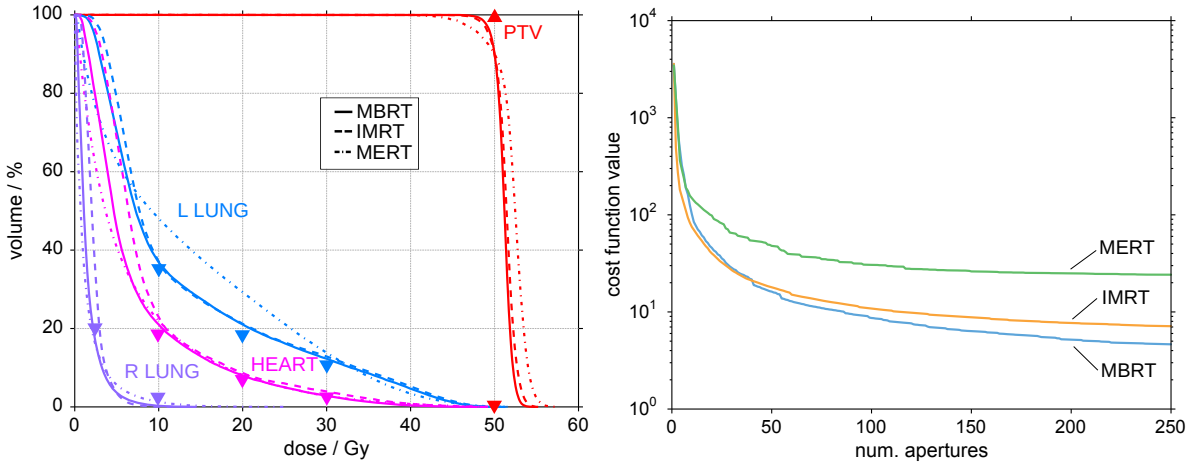

(a) (b)

Figure 6.3 (a) DVH curves for the chest wall plans comparing the MBRT plan, the MERT-only plan and the photon IMRT-only plan. The triangles represent the thresholds of the planning penalties applied during optimisation. All plans are normalized such that 90% of the PTV receives 50 Gy. (b) Evolution of the cost function as more apertures are added to the treatment plan for the plans described in (a).

Figure 6.6 (a) shows a DVH comparison of the two plans illustrating the effect of stricter spillage restrictions on the optimisation structures. For the same target coverage, the strict plan delivers much less dose to the spillage structures. In figures 6.6 (b, c), the stacked area plots show the contribution to the final dose distribution from each modality as apertures are added to the treatment plan. For the slack plan, photons remain a minor component of the plan throughout the aperture generation process. As expected, photons feature more prominently in the strict plan.

6.5.2 Aperture addition schemes

As mentioned in section 6.3.2, the pricing problem yields a deliverable aperture for each beam angle of each modality. Since the bulk of the running time of the algorithm is spent solving the restricted master problem to determine optimal aperture weights at every iteration, it may be preferable to add multiple apertures per iteration if the resulting treatment plan is of comparable quality. In order to quantify the efficiency of different aperture generation schemes,

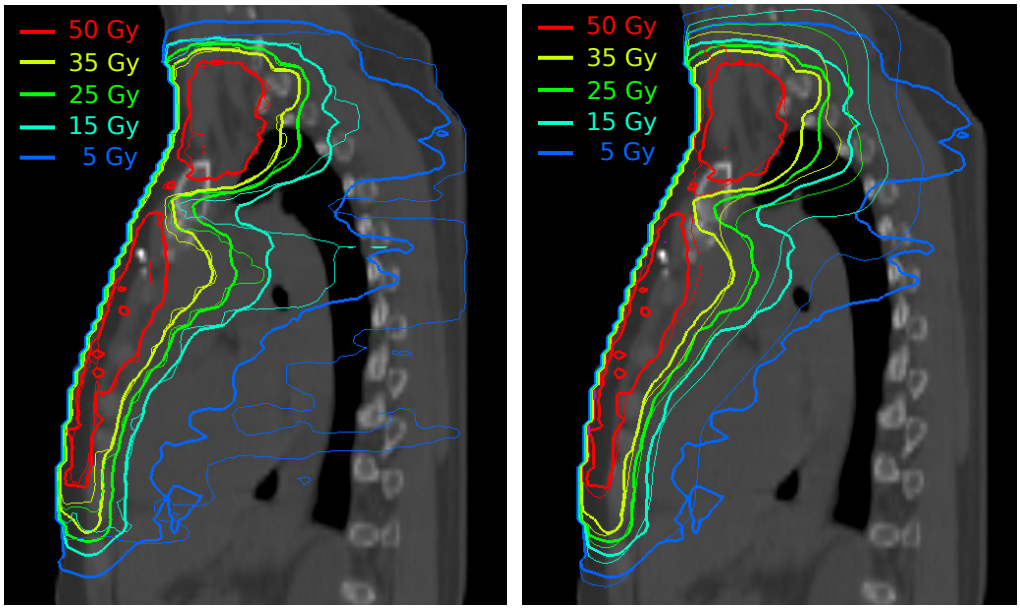

(a) (b)

Figure 6.4 Isodose lines comparing the MBRT plan (thick lines) to (a) the IMRT-only plan and (b) the MERT-only plan for a representative CT slice in the sagittal plane. The IMRT and MERT plan isodose lines are shown as thin lines.

the evolution of the cost function was recorded for four aperture addition schemes (described in section 6.3.2) for the chest wall case.

Figure 6.7 (a) shows that, as expected, schemes which added fewer apertures per iteration had a lower cost function per aperture. However, after approximately 125 apertures, 3 of the 4 schemes had converged to a similar cost function value. Figure 6.7 (b) compares the relative running times of the algorithm during plan creation. The "best per modality" scheme is able to reach 250 apertures in a third of the time compared to the "alternating particle type" scheme, even though the final cost functions were approximately the same.

The different aperture addition schemes may introduce a bias in the final plan components. Figures 6.8 (a, b) compare the distribution of aperture sizes and weight distributions for the different schemes. The distribution of aperture sizes on the final iteration was bimodal for all aperture addition schemes, with electron apertures 50-100% larger than photon apertures on average. The fourth scheme stood out as an outlier both in terms of aperture size and cost function efficiency. Figure 6.8 (b) shows that the weight distribution followed a similar trend in all aperture addition schemes. The dominant components for this case were 6 MV photons

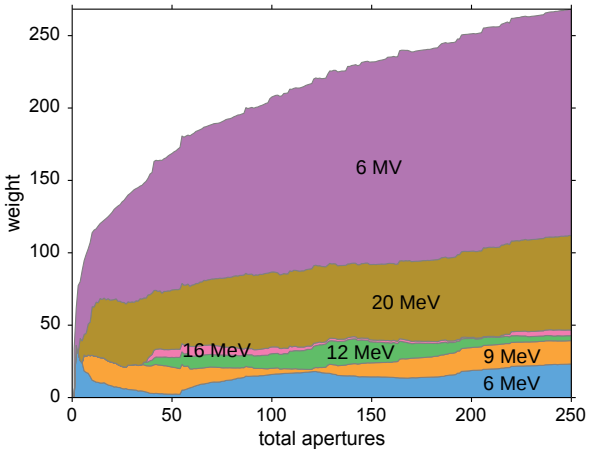

Figure 6.5 Stacked area plots showing the portion of weights being delivered from each energy for the chest wall case. The weights were calculated by binning all apertures by modality and summing their weights. A convergence analysis for up to 3000 apertures showed that largest variations in the relative distribution of weights among modalities occurred in the first 200 apertures

along with 6, 9 and 20 MeV electrons, though different sites are expected to have different dominant components within the available electron energies. For example, sites with no skin involvement are not expected to have such a large 6 MeV component.

6.6 Discussion

6.6.1 Clinical cases

The chest wall case served to examine the creation of multi-modality treatment plans in a complex situation where a large number of competing planning criteria were specified. A subset of the planning criteria are shown as arrows on the DVH curves in figure 6.3. The low-dose region in the DVH illustrates that MBRT is able to reduce the low-dose-bath signature of IMRT due to the OAR-sparing properties of electrons provided through MERT, all the while preserving superior target coverage.

For the sarcoma case, the goal in creating the “slack” and “strict” treatment plans was to artificially fashion a situation where photons were expected to play a larger role, and assess how the algorithm chooses which modality to prioritise as it constructs the treatment plan. The parallel-opposed photon component was placed advantageously to avoid dose spillage

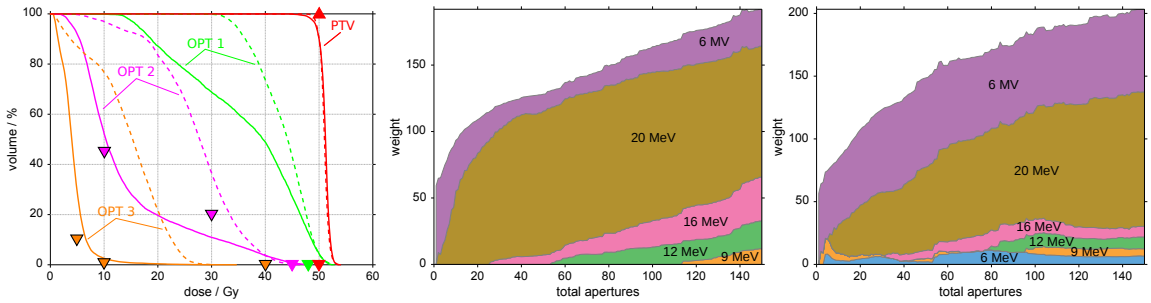

(a) (b) (c)

Figure 6.6 (a) Dose volume histogram comparing the strict plan (full lines) and the "slack" plan (dashed lines) for the optimisation structures used in creating the plans. The planning penalties are indicated as triangles. Penalties common to both plans are borderless. Both plans are normalized such that 90% of the PTV receives 50 Gy. (b, c) Stacked area plots showing the portion of weights being delivered from each energy the "slack" and the strict plan, respectively.

into the optimisation structures. However, given that no bolus was present in the photon dose calculations, the target coverage near the skin was expected to be worse for photons compared to the MERT component. The algorithm was therefore expected to prioritise electrons over photons in the "slack" case where dose spillage is "accepted", whereas the "strict" plan should feature more of the dose being delivered from the photon component, with electrons being used to homogenise the dose near the skin. These scenarios were observed as expected in figures 6.6 (b, c).

For optimisation purposes, beamlet dose calculations are often performed on a less accurate, simplified beam model. The continuing evolution of the treatment plan even as the cost function converges could be undesirable as minimal improvements in the cost function are being obtained from complex combinations of inaccurate aperture dose distributions. The quality of the final dose distribution could degrade when an accurate dose recalculation step is performed. In the interest of limiting plan complexity and maximising aperture size, these results imply that stronger convergence criteria should be imposed on the algorithm. The criteria could simply be based on the relative cost function improvement of the last N iterations, or more ideally on the convergence of clinical plan evaluation parameters. [30]

Note that despite being one of the more common ways to formulate dose volume penalties, the dose volume penalty formulation in this work is in general non-convex [45], hence the cost function could converge to a local minimum instead of the global minimum. Alternative convex methods for allowing fractions of a structure to exceed a dose threshold exist, such as

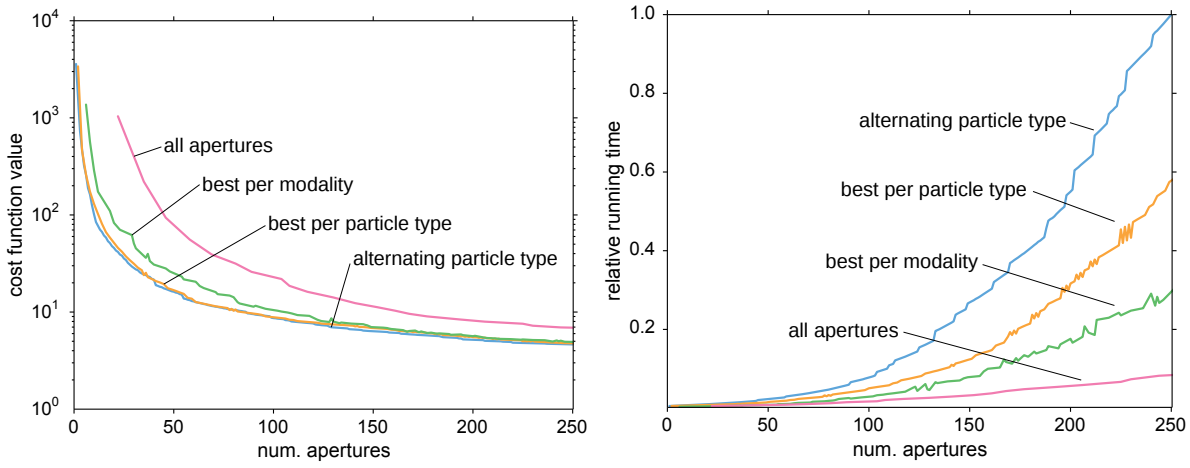

(a) (b)

Figure 6.7 (a) Cost function value as a function of total number of apertures in the treatment plan for the aperture addition schemes described in section 6.3.2. (b) Relative running times for each aperture addition schemes, normalised to the running time of the alternating particle type scheme after 250 apertures (200 minutes on one CPU).

the conditional value at risk (CVaR) formulation [46, 34]. While non-convex DVH criteria have been able to produce high quality treatment plans, recent work by Engberg et al. [47] showed that in a multi-criteria optimisation context, abandoning explicit DVH criteria for CVaR constraints yielded better DVH curves for two cases with three CVaR constraints each. These findings emphasize that plan comparison studies should be carefully designed to account for differences in planning methods.

6.6.2 Aperture addition schemes

As expected, adding multiple apertures per iteration yielded treatment plans with higher cost functions for the same number of apertures. Each candidate aperture is generated using the previous iteration's dose distribution, which does not take into account other apertures that will be added in the same iteration. The inferior performance of the "all apertures" scheme can be understood in this context. A total of 38 apertures were added at every iteration, and too many apertures had a redundant impact on the dose distribution.

The "best per modality" scheme is initially the least efficient of the remaining three schemes as it adds 6 apertures per iteration. However, figures 6.7 (a, b) shows that after 125 apertures, it

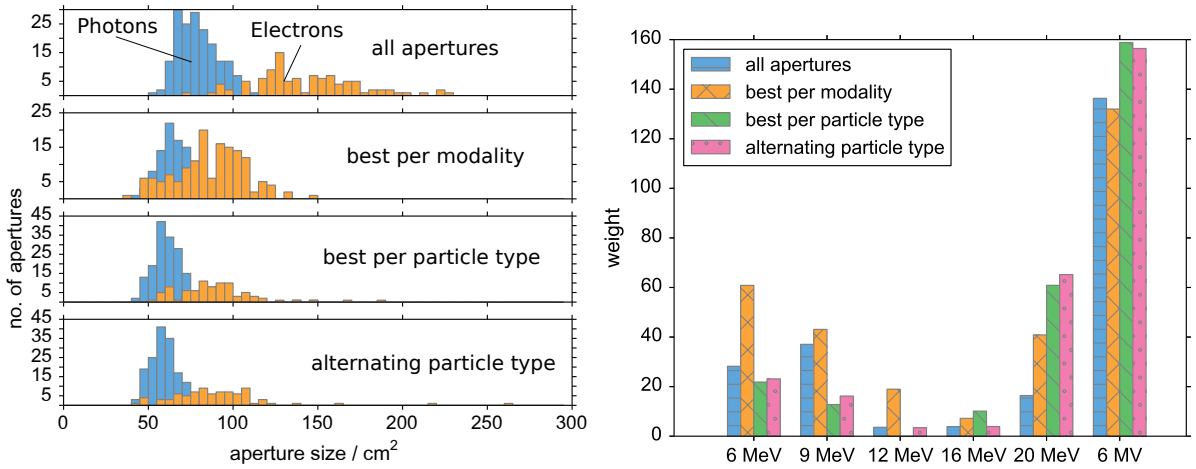

Figure 6.8 (a) Final distribution of aperture sizes for electron and photon components after 250 total apertures for the four aperture addition schemes described in figure 6.7. (b) Final weight distribution for the four aperture addition schemes described in figure 6.7 after 250 apertures.

is able to produce treatment plans with the same efficiency as schemes adding fewer apertures per iteration in a fraction of the running time. On the other hand, figure 6.8 (b) shows that the aperture addition scheme chosen slightly biases the distribution of weights among modalities such that the "best per modality" scheme has the largest electron component. In future studies which include the delivery of mixed modality plans, it may be desirable to limit the number of electron energies for delivery time or QA reasons. The "best per particle type" and "alternating particle type" schemes could provide a way to naturally limit the complexity of mixed beam plans without explicitly penalising the addition of electron energies by artificially lowering their rank in the pricing problem. The relatively poor performance of the "all apertures" scheme emphasizes the need for a heuristic aperture addition scheme for mixed beam planning compared to column generation applied to photon-only IMRT where fewer apertures are generated every iteration.

For all aperture addition schemes, electron apertures were on average 50-100% larger than photon apertures. This feature can be explained by the fact that electrons are able to cover a large portion of the target without depositing dose in normal tissue due to their limited range compared to photons. Photon treatments must combine many smaller apertures from different beam angles to achieve target coverage without depositing large quantities of dose in normal tissue. Another possible explanation for the promotion of larger electron apertures may be due to

the fact that small electron fields lose their characteristic steep dose fall-off. As the MBRT plans were able to match the IMRT-only cost function value with approximately half the apertures, this could offer a way to limit the complexity of aperture shapes and increase the average aperture size compared to conventional treatments. However, as this work did not take into account the different output factors of individual modalities, little can be said about treatment time differences between aperture addition schemes and conventional IMRT treatments. The sum of weights from each modality is roughly constant across aperture addition schemes, therefore differences in treatment time will depend on relative output factors, achievable dose rate and the logistics of changing modalities during treatment.

6.7 Conclusion

In this work, the column generation method was applied to mixed-beam radiotherapy treatment planning. The results demonstrate that the algorithm is able to handle the large increase in degrees of freedom compared to single modality optimisation. Given that the pricing problem can be solved efficiently, the multi-modality column generation-based optimisation model has almost the same runtime as the single modality model for a constant number of desired apertures in the final treatment plan. While this work specifically targeted step-and-shoot mixed electron-photon radiotherapy, applying the heuristic rules explored in the work of Peng et al. [35] could allow the beam components to be optimised as VMAT trajectories instead of a simple step-and-shoot delivery, or a combination of the two.

As expected, the lowest cost function values were obtained by following aperture addition schemes that added the fewest apertures per iteration. Among the aperture addition schemes which did not simply add all generated apertures, the specific choice of scheme did not have a significant impact on the treatment plan quality. However, the final distribution of weights among modalities was slightly biased by the specific scheme. These conclusions are expected to hold for clinical sites with superficial targets.

Rigorous clinical planning studies using this methodology are under way to quantify the advantages of MBRT. The pathway to delivery of MBRT requires research and development on a rigorous QA process including dosimetric QA, the positioning uncertainties related to multiple couch movements due to the shortened SAD and patient-specific aspects such as collision avoidance. Clinical delivery of the MERT component will require a final MC recalculation step with real, MLC-defined apertures as performed by Henzen et al. [23] to obtain the physically

controllable aperture intensities. With the latest generation of linear accelerators providing capabilities for programmable trajectories, the goal is to produce plans which will be deliverable in an automated fashion without requiring user intervention between the photon and the electron components.

6.8 Acknowledgements

The authors acknowledge support from the Fonds de recherche du Québec Nature et Technologie (FRQNT), the CREATE Medical Physics Research Training Network grant (number 432290) of the Natural Sciences and Engineering Research Council (NSERC) and the Canadian Institute of of Health Research Foundation Grant (FDN-143257).

6.9 Disclosure of Conflicts of Interest

The authors have no relevant conflicts of interest to disclose.

References

[1] A. Brahme, J.-E. Roos, and I. Lax, "Solution of an integral equation encountered in rotation therapy," *Physics in medicine and biology*, vol. 27, no. 10, p. 1221, 1982.

[2] C. X. Yu, "Intensity-modulated arc therapy with dynamic multileaf collimation: an alternative to tomotherapy," *Physics in medicine and biology*, vol. 40, no. 9, p. 1435, 1995.

[3] K. Otto, "Volumetric modulated arc therapy: IMRT in a single gantry arc," *Medical physics*, vol. 35, no. 1, pp. 310–317, 2008.

[4] T. Bortfeld, "IMRT: a review and preview," *Physics in Medicine and Biology*, vol. 51, no. 13, p. R363, 2006.

[5] H. E. Romeijn and J. F. Dempsey, "Intensity modulated radiation therapy treatment plan optimization," *Top*, vol. 16, no. 2, pp. 215–243, 2008.

[6] J. L. Bedford, "Treatment planning for volumetric modulated arc therapy," *Medical physics*, vol. 36, no. 11, pp. 5128–5138, 2009.

[7] J. Unkelbach, T. Bortfeld, D. Craft, M. Alber, M. Bangert, R. Bokrantz, D. Chen, R. Li, L. Xing, C. Men, *et al.*, "Optimization approaches to volumetric modulated arc therapy planning," *Medical physics*, vol. 42, no. 3, pp. 1367–1377, 2015.

[8] H. W. Hamacher and K.-H. Küfer, "Inverse radiation therapy planning—a multiple objective optimization approach," *Discrete Applied Mathematics*, vol. 118, no. 1, pp. 145–161, 2002.

[9] D. Craft, D. McQuaid, J. Wala, W. Chen, E. Salari, and T. Bortfeld, "Multicriteria VMAT optimization," *Medical physics*, vol. 39, no. 2, pp. 686–696, 2012.

[10] P. Dong, P. Lee, D. Ruan, T. Long, E. Romeijn, Y. Yang, D. Low, P. Kupelian, and K. Sheng, "4π non-coplanar liver SBRT: a novel delivery technique," *International Journal of Radiation Oncology* Biology* Physics*, vol. 85, no. 5, pp. 1360–1366, 2013.

[11] P. Dong, B. Ungun, S. Boyd, and L. Xing, "Optimization of rotational arc station parameter optimized radiation therapy," *Medical Physics*, vol. 43, no. 9, pp. 4973–4982, 2016.

[12] M. M. Matuszak, J. M. Steers, T. Long, D. L. McShan, B. A. Fraass, H. Edwin Romeijn, and R. K. Ten Haken, "Fusionarc optimization: a hybrid volumetric modulated arc therapy (vmat) and intensity modulated radiation therapy (imrt) planning strategy," *Medical physics*, vol. 40, no. 7, 2013.

[13] R. Li and L. Xing, "An adaptive planning strategy for station parameter optimized radiation therapy (SPORT): Segmentally boosted VMAT," *Medical physics*, vol. 40, no. 5, 2013.

[14] C. Ma, T. Pawlicki, M. Lee, S. Jiang, J. Li, J. Deng, B. Yi, E. Mok, and A. Boyer, "Energy-and intensity-modulated electron beams for radiotherapy," *Physics in medicine and biology*, vol. 45, no. 8, p. 2293, 2000.

[15] M. C. Lee, S. B. Jiang, and C.-M. Ma, "Monte carlo and experimental investigations of multileaf collimated electron beams for modulated electron radiation therapy," *Medical physics*, vol. 27, no. 12, pp. 2708–2718, 2000.

[16] T. Gauer, K. Engel, A. Kiesel, D. Albers, and D. Rades, "Comparison of electron IMRT to helical photon IMRT and conventional photon irradiation for treatment of breast and chest wall tumours," *Radiotherapy and Oncology*, vol. 94, no. 3, pp. 313–318, 2010.

[17] A. Alexander, E. Soisson, T. Hijal, A. Sarfehnia, and J. Seuntjens, "Comparison of modulated electron radiotherapy to conventional electron boost irradiation and volumetric modulated photon arc therapy for treatment of tumour bed boost in breast cancer," *Radiotherapy and Oncology*, vol. 100, no. 2, pp. 253–258, 2011.

[18] A. Alexander, E. Soisson, M.-A. Renaud, and J. Seuntjens, "Direct aperture optimization for FLEC-based MERT and its application in mixed beam radiotherapy," *Medical physics*, vol. 39, no. 8, pp. 4820–4831, 2012.

[19] A. Rodrigues, F.-F. Yin, and Q. Wu, "Dynamic electron arc radiotherapy (DEAR): a feasibility study," *Physics in medicine and biology*, vol. 59, no. 2, p. 327, 2013.

[20] E. E. Klein, M. Mamalui-Hunter, and D. A. Low, "Delivery of modulated electron beams with conventional photon multi-leaf collimators," *Physics in medicine and biology*, vol. 54, no. 2, p. 327, 2008.

[21] F. J. Salguero, B. Palma, R. Arrans, J. Rosello, and A. Leal, "Modulated electron radiotherapy treatment planning using a photon multileaf collimator for post-mastectomized chest walls," *Radiotherapy and Oncology*, vol. 93, no. 3, pp. 625–632, 2009.

[22] D. Henzen, P. Manser, D. Frei, W. Volken, H. Neuenschwander, E. J. Born, D. Vetterli, C. Chatelain, M. Stampanoni, and M. Fix, "Monte carlo based beam model using a photon mlc for modulated electron radiotherapy," *Medical physics*, vol. 41, no. 2, p. 021714, 2014.

[23] D. Henzen, P. Manser, D. Frei, W. Volken, H. Neuenschwander, E. J. Born, A. Joosten, K. Lössl, D. Aebersold, C. Chatelain, *et al.*, "Beamlet based direct aperture optimization for MERT using a photon MLC," *Medical physics*, vol. 41, no. 12, p. 121711, 2014.

[24] M. Blomquist, M. G. Karlsson, B. Zackrisson, and M. Karlsson, "Multileaf collimation of electrons—clinical effects on electron energy modulation and mixed beam therapy depending on treatment head design," *Physics in Medicine and Biology*, vol. 47, no. 7, p. 1013, 2002.

[25] K. Al-Yahya, M. Schwartz, G. Shenouda, F. Verhaegen, C. Freeman, and J. Seuntjens, ""Energy modulated electron therapy using a few leaf electron collimator in combina-

tion with IMRT and 3D-CRT: Monte Carlo-based planning and dosimetric evaluation"," *Medical Physics*, vol. 32, no. 9, pp. 2976–2986, 2005.

[26] Y. Ge and B. A. Faddegon, "Study of intensity-modulated photon–electron radiation therapy using digital phantoms," *Physics in Medicine and Biology*, vol. 56, no. 20, p. 6693, 2011.

[27] B. A. Palma, A. U. Sánchez, F. J. Salguero, R. Arráns, C. M. Sánchez, A. W. Zurita, M. I. R. Hermida, and A. Leal, "Combined modulated electron and photon beams planned by a monte-carlo-based optimization procedure for accelerated partial breast irradiation," *Physics in medicine and biology*, vol. 57, no. 5, p. 1191, 2012.

[28] E. Jimenez-Ortega, C. Miguez-Sanchez, B. Palma, A. Ureba, H. Miras, R. Arrans, A. Barbeiro, J. Baeza, F. Carrasco, and A. L. Plaza, "SU-ET-593: Outcomes and Toxicities From a Clinical Trial of APBI Using MERT+ IMRT with the Same XMLC," *Medical physics*, vol. 42, no. 6, pp. 3472–3472, 2015.

[29] V. Gregoire, T. R. Mackie, W. D. Neve, M. Gospodarowicz, J. A. Purdy, M. van Herk, A. Niemierko, A. Ahnesjö, M. Goitein, N. Gupta, and T. Landberg, "The ICRU Report 83: prescribing, recording and reporting photon-beam intensity-modulated radiation therapy (IMRT)," *J. ICRU*, vol. 10, pp. 1–106, 2010.

[30] C. Men, H. E. Romeijn, Z. C. Taşkın, and J. F. Dempsey, "An exact approach to direct aperture optimization in imrt treatment planning," *Physics in Medicine and Biology*, vol. 52, no. 24, p. 7333, 2007.

[31] S. F. Kry, M. Salehpour, D. S. Followill, M. Stovall, D. A. Kuban, R. A. White, and I. I. Rosen, "The calculated risk of fatal secondary malignancies from intensity-modulated radiation therapy," *International Journal of Radiation Oncology* Biology* Physics*, vol. 62, no. 4, pp. 1195–1203, 2005.

[32] S. Kry, O. Vassiliev, and R. Mohan, "Out-of-field photon dose folllowing removal of the flattening filter from a medical accelerator," *Physics in Medicine and Biology*, vol. 55, pp. 2155–2166, 2010.

[33] D. Shepard, M. Earl, X. Li, S. Naqvi, and C. Yu, "Direct aperture optimization: a turnkey solution for step-and-shoot imrt," *Medical physics*, vol. 29, no. 6, pp. 1007–1018, 2002.

[34] H. E. Romeijn, R. K. Ahuja, J. F. Dempsey, and A. Kumar, "A column generation approach to radiation therapy treatment planning using aperture modulation," *SIAM Journal on Optimization*, vol. 15, no. 3, pp. 838–862, 2005.

[35] F. Peng, X. Jia, X. Gu, M. A. Epelman, H. E. Romeijn, and S. B. Jiang, "A new column-generation-based algorithm for VMAT treatment plan optimization," *Physics in medicine and biology*, vol. 57, no. 14, p. 4569, 2012.

[36] T. Bortfeld, J. Stein, and K. Preiser, "Clinically relevant intensity modulation optimization using physical criteria," in *Proceedings of the XIIth ICCR, Salt Lake City*, vol. 29, 1997.

[37] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty, *Nonlinear programming : theory and algorithms*. Wiley-Interscience series in discrete mathematics and optimization, New Jersey, Hoboken: John Wiley & Sons, 3 ed., 2006.

[38] 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.

[39] Q. Wu and R. Mohan, "Algorithms and functionality of an intensity modulated radiotherapy optimization system," *Medical physics*, vol. 27, no. 4, pp. 701–711, 2000.

[40] M. Renaud and F. DeBlois, "WebTPS: a complete web application for monte carlo treatment plan recalculation," *International Journal of Radiation Oncology, Biology, Physics*, vol. 87, no. 2, p. S625, 2013.

[41] U. Schneider, E. Pedroni, and A. Lomax, "The calibration of CT Hounsfield units for radiotherapy treatment planning," *Physics in medicine and biology*, vol. 41, no. 1, p. 111, 1996.

[42] D. Rogers, I. Kawrakow, J. Seuntjens, B. Walters, and E. Mainegra-Hing, "NRC user codes for EGSnrc," *NRCC Report PIRS-702 (Rev. B)*, 2003.

[43] T. Connell, A. Alexander, M. Evans, and J. Seuntjens, "An experimental feasibility study on the use of scattering foil free beams for modulated electron radiotherapy," *Physics in medicine and biology*, vol. 57, no. 11, p. 3259, 2012.

[44] T. Connell and J. Seuntjens, "Design and validation of novel scattering foils for modulated electron radiation therapy," *Physics in medicine and biology*, vol. 59, no. 10, p. 2381, 2014.

[45] J. Deasy, "Multiple local minima in radiotherapy optimization problems with dose-volume constraints," *Medical Physics*, vol. 24, no. 7, pp. 1157–1161, 1997.

[46] H. E. Romeijn, R. K. Ahuja, J. F. Dempsey, A. Kumar, and J. G. Li, "A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning," *Physics in Medicine and Biology*, vol. 48, no. 21, p. 3521, 2003.

[47] L. Engberg, A. Forsgren, K. Eriksson, and B. Hardemark, "Explicit optimization of plan quality measures in intensity-modulated radiation therapy treatment planning," *Medical Physics*, 2017.

Chapter 7

Robust mixed electron-photon radiation therapy optimisation

Marc-andré Renaud, Monica Serban and Jan Seuntjens

Article published in: *Medical Physics, 2019. (doi: 10.1002/mp.13381)*

7.1 Preface

Following the successful implementation of the column generation method for MBRT treatment planning, we proceeded with a planning study for 3 of each chest wall irradiation, accelerated partial breast irradiation and soft tissue sarcoma cases. However, as a precautionary step, we performed a robustness analysis on the optimised dose distributions by simulating the effects of 5 mm positioning errors in each direction. The preliminary analysis revealed that CTV coverage was not adequately preserved under positioning errors corresponding in magnitude to the CTV-to-PTV expansion margin, despite the target coverage constraints being prescribed to the PTV. For that reason, we modified the column generation-based MBRT optimisation model presented in chapter 6 to produce robust plans by explicitly ensuring that CTV coverage is preserved for positioning errors corresponding to the traditional CTV-to-PTV margin. In this chapter, we introduce the new robust optimisation model and study the characteristics of CTV-based robust plans compared to traditional PTV-based plans.

7.2 Abstract

Purpose: Mixed beam electron-photon radiation therapy (MBRT) is an emerging technique that has the potential to reduce dose to normal tissue while improving target coverage for cancer sites with superficial tumours. Advances in optimisation algorithms and robotic linear accelerators have made the creation and delivery of complex MBRT plans realistic without the need for special additional collimators, devices or re-setup of the patient. However, no study has been performed on the robustness of MBRT dose distributions to patient setup errors. Intensity modulated delivery of other charged particles such as protons have been shown to require robust planning techniques to maintain adequate target coverage under positioning errors. We therefore assess the sensitivity of MBRT treatment plans to positioning uncertainties when created under the traditional PTV-based planning paradigm and present a novel optimisation model for the creation of robust MBRT plans.

Methods: The column generation method was applied to robust MBRT treatment planning by deriving the pricing problem for stochastic and "worst case" minimax optimisation models, two common formulations of robustness. Robust treatment plans were created for two patient cases representative of the cancer sites which stand to benefit from MBRT: soft tissue sarcoma (STS) irradiation and chest wall irradiation with deep-seated internal mammary, axillary and supraclavicular nodes (CW-N). For both patient cases, beamlet dose distributions for electrons and photons were generated for positioning shifts in 6 directions, ± 5 *mm* $(\hat{x}, \hat{y}, \hat{z})$ in addition to a nominal unshifted scenario, for a total of 7 sets of beamlets. Robust plans were created by specifying dose coverage constraints to the CTV, as opposed to the PTV. Comparisons were performed against traditional PTV-based plans created with a single set of unshifted beamlets.

Results: The dose distributions of traditional PTV-based MBRT plans showed significant degradation in target coverage homogeneity when patient positioning errors were considered. For both cancer sites, cold spots below 95% and hot spots above 108% of the prescription dose appeared within the CTV when shifting the patient by 5 mm, corresponding to the margin added to the CTV to form the PTV. In contrast, CTV-based robust plans created with the new optimisation model maintained target coverage within the 95%-108% limits, for all positioning errors.

Conclusion: The quality of MBRT treatment plans created using a traditional PTV-based optimisation model was highly sensitive to patient positioning errors. For both patient cases, positioning errors resulted in perturbations to the nominal dose distributions which would have

rendered PTV-based plans clinically unacceptable. In contrast, CTV-based robust plans were able to maintain adequate target coverage under all positioning error scenarios considered. We therefore conclude that to ensure the fidelity of the dose distribution delivered to the patient, robust optimisation is critical when creating MBRT plans.

7.3 Introduction

Modern radiation therapy treatment planning research focuses on finding novel ways to reduce normal tissue toxicity while preserving target coverage. On the side of delivery techniques, a large amount of work has been dedicated to identifying non-coplanar trajectories that produce superior treatment plans when compared to the standard coplanar trajectories. Examples of such techniques include "4π" optimisation [1], the TORUS technique [2] and non-coplanar SPORT [3]. In addition to improved delivery techniques, advanced treatment planning algorithms such as multicriteria optimisation (MCO) [4] and combined volumetric arc therapy (VMAT) & intensity modulated (IMRT) [5–7] have also contributed to improving the quality of the treatment plans.

For patients with superficial tumours or large tumours with a superficial component, electrons have the potential to greatly reduce the dose delivered to surrounding normal tissue due to the limited range of electrons compared to photons. However, routine clinical adoption of advanced electron techniques such as modulated electron radiation therapy (MERT) has yet to become reality, in part due to logistical issues such as unwieldy electron-specific accessories (e.g., applicators, cerrobend cutouts, etc) and the limited dosimetric gains demonstrated in the literature. Typically, the superior normal tissue sparing of MERT plans comes at the expense of target coverage and homogeneity when compared to photon IMRT plans [8–10].

These issues were addressed in recent studies by demonstrating the viability of delivering modulated electron radiation therapy (MERT) with a photon multi-leaf collimator (MLC) to eliminate the accessory requirements (applicators, cut-outs, electron-specific MLC) of traditional electron RT [11, 10, 12, 13]. By delivering MERT plans with a standard photon MLC, the potential arises for seamless mixed electron-photon delivery without interrupting treatment delivery between modality changes. Indeed, mixed-beam electron-photon radiotherapy (MBRT) plans have been shown to offer additional degrees of freedom to influence the trade-off between normal tissue dose and target coverage [14–18].

However, the topic of plan robustness to patient breathing and positioning (setup) errors has not yet been investigated for MBRT. In conventional photon treatment plans the need for robustness is, in principle, circumvented by prescribing dose to a planning target volume (PTV) which includes margins around clinical target volumes (CTV) to account for setup and delivery uncertainties. The assumption is that even if the CTV moves within the PTV, it will remain covered to an acceptable dose level if the entire PTV is receiving the prescription dose. Nevertheless, it has been shown that for superficial targets such as breast, photon-only treatment plans could benefit from CTV-based robust optimisation [19].

Implicit in the practice of prescribing dose to the PTV is the assumption that the dose distribution itself is not excessively perturbed by patient positioning errors. This assumption does not hold for charged particles, as their dose deposition characteristics are more heavily influenced by small variations in depth caused by positioning errors. For example, intensity modulated proton therapy (IMPT) plans have been shown to degrade significantly when planned using PTV margins and tested for robustness by introducing positioning shifts [20–23]. There is therefore good reason to believe that PTV-based optimisation is also inappropriate for MBRT plans as they involve a large electron component.

In the context of particle therapy planning, two broad categories of robustness modelling have been well-studied: "minimax" and "stochastic programming" [23]. In the minimax formalism, a number of scenarios representing setup and range uncertainties are explicitly modelled and the robust optimisation model seeks to minimise the cost function of the scenario under which the treatment planning objectives are least satisfied [20]. Alternative minimax formulations instead operate on the worst possible dose distribution constructed from combining the worst voxel doses among all scenarios [24]. The stochastic approach instead optimises the expectation value of the cost function with a probability assigned to each scenario [23].

In this work, we extend the column generation-based optimisation model for MBRT presented in Renaud et al. [16] to produce CTV-based robust plans by deriving the pricing problem for the two common mathematical formulations of robustness. We then create CTV-based robust plans for soft tissue sarcoma (STS) and chest wall irradiation with deep-seated internal mammary, axillary and supraclavicular nodes (CW-N) patient cases and compare them to conventional nonrobust PTV-based plans with the goal of determining if robust optimisation is necessary for accurate planning of MBRT.

7.4 Theory

We treat the optimisation problem using the column generation method, a method by which a solution to the master problem (MP) is approximated by repeatedly updating and solving a restricted master problem (RMP) involving only a subset of all possible decision variables. In our case, the solution to the MP is a set of MLC aperture shapes and intensities, along with the resulting dose distribution. At each iteration, the optimal dose distribution under the RMP is then used to construct new apertures to be added to the RMP for the next iteration. The work of Romeijn et al. [25] shows a derivation of the column generation method applied to radiation therapy planning, and in Renaud et al. [16], we applied the method to mixed electron-photon optimisation. In this section, we extend the formalism for MBRT derived in our previous work and we elaborate on the differences between conventional (PTV-based) optimisation and robust (CTV-based) optimisation.

Table 7.1 Symbols used to describe the optimisation problem

K	Set of all possible MLC-defined apertures
$\hat{K}$	Subset of K containing only apertures included in the restricted master problem
S	Set of contoured structures
C	Set of positioning shift scenarios considered
V	Set of all voxels where dose values are specified
V_s	Set of voxels representing the discretisation of structure $s \in S$
A	Set of all beamlets generated
A_k	Set of beamlets included in aperture $k \in K$
D_{njc}	Dose contribution from beamlet $n \in A$ to voxel $j \in V$ for scenario $c \in C$
D_{kjc}	Dose contribution from aperture $k \in \hat{K}$ to voxel $j \in V$ for scenario $c \in C$
D_{jc}	Total dose contribution to voxel $j \in V$ for scenario $c \in C$
y_k	Intensity of aperture $k \in K$
$F_c(\mathbf{y})$	Cost function representing treatment planning goals for scenario $c \in C$
P_{kc}	Price of aperture $k \notin \hat{K}$ for scenario $c \in C$
P_n	Price of beamlet $n \in A$

7.4.1 Stochastic formulation of robustness

Our goal is to obtain a set of aperture shapes and intensities yielding a dose distribution that is robust to a set of delivery uncertainties. Using the stochastic formulation of robustness, we seek to minimise a weighted sum of cost functions among all scenarios. In mathematical terms, if

we consider a general form of the restricted master problem for non-robust optimisation,

$$\begin{aligned} \min_{\boldsymbol{y}} \quad & F(\boldsymbol{y}) \\ \text{subject to} \quad & y_k \geq 0,\ k \in \hat{K} \end{aligned} \tag{7.1}$$

then for robust optimisation, we simply define $F(\boldsymbol{y}) = \sum_{c \in C} \omega_c F_c(\boldsymbol{y})$, where F_c is the cost function for scenario c and ω_c represents the probability of encountering scenario c. As the probability of each scenario is typically unknown, ω_c can also be interpreted as an importance weighting factor on scenario c set by the planner. The RMP for the stochastic formulation is given by

$$\begin{aligned} \min_{\boldsymbol{y}} \quad & \sum_{c \in C} \omega_c F_c(\boldsymbol{y}) \\ \text{subject to} \quad & y_k \geq 0,\ k \in \hat{K} \end{aligned} \tag{7.2}$$

7.4.2 Minimax formulation of robustness

Under the minimax "worst case" formulation, we instead seek a solution to the following problem

$$\begin{aligned} \min_{\boldsymbol{y}} \quad & \max(F_1(\boldsymbol{y}), .., F_{|C|}(\boldsymbol{y})) \\ \text{subject to} \quad & y_k \geq 0,\ k \in \hat{K} \end{aligned} \tag{7.3}$$

where the set C represents the scenarios under which the treatment plan should be robust. This problem can be solved using a standard method by finding the solution to a modification of (7.3) which introduces an additional decision variable, t, and $|C|$ additional constraints:

$$\begin{aligned} \min_{(t,\boldsymbol{y})} \quad & t \\ \text{subject to} \quad & y_k \geq 0,\ k \in K \\ & t \geq F_c(\boldsymbol{y}),\ c \in C \end{aligned} \tag{7.4}$$

7.4.3 Pricing problem

As part of the column generation method, the pricing problem (PP) is a sub-problem solved every iteration after a solution to the RMP has been found. Solving the PP yields one or more apertures to be added to the RMP with the potential of improving the cost function. For both

models presented in this work, we will assume a voxel-based cost function of the form

$$F_c(\boldsymbol{y}) = \sum_{s}^{|S|} \sum_{j}^{|V_s|} F_{sjc}\left(D_{jc}(\boldsymbol{y})\right) \tag{7.5}$$

where the dose in voxel j from scenario c, D_{jc}, is the sum of individual aperture contributions to voxel j in scenario c.

$$D_{jc}(\boldsymbol{y}) = \sum_{k \in K} D_{kjc}\, y_k \tag{7.6}$$

Note that we can extend the sum to be over the entire set of possible apertures K instead of the restricted set $\hat{K}$ since, implicitly, $y_k = 0 \; \forall \; k \notin \hat{K}$.

Stochastic model pricing problem

The pricing problem is derived by applying the Karush-Kuhn-Tucker (KKT) conditions [26] to the MP, assuming a current solution $(\hat{\boldsymbol{y}})$ to the RMP. The stationarity condition is:

$$\nabla \left(\sum_{c \in C} \omega_c F_c(\boldsymbol{y}) - \sum_{k \in K} \rho_k y_k \right) = 0 \tag{7.7}$$

where ρ_k is the KKT multiplier associated with the non-negativity constraint of y_k. Evaluating the stationarity condition on our model yields the following restriction on an optimal solution:

$$\rho_k = \sum_{c \in C} \omega_c \partial_{y_k} F_c(\boldsymbol{y}), k \in K \tag{7.8}$$

this restriction is then combined with the feasibility and complementarity conditions:

$$\begin{aligned} y_k &\geq 0, k \in K \\ \rho_k &\geq 0, k \in K \\ \rho_k y_k &= 0, k \in K \end{aligned} \tag{7.9}$$

If the cost function is convex, then these conditions are necessary and sufficient for optimality. Proceeding as done in the work of Romeijn [25], we assume that we have a solution $\hat{y}$ to the RMP, implicitly setting $y_k = 0 \; \forall \; k \notin \hat{K}$. The only non-trivial optimality condition for $k \notin \hat{K}$ is

$$\rho_k = \sum_{c \in C} \omega_c \partial_{y_k} F_c(\hat{\boldsymbol{y}}) \geq 0 \tag{7.10}$$

Using the voxel-based form of the cost function defined in eq. 7.5, we can rewrite eq. 7.10 as

$$\begin{aligned}\sum_{c\in C}\omega_c\partial_{y_k}F_c(\hat{\boldsymbol{y}}) &= \sum_{c\in C}\omega_c\sum_{j=1}^{|V|}\left(\partial_{D_{jc}}F_{jc}(D_{jc})\right)\left(\partial_{y_k}D_{jc}(\boldsymbol{y})\right)\\ &= \sum_{c\in C}\omega_c\sum_{j=1}^{|V|}\left(\partial_{D_{jc}}F_{jc}(D_{jc})\right)D_{kjc}\geq 0\end{aligned} \tag{7.11}$$

Define

$$P_{kc}(\hat{\boldsymbol{y}}) = \sum_{j=1}^{|V|}\left(\partial_{D_{jc}}F_{jc}(D_{jc})\right)D_{kjc} \tag{7.12}$$

to be the price of a given aperture $k\notin\hat{K}$ for scenario c. Given a set of decision variables (aperture intensities) $\hat{\boldsymbol{y}}$ which solve the RMP, we seek to construct a new aperture ($k\notin\hat{K}$) which yields the lowest (most negative) price P_k:

$$\min_{k\in K}P_k = \min_{k\in K}\quad\sum_{c\in C}\omega_c P_{kc} \tag{7.13}$$

If no aperture yields a price below 0, then from equation 7.10 we conclude that the solution is optimal and no more apertures can improve the cost function. Note that in the single-scenario ($|C|=1$) case, the non-robust pricing problem as derived in Romeijn et al. [25] is recovered.

Minimax pricing problem

Applying the KKT conditions to eq. 7.4 instead of eq. 7.2, the stationarity condition is:

$$\nabla\left(t+\sum_{c\in C}\mu_c(F_c(\boldsymbol{y})-t)-\sum_{k\in K}\rho_k y_k\right)=0 \tag{7.14}$$

where ρ_k is, once again, the KKT multiplier associated with the non-negativity constraint of y_k, and μ_c is the KKT multiplier associated with the $t\geq F_c(\boldsymbol{y})$ constraint in eq. 7.4. The stationarity condition yields the following restrictions on an optimal solution:

$$\begin{aligned}\sum_{c\in C}\mu_c &= 1,\\ \rho_k &= \sum_{c\in C}\mu_c\partial_{y_k}F_c(\boldsymbol{y}),\ k\in K\end{aligned} \tag{7.15}$$

These restrictions are then combined with the feasibility and complementarity conditions

$$\begin{aligned} t - F_c(\boldsymbol{y}) &\geq 0, c \in C \\ \mu_c &\geq 0, c \in C \\ \mu_c(t - F_c(\boldsymbol{y})) &= 0, c \in C \\ y_k &\geq 0, k \in K \\ \rho_k &\geq 0, k \in K \\ \rho_k y_k &= 0, k \in K \end{aligned} \tag{7.16}$$

The only non-trivial optimality condition for $k \notin \hat{K}$ is

$$\rho_k = \sum_{c \in C} \mu_c \partial_{y_k} F_c(\hat{\boldsymbol{y}}) \geq 0 \tag{7.17}$$

Once again using the form of the cost function from eq. 7.5 and the definition of P_{kc} in eq. 7.12, we can rewrite eq. 7.17 as

$$\sum_{c \in C} \mu_c \partial_{y_k} F_c(\hat{\boldsymbol{y}}) = \sum_{c \in C} \mu_c P_{kc} \geq 0 \tag{7.18}$$

Given a set of decision variables $(\hat{t}, \hat{\boldsymbol{y}})$ which solve the RMP, we seek to construct an aperture $k \notin \hat{K}$ which solves the following pricing problem

$$\min_{k \in K} P_k = \min_{k \in K} \sum_{c \in C} \mu_c P_{kc} \tag{7.19}$$

Note that the minimax model pricing problem in eq. 7.19 is identical to the stochastic model pricing problem, with the exception that the probabilities ω_c in the stochastic problem are replaced by the KKT multipliers μ_c which are obtained at every iteration from solving the RMP instead of being user-defined.

The restrictions on μ_c from equations 7.15 and 7.16, $\sum_{c \in C} \mu_c = 1$, $\mu_c \geq 0$ and $\mu_c(t - F_c(\boldsymbol{y})) = 0, c \in C$, imply that only the KKT multipliers associated with the worst (largest) scenario cost functions will be nonzero. Hence, the aperture obtained from the minimax pricing problem is the one with the largest cost function improvement potential to the worst scenario for the current iteration. Note that in the single-scenario ($|C| = 1$) case, the non-robust pricing problem is recovered once again.

Solving the pricing problem using a beamlet-based approach

In a beamlet-based approach, dose distributions are generated for a grid of small beamlets for every beam angle and modality. Beamlets represent fragments of an aperture which, in practice, is collimated using an MLC. The total aperture dose distribution for scenario c is the sum of individual beamlet dose distributions from the beamlets included in the aperture:

$$D_{kjc} = \sum_{n \in A_k} D_{njc}, \tag{7.20}$$

where A_k is the set of beamlets included in aperture k. The price $P(A_k)$ associated with adding an aperture k to the master problem in the stochastic model defined in eq. 7.13 can be rewritten with the beamlet-based aperture dose defined in eq. 7.20,

$$P(A_k) = \sum_{c \in C} \omega_c \sum_{j=1}^{|V|} (\partial_{D_{jc}} F_{jc}(D_{jc}))\, D_{kjc} = \sum_{c \in C} \omega_c \sum_{j=1}^{|V|} (\partial_{D_{jc}} F_{jc}(D_{jc})) \left(\sum_{n \in A_k} D_{njc} \right) \tag{7.21}$$

which, after rearranging the terms, eq. 7.21 can be expressed as a sum over the price contribution from beamlets contained in the aperture:

$$P(A_k) = \sum_{n \in A_k} P_n \tag{7.22}$$

where

$$P_n = \sum_{c \in C} \omega_c \sum_{j=1}^{|V|} (\partial_{D_{jc}} F_{jc}(D_{jc}))\, D_{njc} \tag{7.23}$$

Assuming a grid-based beamlet arrangement where each row is collimated by an MLC leaf pair, the aim is to find the consecutive set of beamlets for which the price is lowest along a given row, and repeat for all beamlet rows to produce an aperture. If the minimum price of a row is positive, then the MLC leaves are completely closed for this row. An efficient algorithm for solving the pricing problem is presented in Romeijn et al. [25]. The pricing problem for the minimax model is solved identically, substituting ω_c by μ_c.

7.5 Methods

7.5.1 Aperture addition schemes

Aperture addition schemes are a concept specific to the column generation method. The pricing problem is solved for every "beam", meaning that one aperture is obtained per beam angle, per modality. Any aperture with a negative price has the potential to improve the cost function of the RMP, therefore we can choose an aperture addition scheme to add a subset of the apertures obtained from the pricing problem to the RMP as a heuristic to limit plan complexity. In previous work [16], we studied differences in the final plan for four aperture addition schemes. In this work, unless stated otherwise, we used the "best per modality" scheme, in which the lowest priced photon aperture along with the lowest priced aperture of each individual electron energies are incorporated to the RMP at every iteration.

7.5.2 Cost function

As indicated in equation 7.5, we employed voxel-based cost functions that are defined based on the contoured structure that the voxel belongs to. Individual voxels were forced to belong only to a single structure. A user-defined priority was defined for overlapping structures such that voxels in the overlapping regions would be assigned to the highest priority structure. Upper and lower dose limits were implemented as one-sided quadratic penalty functions:

$$F^{-}_{sjc}(D_{jc}) = \frac{\alpha^s}{|V_s|}\max(0, T^{s-} - D_{jc})^2 \tag{7.24}$$

$$F^{+}_{sjc}(D_{jc}) = \frac{\alpha^s}{|V_s|}\max(0, D_{jc} - T^{s+})^2 \tag{7.25}$$

where α^s is a user-defined weight for the penalty function of structure s, and T is the threshold dose before a penalty is applied.

In addition to the standard one-sided quadratic penalty functions, we also included dose volume penalties implemented as described by Bortfeld [27] and Wu [28]. The formulation $V^s(>T) < V_m$, which is interpreted as "the volume of structure s receiving dose greater than T must be less than V_m", is represented in the optimisation model as

$$F^{DV}_{sjmc}(D_{jc}) = \frac{\alpha^s_m}{|V_s|}\Theta\left(D_{V_m} - D_{jc}\right) \cdot \max(0, D_{jc} - T_m)^2 \tag{7.26}$$

where D_{V_m} is the current point on the dose volume histogram such that $V(D_{V_m}) = V_m$. Θ is the Heaviside function. This formulation ensures that only voxels with dose values between T and D_{V_m} are penalised. The subscript m is added to emphasise that multiple dose volume constraints can be specified for the same structure. All other symbols preserve the same definition as in table 7.1 and equations 7.24 and 7.25. Note that this formulation of DVH penalty functions is non-convex and therefore the algorithm is not guaranteed to converge to a global minimum. Convex formulations of partial volume constraints exist and have been shown to be superior in a multicriteria optimisation setting [29]. We nonetheless chose to use the non-convex formulation to keep the planning process intuitive, and we expect the conclusions reached in this work to be unaffected by the choice of partial volume objective formulation.

The total cost function for a scenario is given in equation 7.27

$$F_c(\boldsymbol{y}) = \sum_{s}^{|S|} \left(\sum_{j}^{|V_s|} \left(F_{sjc}^{-}\left(D_{cj}(\boldsymbol{y})\right) + F_{sjc}^{+}\left(D_{cj}(\boldsymbol{y})\right) \right) + \sum_{m} \sum_{j}^{|V_s|} F_{sjmc}^{DV}\left(D_{cj}(\boldsymbol{y})\right) \right) \tag{7.27}$$

with the understanding that some structures may have only a subset of the possible penalty functions defined. For example, normal tissue structures would not have any lower bound (F_{sjc}^{+}) penalties assigned to them.

To solve the RMP at each iteration of the column generation method, an open-source C++ large scale nonlinear optimisation solver based on the interior point method (IPOPT) was used [30].

7.5.3 Example cases

The robust optimisation models were compared to the PTV-based approach on a STS case and a CW-N case. For the robust plans, target dose coverage penalties were specified to the CTV, whereas for the non-robust plans, they were applied to the PTV. The comparison and dosimetric analysis between robust and non-robust plans is always performed on the CTV, as the goal is to determine if the CTV remains adequately covered under setup errors.

For the example cases, a total of seven independent positioning shift scenarios were considered: patient positioning shifts of ± 5 *mm* $\hat{\boldsymbol{x}}$ (right-left), ± 5 *mm* $\hat{\boldsymbol{y}}$ (anterior-posterior), $\pm$ 5 *mm* $\hat{\boldsymbol{z}}$ (cranio-caudal) and a nominal, unperturbed scenario. The shifts corresponded to the margin

added to the CTV to create the PTV used clinically to treat these patients. In the interest of reducing computational time, combined shifts were not considered, as each scenario required the generation of a full set of beamlets for both photons and electrons. We also assigned equal weight to each scenario in the stochastic model.

For both the STS and CW-N cases, the PTV was constructed by extending the CTV with a 5 mm margin and further cropping 2 mm from the skin contour. For the CW-N case, the PTV was additionally cropped at the lung interface. For planning purposes, the clinically drawn CTV was also cropped 2 mm from the skin contour.

Photon beamlets were generated with a 6 MV beam model using an in-house collapsed cone convolution superposition (CCCS) engine. Electron beamlets were generated with energies of 6, 9, 12, 16 and 20 MeV using DOSXYZnrc [31] with phase space files generated from a BEAMnrc [32] electron Monte Carlo beam model of the Varian TrueBeam accelerator (Varian Medical Systems, Palo Alto, CA). The MC beam model was modified from a previously commissioned Varian Clinac 2100 EX model [33] to match TrueBeam electron measured data. The phase space files were generated for a 40×40 cm^2 field and manually split into 1×1 cm^2 subfields to produce the electron beamlets [10].

Case	**Photon beams**
STS	Arc [-160, 40]°, 100 cm SAD
CW	IMRT [-60, 200]° every 20°, 100 cm SAD
Case	**Electron beams**
STS	-60, -30, 0, 30, 60°, $\approx$ 80 cm SSD
CW	0, 22.5, 45, 67.5, 90°, $\approx$ 80 cm SSD

Table 7.2 Photon and electron beam arrangement and gantry angles for the CW-N and STS cases

The CW-N and STS beam angles are indicated in table 7.2. Only beamlets falling within the PTV from a beam's eye view (BEV) perspective were generated, with additional beamlets along the perimeter of the BEV projection of the PTV to ensure that the target is entirely covered. For both cases, 0.5 x 0.5 cm^2 photon beamlets were generated, while electron beamlets were kept at 1 x 1 cm^2. For the STS case, a single isocenter was used for all electron beam angles to simplify a planned future delivery and quality assurance process. The photon component of the STS case was planned as a VMAT arc, taking into account MLC leaf travel constraints, whereas the CW-N photon component was planned as a step-and-shoot IMRT delivery. The electron component was planned as step-and-shoot MERT delivery for both cases. Tables of

optimisation constraints on anatomical structures are included in the supplementary material, tables B1 and B2 for the STS and CW-N cases, respectively.

The final dose distributions for both example cases were normalised such that 95% of the volume of the PTV received 50 Gy in the PTV-based optimisation, and 95% of the CTV volume received 50 Gy in the nominal scenario for the CTV-based robust optimisation. The robustness analysis is performed on dose distributions obtained from beamlet-based apertures. For an analysis of robust MBRT delivery, which will be the topic of future work, a final MC recalculation and re-weighting of each aperture will be necessary to obtain the monitor units for each aperture [10].

The workflow for producing a treatment plan proceeds as follows

---

Require: Generated beamlets for all beam angles, modalities and positioning shift scenarios

1: Initialise $\hat{K} = \varnothing$, $D_{jc} = 0 \; \forall j \in V$
2: **while** $|\hat{K}|$ is below the user-specified number of apertures desired **do**
3: Calculate $\partial_{D_{jc}} F_{jc}(D_{jc}) \; \forall c \in C$ for the current optimal dose distribution
4: Solve the PP for each beam angle and each modality
5: **if** PP $\geq 0 \; \forall \; K \notin \hat{K}$ **then return** $(\hat{K}, \boldsymbol{y})$
6: **end if**
7: Add candidate apertures to $\hat{K}$ depending on aperture addition scheme
8: Solve the RMP to obtain a new set of y_k weights
9: **end while**
10: **return** $(\hat{K}, \boldsymbol{y})$

7.6 Results

Plan evaluation and comparison between robust CTV-based plans and PTV-based plans was done through DVH bands and 3D dose distribution. For a given contoured structure, the thickness of the DVH band was determined by the extrema of DVH values among all positioning shift scenarios, with thinner bands indicating a more robust plan. All CTV-based robust optimisations were performed using the stochastic optimisation model. A brief discussion of stochastic optimisation versus minimax models is presented in section 7.6.4.

7.6.1 Soft tissue sarcoma

Apertures for the photon arc were determined at every 2 degree increment throughout the arc. The total number of electron apertures was limited to 50 for both PTV-based and CTV-based robust optimisations. Figure 7.1a compares DVH curves for the target and OARs for CTV-based robust and PTV-based plans. Figure 7.1b shows superimposed isodose lines for the CTV-based robust and PTV-based plans for the 95% prescription dose level to illustrate the differences in coverage between the plans. Figure 7.2 highlights differences in coverage between the

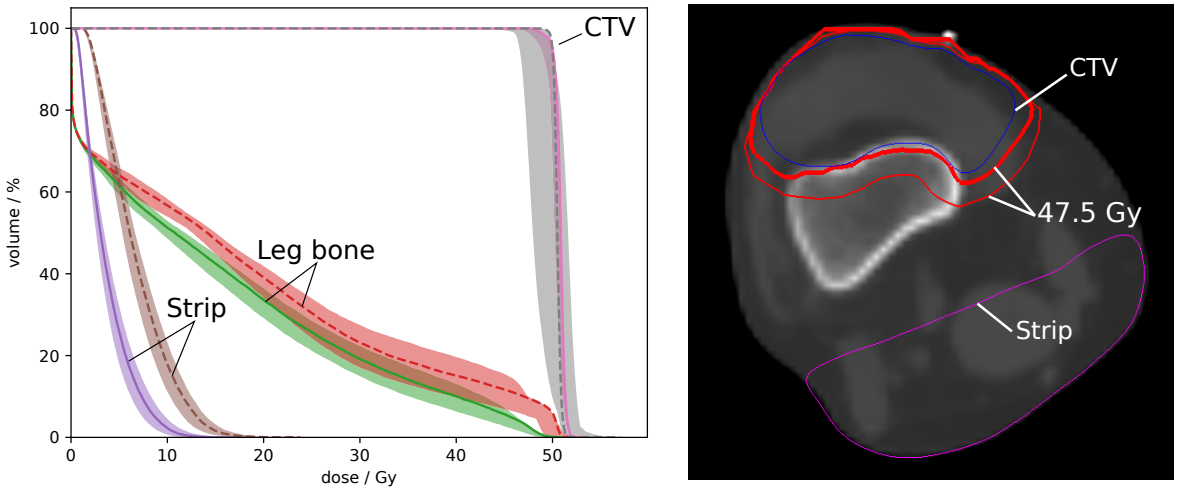

Figure 7.1 (a) DVH bands for the target and OARs comparing the PTV-based plan (dashed lines) and the CTV-based robust plan (full lines). (b) Isodose comparison between the total dose distribution of the CTV-based robust plan and the PTV-based plan. The 95% (47.5 Gy) isodose line for both plans is shown as a thick and thin red line for the CTV-based robust plan and the PTV-based plan, respectively. The strip of normal tissue indicated in (b) was used to reduce the dose to that area and control the incidence of lymphoedema.

CTV-based robust plans and the PTV-based plans under positioning shifts. Figures 7.2a and 7.2b show the dose colourwash for the PTV-based optimisation for the scenarios involving lateral shifts of $\pm$ 5 mm, whereas Figures 7.2c and 7.2d show the equivalent scenario dose colourwash for the CTV-based robust optimisation. The colourwash scale in all panels of Figure 7.2 is set to [95%, 108%] of the prescription dose to emphasise areas of under- and over-dosing induced by the positioning shifts. Figure 7.3 shows line profiles comparing the composition of the total dose distribution for the PTV-based and CTV-based robust plans at different locations in the target, along the simulated shift directions. Figure 7.3a displays the two lines, P1 and

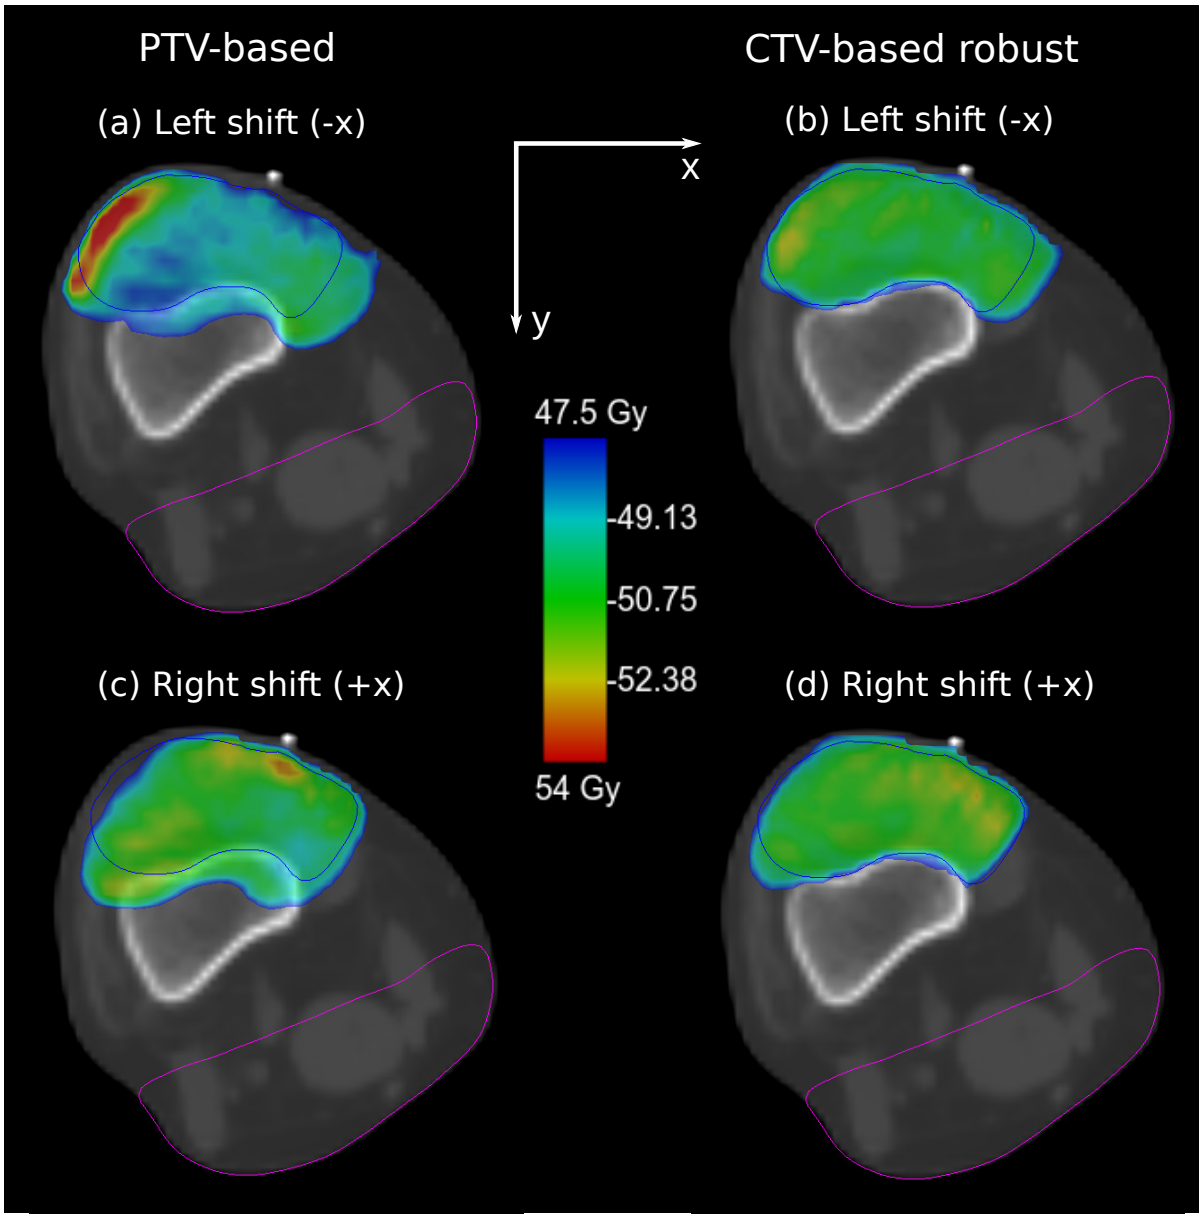

Figure 7.2 Colourwash comparison between the total dose distributions of the STS PTV-based and CTV-based robust plans for the indicated shifts. Note that the colourwash scale is set to [95%, 108%] to emphasise cold and hot spots. The shifts represent patient shifts relative to a fixed isocenter position.

P2, along which the line profiles are taken, and the 15 Gy dose sleeve representing the electron contribution to the total dose with both optimisation approaches. The line profiles for lines P1 and P2 are shown in Figures 7.3b and 7.3c, respectively.

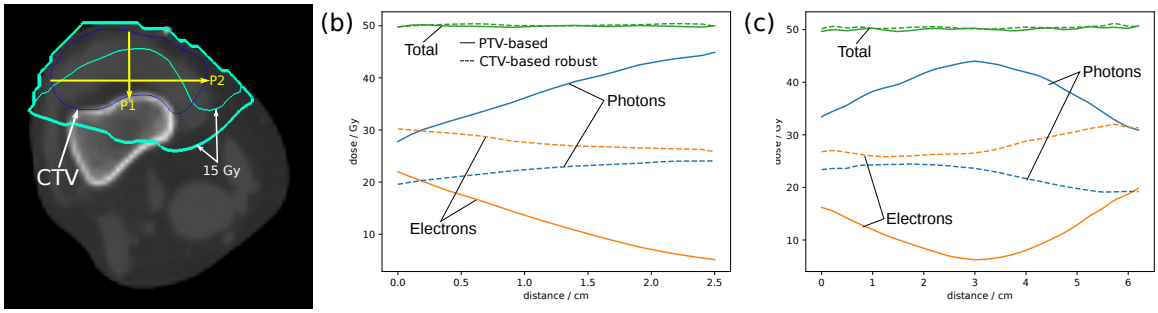

Figure 7.3 (a) Axial CT slice showing the position of lines P1 and P2 along which line profiles were calculated. An isodose comparison of the electron component of the PTV-based and CTV-based robust plans is shown on the same slice. The thick and thin 15 Gy isodose lines are shown for the electron component of the CTV-based robust and PTV-based plans, respectively. Line profiles for PTV-based (full lines) and CTV-based robust (dashed lines) plans decomposed by particle type for (b) line P1 and (c) line P2.

7.6.2 Chest wall and nodal irradiation

Unlike the STS case, where the target is fairly uniform in shape, the CW-N case features a long ($>$ 25 cm) target with distinct regions. For example, the chest wall region is relatively shallow and uniform in depth, whereas the internal mammary, axilla and supraclavicular nodal parts of the target are deep seated and irregular in shape. For both the PTV-based and CTV-based robust optimisations, the algorithm was stopped when 200 total apertures were included in the treatment plan.

Figure 7.4a compares DVH bands for the CW-N target volume and OARs. Figure 7.4b-d shows the 95% isodose lines for both plans at three distinct positions in the target: CW (panel b), axilla (panel c) and supraclavicular (panel d) nodes. Figure 7.5 illustrates how the 95% PTV dose coverage is compromised or how hot spots are generated under certain positioning shifts. Each panel is contrasted with the CTV-based robust plan dose distribution for the same positioning shifts showing that coverage is maintained. The colourwash scale spans [95%, 108%] of the prescription dose to emphasise areas of under- and over-dosing. Figure 7.6 shows line profiles comparing the individual contribution of electrons and photons to the total dose

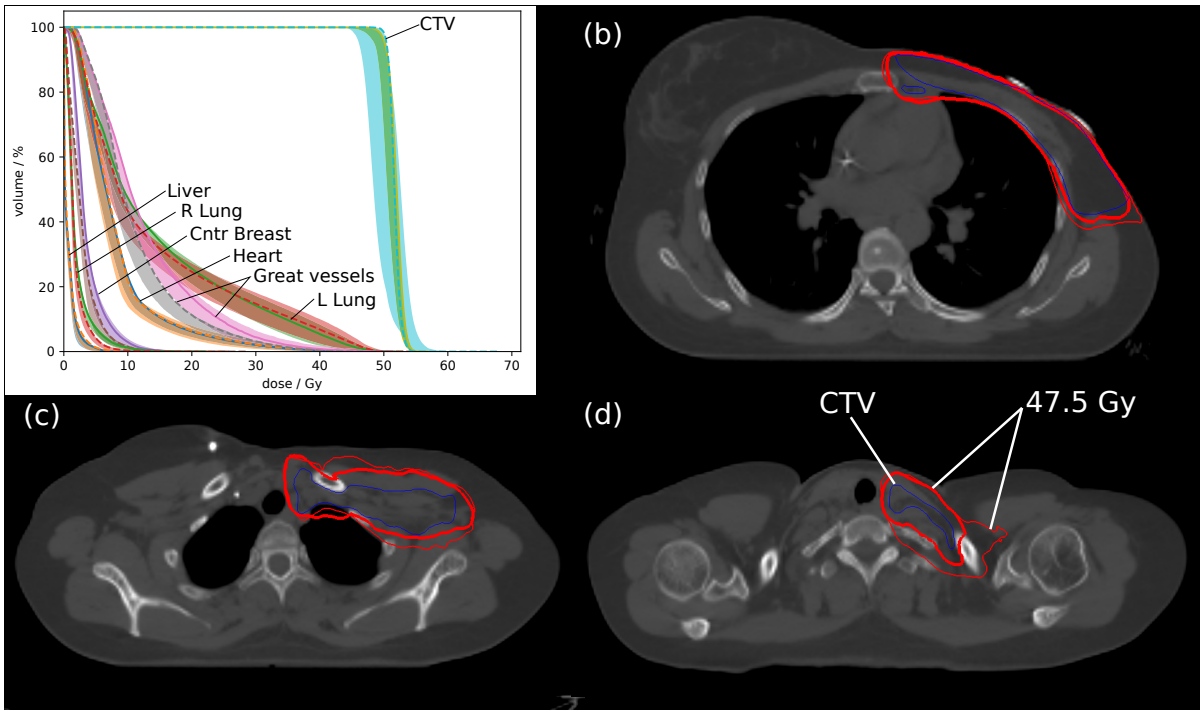

Figure 7.4 (a) DVH bands for the target volume and OARs. Dashed lines and full lines represent the nominal (no shift) scenario DVH curve for the PTV-based and CTV-based robust plans, respectively. (b-d) 95% (47.5 Gy) isodose lines of the total dose distributions in the (b) chest wall, (c) axilla and (d) supraclavicular node regions within the target volume. The thick and thin lines correspond to the CTV-based robust plan and the PTV-based plan, respectively. The CTV contour is indicated in blue.

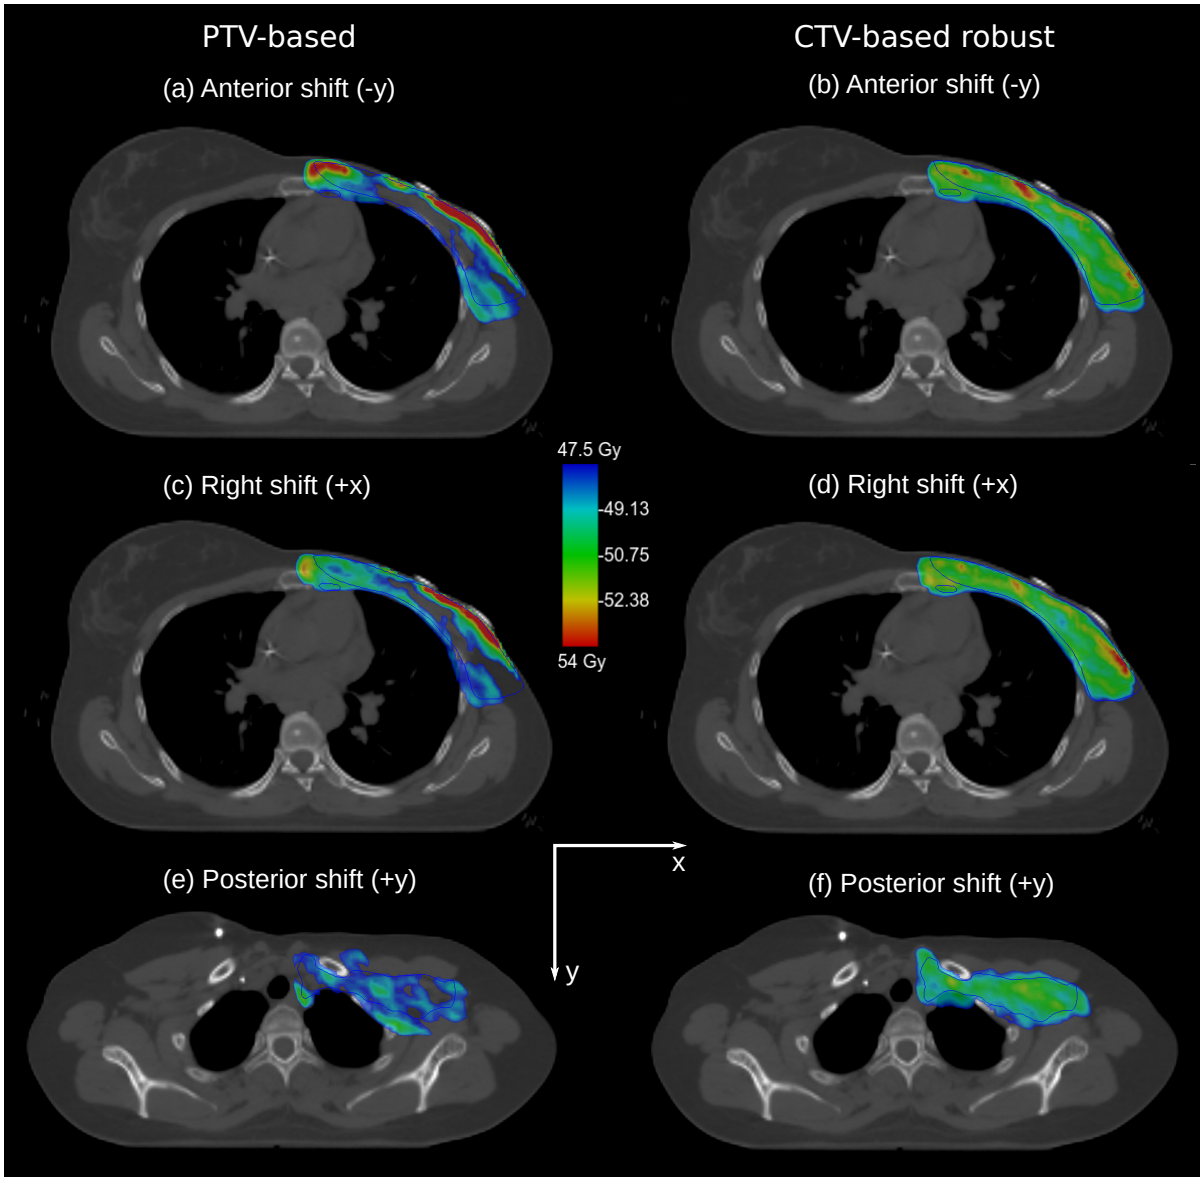

Figure 7.5 Colourwash comparison between the CW PTV-based and CTV-based robust plans for the indicated shifts. The colourwash scale is set to [95%, 108%] to emphasise cold and hot spots. The CTV contour is indicated in blue. The shifts represent patient shifts relative to a fixed isocenter position.

between PTV-based and CTV-based robust plans. Note that the electron dose distribution is the composite distribution of all electron energies. To further illustrate the differences between CTV-based robust and PTV-based plans, 20 Gy and 15 Gy isodose lines for the electron component of each plan are shown in Figures 7.6a,d.

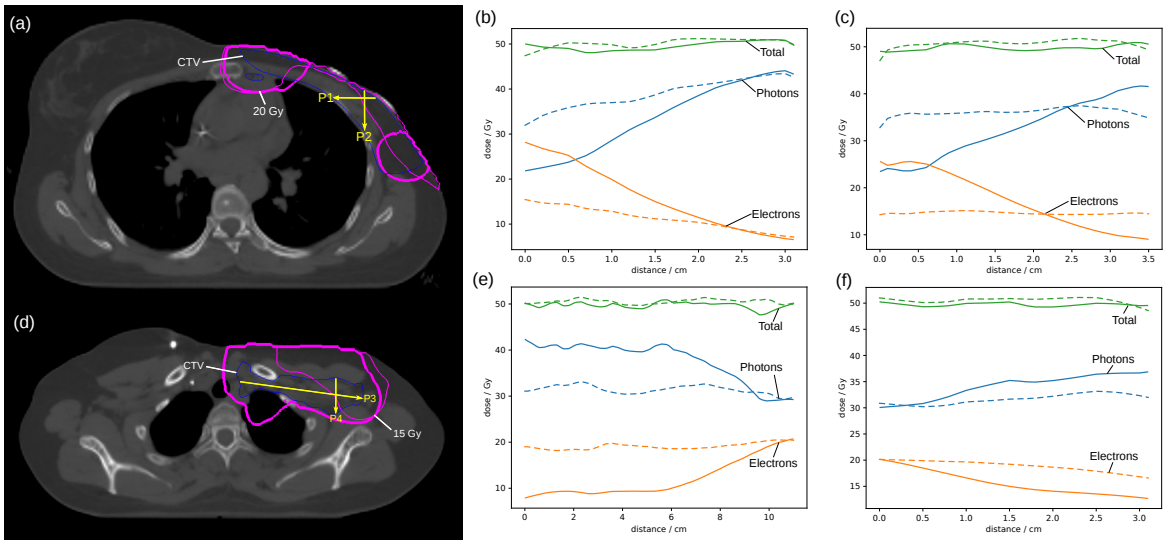

Figure 7.6 (a), (d) Axial CT slices showing lines P1-P4 along which line profiles were calculated. A 20 Gy and 15 Gy isodose curve comparison of the electron-only component of the CW-N plans is shown on the same slices. The thick and thin lines correspond to the CTV-based robust plan and the PTV-based plan, respectively. (b,c) Line profiles showing the contribution of electrons and photons to the total dose along lines P1 and P2, respectively. (e,f) Line profiles showing the contribution of electrons and photons to the total dose along lines P3 and P4, respectively.

7.6.3 Aperture addition schemes

In our previous work describing the PTV-based optimisation method for MBRT [16], the quality of plans generated using a series of different aperture addition schemes was investigated. The aperture addition scheme determines which types of apertures are added per iteration and the evolution of the cost function reduction is dependent on the addition scheme chosen[16]. For example, the "*best per modality*" scheme will include the highest quality aperture from each modality (i.e., five electron and one photon aperture per iteration), whereas the "*best per particle*" scheme will include both the best photon and the best electron apertures (two apertures per iteration). We briefly revisited the topic by looking at the robustness of plans

created under different aperture addition schemes. Figures 7.7a) and 7.7b) show DVH bands for STS plans created using both the "best per modality" and "best per particle" addition schemes with the PTV-based and CTV-based robust optimisation models, respectively. For PTV-based optimisation, the best per modality scheme resulted in less robust CTV coverage compared to the best per particle scheme. Under the CTV-based robust optimisation model, both schemes produced equally robust plans.

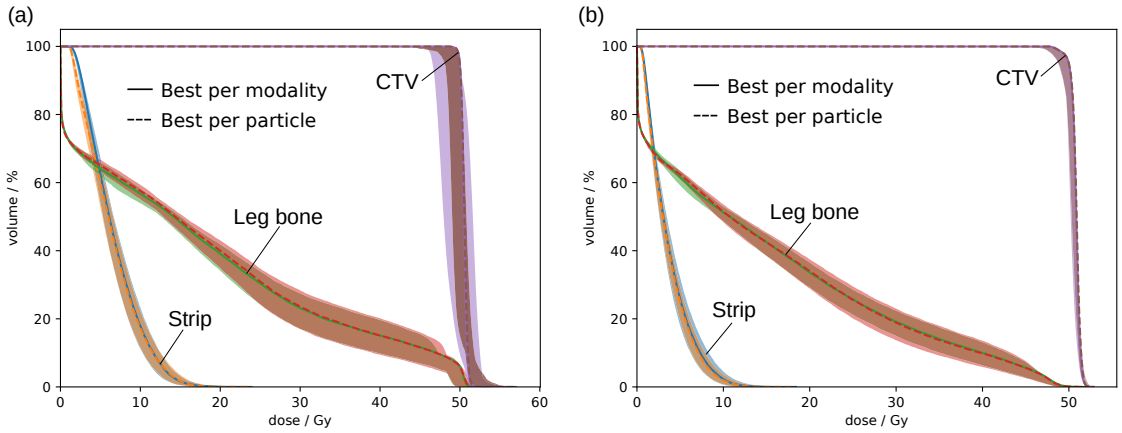

Figure 7.7 Comparison of STS DVH bands for (a) PTV-based non-robust and (b) CTV-based robust MBRT plans generated using the "best per modality" (full lines) and "best per particle" (dashed lines) aperture addition schemes described in [16]. The "best per particle" scheme added only two apertures per iteration, the lowest priced photon and electron apertures, as opposed to the "best per modality" which added the lowest priced aperture of each modality (6 apertures / iteration).

7.6.4 Minimax and stochastic optimisation models

The RMPs for both minimax and stochastic models have similar computational complexity. A nonlinear solver must evaluate the cost function, constraints, gradient, constraint Jacobian and Lagrangian Hessian at each step [30], and the aggregate computational cost of these evaluations is the same for both models. For example, although the cost function is simply the value of the t decision variable in the minimax case, each $F_c(\boldsymbol{y})$ must still be evaluated as part of the constraint calculations. The pricing problem can be solved more efficiently for the minimax model as only scenarios with non-zero values of μ_c contribute, however, the majority of the running time is spent solving the RMP at each iteration, rather than the PP.

Each scenario has an identical linear contribution to the models, either as an additional constraint in the case of the minimax model or simply as another component of the total cost function for the stochastic model. As additional scenarios do not affect the total number of decision variables in the problems, the overall running time of the algorithm simply increases linearly with the number of scenarios considered.

With respect to the quality of plans created using stochastic versus minimax modelling for robust optimisation, Unkelbach et al. state that "there is no comprehensive evidence that one method is generally superior" [23]. As shown in the Supplementary Material, Figure A1, the robust CTV-based plans created with the two models are near-identical for the cases presented in this work. The purpose of this paper is to present a robust optimisation methodology for MBRT and gain an intuition for how robustness is achieved with this delivery technique. Since both models produced near-identical plans for the cases studied, in the remainder of the paper, we opted to only discuss our results comparing the stochastic modelling approach to the PTV-based approach.

7.7 Discussion

7.7.1 Soft tissue sarcoma

In Figure 7.1a, we observe that the CTV-based robust optimisation resulted in a much tighter DVH band for the CTV, indicating that coverage does not degrade under positioning shifts chosen to correspond to the clinically used margin for the CTV to PTV expansion (5 mm). The CTV-based robust and PTV-based plans have comparable target coverage in the nominal scenario, but the DVH band for the PTV-based plan shows that positioning shifts resulted in a significant loss of dose homogeneity in the CTV. For the PTV-based plan, the positioning shift which resulted in the worst degradation in CTV coverage had a homogeneity index (HI = $D_{2\%}/D_{98\%}$) of 1.11 compared to 1.03 for the nominal, unshifted scenario. For the CTV-based plan, the worst HI was 1.08 compared to 1.05 in the nominal scenario.

The leg bone and the strip of normal tissue received less dose in the CTV-based robust plan due to the fact that achieving robust CTV coverage did not require a margin as large at the distal edge of the PTV. By taking advantage of the limited range of electrons, a significant component of the total dose can be delivered with an anterior-posterior (AP) beam. The dose delivered

from this angle is inherently robust to small positioning uncertainties along that beam's axis. For that reason, CTV coverage can be maintained without the nominal 5 mm margin.

The 95% isodose line in Figure 7.1b can be used to study the behaviour of the robust algorithm. In the PTV-based plan a high PTV conformality was achieved, as anticipated from the optimisation constraints. The 95% isodose line for the CTV-based robust plan extends to the PTV in the lateral edges of the tumour, but matches the CTV contour at depth, which is expected as electrons are naturally robust to positioning shifts in the beam direction. The robust algorithm, in essence, produced a customised margin which took into account the sensitivity of the dose distribution to the positioning shifts considered.

Figures 7.2a,c show that with PTV-based optimisation, a hot spot or loss of coverage occurred when the patient was shifted laterally to the left or laterally to the right, respectively. Figures 7.2b,d show that under the same shifts, the CTV-based robust optimisation maintained target coverage and homogeneity. This is in part due to the fact that there was no PTV margin available in the areas where the CTV extended close to the skin surface. Moreover, the CTV-based robust optimisation not only leads to larger apertures to ensure CTV coverage under setup errors, but also leads to a significant re-weighting of each modality within the plan such that dose profiles along shift directions are flatter, as shown in Figs. 7.3b, c).

7.7.2 Chest wall and nodal irradiation

Similarly to the STS case, Figure 7.4a shows that for the CW-N case, the CTV coverage was more robust when the effect of positioning shifts on the CTV were explicitly taken into account, as opposed to simply covering the PTV. The 95% isodose line of the CTV-based robust optimisation generally extends as far as the PTV in the lateral direction while remaining closer to the CTV at depth. The CTV-based robust plan eliminated scenarios where over-dosage above 108% of the prescription dose occur, and achieved improved coverage. Similar doses to OARs and DVH bands were obtained with both optimisation approaches.

Figures 7.5a,c show both under- and over-dosing within the target when 5 mm anterior and lateral right shifts were applied. Figure 7.5e shows a general degradation of coverage in the axilla region for a 5 mm posterior shift. The robust dose distributions corresponding to the same shifts, shown in Figures 7.5b,d,f, exhibit adequate target dose coverage and homogeneity.

7.7.3 Differences between non-robust and robust plans

To gain an intuition for the process by which the CTV-based algorithm achieves robustness, the total dose distributions were split into individual modality contributions. In Figures 7.3 and 7.6, line profiles were plotted to compare PTV-based and CTV-based robust plans in the nominal scenario. All electron energies were combined to allow one to contrast trends between photons and electrons. The profiles show that while the total dose remained identical between the two plans, the electron and photon contributions to the total dose were drastically different.

The robustness of the overall plan can be understood by considering the contribution to the total dose of different modalities summed over all beam angles, and studying the sensitivity to applied shifts. In the CTV-based robust plans, dose profiles for individual modalities taken along the direction of the applied shifts (Figures 7.3 and 7.6) show lower dose gradients within target boundaries compared to the PTV-based plans. That is, the robust optimiser attempted to cover the CTV with a uniform dose from each modality individually. This observation can be used to explain the hot spot created in the PTV-based plan when subjected to the 5 mm $+\hat{x}$ shift shown in Figure 7.2a.

In photon-only plans, achieving target homogeneity while sparing normal tissue requires multiple beam directions. By promoting electrons at the expense of photons, the optimiser is able to take advantage of the limited range of electrons to deliver a significant component of the total dose from a narrow range of beam angles, which renders the dose distribution robust to small positioning shifts along those beam axes. Figures 7.3b,c show the overall electron component going from less than 10% of the total dose in the PTV-based plan to over 50% in the CTV-based robust plan. To produce a uniform electron dose contribution at depth, higher electron energies were preferentially used in the CTV-based robust plans compared to the PTV-based plans.

Due to the magnitude of dose perturbations observed in PTV-based plans, we recommend including more positioning uncertainty scenarios such as diagonal shifts combining multiple axes in the robust optimisation process. For cases where the electron beams must penetrate through a significant length of bone material before reaching the target, range uncertainties originating from the CT-to-density and material assignment process may need to be taken into account as part of the robust optimisation process [34, 35]. However, in the cases presented in this work, the presence of bone is not expected to interfere with CTV coverage as it is typically downstream from the target.

Since the robustness analysis for this work has been performed on dose distributions obtained from beamlet-based apertures, an assessment of robust MBRT delivery will require a final MC recalculation and re-weighting of each aperture will be necessary to obtain proper monitor units [10]. However, this re-optimisation does not significantly affect the conclusions of this work with respect to robustness.

7.7.4 Aperture addition schemes

In Renaud et al. [16], we concluded that the choice of aperture addition scheme had little impact on the quality of treatment plans generated using the column generation method assuming enough apertures were included in the treatment plan. In other words, the cost function values of plans generated using different aperture addition schemes converged to a similar value. The recommendation was therefore to use the scheme which resulted in the lowest running time for the algorithm to produce a plan.

However, when studying the robustness of the PTV-based plans, a more nuanced picture emerged. Figure 7.7a shows that despite having indistinguishable DVH curves in the nominal positioning setup (the solid and dashed lines), the "best per modality" plan was less robust than the "best per particle" plan. There is an effective degeneracy in the modality composition of MBRT plans producing near-identical nominal treatment plans but with varying degrees of robustness. As demonstrated in Figure 8b of Renaud et al. [16], the best per modality aperture addition scheme resulted in a larger fraction of the dose delivered from lower energy electrons, which are more perturbed by positioning shifts.

This finding motivates the use of an algorithm which explicitly accounts for robustness when producing a MBRT treatment plan, such as the one presented in this work. Indeed, Figure 7.7b shows that when optimised robustly, the STS plans created with different aperture addition schemes were once again identical in quality, including robustness.

7.8 Conclusion

In this work we have extended our previous column generation-based MBRT optimisation model to produce plans which explicitly minimised the perturbation on the dose distribution induced by positioning errors. We have shown that both STS and CW-N plans were made robust using our optimisation models. In the STS case, the CTV-based robust plan resulted in less

dose being delivered to healthy tissue due to the fact that the CTV could be reliably covered under 5 mm positioning shifts without needing to cover the entire PTV.

The contribution to the total dose from electrons and photons differed greatly between non-robust and robust plans. For the same total dose, the robust plans had more uniform electron and photon dose profiles within the target volume. The fraction of total dose delivered by electrons was typically larger in the robust plans, in many cases overtaking the photon component.

Finally, when considering DVH curves, two PTV-based (nonrobust) plans for the same patient created with different aperture addition schemes can appear identical in the nominal scenario but have varying levels of robustness when all positioning shifts are considered. The difference in robustness disappeared when CTV-based robust optimisation was performed.

Similarly to the conclusion reached by studies investigating the robustness of IMPT plans, we conclude that PTV-based nonrobust optimisation is inferior to CTV-based robust optimisation at maintaining target coverage under positioning errors for MBRT.

Acknowledgments

The authors acknowledge support from the Fonds de recherche du Québec Nature et Technologie (FRQNT), the CREATE Medical Physics Research Training Network grant (number 432290) of the Natural Sciences and Engineering Research Council (NSERC) and the Canadian Institute of of Health Research Foundation Grant (FDN-143257).

Conflicts of interest

The authors have no conflicts to disclose.

References

[1] P. Dong, P. Lee, D. Ruan, T. Long, E. Romeijn, Y. Yang, D. Low, P. Kupelian, and K. Sheng, "4π non-coplanar liver SBRT: a novel delivery technique," *International Journal of Radiation Oncology *Biology* Physics*, vol. 85, no. 5, pp. 1360–1366, 2013.

[2] C. B. Locke and K. K. Bush, "Trajectory optimization in radiotherapy using sectioning (TORUS)," *Medical physics*, 2017.

[3] P. Dong, H. Liu, and L. Xing, "Monte Carlo tree search -based non-coplanar trajectory design for station parameter optimized radiation therapy (SPORT)," *Physics in Medicine & Biology*, vol. 63, no. 13, p. 135014, 2018.

[4] D. Craft, D. McQuaid, J. Wala, W. Chen, E. Salari, and T. Bortfeld, "Multicriteria VMAT optimization," *Medical physics*, vol. 39, no. 2, pp. 686–696, 2012.

[5] J. L. Robar and C. Thomas, "Hybridarc: A novel radiation therapy technique combining optimized dynamic arcs and intensity modulation," *Medical Dosimetry*, vol. 37, no. 4, pp. 358 – 368, 2012.

[6] R. Li and L. Xing, "An adaptive planning strategy for station parameter optimized radiation therapy (SPORT): Segmentally boosted VMAT," *Medical physics*, vol. 40, no. 5, 2013.

[7] A. W. M. Sharfo, M. L. Dirkx, S. Breedveld, A. M. Romero, and B. J. Heijmen, "Vmat plus a few computer-optimized non-coplanar imrt beams (vmat+) tested for liver sbrt," *Radiotherapy and Oncology*, vol. 123, no. 1, pp. 49–56, 2017.

[8] C.-M. Ma, M. Ding, J. S. Li, M. C. Lee, T. Pawlicki, and J. Deng, "A comparative dosimetric study on tangential photon beams, intensity-modulated radiation therapy (IMRT) and modulated electron radiotherapy (MERT) for breast cancer treatment," *Physics in Medicine & Biology*, vol. 48, no. 7, p. 909, 2003.

[9] A. Alexander, E. Soisson, M.-A. Renaud, and J. Seuntjens, "Direct aperture optimization for FLEC-based MERT and its application in mixed beam radiotherapy," *Medical physics*, vol. 39, no. 8, pp. 4820–4831, 2012.

[10] D. Henzen, P. Manser, D. Frei, W. Volken, H. Neuenschwander, E. J. Born, A. Joosten, K. Lössl, D. Aebersold, C. Chatelain, *et al.*, "Beamlet based direct aperture optimization for MERT using a photon MLC," *Medical physics*, vol. 41, no. 12, p. 121711, 2014.

[11] F. J. Salguero, B. Palma, R. Arrans, J. Rosello, and A. Leal, "Modulated electron radiotherapy treatment planning using a photon multileaf collimator for post-mastectomized chest walls," *Radiotherapy and Oncology*, vol. 93, no. 3, pp. 625–632, 2009.

[12] L. S. A. M., G. I. M., B. Magdalena, and Z. Sergei, "Validation of Varian TrueBeam electron phase-spaces for Monte Carlo simulation of MLC-shaped fields," *Medical Physics*, vol. 43, no. 6, pp. 2894–2903, 2016.

[13] S. Mueller, M. K. Fix, D. Henzen, D. Frei, D. Frauchiger, K. Loessl, M. F. Stampanoni, and P. Manser, "Electron beam collimation with a photon MLC for standard electron treatments," *Physics in Medicine & Biology*, vol. 63, no. 2, p. 025017, 2018.

[14] B. A. Palma, A. U. Sánchez, F. J. Salguero, R. Arráns, C. M. Sánchez, A. W. Zurita, M. I. R. Hermida, and A. Leal, "Combined modulated electron and photon beams planned by a Monte-Carlo-based optimization procedure for accelerated partial breast irradiation," *Physics in Medicine & Biology*, vol. 57, no. 5, p. 1191, 2012.

[15] C. Miguez, E. Jimenez-Ortega, B. A. Palma, H. Miras, A. Ureba, R. Arrans, F. Carrasco-Pena, A. Illescas-Vacas, and A. Leal, "Clinical implementation of combined modulated electron and photon beams with conventional mlc for accelerated partial breast irradiation," *Radiotherapy and Oncology*, vol. 124, no. 1, pp. 124 – 129, 2017.

[16] M.-A. Renaud, M. Serban, and J. Seuntjens, "On mixed electron-photon radiation therapy optimisation using the column generation approach," *Medical Physics*, vol. 44, no. 8, pp. 4287–4298, 2017.

[17] S. Müller, M. K. Fix, A. Joosten, D. Henzen, D. Frei, W. Volken, R. Küng, D. M. Aebersold, M. F. Stampanoni, and P. Manser, "Simultaneous optimization of photons and electrons for mixed beam radiotherapy," *Physics in Medicine & Biology*, vol. 62, no. 14, p. 5840, 2017.

[18] S. Mueller, P. Manser, W. Volken, D. Frei, R. Kueng, E. Herrmann, O. Elicin, D. M. Aebersold, M. F. M. Stampanoni, and M. K. Fix, "Part 2: Dynamic mixed beam radiotherapy (DYMBER): Photon dynamic trajectories combined with modulated electron beams," *Medical Physics*, vol. preprint, no. 0, 2018.

[19] M. Byrne, Y. Hu, and B. Archibald-Heeren, "Evaluation of raystation robust optimisation for superficial target coverage with setup variation in breast imrt," *Australasian physical & engineering sciences in medicine*, vol. 39, no. 3, pp. 705–716, 2016.

[20] F. Albin, F. Anders, and H. Björn, "Minimax optimization for handling range and setup uncertainties in proton therapy," *Medical Physics*, vol. 38, no. 3, pp. 1672–1684, 2011.

[21] W. Liu, X. Zhang, Y. Li, and R. Mohan, "Robust optimization of intensity modulated proton therapy," *Medical physics*, vol. 39, no. 2, pp. 1079–1091, 2012.

[22] Z. Maryam, C. Wenhua, L. Wei, K. Laleh, R. Sharmalee, M. Radhe, and L. Gino, "Comparison of linear and nonlinear programming approaches for "worst case dose" and "minmax" robust optimization of intensity-modulated proton therapy dose distributions," *Journal of Applied Clinical Medical Physics*, vol. 18, no. 2, pp. 15–25, 2017.

[23] J. Unkelbach and H. Paganetti, "Robust proton treatment planning: Physical and biological optimization," *Seminars in Radiation Oncology*, vol. 28, no. 2, pp. 88 – 96, 2018. Proton Radiation Therapy.

[24] D. Pflugfelder, J. J. Wilkens, and U. Oelfke, "Worst case optimization: a method to account for uncertainties in the optimization of intensity modulated proton therapy," *Physics in Medicine & Biology*, vol. 53, no. 6, p. 1689, 2008.

[25] H. E. Romeijn, R. K. Ahuja, J. F. Dempsey, and A. Kumar, "A column generation approach to radiation therapy treatment planning using aperture modulation," *SIAM Journal on Optimization*, vol. 15, no. 3, pp. 838–862, 2005.

[26] M. S. Bazaraa, H. D. Sherali, and C. M. Shetty, *Nonlinear programming : theory and algorithms*. Wiley-Interscience series in discrete mathematics and optimization, New Jersey, Hoboken: John Wiley & Sons, 3 ed., 2006.

[27] T. Bortfeld, J. Stein, and K. Preiser, "Clinically relevant intensity modulation optimization using physical criteria," in *Proceedings of the XIIth ICCR, Salt Lake City*, vol. 29, 1997.

[28] Q. Wu and R. Mohan, "Algorithms and functionality of an intensity modulated radiotherapy optimization system," *Medical physics*, vol. 27, no. 4, pp. 701–711, 2000.

[29] L. Engberg, A. Forsgren, K. Eriksson, and B. Hårdemark, "Explicit optimization of plan quality measures in intensity-modulated radiation therapy treatment planning," *Medical Physics*, 2017.

[30] 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.

[31] D. Rogers, I. Kawrakow, J. Seuntjens, B. Walters, and E. Mainegra-Hing, "NRC user codes for EGSnrc," *NRCC Report PIRS-702 (Rev. B)*, 2003.

[32] D. Rogers, B. Faddegon, G. Ding, C.-M. Ma, J. We, and T. Mackie, "BEAM: a Monte Carlo code to simulate radiotherapy treatment units," *Medical physics*, vol. 22, no. 5, pp. 503–524, 1995.

[33] T. Connell, A. Alexander, P. Papaconstadopoulos, M. Serban, S. Devic, and J. Seuntjens, "Delivery validation of an automated modulated electron radiotherapy plan," *Medical physics*, vol. 41, no. 6Part1, 2014.

[34] F. Verhaegen and S. Devic, "Sensitivity study for ct image use in monte carlo treatment planning," *Physics in Medicine & Biology*, vol. 50, no. 5, p. 937, 2005.

[35] B. Vanderstraeten, P. W. Chin, M. Fix, A. Leal, G. Mora, N. Reynaert, J. Seco, M. Soukup, E. Spezi, W. De Neve, *et al.*, "Conversion of ct numbers into tissue parameters for monte carlo dose calculations: a multi-centre study," *Physics in Medicine & Biology*, vol. 52, no. 3, p. 539, 2007.

Supplementary material

Robust mixed electron-photon radiation therapy optimisation
Supplementary data

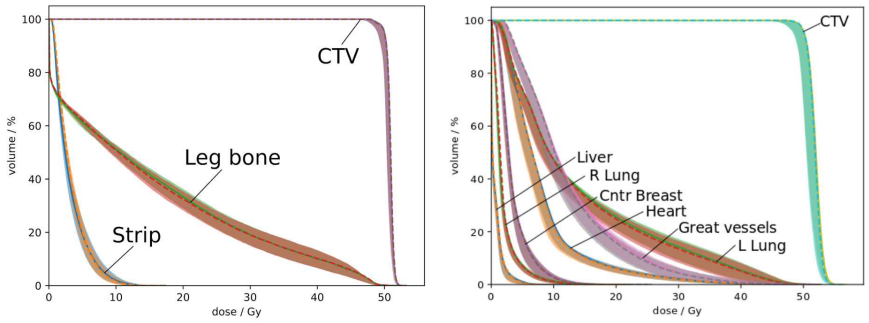

Fig A1: DVH band comparison between stochastic (full lines) and minimax (dashed lines) robust models for the STS case and the CWI case, respectively. The lines represent the DVH curve under the nominal scenario with no positioning shift. The plans are near identical for both cases presented in this work, hence we perform our comparisons to PTV-based plans with the stochastic model only.

The optimisation constraints used in the creation of treatment plans for the two cases presented in this work are given in the following tables. Only constraints on anatomical structures were are shown. In addition to anatomical structures, optimisation-specific structures were created to refine the dose distribution in areas where the anatomical structures did not reach.

STS case			
Structure	Bound	Dose limit	% volume
PTV	Lower	50	100
	Upper	50	0
RT Leg bone	Upper	45	12
	Upper	5	58
	Upper	10	45
Strip	Upper	20	0

Table B1: Optimisation constraints on anatomical structures used to produce the STS plans. For the CTV-based robust plan, the PTV constraint was replaced with an identical CTV constraint instead.

Robust mixed electron-photon radiation therapy optimisation
Supplementary data

Chestwall case			
Structure	Bound	Dose limit	% volume
PTV	Lower	50	100
	Upper	50	0
Contralat. breast	Upper	10	0.5
	Upper	5	5
	Upper	2.5	40
Heart	Upper	30	2
	Upper	20	4
	Upper	10	18
Left lung	Upper	30	10
	Upper	20	20
	Upper	10	38
Right lung	Upper	10	0
	Upper	5	2
	Upper	2.5	10
Greast vessels	Upper	30	2
	Upper	25	7
	Upper	20	23
	Upper	15	45
LAD	Upper	30	2
	Upper	20	24
	Upper	15	35
	Upper	10	65
Nerve Structures	Upper	50	0
	Upper	45	86.5
	Upper	35	93
Spinal Cord	Upper	34	0.5
Thyroid	Upper	40	7
	Upper	25	10
	Upper	15	15
Larynx	Upper	5	18
Esophagus	Upper	30	10
	Upper	20	15

Robust mixed electron-photon radiation therapy optimisation
Supplementary data

Liver	Upper	2	10

Table B2: Optimisation constraints on anatomical structures used to produce the CW-N plans. For the CTV-based robust plan, the PTV constraint was replaced with an identical CTV constraint instead.

Chapter 8

Delivery and quality assurance of mixed electron-photon radiation therapy plans

Marc-andré Renaud, Monica Serban and Jan Seuntjens

Article in preparation for submission to Medical Physics

8.1 Preface

The work presented in chapters 6 and 7 established that we successfully produced an algorithm capable of creating complex MBRT plans despite the large increase in degrees of freedom compared to photon-only RT. However, until the work presented in this chapter, the studies had been entirely simulated and were not experimentally delivered on real clinical accelerators. In this chapter, the aim was to validate that the MBRT algorithm produced deliverable treatment plans by comparing the planned doses to measurements in quality assurance (QA) conditions. In addition, to begin the clinical implementation process, we also set out to define the steps of a patient-specific QA procedure for clinical MBRT treatments.

8.2 Abstract

Purpose: In previous work, we have demonstrated that a mixed electron-photon radiation therapy (MBRT) planning algorithm can produce treatment plans which improve normal tissue

sparing without compromising target coverage, when compared to photon-only plans. The purpose of this work is to validate the MBRT delivery process on a Varian TrueBeam accelerator and define the steps for a patient-specific quality assurance (QA) protocol based on ion chamber point measurements.

Methods: To validate the MC beam models used to calculate the MBRT dose distributions for each modality in the planning process, a simple, single beam angle MBRT treatment plan was created and delivered on a slab of Solid Water with a film positioned at a depth of 3 cm. The measured film absorbed dose was compared to the calculated dose.

For patient-specific QA, a polymethyl methacrylate (PMMA) cylinder was machined and holes were made to fit an ionisation chamber. A complex MBRT plan involving a photon arc and three electron delivery angles was created with the aim of reproducing a clinically realistic dose distribution for a typical soft tissue sarcoma of the extremities. The treatment plan was delivered on the PMMA cylinder with an Exradin A1SL chamber placed at a nominal depth of 1.4 cm.

An existing EGSnrc user-code, SPRRZnrc, was modified to calculate stopping power ratios between any material in the same voxelised geometry used for dose calculation purposes. This modified code, called SPRXYZnrc, was used to calculate a k_{MBRT} correction factor accounting for the differences in electron fluence spectrum at the measurement point compared to that at reference conditions. With this approach, potential fluence perturbation correction factors at low electron beam energies were neglected.

Results: The film measurement from the Solid Water phantom treatment plan was in good agreement with the simulated dose distribution, with a gamma pass rate of 96.4 % for a 2 %/2 mm criteria. For the PMMA phantom delivery, the ion chamber measurements for each modality agreed to better than 2.7 %, and the total delivered dose agreed with the MC-simulated dose within 0.2 %. The beam quality correction factors amounted to, at most, a 4 % correction on the ion chamber measurement. The electron specific delivery of the ion chamber measurement procedure was realisable in less than 5 minutes, and the entire QA process including positioning was performed in less than 15 minutes.

Conclusion: The agreement between measured and simulated MBRT doses indicates that the dose distributions obtained from the MBRT treatment planning algorithm are realistically achievable. The SPRXYZnrc MC code allowed for convenient calculations of k_{MBRT} simultaneously with the dose distributions, making the proposed patient-specific QA protocol feasible for clinical implementation. However, further investigation is necessary to establish the accuracy

of our approach in calculating the ionisation chamber correction factors k_{MBRT} at low electron energies.

8.3 Introduction

The majority of patients undergoing external beam radiation therapy (EBRT) are treated with photons, despite medical linear accelerators also being capable of delivering electrons. All available photon and electron energies constitute modalities that could potentially be combined into a mixed beam (MBRT) plan to produce a superior treatment plan compared to single-energy (and single-particle) plans, provided that MBRT delivery is logistically feasible within the clinical workflow.

The modulated electron radiation therapy (MERT) literature has historically focused between studying MERT delivery using tertiary electron-specific collimators such as the eMLC [1–5] or the few-leaf electron collimator (FLEC) [6–8], and utilising the photon MLC (pMLC) already present in modern linacs [1, 9–12]. Early studies showed that a shortened source-to-surface distance (SSD), typically 70 cm, was necessary to produce clinically acceptable electron dose distributions due to the degradation of electron field penumbras in air [13]. Traditional electron RT remains cumbersome to deliver in comparison to photon RT, requiring custom patient-specific cut-outs, bolus and more time consuming setup. MLC-based MERT has seen limited adoption in the clinic due to additional time consuming tasks related to set-up and commissioning compared to conventional photon RT, and also due to the high plan quality of modern intensity modulated (IMRT) and volumetric modulated arc therapy (VMAT) dose distributions.

However, there remains a subset of patients with tumours with a superficial component that would benefit significantly from the limited range of electron radiation to spare organs at risk (OAR) downstream from the tumour. In recent years, there has been a renewed interest in pMLC-based MERT delivery in the context of mixed electron-photon beam treatments. Previous planning studies have shown that, while electron-only MERT treatment plans typically deliver lower doses to normal tissue compared to photon plans, they are unable to provide the same level of dose homogeneity within the target [14, 12]. On the other hand, pMLC-based MBRT plans have recently been shown to provide superior OAR sparing compared to IMRT or VMAT plans without sacrificing target coverage [11, 15–17].

MBRT plans delivered using a pMLC as the sole collimation device would be the simplest to integrate into the clinical workflow as they do not, in principle, require staff intervention when switching modalities. Miguez et al. have demonstrated that pMLC-based MBRT for accelerated partial breast irradiation (APBI) can be implemented safely in the clinic. However, the treatment deliveries typically involved only three gantry angles and fewer than 10 fields [18]. Recently, Mueller et al. have delivered one brain and two head & neck MBRT plans with a non-coplanar photon component onto an anthropomorphic Alderson head phantom with films and shown 2%/2 mm gamma pass rates above 99.2% for all cases when compared to the expected simulated dose distribution [17], supporting the notion that MBRT can be delivered accurately. However, while the photon delivery was more complex than in the work of Miguez et al., the electron component remained simple, with one or two apertures per energy per beam angle. In addition, the delivered treatment plans did not contain low energy (6 or 9 MeV) components, which are likely to be the most challenging to model accurately.

In this work, we present comparisons between simulated and measured MBRT dose distributions for a simple, inherently robust delivery and a complex delivery. The aim was to 1) validate that the MC beam model used to produce MBRT treatment plans could accurately determine the number of MUs necessary to produce a desired dose distribution from each modality, 2) define the steps of a possible a patient-specific QA protocol based on point measurements with a single ionisation chamber position and 3) present a methodology for calculating beam quality correction factors for ionisation chamber measurements in MERT and MBRT fields.

8.4 Methods

8.4.1 Reference dose measurements for MLC-defined electron fields

Reference dose calibration for the applicator-less electron beams was performed on a Varian TrueBeam linear accelerator using an IBA Blue Phantom 2 water tank and an Exradin A1SL ionisation chamber. The reference conditions were defined to be 80 cm source-to-surface distance (SSD), with the the MLC leaves positioned to define a 10 x 10 cm^2 field when projected at the machine isocenter. This choice of non-standard reference conditions was made to more closely align the reference conditions with the delivery conditions for the electron component of MBRT plans. The jaws were set to 35 x 35 cm^2, the largest allowable size for our

MBRT planning algorithm. Reference dose measurements were performed at d_{ref} following the AAPM Task Group 51 (TG-51) protocol [19].

The charge measured in the ionisation chamber was converted to absorbed dose to water using equation 8.1,

$$D = M_c \, k_{Q,ecal} \, k'_Q \, N^{Co}_{D,w} \tag{8.1}$$

where M_c is the ionisation chamber reading, corrected for environmental conditions, ion recombination and polarity, $k_{Q,ecal}$ and k'_Q are the beam quality conversion factors described in Muir et al, (2014) [20]. The A1SL $N^{Co}_{D,w}$ coefficient used in this work was traceable to national primary absorbed dose standards.

The conversion factors are similar to $k_{R_{50}}$ and k_{ecal} described in the TG-51 report [19] but explicitly take into account P_{gr}, the gradient correction for the ionisation chamber used in this work. The values for the beam quality conversion factors were obtained from the Monte Carlo work by Muir et al, (2014) [20]. While these conversion factors were calculated in standard reference conditions at 100 cm SSD rather than the reference conditions used in this work, they are specified in terms of R_{50} which we assume remains a faithful representation of the electron fluence spectrum at the reference point for an 80 cm SSD setup.

8.4.2 Absorbed dose measurements in MBRT fields

Ionisation chambers are calibrated in terms of dose to water at the reference depth for a specific beam quality. The k'_Q beam quality correction factors used in eq. 8.1 are therefore only valid for the reference conditions described in section 8.4.1. When attempting to perform measurements in MBRT fields, we must correct the ionisation chamber response for the exact electron fluence spectrum at the measurement point in the MBRT field, which can vary greatly from the electron fluence spectrum at reference conditions. Differences in electron fluence spectrum can be caused by differences in measurement depth as well as by intensity modulation and delivery from multiple angles. The electron apertures for a given energy may not deliver radiation directly aimed at the measurement point, as shown in figure 8.1, leading to a potentially different electron fluence spectrum compared to reference conditions.

The beam quality conversion factor is defined as the ratio of absorbed dose to water, D_w to the absorbed dose in the air cavity of an ionisation chamber, D_{ch} between a beam quality Q and cobalt-60, which can be calculated with Monte Carlo methods assuming a fully characterised model of the ionisation chamber is included in the calculation [20–22]. In the methodology

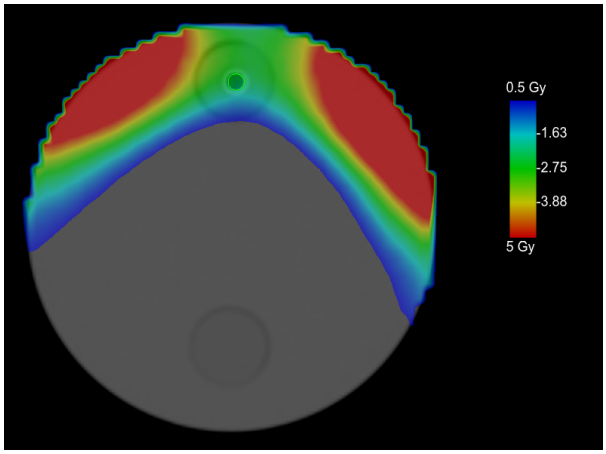

Figure 8.1 Dose colourwash for the 6 MeV component of a MBRT plan delivered on a PMMA phantom. The beam quality of particles inside the chamber air cavity (shown as a green contour) will differ significantly from the beam quality at reference conditions.

underlying AAPM's TG-51 protocol, however, this ratio was approximated as a ratio of Spencer-Attix stopping power ratios corrected for fluence perturbations,

$$k_Q = \left(\frac{D_w}{D_{ch}}\right)^Q_{Co} = P^Q_{gr} \frac{\left[\left(\frac{\bar{L}}{\rho}\right)^{water}_{air} P_{cel} P_{fl} P_{wall}\right]^Q}{\left[\left(\frac{\bar{L}}{\rho}\right)^{water}_{air} P_{cel} P_{fl} P_{wall}\right]^{Co}} \tag{8.2}$$

In this work we followed the latter methodology to apply a conversion factor between the beam quality in reference condition and the MBRT fields,

$$k_{MBRT} = \left(\frac{D_w}{D_{ch}}\right)^{MBRT}_Q \tag{8.3}$$

While, as indicated above, an evaluation of 8.3 requires a MC simulation of the local electron fluence with inclusion of the full chamber geometry, we assume that the first order contribution to k_{MBRT} is due to stopping power-ratio differences between water and air for the MBRT beam quality and the reference beam quality and sufficiently accurately corrects for the difference in chamber response between these two situations. This approximation ignores electron fluence *perturbation* by the presence of the ionisation chamber, but does capture the differences in energy response of the detector between reference conditions and MBRT conditions. A similar

approach has been used in the work by Al-Yahya et al to calibrate nonstandard electron fields created by the FLEC. [6]

$$k_{MBRT} \approx \frac{\left[\left(\frac{\bar{L}}{\rho}\right)^{w}_{air}\right]^{MBRT}}{\left[\left(\frac{\bar{L}}{\rho}\right)^{w}_{air}\right]^{Q}} \tag{8.4}$$

where $\left(\frac{\bar{L}}{\rho}\right)^{w}_{air}$ is the Spencer-Attix stopping power ratio (SPR) between water and air [23]. With this correction factor, the dose measured by an ionisation chamber in MBRT fields is given by:

$$D = M_c \; k_{Q,ecal} \; k'_Q \; k_{MBRT} \; N^{Co}_{D,w} \tag{8.5}$$

To obtain the SPRs for each electron component, we have modified the SPRRZnrc EGSnrc user-code [24] to score stopping power ratios in the same voxelised geometry as DOSXYZnrc, the code used to obtain the final patient dose distributions. This modified code, SPRXYZnrc, allows a SPR distribution to be obtained in the same geometry with the same input file as the one used for DOSXYZnrc. For each electron beam energy, the SPR values between water and air for voxels inside the chamber volume of the phantom were averaged based on the weight of each field and used to determine k_{MBRT}.

8.4.3 Phantom simulation and planning

A cylindrical polymethyl methacrylate (PMMA) phantom, shown in figure 8.2, was machined for the purposes of this work. The aim was to produce a phantom with a geometry similar to a patient extremity to perform QA measurements on MBRT plans created for soft tissue sarcomas of the leg. For ionisation chamber measurements, two holes were drilled with centres at depths of 1.4 cm and 2.1 cm from the surface of the cylinder, respectively, using physical measurements to determine the distances. The depth of the holes was chosen such that the active volume of an ionisation chamber would align with the centre of the phantom, lengthwise.

CT simulation scans of Solid Water (Sun Nuclear Corporation, Melbourne, Florida) slabs and the PMMA cylinder were obtained. In the case of the PMMA phantom, an Exradin A1SL (Standard Imaging, Madison, Wisconsin) ionisation chamber was present in the 1.4 cm hole in order to contour the active volume of the chamber for SPR calculation purposes. CT markers were placed using the in-room lasers for reproducible positioning. The MBRT treatment

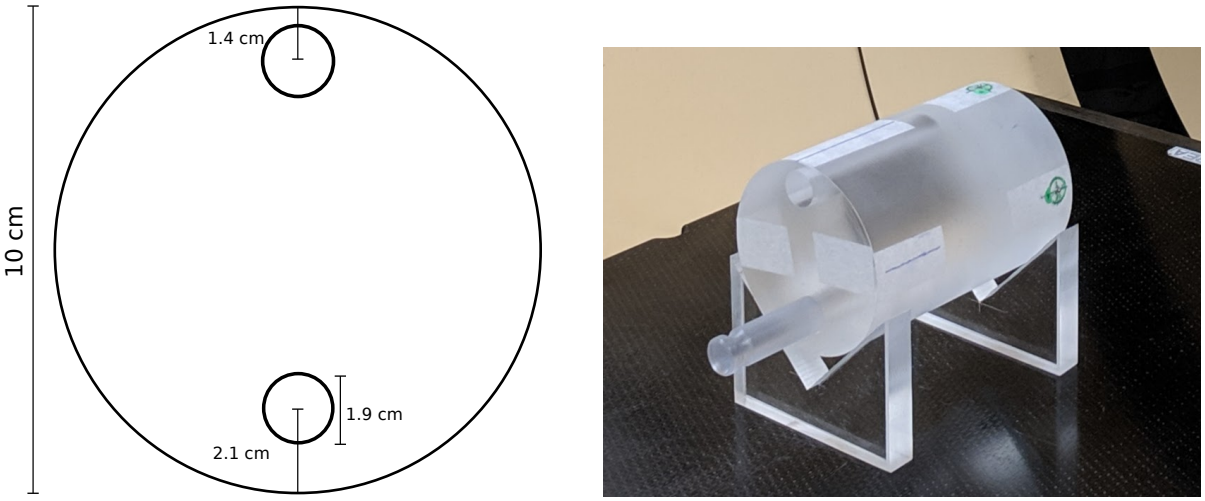

Figure 8.2 (a) Schematic diagram of the PMMA cylinder machined for ionisation chamber measurements. (b) Setup of the PMMA cylinder during the CT scanning process.

planning algorithm described in Renaud et al. [15] was used to create treatment plans for the Solid Water slab and the PMMA cylinder.

For the Solid Water slab phantom, a simple MBRT plan was created with a single electron beam angle and an arrangement of one photon beam angle every 45°. A total of 50 apertures spread across 6, 9, 12, 16 MeV electrons and 6 MV photons were included in the treatment plan. The aim was to ensure that many modalities participated in the plan while maintaining a simple plan delivery rather than producing a clinically realistic plan. In addition, to further simplify the delivery, the photon beam angles were collapsed at a gantry angle of 0°. All apertures were delivered as a step-and-shoot delivery.

The plan created for the PMMA cylinder aimed to reproduce a clinically realistic plan for a superficial target on the top half of the cylinder. The electron beam angles were (-30, 0, 30)° and the photon component was an arc from -110° to 110°. For this plan, the physical delivery was not simplified in any way compared to the planned delivery. The electron component included a total of 40 electron apertures and was delivered as a step-and-shoot delivery. The beam delivery parameters for each plan are summarised in table 8.1.

Figure 8.3 shows dose colourwashes for a representative axial slice for both plans. For both phantoms, the target is contoured in red. In both cases, the dose distribution is normalised so that 95% of the PTV volume receives 50 Gy.

	Solid water plan	PMMA cylinder plan
Photon angles (°)	$(22.5, 67.5, ..., 337.5)$	arc from -110 to 110
Electron angles (°)	0	$(-30, 0, 30)$
Photon SAD (cm)	100	100
Electron SAD (cm)*	80	80

Table 8.1 Summary of the planned beam delivery angles and phantom positions. For the Solid Water plan delivery, the photon beam angles were collapsed to 0°. * The shortened electron SAD is realised in practice as a virtual SAD, meaning that the couch is moved such that the centre of the target is 80 cm from the source for each beam angle.

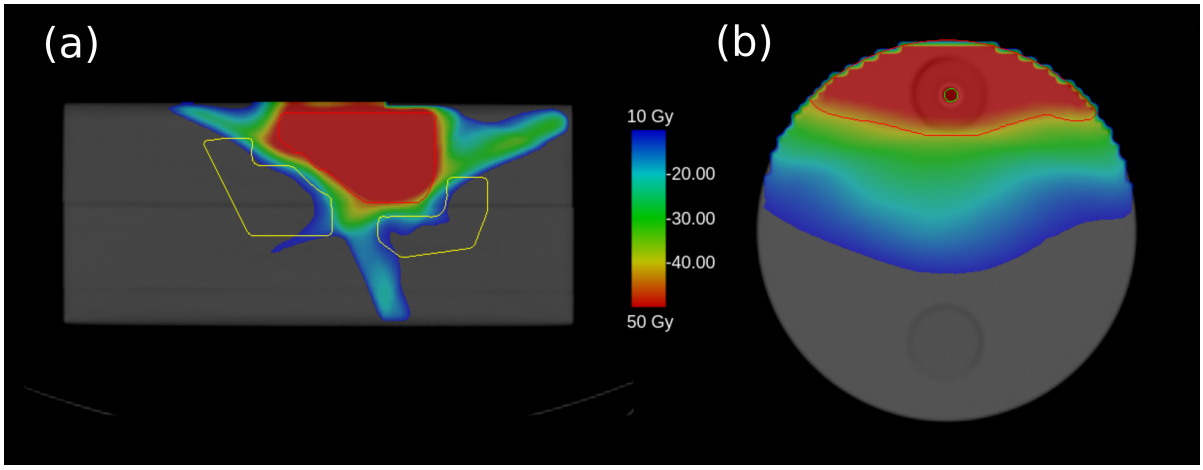

Figure 8.3 Dose colourwash for a representative axial slice of the (a) simplified MBRT plan delivered on Solid Water slabs and (b) complex MBRT plan delivered on the PMMA cylinder. In both cases, the target is shown as a red contour.

8.4.4 Calculated dose distributions

In previous work, we have demonstrated PTV-based and CTV-based robust treatment planning algorithms to produce MBRT plans from beamlet-based apertures [15, 25]. However, beamlet-based apertures require a final MC recalculation to account for the effects of MLC leaves and jaw position on the dose distribution and the relative output of each aperture. In this work, each aperture of the treatment plans we created for QA purposes was recalculated using a validated MC beam model, and the MC dose distributions were renormalised from dose per primary particle to dose per monitor unit.

MC simulations were performed using the phase space files distributed by Varian for the TrueBeam linear accelerator as the particle source (Virtual Linac) [26]. The particles sampled

from the phase space files were transported through a BEAMnrc model consisting of the jaws, the base plate, the Millenium 120 MLC and the exit window [27]. Particles were further transported in a voxelised geometry by using DOSXYZnrc [28]. The electron transport cutoff (ECUT) was 0.7 MeV, and the photon cutoff (PCUT) was 0.01 MeV. The EXACT boundary crossing algorithm was used, with a skin depth of 3 mean free paths. The electron stepping algorithm was PRESTA-II. The voxel sizes used in the MC simulations were 2.5 mm^3 and 2 mm^3 for the Solid Water slabs and PMMA cylinder, respectively. For SPR calculations, an ECUT of 0.521 MeV was used instead of 0.7 MeV.

To convert MC dose values into absorbed dose per Monitor Unit (Gy / MU), the MC reference dose calibration factors were obtained by reproducing the conditions described in section 8.4.1 in a Monte Carlo simulation, and noting the MC dose per primary particle value at d_{ref} for each energy. The MC aperture dose distributions were then renormalised as

$$D_{MU}^{MC} = D^{MC} \frac{D_{ref}^{meas}}{D_{ref}^{MC}} \tag{8.6}$$

where D_{ref}^{meas} was the measured dose per MU in reference conditions, and D_{ref}^{MC} was the MC dose per primary particle calculated in the same reference conditions [29]. A monitor chamber backscatter correction was not applied as it has been shown that the correction is negligible when the jaw opening is kept large and MLC leaves are used to collimate the field [30]. The same procedure was performed to renormalise photon MC aperture dose distributions, however the reference conditions were taken as standard TG-51 conditions rather than the 80 cm SSD setup done for electrons.

The same treatment planning optimisation criteria were then used to re-optimise the relative weight of each aperture using the MC-calculated aperture dose distributions and obtain the monitor units for each aperture. The number of monitor units delivered from each modality for both plans is given in table 8.2.

	Monitor units per modality					
Plan	6 MeV	9 MeV	12 MeV	16 MeV	20 MeV	6 MV
Solid Water	11932	6528	7052	3385	0	6800
PMMA	8194	1940	260	1514	1083	4181

Table 8.2 Distribution of monitor units per modality for the two plans delivered in this work, normalised to deliver 50 Gy to 95% of the contoured target volume.

8.4.5 Phantom setup and delivery

The phantoms were positioned such that the distance between the geometric centre of the PTV and the source was 80 cm for the electron component (i.e. a virtual 80 cm SAD) and 100 cm for the photon component. The plans were delivered on a Varian TrueBeam linear accelerator, which has a nominal SAD of 100 cm, hence shortened SAD deliveries require a different couch position for each beam delivery angle.

The treatment plans were delivered using the TrueBeam developer mode which allows the couch to be moved dynamically during treatment. The setup and delivery process was as follows:

1. The phantom was positioned on the treatment couch and the CT markers were aligned with the in-room lasers.

2. The couch positions (lat, lng, vrt) displayed on the treatment console were noted.

3. Using our in-house treatment planning system (TPS), the position of the machine isocenter in the CT coordinate system was identified, as shown in Fig. 8.4 (a), to establish a transformation between the couch coordinate system and the CT coordinate system.

4. The treatment plan was exported as an XML file using our TPS by supplying the information shown in Fig. 8.4 (b). The TPS automatically creates the XML files necessary for delivery using the TrueBeam developer mode and determines the couch position of each control point based on the transformation between the couch and CT coordinate systems shown in eq. 8.7.

5. The treatment plan was delivered on a per-modality basis, as the TrueBeam developer mode does not yet support changing between photon or electron energies within a single XML file.

The couch positions for each control point were determined using a simple translation of the couch based on the difference between the machine isocenter position at each control point and the machine isocenter in the setup position,

$$\boldsymbol{p}_{cpt}^{couch} = \boldsymbol{p}_{setup}^{couch} + (\boldsymbol{p}_{cpt}^{CT} - \boldsymbol{p}_{setup}^{CT}) \tag{8.7}$$

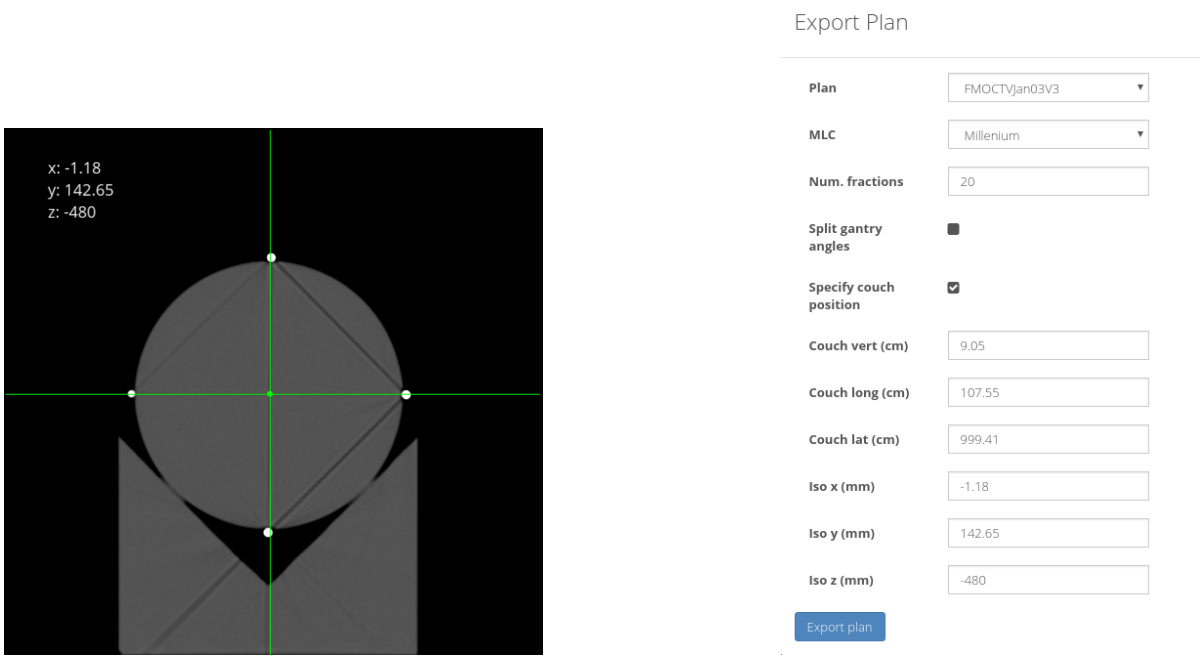

(a) (b)

Figure 8.4 (a) Identification of the machine isocenter position in the CT coordinate system when the phantom BBs are aligned with the in-room lasers. (b) Example of the information supplied by the user when exporting a plan as developer mode XML files from our in-house TPS.

8.4.6 Measurement setup

For the Solid Water slab phantom, a Gafchromic EBT3 film was placed at a depth of 3 cm and irradiated with all components of the treatment plan. The film was scanned 18 hours after irradiation using an Epson Expression 1100XL flatbed scanner (Epson America, Inc., Long Beach, CA). The film was then read into FilmQA Pro 2015 (Ashland Advanced Materials, Bridgewater, New Jersey) and the red colour channel was converted to dose using calibration films obtained on the same day.

The ionisation chamber measurements were performed using an Exradin A1SL chamber inside a Solid Water plug inserted into the 1.4 cm hole of the PMMA cylinder. The MBRT plan MC doses for all modalities were recalculated to account for the differences in materials between the planning and measurement conditions. The active volume of the ionisation chamber was converted from air to water in the MC simulation as the chamber is calibrated in terms of absorbed dose to water and thus nominally reports absorbed dose to water. Figure 8.5 shows

the materials and densities used for the MC dose calculation. The same phantom was used to calculate SPRs inside the chamber volume.

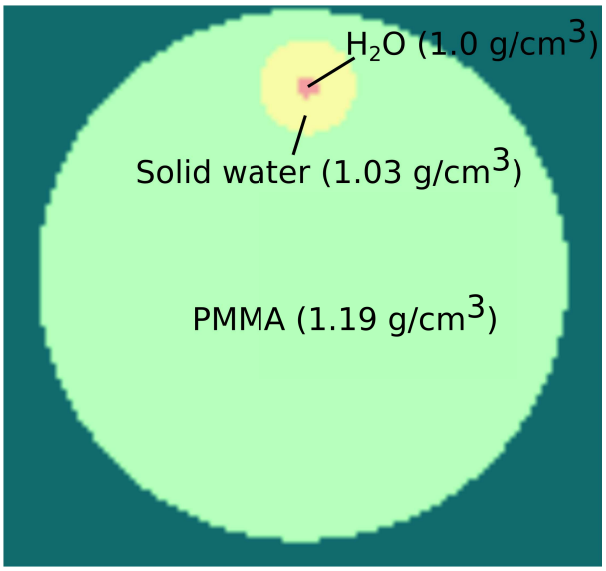

Figure 8.5 MC phantom materials and densities used to calculate the dose inside the ionisation chamber.

8.5 Results

8.5.1 Reference dose measurements

Table 8.3 shows the calibration depths and values measured in the applicator-less electron radiation therapy reference conditions described in section 8.4.1. The R_{50} beam quality specifiers were obtained from percent depth dose curves measured using an IBA RFD 3G diode detector (Ion Beam Applications, Louvain-la-neuve, Belgium) in the same reference conditions. The Type-A uncertainty on the MC-calculated SPRs was less than 0.2 % for all modalities. All uncertainties stated are k = 1.

8.5.2 Film measurement on Solid Water slab

The total plan delivery was divided into 25 fractions and a single fraction was delivered on the film. The number of fractions was chosen such that the maximum dose on the film was

Energy (MeV)	R_{50} (cm)	d_{ref} (cm)	k'_Q	$(SPR^w_{air})_Q$	D_{ref} (cGy / MU)
6	2.46	1.38	1.0318	1.079	1.498 ± 0.016
9	3.70	2.12	1.0197	1.062	1.508 ± 0.016
12	5.10	2.96	1.0104	1.045	1.518 ± 0.016
16	6.73	3.94	1.0024	1.034	1.487 ± 0.016
20	8.31	4.89	0.9963	1.023	1.426 ± 0.015

Table 8.3 Beam quality specifiers and conversion factors for each electron beam energy for the Exradin A1SL ion chamber [20], measured in a water tank using a 10x10 cm^2 MLC-defined field at 80 cm SSD. $k_{Q,ecal}$ for the A1SL was taken to be 0.914. The SPRs were calculated using SPRXYZnrc with a Type A uncertainty of less than 0.2%. An uncertainty of 1.1 % on the reference absorbed dose per monitor unit was assessed using the same uncertainty budget approach as detailed in McEwen et al. [22].

approximately 70% of the maximum calibration dose. Fig. 8.6 presents the results of the comparison between the film measurement and the planned dose distribution. A 2%/2mm gamma analysis was performed with a global 10% dose threshold. The pass rate was found to be 96.4%. The MC uncertainty on voxel doses was 2%. It should be noted that MC uncertainty can artificially inflate the gamma pass rate.

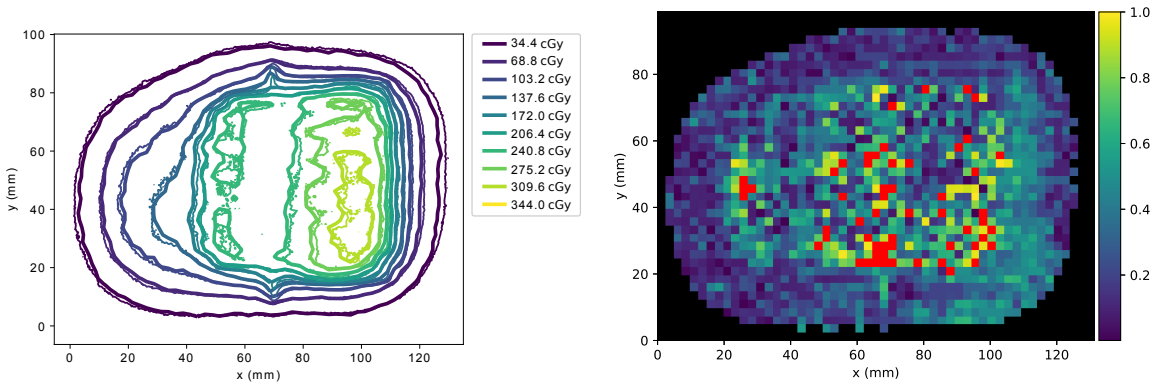

(a) (b)

Figure 8.6 (a) Isodose comparison between measured film dose (thin lines) and simulated dose (thick lines) for the Solid Water plan. (b) Gamma map for a 2% / 2 mm passing criteria. Pixels in red have $\gamma >= 1$.

8.5.3 Ionisation chamber measurements on PMMA cylinder

The total dose from the PMMA treatment plan was divided into 10 fractions ($N_{frac} = 10$) and a single fraction was delivered in the measurement setup described in section 8.4.6. The number of fractions was chosen to ensure that the measured signal from each modality was large compared to background leakage. In practice, the type of clinical plan calculated on the PMMA cylinder would more likely be split into 20 or 25 fractions.

Table 8.4 shows the measured dose values in the ionisation chamber compared to Monte Carlo-calculated doses for each modality. For each modality, $\left[\left(\frac{\bar{L}}{\rho}\right)^{w}_{air}\right]^{MBRT}$ was Monte Carlo-calculated using the SPRXYZnrc code with the same geometry and particle source as for the dose calculation performed in DOSXYZnrc. The MBRT beam quality correction factor, k_{MBRT}, was then determined using eq. 8.3 using the reference SPRs in table 8.3. The absorbed dose to water was then calculated using eq. 8.5.

	20 MeV	16 MeV	12 MeV	9 MeV	6 MeV	6 MV
$(SPR^{w}_{air})_{MBRT}$	0.982	0.996	1.017	1.049	1.087	1.119
k_{MBRT}	0.960	0.966	0.973	0.988	1.007	0.998
D_{meas} (Gy)	8.91 ± 0.19	14.87 ± 0.44	3.16 ± 0.07	4.59 ± 0.30	11.63 ± 0.56	11.86 ± 0.13
D_{MC} (Gy)	8.86 ± 0.05	14.99 ± 0.09	3.19 ± 0.01	4.47 ± 0.03	11.40 ± 0.07	12.02 ± 0.02
ΔD $(\%)$	0.52 ± 2.2	-0.8 ± 3.0	-0.82 ± 2.3	2.7 ± 6.4	2.0 ± 4.8	-1.4 ± 1.1
	Electrons				Total	
D_{meas} (Gy)	43.16 ± 1.55				55.02 ± 1.68	
D_{MC} (Gy)	42.91 ± 0.25				54.93 ± 0.27	
ΔD $(\%)$	-0.6 ± 3.6				-0.2 ± 3.1	

Table 8.4 Measurement data from an Exradin A1SL ionisation chamber placed in the 1.4 cm hole of the PMMA phantom compared to simulated MC doses. The measurement doses were corrected for the beam quality at the measurement point using k_{MBRT}. The uncertainty on k_{MBRT} is estimated at 1%, dominated by the type B uncertainty associated with neglecting the cavity fluence perturbation.

The agreement between measurement and calculation for each modality was within 2.7 %. For the sum of all electrons, the measured dose and calculated dose agreed to within 0.6 % and the total dose at the measurement point agreed to within 0.2 %.

In determining the uncertainty associated with the measured dose, D_{meas}, we combined the standard uncertainty on the reference dose determination (1.1%) with the uncertainties associated with the determination of k_{MBRT}, most notably by estimating the uncertainty associated with omitting the cavity fluence perturbation correction (p_{cav}) in the determination of

k_{MBRT}, and a dose non-uniformity uncertainty based on the heterogeneity of the dose inside the chamber volume calculated using the MC doses.

The uncertainty associated with neglecting the fluence perturbation correction in the MBRT field was estimated by first assigning a hypothetical beam quality to each energy of the MBRT delivery ($R_{50,MBRT}$). This specifier was determined by inverting the $\left(\frac{\bar{L}}{\rho}\right)^{w}_{air}$ to R_{50} relationship given in the IAEA TRS-398 report, Appendix B, based on the $\left[\left(\frac{\bar{L}}{\rho}\right)^{w}_{air}\right]^{MBRT}$ values calculated using SPRXYZnrc. This hypothetical $R_{50,MBRT}$ was then used to obtain a value for $p_{cav,MBRT}$ using the equation for p_{cav} for cylindrical chambers as a function of R_{50} provided by TRS-398 Appendix B [31], which is based on a broad set of experimental data. $p_{cav,ref}$ was determined from the same formula but using the reference beam R_{50} instead. The relative difference between $p_{cav,MBRT}$ and $p_{cav,ref}$ (i.e., $1.0 - p_{cav,MBRT}/p_{cav,ref}$) does not exceed 1% and was treated as an uncertainty factor.

The uncertainty due to the non-uniformity of the dose inside the chamber volume was determined based on the minimum and maximum dose values inside the volume, and assuming a triangular distribution (i.e., $(D_{max} - D_{min})/D_{avg}/\sqrt{(6)}$).

Figure 8.7 b) shows the SPR between the phantom material and air as a function of depth along the line shown in Fig. 8.7 a). The effect of the different material compositions on the SPR is clearly visible.

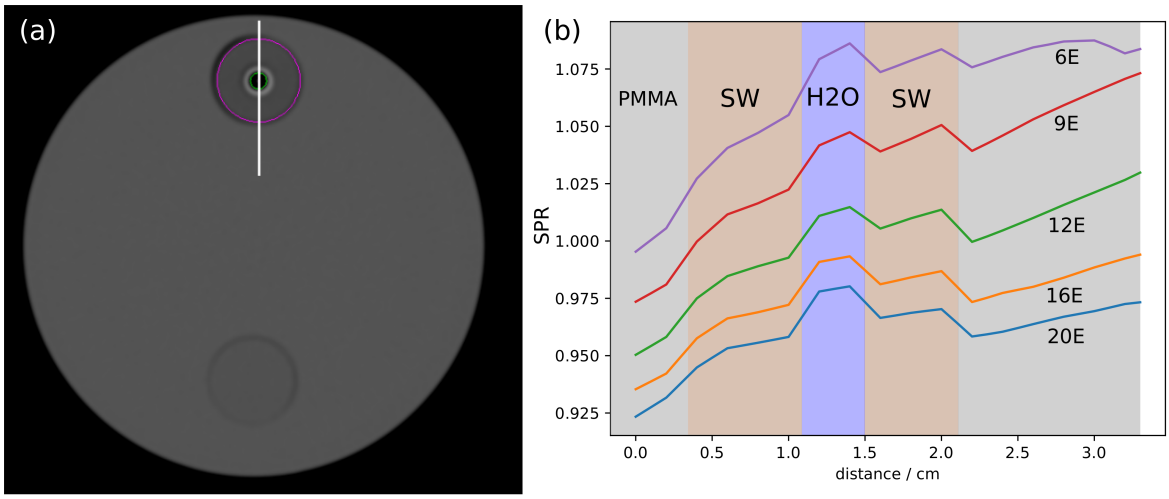

Figure 8.7 (a) Axial CT slice of the PMMA phantom showing the line along which the SPRs are plotted for each modality. (b) SPR between the phantom medium and air. The active volume of the ionisation chamber was modelled as water.

8.6 Discussion

The purpose of the film delivery on slabs of Solid Water phantom was to identify large errors in the planning, simulation and delivery process. The setup was inherently robust to positioning errors as the entire plan was delivered at a normal incidence to the flat phantom, therefore good agreement was expected between the film measurement and the simulated dose. The gamma pass rate was found to be 96.4 % for a 2 %/2 mm criteria, with many of the failing voxels distributed with no discernible pattern.

The cylindrical phantom delivery was set up to closely resemble the types of dose distributions obtainable for MBRT applications to soft tissue sarcomas of the leg. Electron dose distributions are known to be perturbed significantly more than photon dose distributions when delivered at oblique incidences; therefore, this delivery can be seen as a particularly challenging case for the electron MC beam models.

The measured doses agreed to better than 2.7 % when individual modalities were compared. When all electron modalities were combined, the measurement was within 0.6 % of the planned dose, and the difference between the total measured and simulated MBRT doses was 0.2 %, well within the uncertainty of the measurement. As expected, the difference between measured and simulated absorbed doses was largest for the lowest energy electrons, 6 MeV and 9 MeV. In addition to the accuracy of the MC models, the accuracy of the setup was expected to have a significant impact on the delivered dose distribution. Although a mini ion chamber (A1SL, 0.053 cc) was used to perform all measurements, there remained large dose gradients inside the chamber volume, contributing uncertainty to the measurement as the effective point of measurement is not well defined.

By calculating k_{MBRT} as a simple ratio of stopping power-ratios between measurement and reference conditions, we have implicitly assumed that P_{gr} was constant, which may explain part of the differences between measured and calculated doses. In addition, effects of electron fluence perturbation are ignored in our approach, which would affect the results predominantly at low electron energies (6 MeV and 9 MeV). The overall agreement between planned and delivered dose, however, confirms that for this situation the effects are limited.

The k_{MBRT} correction factor resulted in a 4 % correction in measured dose for the highest electron energies. Despite k_{MBRT} being necessary for accurate measurements in MBRT conditions, we conclude that the measurement procedure followed in this work consists of a viable procedure for MBRT plan verification using an ionisation chamber. For patient-specific QA, the

same procedure as described in sections 8.4.5 and 8.4.6 would be followed, with the addition of a preliminary step consisting of producing a *hybrid plan* on the measurement phantom by transposing a patient plan and recalculating the dose distribution on the phantom.

Both MBRT plans were deliverable with a single setup procedure as the TrueBeam developer mode allows dynamic couch positioning. Changing between modalities was the longest overhead associated with MBRT compared to conventional photon radiation therapy. In both cases, the electron component of the plan was deliverable in less than 5 minutes, including modality changes, but not including setup time. In terms of beam-on time, all electron components can be delivered at a rate of 1000 MU / min, which speeds up delivery over the 600 MU / min maximum of photon beams with a flattening filter on the TrueBeam accelerator. The complete QA procedure was performed in under 15 minutes.

As can be seen from table 8.2, the number of monitor units per component greatly increases for the lower electron energies due to the fact that output factors degrade quickly for the combination of low energies and small fields. For example, we measured the 3×3 cm^2 MLC-defined field output factor to be 0.251 for 6 MeV, compared to 0.855 for 20 MeV. While the deliveries of these components is still accomplished rapidly due to the high dose rates achievable for electrons, Connell et al. found that electron output factors were highly sensitive to minute (≈ 2 mm) changes in collimation when the field sizes were on the order of 3×3 cm^2 [8], which could lead to large discrepancies between simulated and delivered doses if the linac jaws were miscalibrated. We therefore recommend that the MC-simulated output factors be compared to measured output factors regularly if MBRT plans are delivered with a significant low energy electron component. By default, some manufacturers allow tolerances on MLC leaf positioning during treatment which could allow the beam to be enabled if the leaf is within, e.g., 2 mm of its intended position. For low electron beam energies and smaller MLC apertures, such tolerances could also lead to large differences in output factors. If possible, lower tolerances on MLC leaf positions should be used during MBRT delivery.

As previously mentioned, the virtual SAD delivery of the electron components requires a different couch position for each beam angle, which is not currently part of routine clinical practice. In an effort to produce plans which required fewer changes to current practice, we attempted to deliver the MBRT plans with a single couch position, hence a varying shortened SSD. However, with such a setup, the target volume is typically not on the beam central axis, leading to off-axis irradiation. The transition between beam angles therefore required large movements of the collimators, during which the beam had to be in the *beam hold* state. Beam

holds on Varian accelerators are created by adjusting the grid voltage in the electron gun; during the beam hold state the accelerator is active with RF in the wave guide but the electron source is in a hold state. However, during the course of delivering these plans, we have discovered the presence of a persistent, low level of leakage radiation during the beam hold states which is high enough to trigger a machine interlock and interrupt the delivery.

A discussion with a representative from Varian Medical Systems confirmed that the leakage is due to the absence of attenuating material in the path of the beam compared to photon beams, which means that low electron currents in the waveguide can reach to appreciable doses to the patient even with the beam in the beam hold state. Virtual SAD deliveries result in less collimator movement which reduces the need for beam holds. However, long beam holds could still occur in virtual SAD deliveries if large collimator movements were needed, for example, when multiple physically separated small lesions are treated.

8.7 Conclusion

In this work, the aim was to confirm that the treatment plans produced by our MBRT treatment planning algorithm [15, 25] could be delivered and measured accurately on a Varian TrueBeam accelerator. In addition, we sought to define the steps for a patient-specific QA protocol based on ionisation chamber measurements.

These goals were first achieved by delivering a simple plan on slabs of Solid Water with a film placed at a depth of 3 cm, and comparing the results to the simulated delivery. The resulting gamma pass rate of 96.4 % for a 2 %/2 mm criteria confirmed that the MC beam models used in this work performed accurately.

A patient-specific QA procedure was evaluated by delivering a complex plan on a PMMA cylinder specifically constructed for ionisation chamber measurements and performing dose measurements at a single point. The measurement from each modality was corrected by a beam quality correction factor calculated using a MC code specifically created to obtain both dose distributions and correction factors with the same input in order to facilitate the QA process. The complete QA measurement procedure was realised in under 15 minutes, and the agreement between measured and simulated dose agreed to within 0.2 %, leading us to conclude that the procedure can be applied for general patient-specific QA. The QA procedure on the cylindrical phantom needs to be further extended to include relative film measurements analysed using the dose conversion procedure described in this work.

Future work will assess the sensitivity of the QA measurement process to positioning errors. MBRT plans created using robust optimisation will no longer be robust to positioning errors when transposed onto measurement phantoms. Therefore, a study is necessary to determine the variation in agreement between simulated and measured doses from clinical robust plans delivered on QA phantoms and define realistic tolerances for QA setups of robust plans for the clinical site of interest.

The fact that beam holds are unreliable for electron beams presents an additional constraint to the delivery of MBRT but this is an engineering problem that could likely be addressed relatively easily by the manufacturer.

8.8 Acknowledgements

The authors would like to acknowledge Dr. Tanner Connell for helpful discussions on film measurements, Dr. James Renaud for assistance with reference dosimetry, and Joe Larkin for machining the phantom used in this work. The authors acknowledge support from the Fonds de recherche du Québec Nature et Technologie (FRQNT), the CREATE Medical Physics Research Training Network grant (number 432290) of the Natural Sciences and Engineering Research Council (NSERC) and the Canadian Institute of of Health Research Foundation Grant (FDN-143257).

References

[1] M. C. Lee, S. B. Jiang, and C. M. Ma, "Monte Carlo and experimental investigations of multileaf collimated electron beams for modulated electron radiation therapy," *Medical Physics*, vol. 27, no. 12, pp. 2708–2718, 2000.

[2] C. M. Ma, T. Pawlicki, M. C. Lee, S. B. Jiang, J. S. Li, J. Deng, B. Yi, E. Mok, and A. L. Boyer, "Energy- and intensity-modulated electron beams for radiotherapy," *Physics in Medicine and Biology*, vol. 45, no. 8, pp. 2293–2311, 2000.

[3] K. R. Hogstrom, R. A. Boyd, J. A. Antolak, M. M. Svatos, B. A. Faddegon, and J. G. Rosenman, "Dosimetry of a prototype retractable eMLC for fixed-beam electron therapy," *Medical Physics*, vol. 31, no. 3, pp. 443–462, 2004.

[4] T. Gauer, D. Albers, F. Cremers, R. Harmansa, R. Pellegrini, and R. Schmidt, "Design of a computer-controlled multileaf collimator for advanced electron radiotherapy," *Physics in Medicine and Biology*, vol. 51, no. 23, pp. 5987–6003, 2006.

[5] T. Gauer, J. Sokoll, F. Cremers, R. Harmansa, M. Luzzara, and R. Schmidt, "Characterization of an add-on multileaf collimator for electron beam therapy," *Physics in Medicine and Biology*, vol. 53, no. 4, pp. 1071–1085, 2008.

[6] K. Al-Yahya, F. Verhaegen, and J. Seuntjens, "Design and dosimetry of a few leaf electron collimator for energy modulated electron therapy," *Medical Physics*, vol. 34, no. 12, pp. 4782–4791, 2007.

[7] A. Alexander, F. Deblois, and J. Seuntjens, "Toward automatic field selection and planning using Monte Carlo-based direct aperture optimization in modulated electron radiotherapy," *Physics in Medicine and Biology*, vol. 55, no. 16, pp. 4563–4576, 2010.

[8] T. Connell, A. Alexander, P. Papaconstadopoulos, M. Serban, S. Devic, and J. Seuntjens, "Delivery validation of an automated modulated electron radiotherapy plan," *Medical Physics*, vol. 41, no. 6, 2014.

[9] F. C. Du Plessis, A. Leal, S. Stathakis, W. Xiong, and C. M. Ma, "Characterization of megavoltage electron beams delivered through a photon multi-leaf collimator (pMLC)," *Physics in Medicine and Biology*, vol. 51, no. 8, pp. 2113–2129, 2006.

[10] E. E. Klein, M. Vicic, C. M. Ma, D. A. Low, and R. E. Drzymala, "Validation of calculations for electrons modulated with conventional photon multileaf collimators," *Physics in Medicine and Biology*, vol. 53, no. 5, pp. 1183–1208, 2008.

[11] F. J. Salguero, B. Palma, R. Arrans, J. Rosello, and A. Leal, "Modulated electron radiotherapy treatment planning using a photon multileaf collimator for post-mastectomized chest walls," *Radiotherapy and Oncology*, vol. 93, no. 3, pp. 625–632, 2009.

[12] D. Henzen, P. Manser, D. Frei, W. Volken, H. Neuenschwander, E. J. Born, a. Joosten, K. Lössl, D. M. Aebersold, C. Chatelain, M. F. M. Stampanoni, and M. K. Fix, "Beamlet based direct aperture optimization for MERT using a photon MLC.," *Medical physics*, vol. 41, p. 121711, dec 2014.

[13] E. E. Klein, Z. Li, and D. A. Low, "Feasibility study of multileaf collimated electrons with a scattering foil based accelerator," *Radiotherapy and Oncology*, vol. 41, no. 2, pp. 189–196, 1996.

[14] A. Alexander, E. Soisson, M.-A. Renaud, and J. Seuntjens, "Direct aperture optimization for FLEC-based MERT and its application in mixed beam radiotherapy.," *Medical physics*, vol. 39, pp. 4820–31, aug 2012.

[15] M.-A. Renaud, M. Serban, and J. Seuntjens, "On mixed electron-photon radiation therapy optimization using the column generation approach," *Medical Physics*, vol. 44, no. 8, pp. 4287–4298, 2017.

[16] S. Mueller, M. K. Fix, A. Joosten, D. Henzen, D. Frei, W. Volken, R. Kueng, D. M. Aebersold, M. F. Stampanoni, and P. Manser, "Simultaneous optimization of photons and electrons for mixed beam radiotherapy," *Physics in Medicine and Biology*, vol. 62, no. 14, pp. 5840–5860, 2017.

[17] S. Mueller, M. K. Fix, D. Henzen, D. Frei, D. Frauchiger, K. Loessl, M. F. M. Stampanoni, and P. Manser, "Electron beam collimation with a photon MLC for standard electron treatments," *Physics in Medicine & Biology*, vol. 63, no. 2, p. 025017, 2018.

[18] C. Míguez, E. Jiménez-Ortega, B. A. Palma, H. Miras, A. Ureba, R. Arráns, F. Carrasco-Peña, A. Illescas-Vacas, and A. Leal, "Clinical implementation of combined modulated electron and photon beams with conventional MLC for accelerated partial breast irradiation," *Radiotherapy and Oncology*, vol. 124, no. 1, pp. 124–129, 2017.

[19] P. R. Almond, P. J. Biggs, B. M. Coursey, W. F. Hanson, M. S. Huq, R. Nath, and D. W. Rogers, "AAPM's TG-51 protocol for clinical reference dosimetry of high-energy photon and electron beams," *Medical Physics*, vol. 26, no. 9, pp. 1847–1870, 1999.

[20] B. R. Muir and D. W. Rogers, "Monte Carlo calculations of electron beam quality conversion factors for several ion chamber types," *Medical Physics*, vol. 41, no. 11, p. 111701, 2014.

[21] B. R. Muir and D. W. Rogers, "Monte Carlo calculations of kQ, the beam quality conversion factor," *Medical Physics*, vol. 37, no. 11, pp. 5939–5950, 2010.

[22] M. McEwen, L. Dewerd, G. Ibbott, D. Followill, D. W. Rogers, S. Seltzer, and J. Seuntjens, "Addendum to the AAPM's TG-51 protocol for clinical reference dosimetry of high-energy photon beams," *Medical Physics*, vol. 41, no. 4, 2014.

[23] P. Andreo, D. T. Burns, A. E. Nahum, J. Seuntjens, and F. H. Attix, *Fundamentals of ionizing radiation dosimetry*. John Wiley & Sons, 2017.

[24] D. W. O. Rogers, I. Kawrakow, J. P. Seuntjens, B. R. B. Walters, and E. Mainegra-Hing, "NRC User Codes for EGSnrc," tech. rep., 2003.

[25] M.-A. Renaud, M. Serban, and J. Seuntjens, "Robust mixed electron-photon radiation therapy optimisation," *Medical Physics*, vol. 0, no. ja, 2019.

[26] A. Rodrigues, D. Sawkey, F. F. Yin, and Q. Wu, "A Monte Carlo simulation framework for electron beam dose calculations using Varian phase space files for TrueBeam Linacs," *Medical Physics*, vol. 42, no. 5, pp. 2389–2403, 2015.

[27] D. W. O. Rogers, B. Walters, and I. Kawrakow, "BEAMnrc Users Manual," tech. rep., National Research Council of Canada, 2011.

[28] B. Walters, I. Kawrakow, and D. W. O. Rogers, "DOSXYZnrc Users Manual," tech. rep., National Research Council of Canada, 2016.

[29] C. M. Ma, R. A. Price, J. S. Li, L. Chen, L. Wang, E. Fourkal, L. Qin, and J. Yang, "Monitor unit calculation for Monte Carlo treatment planning," *Physics in Medicine and Biology*, vol. 49, no. 9, pp. 1671–1687, 2004.

[30] S. A. Lloyd, I. M. Gagne, M. Bazalova-Carter, and S. Zavgorodni, "Measured and Monte Carlo simulated electron backscatter to the monitor chamber for the Varian TrueBeam Linac," *Physics in Medicine and Biology*, vol. 61, no. 24, pp. 8779–8793, 2016.

[31] INTERNATIONAL ATOMIC ENERGY AGENCY, *Absorbed Dose Determination in External Beam Radiotherapy*. No. 398 in Technical Reports Series, Vienna: INTERNATIONAL ATOMIC ENERGY AGENCY, 2000.

Chapter 9

Summary and outlook

9.1 Summary

Modern radiotherapy critically hinges on the process of treatment planning, where computer algorithms are used to optimise radiation dose to the target while keeping dose to normal tissues and critical organs as low as possible. However, while current external beam treatment planning algorithms focus on finding optimal gantry, couch and collimator settings to achieve dose prescriptions, they are only using a single modality, i.e., typically photons. For patients with superficial tumours, this has meant that a major potential improvement in normal tissue sparing has been overlooked. The aim of this thesis was to incorporate the particle type degree of freedom into the modern treatment planning workflow and to produce mixed external beam electron-photon plans deliverable using only a photon MLC.

As reviewed in chapter 4, mixed beam radiation therapy is not a new research topic. However, prior to our work, most attempts featured elements of forward planning or sequential optimisation, rather than fully inversely planned, simultaneous optimisation. A true simultaneous MBRT optimisation algorithm should ideally never produce plans that are inferior to single-modality plans, as they are simply a subset of the MBRT solution space. However, due to the added computational burden of optimising multiple modalities simultaneously, previous attempts at MBRT optimisation had typically only sampled a subset of the available single-modality degrees of freedom, such as a limited number of beam delivery angles. One of the major accomplishments of this work was to develop an algorithm that can optimise all the degrees of freedom of modern photon planning algorithms while also being able to accept a large increase in computational complexity due to the added electron components.

9.1.1 Development of a framework for MBRT treatment planning

The data requirements for MBRT planning precluded the use of commercial TPS software. Hence, prior to investigating optimisation algorithms, a web-based research dose calculation and treatment planning platform called Radify was created as a general tool to visualise patient data, coordinate the creation of beamlet dose distributions and control the submission of optimisation jobs. Treatment plan recalculation using an MC dose calculation engine was also included as a functionality. The design of Radify followed a standard back-end web server arrangement where a database and an API were set up to respond to requests and return data for a front-end GUI to interpret and display. The API was also set up to respond to requests originating from other sources than the front-end GUI, in order to accommodate user-created scripts for large scale data processing tasks.

A collapsed cone convolution superposition (CCCS) photon dose calculation engine was developed for the purpose of generating photon beamlets. A simplified single source 6 MV beam model was used to ensure fast calculation times. On average, the engine was capable of calculating 2400 beamlet dose distributions per minute using an Nvidia GTX 1080 Ti graphics card. For electrons, a previously validated electron MC beam model was used to create a phase space file at the level of the MLC plane. The phase space was then split into individual 1 x 1 cm^2 smaller phase space files which were used as particle sources to calculate beamlet dose distributions in patient geometries.

The optimisation algorithms were implemented as modules of a C++ library. The IPOPT general purpose nonlinear solver was used to solve the master problems for all the optimisation models. While the major focus of this work has been on developing the optimisation models for MBRT planning, modules were also created for standard photon-only IMRT and VMAT, rotating shield brachytherapy and experimental kV arc therapy optimisation as well.

This work served as the foundation to enable the investigation of the column generation method applied to MBRT planning. In addition, the framework has served numerous researchers in developing new radiation therapy delivery techniques and performing large scale data processing tasks.

9.1.2 The column generation method applied to MBRT planning

Following the implementation of the basic building blocks for treatment planning, the primary objective was to implement the column generation (CG) method and apply it to MBRT planning.

CG was initially introduced in radiation therapy by Romeijn et al. as an alternative to simulated annealing for direct aperture optimisation of photon IMRT plans. In chapter 6, we expanded the optimisation model presented in Romeijn et al. to include electron control points as well as photon control points. CG was chosen for MBRT planning as the pricing problem for beamlet-based DAO can be solved efficiently for each control point, meaning that the algorithm scales favourably with increasing degrees of freedom.

Based on soft tissue sarcoma and chest wall patient cases, we confirmed that, as expected, MBRT plans resulted in lower cost function values than IMRT or MERT. We then confirmed that treatment planning constraints which advantaged photons over electrons resulted in a larger photon dose contribution to the final optimised plan, and vice versa. These results were expected assuming enough iterations of the column generation method were allowed to occur, but it remained to be verified that high-quality MBRT plans could be generated with a clinically viable number of apertures.

In the final section of chapter 6, we explored the performance of multiple aperture addition schemes. The pricing problem produces the best candidate aperture from each control point and modality, and any aperture with a negative price has a potential to improve the cost function. However, in radiation therapy, it is reasonable to seek to obtain the highest quality plan with the fewest number of apertures. We found that for all but one especially greedy aperture addition scheme, the cost function and plan quality converged to a similar value provided enough apertures were included in the treatment plan. There were minor differences in the weight distribution among modalities for each addition scheme reflecting the biases of each scheme, e.g., a scheme which added 1 photon aperture and 5 electron apertures at every iteration featured slightly more dose delivered from electrons in the final plan, compared to a scheme which added 1 photon and 1 electron aperture.

Based on our results, we concluded that the CG method is well-suited for MBRT treatment planning. The choice of the aperture mixing scheme did not heavily bias the resulting treatment plan, but a greedier scheme which added more apertures per iteration resulted in a lower running time for the algorithm.

9.1.3 Robust MBRT treatment plan optimisation

In the work presented in chapter 7, we investigated the possibility of making MBRT plans robust to positioning errors. Based on previous studies of intensity modulated proton radiation

therapy, there was ample evidence that dose distributions involving charged particle beams would be less robust than photon-only radiation therapy. This hypothesis was confirmed by assessing the degradation of CTV coverage and OAR sparing in chest wall and sarcoma plans created using the PTV-based optimisation model derived in chapter 6.

Two new CTV-based robust optimisation models were constructed to explicitly ensure that CTV coverage is preserved for a series of positioning error *scenarios*: a *minimax* model to ensure that the worst scenario is still clinically acceptable, and a *stochastic* model which minimises the expectation value of the cost function based on user-specified probabilities for each scenario to happen. The pricing problem, essential for the application of the CG method, was derived for both models. We found that the two models produced nearly identical treatment plans and used the stochastic optimisation model for the remainder of the work.

The optimisation models were tested on the soft-tissue sarcoma and chest wall irradiation cases presented previously in chapter 6. The resulting CTV-based robust treatment plans were able to preserve CTV coverage under the positioning error scenarios included in the optimisation model, unlike the PTV-based treatment plans. We decomposed the CTV-based robust and PTV-based dose distributions into electron and photon contributions to the total dose, and analysed the differences in dose homogeneity using line profiles. We found that for an identical total dose, the robust dose distributions featured a significantly more even distribution between modalities as a function of depth. In other words, the fraction of electron to photon dose contribution to the total dose remained constant throughout the line profile.

Previously, for the PTV-based model, we found that the choice of aperture addition scheme did not have a significant impact on the quality of the treatment plan. However, after performing a robustness analysis on the dose distributions, we found that the different aperture schemes resulted in different levels of robustness due to the minor differences in the plan modality distributions. This difference between aperture addition schemes disappeared when the CTV-based robust model was used, allowing us to conclude once again that the choice of aperture addition scheme can be made entirely from the perspective of minimising the running time, provided that the robust optimisation model is used.

Based on our results, we concluded that MBRT plans are sensitive to positioning uncertainties and PTV-based optimisation can lead to significant under- and over-dosing of the CTV for positioning errors on the same order as the CTV-to-PTV margin. However, we found that a CTV-based robust optimisation model is able to produce plans which maintain target coverage under the same positioning error scenarios.

9.1.4 Delivery and quality assurance of mixed electron-photon radiation therapy plans

The work presented in chapter 8 was aimed at providing confidence in the results of the MBRT optimisation algorithm developed in the previous chapters by comparing the simulated dose distributions to measurements. Two MBRT treatment plans were created, one aimed at a simple, inherently robust delivery on solid water slabs, and another aimed at validating a complex delivery on a PMMA cylinder.

For the solid water delivery, a film was placed at a depth of 3 cm during delivery, and the measured dose agreed with the simulated dose with a 96.4% gamma pass rate for a 2 %/2 mm criteria.

A complex plan involving a photon arc and 3 electron delivery angles was delivered on the PMMA cylinder, and point dose measurements were taken using an Exradin A1SL ion chamber at a single nominal depth of 1.4 cm. Due to the large differences between the measurement point and the reference conditions in which the ion chamber was calibrated for each modality, a MC-derived beam quality conversion factor was calculated for each modality to correct the measurement. The measured dose from individual modalities agreed within 2.7% with simulated doses, and the total dose measured at the point agreed within 0.2%.

This film measurement results showed that the dose distributions obtained by the MBRT planning algorithm were realistic, deliverable and verifiable in the case of a simple delivery. The ion chamber measurement procedure was performed in under 15 minutes, making it viable as a clinical patient-specific QA protocol.

9.2 Concurrent work

In parallel with this work, Henzen et al. showed that a beamlet-based approach can be used to create photon MLC apertures for MERT using simulated annealing. The effects of MLC leaves on aperture dose distributions were found to be minimal if a full MC recalculation including a model of the MLC was performed for each aperture, followed by weight reoptimisation [1].

Mueller et al. then developed a simultaneous MBRT planning algorithm based on the same simulated annealing approach [2]. Similarly to the results presented in this work, they were able to obtain superior organ at risk sparing while maintaining PTV coverage compared to photon-only plans.

Any advancements in photon-only radiation therapy, such as so-called 4π or non-coplanar deliveries [3, 4], can be adopted into an MBRT context. Mueller et al. have already demonstrated that combined 4π photon trajectories and step-and-shoot MERT can further improve normal tissue sparing compared to MBRT delivered with a coplanar photon component and photon-only plans [5]. They were also able to demonstrate the deliverability of these MBRT plans on modern linear accelerators by obtaining excellent agreement between film measurements and simulated dose distributions.

The work in this thesis, and the concurrent work performed by Henzen et al. and Muller et al. adds confidence to the clinical viability of MBRT. The authors were able to create complex MERT and MBRT plans and verify that the delivery was accurate. However, the plans created tended to include a limited number of electron beam apertures, which leaves room for more validation work to be performed on plans where a large number of electron apertures are needed. In addition, robustness analysis was not performed on the plans shown in their work, and line profiles showing the electron and photon contributions to the total dose had similar high-gradient features as the PTV-based plans that were shown to be non-robust in this thesis.

In 2017, Miguez et al. reported on their clinical implementation of MBRT applied to 7 patients undergoing accelerated partial breast irradiation. To our knowledge, the authors are the only group to have incorporated MBRT as part of clinical practice. The plans featured between 2 and 4 beam angles, and at most two electron energies were used for the electron component, meaning that the treatment plans were considerably simpler than the ones created in this work or the work of Mueller et al. The authors concluded that MBRT is a feasible and secure technique, with no moderate or severe grade of toxicity reported.

9.3 Conclusion and future directions

The good agreement between measured and simulated doses for the deliveries performed in chapter 8 leads us to conclude that the goals of this thesis have largely been achieved, paving the way towards rigorous clinical investigations of MBRT deliveries such as robustness analysis in both patient and QA deliveries and collision detection.

Currently, the beamlet generation process for electron is the most time consuming part of the optimisation workflow. Considerations about the robustness of MBRT deliveries will require more uncertainty scenarios to be considered, which will lead to an even larger number of

electron beamlets to be generated. The use of fast *macro* MC engines could drastically reduce the time needed for plan creation [6–8].

A majority of the work performed on MBRT remains in the realm of simulations rather than physical deliveries, with the notable exception of MBRT plans delivered for accelerated partial breast irradiation treatments [9–11]. As mentioned above, the issues remaining before clinical implementation include a rigorous assessment of the robustness of MBRT deliveries and, in the case of virtual SAD deliveries, a patient-specific collision detection model. The primary focus of future work should be directing at achieving this transition between simulations and clinical delivery on modern linear accelerators.

References

[1] D. Henzen, P. Manser, D. Frei, W. Volken, H. Neuenschwander, E. J. Born, a. Joosten, K. Lössl, D. M. Aebersold, C. Chatelain, M. F. M. Stampanoni, and M. K. Fix, "Beamlet based direct aperture optimization for MERT using a photon MLC.," *Medical physics*, vol. 41, p. 121711, dec 2014.

[2] S. Mueller, M. K. Fix, A. Joosten, D. Henzen, D. Frei, W. Volken, R. Kueng, D. M. Aebersold, M. F. Stampanoni, and P. Manser, "Simultaneous optimization of photons and electrons for mixed beam radiotherapy," *Physics in Medicine and Biology*, vol. 62, no. 14, pp. 5840–5860, 2017.

[3] P. Dong, P. Lee, D. Ruan, T. Long, E. Romeijn, Y. Yang, D. Low, P. Kupelian, and K. Sheng, "4π non-coplanar liver SBRT: A novel delivery technique," *International Journal of Radiation Oncology Biology Physics*, vol. 85, no. 5, pp. 1360–1366, 2013.

[4] P. Dong, H. Liu, and L. Xing, "Monte Carlo tree search -based non-coplanar trajectory design for station parameter optimized radiation therapy (SPORT)," *Physics in Medicine and Biology*, vol. 63, no. 13, 2018.

[5] S. Mueller, M. K. Fix, D. Henzen, D. Frei, D. Frauchiger, K. Loessl, M. F. M. Stampanoni, and P. Manser, "Electron beam collimation with a photon MLC for standard electron treatments," *Physics in Medicine & Biology*, vol. 63, no. 2, p. 025017, 2018.

[6] H. Neuenschwander, T. R. MacKie, and P. J. Reckwerdt, "MMC-a high-performance Monte Carlo code for electron beam treatment planning," *Physics in Medicine and Biology*, vol. 40, no. 4, pp. 543–574, 1995.

[7] K. Jabbari, P. Keall, and J. Seuntjens, "Considerations and limitations of fast Monte Carlo electron transport in radiation therapy based on precalculated data," *Medical Physics*, vol. 36, no. 2, p. 530, 2009.

[8] M. A. Renaud, D. Roberge, and J. Seuntjens, "Latent uncertainties of the precalculated track Monte Carlo method," *Medical Physics*, vol. 42, no. 1, pp. 479–490, 2015.

[9] B. A. Palma, A. U. Sánchez, F. J. Salguero, R. Arráns, C. M. Sánchez, A. W. Zurita, M. I. R. Hermida, and A. Leal, "Combined modulated electron and photon beams planned by a Monte-Carlo-based optimization procedure for accelerated partial breast irradiation," *Physics in Medicine and Biology*, vol. 57, no. 5, pp. 1191–1202, 2012.

[10] E. Jimenez-Ortega, C. Miguez-Sanchez, B. Palma, A. Ureba, H. Miras, R. Arrans, A. Barbeiro, J. Baeza, F. Carrasco, and A. L. Plaza, "SU-E-T-593: Outcomes and Toxicities From a Clinical Trial of APBI Using MERT+IMRT with the Same XMLC," *Medical Physics*, vol. 42, no. 6, pp. 3472–3472, 2015.

[11] C. Míguez, E. Jiménez-Ortega, B. A. Palma, H. Miras, A. Ureba, R. Arráns, F. Carrasco-Peña, A. Illescas-Vacas, and A. Leal, "Clinical implementation of combined modulated electron and photon beams with conventional MLC for accelerated partial breast irradiation," *Radiotherapy and Oncology*, vol. 124, no. 1, pp. 124–129, 2017.