LANCASTER UNIVERSITY

Managing radiotherapy treatment trade-offs using multi-criteria optimisation and data envelopment analysis

by

Kuan-Min Lin

A thesis submitted in partial fulfillment for the degree of Doctor of Philosophy

in the

Management School

Department of Management Science

March 2016

Declaration of Authorship

I, Kuan-Min Lin, declare that this thesis titled, "Managing radiotherapy treatment trade-offs using multi-criteria optimisation and data envelopment analysis" and the work presented in it are my own. The thesis has not been submitted in substantially the same form for the award of a higher degree elsewhere.

Research conducted in this thesis was supervised by Matthias Ehrgott and the work presented in Chapter 4, 5 and 6 was jointly supervised by Andrea Raith. The study in Chapter 4 was a collaborative work with industrial advisors John Simpson, who performed plan re-optimisation, and Giuseppe Sasso, who conducted plan evaluation. The study in Chapter 4 was published in Lin et al. (2013) and the materials presented in both Chapter 5 and 6 have been submitted for publication.

Abstract

Techniques for managing trade-offs between tumour control and normal tissue sparing in radiotherapy treatment planning are reviewed and developed.

Firstly, a quality control method based on data envelopment analysis is proposed. The method measures the improvement potential of a plan by comparing the plan to other reference plans. The method considers multiple criteria, including one that represents anatomical variations between patients. An application to prostate cases demonstrates the capability of the method in identifying plans with further improvement potential.

A multi-criteria based planning technique that considers treatment delivery is then proposed. The method integrates column generation in the revised normal boundary intersection method, which projects a set of equidistant reference points onto the non-dominated set to form a representative set of non-dominated points. The delivery constraints are considered in the column generation process. Essentially, the method generates a set of deliverable plans featuring a range of treatment trade-offs. Demonstrated by a prostate case, the method generates near-optimal plans that can be delivered with dramatically lower total fluence than the optimal ones post-processed for treatment delivery constraints.

Finally, a navigation method based on solving interactive multi-objective optimisation for a discrete set of plans is developed. The method sets the aspiration values for criteria as soft constraints, thus allowing the planner to freely express his/her preferences without causing infeasibility. Navigation is conducted on planner-defined clinical criteria, including the non-convex dose-volume criteria and treatment delivery time. Navigation steps on a prostate case are demonstrated with a prototype implementation. The prostate case shows that optimisation criteria may not correctly reflect plan quality and can mislead a planner to select a "sub-optimal" plan. Instead, using clinical criteria provides the most relevant measure of plan quality, hence allowing the planner to quickly identify the most preferable plan from a representative set.

Acknowledgements

First of all, I express my gratitude to my supervisors, Professor Matthias Ehrgott and Dr Andrea Raith, for their dedicated support on my academic development. Knowledge offered by Matthias was invaluable. It was under his expertise, patience and friendliness I were able to complete this study. Andrea has been a great co-supervisor. Her dedication, enthusiasm and positive attitude motivated me throughout the study. Without their guidance this thesis would not have been possible.

I also thank my industrial advisors, Dr John Simpson and Dr Giuseppe Sasso, for offering knowledge and insights of radiotherapy.

I am grateful to the departmental members of Management Science, Lancaster University Management School. Thank you to Ivan Svetunkov, Oliver Schaer, Patrick Saoud, Xin Liu and Qiulin Yang for making all the PhD activities more enjoyable. A big thanks to Ethan Liu for being a great officemate and for many interesting academic discussions. I also thank Ms Gay Bentinck who I have troubled several times with all the administrative matters.

I express my gratitude to all who I got the chance to meet and work with. I thank the members of the Department of Mathematics at the University of Kaiserslautern for their warm hospitality during a three month research visit. Special thanks to Professor Horst W. Hamacher for his valuable comments. Additionally, I thank Dr Daniele Manerba, who I met at IFORS 2014, for his practical tips on applying column generation.

Heartfelt thanks to all the friends who supported me in different ways. Special thanks to Ming Lin, Yaning Chang, Peiwen Yeh, Claude Yang, Meiyin Lu, Wanju Lee and Vararin Charoenchaikorn for going through the ups and downs of research life with me.

Last, but not least, I extend my deepest gratitude to my parents, for their unconditional love, and to my brothers, for their support and understanding.

Contents

List of Figures

List of Tables

Abbreviations

3DCRT	three-dimensional conformal radiation therapy
CI	conformity index
CPU	central processing unit
CRS	constant returns-to-scale
CT	computed tomography
DAO	direct aperture optimisation
DEA	data envelopment analysis
DMU	decision making unit
DTH	distance to target histogram
DV(H)	dose-volume (histogram)
FDH	free disposal hull
FMO	fluence map optimisation
(g)EUD	(generalised) equivalent uniform dose
GPU	graphical processing unit
Gy	Gray
HI	homogeneity index
IMPT	intensity modulated proton therapy
IMRT	intensity modulated radiation therapy
LFH	left femur head
linac	linear accelerator
MCO	multi-criteria optimisation
MIP	mixed integer programme

MLC	multileaf collimator
MOLP	multi-objective linear programme
MOP	multi-objective programme
MP	master problem
MRI	magnetic resonance imaging
MU	monitor units
NTCP	normal tissue complication probability
OAR	organ at risk
OVH	overlap volume histogram
PET	positron emission tomography
PTV	planning target volume
RFH	right femur head
RMP	restricted master problem
RNBI	revised normal boundary intersection
SP	subproblem
TPC	tumour control probability
VMAT	volumetric-modulated arc therapy
VRS	variable returns-to-scale

Notation

$y^1 \leqq y^2$	$y^1_k \leqq y^2_k$ for $k = 1, \ldots, p$
$y^1 \leq y^2$	$y^1_k \leqq y^2_k$ and $y^1 \neq y^2$
$y^1 < y^2$	$y^1_k < y^2_k$ for $k = 1, \ldots, p$
$\mathbb{R}^p_{\leqq}$	$\{y \in \mathbb{R}^p : y \geqq 0\}$
$\mathbb{R}^p_{>}$	$\{y \in \mathbb{R}^p : y > 0\}$
$\mathbb{R}^p_{\geq}$	$\{y \in \mathbb{R}^p : y \geq 0\}$
i, j, k	dummy indices
e	vector of ones
ε	error vector
ϵ	error value
q	reference point
λ	weight vector
A	constraint matrix; dose deposition matrix
a^j	jth column of matrix A
C	objective matrix for an MOLP
x	variable vector; vector of bixel intensities
X	feasible set
y	vector of objective function values
Y	feasible set in objective space
r	input vector
R	input matrix
w	output vector

W	output matrix
z	vector of environmental variables
Z	matrix of environmental variables
s^-, s^+, s^z	vector of slack variables
D_v	dose-at-volume parameter
V_d	volume-at-dose parameter

Chapter 1

Introduction

Radiation therapy, or radiotherapy, is a medical treatment modality that uses external ionising radiation to control or eliminate cancerous cells. It is the most effective non-surgical cancer treatment modality (Department of Health Cancer Policy Team, 2012) with 50% of all cancer patients recommended with radiotherapy as part of their cancer treatment (Delaney et al., 2005). Around 40% of cured cancer cases receive radiotherapy as part of the cancer treatment (Baskar et al., 2012). Radiotherapy is also cost effective, accounting for only 5% of the total cost of cancer care (Ringborg et al., 2003).

Radiotherapy planning involves managing conflicting goals between irradiating cancerous tissues and sparing of surrounding critical organs. Conventional treatment planning practice is conducted in a trial-and-error manner. As a consequence, the treatment planning process is highly inefficient and the resulting plan quality is not guaranteed. Techniques for managing treatment trade-offs have developed considerably over the last decade. In particular, a multi-criteria optimisation (MCO) based planning approach has been deployed clinically (Craft and Richter, 2013). The multi-criteria planning approach generates a set of plans capturing potentially preferable treatment trade-offs. Given the set of plans, the planning practice becomes an informed decision making process in which the planner compares and evaluates treatment plans in order to find the most preferable plan from the set.

While multi-criteria based planning allows a more informed decision making process, it is not without limitations. This thesis addresses the limitations and further develops the planning practice, taking into consideration other planning approaches in managing treatment trade-offs. In particular, we focus on the following aspects:

- Plan quality control, in which the improvement potential of a given plan is measured quantitatively, thus ensuring only top quality plans are approved for treatments.
- Multi-criteria based plan generation incorporating delivery constraints, such that plans reviewed by the oncologists are directly deliverable without the need for further adjustments that can degrade plan quality.
- Plan evaluation and comparison with clinically relevant criteria, such that given a set of potential plans, one can quickly identify the most preferable plan using any user-defined criteria.

The thesis is organised as follows. The rest of this chapter provides relevant information on radiotherapy treatment planning. In Chapter 2 we provide background material on MCO and data envelopment analysis (DEA). In Chapter 3, we provide a literature review on radiotherapy planning approaches, emphasising techniques for managing treatment trade-offs. We categorise the planning approaches based on how the planners interact with the planning software during the planning process. In Chapter 4, we propose a quality assessment method based on DEA. We apply the method to a set of prostate plans and demonstrate its efficacy. In Chapter 5, an MCO based plan generation method is proposed. The plan generation method uses the column generation technique to incorporate treatment delivery constraints. As a result, the generated plans are directly deliverable, thus avoiding subsequent modification that can degrade plan quality. Chapter 7 considers how a planner can effectively identify the most preferred plan from a given set of candidates. In contrast to previously proposed methods, our technique allows plan comparison based on any user defined criteria, regardless of the mathematical form of the criteria. In particular, convexity is not assumed.

1.1 Radiotherapy

Radiotherapy is based on the fact that cancerous cells are incapable of reproducing themselves if they are damaged by radiation while slightly damaged non-cancerous cells are capable of doing so. Radiotherapy exploits this therapeutic advantage by focusing radiation to a targeted region to kill the cancerous cells while sparing surrounding healthy structures. The two major ways to deliver radiation are external beam radiotherapy and brachytherapy. Brachytherapy is a form of treatment where a radiation source is placed inside the patient's body. In external beam radiotherapy, the radiation, in the form of x-rays, electrons, protons or gamma rays, is generated by a linear accelerator and delivered from a gantry to the tumour. In this thesis we consider external beam radiotherapy.

The goal of radiotherapy is to maximise tumour control without causing unacceptable complications in the normal tissue. This goal is usually expressed in terms of dose to relevant structures, or volumes, of interests, i.e., to deliver a high and uniform dose of a certain level to the target volume while limiting the dose to the surrounding healthy structures to certain acceptable levels. However, achieving the treatment goal can be challenging since as target dose increases, dose to the surrounding healthy structures typically increases. Treatment techniques such as three-dimensional conformal radiation therapy (3DCRT) and intensity modulated radiation therapy (IMRT) were developed to address this challenge.

The 3DCRT is a treatment technique which shapes the radiation beams to be conformal to the beam's-eye-view of the tumour volume. The radiation beam is shaped by simple modifying devices such as wedges or compensating filters. IMRT, first proposed by Brahme (1988), is an advanced form of 3DCRT. In IMRT, radiation fields are shaped by a device called multileaf collimator (MLC) that is attached to the gantry. An MLC consists of a number of pairs of metal leaves that can move in and out of the path of the radiation independently (see Figure 1.1 for an illustration). When radiation passes through the MLC, part of the radiation is blocked by the metal leaves, resulting in a shaped radiation field. The radiation fields of IMRT may not conform to the contour of the target volume, but the cumulative effects of these shaped radiation fields result in a modulated fluence

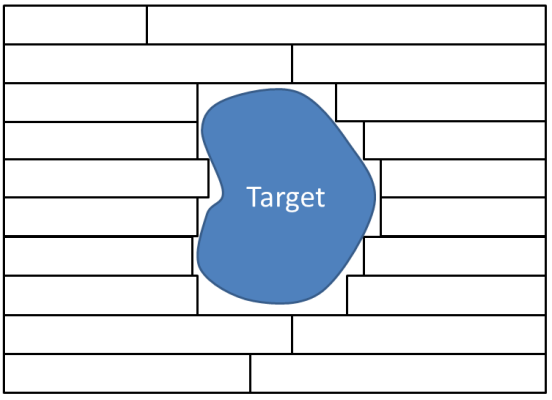

Figure 1.1: An illustration of MLC leaves.

map with a high dose to the tumour while sparing doses to surrounding healthy structures.

1.2 Radiotherapy treatment process

In this section we provide an overview of the radiotherapy treatment process, which consists of image acquisition and simulation, treatment planning, quality assurance and treatment delivery. Readers are referred to specialised textbooks, e.g., Webb (2001) and Schlegel and Mahr (2007) for further detail.

Image acquisition and simulation The first step of radiotherapy planning is image acquisition and simulation. The patient is set up in the treatment position while being scanned by a computed tomography (CT) scanner. Custom-designed body moulds and immobilization devices are usually used to immobilize the patient. The use of these devices should also assist in positioning the patient accurately and consistently for the treatments. Imaging technologies such as CT imaging, magnetic resonance imaging (MRI), positron emission tomography (PET) and other functional imaging studies can be used to facilitate structure definition. The imaging process produces 2D transverse scans of the treatment site. The scans are then reviewed and the borders of the essential structures are contoured manually. The essential structures include the tumour volume, planning target volume (PTV), which is an extension of the tumour volume that is assumed to contain

microscopic extensions of the the cancer, and the organs at risk (OARs). By calibrating the 2D scans along with the contouring of each structure, the treatment software produces a 3D simulation of the treatment anatomy.

Planning The treatment planning proceeds once the 3D simulation of the anatomy is completed. The oncologist determines a set of plan requirements, for example, the prescription dose to the tumour and the tolerable dose-to-volume parameters for the OARs (see Section 1.5 for details). A treatment planner then inputs appropriate planning parameters, which reflect the planning requirements, to a computerised treatment planning system in order to generate a plan that satisfies the plan requirements. Essentially the planning parameters are used to construct planning optimisation problems (as described in Section 1.4). By solving the optimisation problems, one obtains a treatment plan to be evaluated against the plan requirements. Conventional treatment planning practice requires a planner to iteratively adjust the planning parameters in a trial-and-error manner until a satisfactory plan is generated. Advanced planning practices avoid some of the trial-and-error planning practice through techniques such as MCO and knowledge-based planning (see Chapter 3 for details).

Quality assurance IMRT requires exact delivery of radiation to specified regions, thus it is important to assure that the selected treatment plan can be delivered precisely. Quality assurance verifies that the plan can be delivered by the mechanical hardware accurately to achieve the desired dose distribution. Examples of quality assurance include dosimetric verification using phantoms, verification of leaf positions and verification of patient setup.

Treatment delivery During a treatment, radiation is generated by a linear accelerator (*linac*) and delivered through a gantry that can rotate along a central axis. The patient is positioned on a couch that allows vertical, longitudinal, lateral and rotational movement. Through the positioning of the couch and gantry, radiation can be delivered to the target volume from almost any angle, provided that the couch and the gantry do not collide. During the delivery, the gantry moves to one beam direction, delivers the radiation intensities assigned to that

beam direction and proceeds to the next beam direction. During this process an MLC is used to modulate the fluence. The MLC can be operated in two modes: *static* mode and *dynamic* mode. In static mode, radiation is delivered when the MLC leafs remain static. During the transition from one segment to the other, radiation is turned off. The static mode is also referred to as *step-and-shoot* or *stop-and-shoot*. Alternatively, in dynamic mode, which is also referred to as the *sliding window* technique, radiation is delivered continuously while the MLC leafs sweep across the leaf positions unidirectionally to achieve the desirable fluence modulation. While the dynamic MLC mode enables more efficient delivery than the static MLC mode, the mechanism is more complicated and requires a more complicated planning and quality assurance. In this thesis we primarily consider MLC in static mode.

1.3 Mathematical modelling of radiotherapy treatments

To mathematically model the beam intensity, the radiation fields at each beam direction are discretised into a grid of rectangular fields called *bixels*. Bixels are used to represent the accumulated radiation intensities that are delivered from a specific position of a radiation field. The size of the bixels is determined by the width of the individual MLC leaf and the positions along the MLC leaf path, i.e., where the leaf can stop. To evaluate the anatomical dose, the treatment site is discretised into m 3D cubic volumes called *voxels*. Dose is calculated at a certain point within a voxel, called a dose point. It is assumed that the dose value throughout the voxel is the same as the dose value at the dose point.

Given a set of beam directions $k = \{1, \ldots, o\}$ for a treatment, let the radiation fields of all beams be represented by bixels $j = 1, \ldots, n$ and the patient volume be represented by voxels $i = 1, \ldots, m$. The radiation dose distributed to voxel i from bixel j under unit intensity, is represented by a_{ij}. This value is calculated based on the physical behaviour of radiation as it travels through the body. Different dose calculation models, such as pencil beam models (Jeleń and Alber, 2007), the superposition algorithm (Keall and Hoban, 1996) and Monte Carlo simulation

(Reynaert et al., 2007), exist for dose calculation, with simulation models being the gold standard. The a_{ij} can be grouped into a dose deposition matrix A by indexing rows by i and columns by j. The mapping between the intensity and dose is (at least approximately) linear (Kolmonen et al., 1998), hence the dose distribution d is calculated as

$$d = Ax, \tag{1.1}$$

where $d \in \mathbb{R}^m$ is the dose vector in which d_i represents the dose delivered to voxel i. Vector $x \in \mathbb{R}^n$ is the radiation intensity vector in which x_j describes the radiation intensity for bixel j. For convenience, A can be partitioned and re-ordered into sub-matrices according to the structure type of the voxel, i.e., $A_T \in \mathbb{R}^{m_T \times n}$, $A_C \in \mathbb{R}^{m_C \times n}$ and $A_N \in \mathbb{R}^{m_N \times n}$ for the tumour T with m_T voxels, for critical organs C with m_C voxels and for normal tissue N with m_N voxels, respectively.

1.4 Treatment planning optimisation problems

IMRT planning can be considered as three sequential optimisation problems: the beam angle optimisation problem, the fluence map optimisation problem and the realisation (or leaf sequencing) problem (Ehrgott et al., 2010).

Beam angle optimisation The starting point of the planning process is to select a number of beam angles from which the radiation will be delivered to the treatment site. This is referred to as the beam angle optimisation problem. The beam angles selected for a treatment determine the dose deposition matrix A and hence directly affect the quality of plans produced in the subsequent optimisation. The beam angle optimisation problem is a complex non-convex problem (Censor and Unkelbach, 2012). In a clinical environment, the beam directions are usually determined based on the geometry of the treatment site and the experience of the planner. An IMRT treatment often utilises 4 to 9 beam angles (Lim et al., 2014). An overview of the beam angle optimisation problem can be found in Ehrgott et al. (2008c).

Fluence map optimisation (FMO) After the beam directions are decided, the dose distribution matrix is derived and equation (1.1) is incorporated in an optimisation model to find a set of optimal bixel intensities that result in a satisfactory treatment outcome. These bixel intensities are also referred to as the intensity pattern or fluence map. Due to physical limitations, it is impossible to irradiate cancerous tissues without irradiating surrounding healthy tissues. Rather, the planner needs to make trade-offs on the degree of tumour control and healthy tissue sparing. These treatment trade-offs are primarily considered in FMO. Further information on FMO is provided in Section 1.6.

Realisation problem The final step focuses on the realisation of the intensity pattern obtained from FMO. In the delivery of a treatment, an MLC device is used to shape the radiation fields delivered to the treatment site. The shapes that are formed by the MLC during part of the treatment are referred to as *apertures* or *segments*. The radiation output assigned to a segment is referred as the segment weight, measured in monitor units (MU), which is the unit of output measure for the linac. The goal of the realization problem is to find a sequence of segments, each with an associated segment weight, that efficiently reproduces the intensity pattern, by for example, minimising total beam-on time or minimising the number of segments (Baatar et al., 2005). This process is also referred to as *segmentation*. The segments must comply with the physical constraints of the MLC leaves. The elementary ones are collision constraints, that prevent opposing leaves to overlap and constraints that ensure the opening in any MLC row is continuous, i.e., all open bixels in a row are consecutive. A survey of the realisation problem can be found in Ehrgott et al. (2008b).

1.5 Plan evaluation

In this section, we present common criteria used to evaluate the quality of a treatment plan. Throughout this thesis, we will refer to these criteria as *clinical evaluation criteria*. In contrast, we will refer to criteria used in an optimisation model as *optimisation criteria* (reviewed in Section 1.6). Note that a criterion can be both a clinical evaluation criterion and an optimisation criterion. These terms are

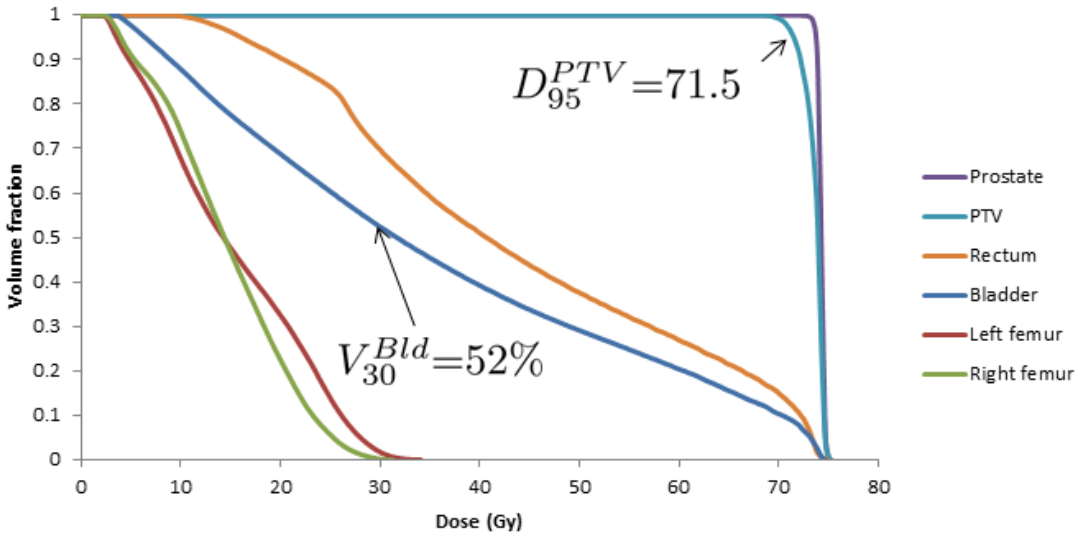

Figure 1.2: DVH of a sample prostate treatment plan.

used to distinguish whether a criterion is used for evaluation or for optimisation during a treatment planning process.

The primary tool for plan evaluation is often based on assessing the so-called *dose-volume histogram* (DVH). A DVH depicts the cumulative volume of a structure that receives at least a given dose value. An example DVH is illustrated in Figure 1.2. Certain points on the DVH are often selected as planning requirements, specified in terms of

- a *dose-at-volume* (D_v) parameter, which is the dose of which the associated iso-dose volume contains $v\%$ of the volume of a given structure.
- a *volume-at-dose* (V_d) parameter, which is the percentage volume of a given structure that is contained by the d Gray (Gy, unit for radiation dose) iso-dose volume.

Example treatment requirements for prostate cancer cases, with the prescription dose denoted as d_p, are shown in Table 1.1.

Other common evaluation criteria include *homogeneity index* (HI) and *conformity index* (CI). HI measures the uniformity of dose distribution within a target volume. Ideally, one would like the voxels within a target volume to receive a uniform dose.

Table 1.1: An example of DV requirements for prostate treatment planning.

Structure	DV parameter		DV requirement
Tumour volume	D_{99}	$\geq$	99% $\times$ d_p Gy
PTV	D_{95}	$\geq$	95% $\times$ d_p Gy
Rectum	V_{60}	$\leq$	40%
	V_{70}	$\leq$	10%
Bladder	D_{50}	$\leq$	40 Gy
	D_{25}	$\leq$	75 Gy
Femur heads	D_{10}	$\leq$	50 Gy

Various formulations for HI have been proposed (see, e.g., Kataria et al. (2012)). The formulation proposed by Wu et al. (2003a), as shown in (1.2), is the most common HI index used in the literature,

$$\text{HI} = \frac{D_2 - D_{98}}{d_p} \times 100\%. \tag{1.2}$$

The use of D_2 and D_{98}, instead of the exact minimum and maximum dose, better represents the dose received by a volume because the calculation of the true maximum and minimum is typically unreliable (Wu et al., 2003b). CI measures how well the dose distribution conforms to the target volume. Various formulations of CI exist (see Feuvret et al. (2006) for a review). Here we consider the CI formulation proposed by Lomax and Scheib (2003) ,

$$\text{CI} = \frac{V_{(T,d)}}{V_d}, \tag{1.3}$$

where $V_{(T,d)}$ is the d Gy isodose volume of the target and V_d is the d Gy isodose volume of the total patient volume. Hence a good target coverage by the d Gy isodose volume is indicated by a CI value close to 1.

Criteria based on biological responses can also be used for plan evaluation. Common biological criteria include tumour control probability (TCP), normal tissue complication probability (NTCP) and (generalised) equivalent uniform dose ((g)EUD). Variations of these criteria exist due to different modelling approaches (mechanistic or phenomenological) and different modelling factors are being considered (e.g., effects of repopulation, redistribution and/or reoxygenation). Readers are referred to Niemierko (2006) for an overview and to Romeijn et al. (2004) and

Hoffmann et al. (2008) for a review of commonly used biological criteria. In this thesis, we consider the gEUD formulation proposed by Niemierko (1999):

$$\text{gEUD} = \left(\frac{1}{M} \sum_{i=1}^{M} d_i^a \right)^{\frac{1}{a}}, \tag{1.4}$$

where M is the number of the voxels of a given structure and a is a structure specific parameter that is negative for tumours and positive for non-tumourous tissues.

1.6 Fluence map optimisation

The goal of FMO is to find a fluence map that maximises the probability of tumour control while avoiding unacceptable complications in healthy structures. In treatment planning, this goal is often translated to dose and volume requirements derived from evidence-based knowledge on the outcomes of treatments (Marks et al., 2010). These requirements naturally lead to *physical* optimisation models, which are models formulated based on dose and volume. In contrast, the *biological* optimisation models, which are models based on biological criteria such as TCP and NTCP, are not commonly used due to uncertainties in the model parameters (Niemierko, 2006). In this section we introduce FMO in terms of physical optimisation models due to their active use in research and practice (Bortfeld, 1999, Ehrgott et al., 2008a, Shepard et al., 1999).

Consider the following abstract FMO model:

$$\begin{aligned} \min \quad & f_0(x) \\ \text{s.t.} \quad & f_i(x) \leqq 0 \quad \text{for} \quad i = 1, \ldots, r, \\ & x \geqq 0. \end{aligned} \tag{1.5}$$

Here x is a vector representing bixel intensities, which is required to be non-negative. The objective function $f_0(x)$ provides a quantitative measure on the quality of the intensity pattern. Constraints $f_i(x) \leqq 0$ specify the planning requirements as a function of the intensity pattern. The goal of the optimisation

problem is to find an optimal intensity pattern that results in a minimal objective function value. Here minimisation is considered without loss of generality since maximisation of a function f is equivalent to minimisation of $-f$.

To find a satisfactory intensity pattern, the choice of objective function and constraints should reflect the desirable characteristics of an ideal clinical outcome. To do so, the planner specifies a prescription dose for the target and tolerance dose levels for the critical organs and normal tissues. In particular we consider dose lower bound $LB_T \in \mathbb{R}^{m_T}$ for the tumour and dose upper bounds $UB_T \in \mathbb{R}^{m_T}$, $UB_C \in \mathbb{R}^{m_C}$ and $UB_N \in \mathbb{R}^{m_N}$ for the tumour, critical organs and normal tissue, respectively.

A majority of physical functions are of the form of an $l_p - norm$ which essentially measures the distance between a dose vector d to a reference dose vector d_r

$$f_0(d) = c\|d - d_r\|_p, \tag{1.6}$$

where c is a scaling parameter and the anatomical dose d is related to the bixel intensities through equation (1.1). For healthy structures, under-dosing can be neglected using

$$f_0(d) = c\| \left(d - d_r\right)_+ \|_p, \tag{1.7}$$

with $(\cdot)_+ = \max\{0, \cdot\}$. In conventional planning, trade-offs between target coverage and OAR sparing are usually managed by assigning a non-negative *importance scale* or *weighting factor* for each objective function, e.g., λ_T for the turmour, λ_C for critical organs and λ_N for normal tissue. The weighted objectives are then aggregated to form a single objective function, which is being minimised. As an example, a min-max formulation by Lim et al. (2007), as shown in (1.8), minimises the weighted sum of maximum deviation on the tumour and maximum overdose on healthy structures using the $l_\infty - norm$

$$\min \quad \lambda_T\|A_T x - d_p\|_\infty + \lambda_C\|(A_C x - UB_C)_+\|_\infty + \lambda_N\|(A_N x - UB_N)_+\|_\infty. \tag{1.8}$$

In clinical practice, a weighted least-squares model (which is essentially a squared $l_2 - norm$) has become an accepted standard (Ehrgott et al., 2010).

The optimisation model can be made more clinically relevant by the use of constraints. One obvious choice for constraints are the *DV constraints*, which incorporate the DV requirements into the model. Precise formulation of DV constraints requires the use of binary variables b_i which turn the model into a mixed integer programme (MIP). For example, the DV constraint for critical organs C can be expressed as

$$\begin{aligned} A_i x &\leqq (1 + b_i F) U B_{C_i} \quad \forall i \in C \\ \sum_{i \in C} b_i &\leqq P m_C, \end{aligned} \tag{1.9}$$

where F is the overdose fraction and P is the volume fraction. Alternatively, the DV constraints can be incorporated in the model implicitly through an iterative post-optimisation process, e.g., by iteratively introducing a penalty objective for voxels violating the DV constraints (Spirou and Chui, 1998). Note that DV constraints are non-convex (Deasy, 1997) with existence of local minima (see, e.g., Llacer et al. (2003) which can have major differences from the global minimum (Jeraj et al., 2003, Wu et al., 2003a).

An alternative form of constraints are *dose constraints*, which set the threshold values for the dose received by the voxels of a certain structure, for example, $A_T x \geqq LB_T$ and $A_c x \leqq UB_C$. However, it can be non-trivial to set the threshold values, since if the constraints are too loose, plan improvement potential is sacrificed while if the constraints are too tight, the model becomes infeasible.

To address the infeasibility issue, Holder (2003, 2005, 2006) proposes to use elastic constraints, which allow violation of threshold values while violations are penalised:

$$
\begin{array}{lrcll}
\min & \omega l^T\alpha + u_C^T\beta + u_N^T\gamma & & & \\
\text{s.t.} & LB_T - L\alpha \leqq & A_T x & \leqq UB_T & \\
& A_C x \leqq & UB_C + U_C\beta & & \\
& A_N x \leqq & UB_N + U_N\gamma & & \\
& 0 \leqq & L\alpha & \leqq LB_T & \\
& -UB_C \leqq & U_C\beta & & \\
& 0 \leqq & U_N\gamma & & \\
& 0 \leqq & x, & &
\end{array}
\tag{1.10}
$$

where ω is the weighting factor for achieving the dose lower bound for the tumour and $\alpha \in \mathbb{R}^{qT}, \beta \in \mathbb{R}^{qC}, \gamma \in \mathbb{R}^{qN}, l \in \mathbb{R}^{qT}, u_C \in \mathbb{R}^{qC}, u_N \in \mathbb{R}^{qN}, L \in \mathbb{R}^{m_T \times qT}, U_C \in \mathbb{R}^{m_C \times qC}$ and $U_N \in \mathbb{R}^{m_N \times qN}$. Let e be a vector of ones with an appropriate dimension. The author proposes two types of solution analysis. One is average analysis, where $l = \frac{1}{m_T}e, u_C = \frac{1}{m_C}e$ and $u_N = \frac{1}{m_N}e$ and L, U_C and U_N are identity matrices (thus $qT = m_T, qC = m_C, qN = m_N$). The other is absolute analysis, where $l = 1, u_C = 1, u_N = 1, L = e, U_C = e$ and $U_N = e$ (thus $qT = qC = qN = 1$). A multi-objective formulation of Holder's average-analysis model is used in Chapter 6.

1.7 Extension: Volumetric-modulated arc therapy

Volumetric-modulated arc therapy (VMAT) is a further development of IMRT in which radiation is delivered continuously with an MLC in dynamic mode while the gantry rotates around the target to form one or more radiation arcs. Since VMAT can utilize many more angles to deliver the radiation than fixed-gantry IMRT in a given treatment time, it provides better flexibility in shaping the dose distribution (Bedford, 2009). It is acknowledged that the dosimetric quality of VMAT plans is comparable to that of IMRT plans while the delivery efficiency is better due

to continuous delivery and variable dose rate (see Studenski et al. (2013) and the references within).

Although VMAT is delivered continuously, the dose distribution of VMAT is not calculated from a continuous model. Instead, the dose distribution delivered by VMAT is approximated with discrete static beams that spread evenly and tightly over the arc interval, each with an assigned segment and MU. During treatment delivery, the machine moves the gantry and the MLC leaves from one planned position to the next planned position, with variable movement speed.

Since the dose distribution is calculated from discrete static beam samples, the FMO of VMAT is almost identical to that of IMRT, except that a large set of beams is used. However, in terms of realisation, a VMAT plan needs to consider an additional constraint when compared to an IMRT plan: the inter-leaf connectivity constraint. That is, the leaf displacements from one planned beam angle to the next must be attainable given the maximum leaf movement speed and the planned gantry rotational speed. Optimisation methods in VMAT primarily consider the leaf sequencing aspect of the planning problem (Unkelbach et al., 2015). A review of VMAT is available in Yu and Tang (2011).

Chapter 2

Multi-criteria optimisation and data envelopment analysis

In this chapter, we present basic materials of MCO (Section 2.1) and DEA (Section 2.2).

2.1 Introduction to multi-criteria optimisation

In this section we present concepts, definitions and solution approaches of MCO. Note that the terms "multi-criteria optimisation" and "multi-objective optimisation" are used interchangeably to refer to the field of optimisation with more than one objective function. In the literature, "multi-criteria optimisation" is typically used by studies in the field of radiotherapy treatment planning whereas "multi-objective optimisation" is usually used by studies in the field of operational research. A comprehensive introduction of MCO materials is available in Ehrgott (2005).

2.1.1 Definitions

Consider a general multi-objective programme (MOP)

$$\min \{f(x) : x \in X\}, \tag{2.1}$$

where $x \in \mathbb{R}^n$ is a vector of decision variables and X is the feasible set of x. $f(x) = (f_1, \ldots, f_p)^T$ is a vector of criteria or objective functions. We denote Y as the feasible set in objective space,

$$Y = \{f(x) : x \in X\}. \tag{2.2}$$

For multiobjective linear programmes (MOLPs), we express $f(x)$ as Cx where $C \in \mathbb{R}^{p \times n}$ is the coefficient matrix consisting of row vectors $c^k \in \mathbb{R}^n$ for $k = 1, \ldots, p$.

We use the following notation for comparison of vectors.

Definition 2.1. Let $y^1, y^2 \in \mathbb{R}^p$, we say

$$\begin{aligned} y^1 \leqq y^2 &\Leftrightarrow y^1_k \leqq y^2_k \quad \text{for} \quad k = 1, \ldots, p \\ y^1 \leq y^2 &\Leftrightarrow y^1_k \leqq y^2_k \quad \text{and} \quad y^1 \neq y^2 \\ y^1 < y^2 &\Leftrightarrow y^1_k < y^2_k \quad \text{for} \quad k = 1, \ldots, p \end{aligned}$$

Furthermore, we define the cones $\mathbb{R}^p_{\geqq} = \{y \in \mathbb{R}^p : y \geqq 0\}$, $\mathbb{R}^p_{>} = \{y \in \mathbb{R}^p : y > 0\}$ and $\mathbb{R}^p_{\geq} = \{y \in \mathbb{R}^p : y \geq 0\}$.

Definition 2.2 (Dominance)**.** For two vectors $y^1, y^2 \in \mathbb{R}^p$, we say

- y^1 *dominates* y^2 if $y^1 \leq y^2$
- y^1 *weakly dominates* y^2 if $y^1 \leqq y^2$
- y^1 *strictly dominates* y^2 if $y^1 < y^2$.

In MCO, as more than one objective is considered, a single solution that simultaneously optimises all criteria generally does not exist. Instead, MCO seeks for solutions that cannot improve in any single criterion without deteriorating at least

one other criterion. Solutions with this property are referred to as *efficient* solutions and the points obtained by mapping the efficient solutions to the objective space are referred to as *non-dominated* points (Definition 2.3).

Definition 2.3 (Efficiency, Non-dominance)**.** A feasible solution $\hat{x} \in X$ of (2.1) is called efficient if there exists no $x \in X$ such that $f(x) \leq f(\hat{x})$. The set of all efficient solutions of (2.1) is called the efficient set in decision space and is denoted by X_E. The corresponding point $\hat{y} = f(\hat{x})$ in criterion space is called a non-dominated point and $Y_N = \{f(x) : x \in X_E\}$ is the non-dominated set in criterion space of (2.1).

Definition 2.4 (Weak efficiency, Weak non-dominance)**.** A feasible solution $\hat{x} \in X$ of (2.1) is called *weakly efficient* if there exists no $x \in X$ such that $f(x) < f(\hat{x})$. The set of all weakly efficient solutions of (2.1) is called the weakly efficient set in decision space and is denoted by X_{WE}. The corresponding point in criterion space $\hat{y} = f(\hat{x})$ is called a *weakly non-dominated* point and $Y_{WN} = \{f(x) : x \in X_{WE}\}$ is the weakly non-dominated set in criterion space of (2.1).

Definition 2.5 (ε-efficiency, ε-nondominance)**.** Let $\varepsilon \in \mathbb{R}^p_{\geqq}$.

- A feasible solution $\hat{x} \in X$ of (2.1) is called *ε-efficient* if there exists no $x \in X$ such that $f(x) \leq f(\hat{x}) - \varepsilon$. The corresponding point in criterion space $\hat{y} = f(\hat{x})$ is called an *ε-nondominated point*.
- A feasible solution $\hat{x} \in X$ of (2.1) is called *weakly ε-efficient* if there exists no $x \in X$ such that $f(x) < f(\hat{x}) - \varepsilon$. The corresponding point in criterion space $\hat{y} = f(\hat{x})$ is called a *weakly ε-nondominated point*.

2.1.2 Obtaining non-dominated points

In this subsection we introduce common MCO techniques that can be used to obtain one or more non-dominated points.

Weighted-sum method Let $\lambda^T = (\lambda_1, \lambda_2, \ldots, \lambda_p)$ be a weight vector with entries representing the relative weighting for each criterion. By assign weightings

to the criteria, (2.1) is transformed to a single criterion weighted-sum problem:

$$\min\{\lambda^T f(x) : x \in X\}. \tag{2.3}$$

Essentially, λ specifies the marginal rate of substitution among the criteria and can be visualised as a hyperplane with a normal λ in criterion space where any point y on the hyperplane has the same value of $\lambda^T y$. Thus, to obtain an optimal solution of (2.3), one moves the hyperplane toward the non-dominated set of Y to the greatest extent while ensuring at least one point on the hyperplane is feasible. Let an optimal solution of (2.3) be $\hat{x}$. At optimum, λ and $f(\hat{x})$ then define a supporting hyperplane of Y. Optimal solutions of (2.3) with positive/non-negative weights are always efficient/weakly efficient (Ehrgott, 2005, Proposition 3.9).

Constraint methods Another way to obtain non-dominated points is through imposing constraints to the model. The constraints limit the solution space such that (weakly) efficient solutions can be found by solving a single criterion problem. One of the most well known constraint methods is the ϵ-constraint method in which only one of the original objectives is minimised while others are transformed into constraints with upper bounds ε:

$$\begin{aligned} \min \quad & f_i(x) \\ \text{s.t.} \quad & f_j(x) \leqq \varepsilon_j \qquad \text{for} \quad j \in \{1, \ldots, p\} \setminus \{i\} \\ & x \in X. \end{aligned} \tag{2.4}$$

It can be non-trivial to determine an appropriate ε. If the constraints are too tight, (2.4) may be infeasible. Instead, the elastic constraint method (Ehrgott and Ryan, 2002) relaxes these constraints and penalises violation of the upper bounds:

$$\begin{aligned} \min \quad & f_i(x) + \sum_{j \in \{1,\ldots,p\} \setminus \{i\}}^{p} \mu_j s_j \\ \text{s.t.} \quad & f_j(x) - s_j \leqq \varepsilon_j \quad \text{for} \quad j \in \{1,\ldots,p\} \setminus \{i\} \\ & s_j \geqq 0 \quad \text{for} \quad j \in \{1,\ldots,p\} \setminus \{i\} \\ & x \in X, \end{aligned} \tag{2.5}$$

where $\mu_j \geq 0, j = \{1,\ldots,p\} \setminus \{i\}$. An optimal solution of (2.4) and (2.5) is weakly efficient (Ehrgott, 2005, Proposition 4.3 and 4.8) but with appropriate j, ε and μ, one obtains efficient solutions by solving (2.5) (Ehrgott, 2005, p. 104). Furthermore, given a weakly efficient solution $\bar{x}$, a corresponding efficient solution can be obtained by solving a hybrid model of the weighted-sum method and the ε-constraint method:

$$\begin{aligned} \min \quad & \sum_{i}^{p} f_i(x) \\ \text{s.t.} \quad & f_i(x) \leqq f_i(\bar{x}) \quad \text{for} \quad i = 1,\ldots,p \\ & x \in X. \end{aligned} \tag{2.6}$$

Reference-point based methods Reference-point based methods transform an MOP to a single criterion problem that minimises a certain metric between the original criteria $f(x)$ and a reference point $q \in \mathbb{R}^p$. One example of the reference-point methods is the weighted Tchebycheff method (Steuer and Choo, 1983):

$$\begin{aligned} \min \quad & \max_{i=1,\ldots,p} \lambda_i \left(f_i(x) - q_i \right) \\ \text{s.t.} \quad & x \in X. \end{aligned} \tag{2.7}$$

Prioritised goal programming In prioritised goal programming, the multiple criteria are prioritised and each assigned a goal value. Single-criterion optimisations are conducted in the prioritised sequence to minimise the deviation between criterion value $f_j(x)$ and its goal value q_j. The achieved criterion value $\bar{q}_i$ from previous higher-priority ($i = 1, \ldots, j-1$) optimisations are turned into hard constraints for subsequent (lower-priority) optimisations:

$$\begin{aligned} \min \quad & f_j(x) - q_j \\ \text{s.t.} \quad & f_i(x) \leqq \bar{q}_i \qquad \text{for} \quad i = 1, \ldots, j-1 \\ & x \in X. \end{aligned} \tag{2.8}$$

A variation of prioritised goal programming is lexicographic programming where the high-priority criteria are optimised (i.e., $\min_{x \in X} f_j(x)$) and the optimised values are turned into constraints for subsequent optimisations. Solutions produced from lexicographic optimisation are efficient (Ehrgott, 2005, Lemma 5.2) whereas this may not be true for prioritised goal programming.

2.1.3 Solution approaches to MCO

The goal of MCO is to find a most preferable solution for the decision maker. Hence the preferences of the decision maker influence the solution approaches of MCO. In this subsection, we classify the solution approaches of MCO subject to the availability of the decision maker's preferences during the solution process. Further detail of the classification scheme is available in Miettinen (1999).

- *No-preference* methods, in which the decision maker has no preference on the solution.
- *A priori* methods, in which the decision maker's preferences are specified before the solution process.
- *A posteriori* methods, in which a set of solutions that capture all decision trade-offs are obtained with absence of decision maker's preferences. Given

the set of solutions, the decision maker can learn the decision trade-offs and then select the most preferred solution out of the solution set.

- *Interactive* methods, in which the decision maker's preferences are gradually articulated by iteratively exploring potential solutions during the solution process.

In the no-preference methods, the decision maker is satisfied as long as a solution is efficient. The a priori methods seek to find an efficient solution that best matches the decision maker's pre-specified preferences. The interactive methods can be considered as iteratively conducting an a priori method subject to changing preferences. Both the a priori and the interactive methods do not require a decision maker to obtain the non-dominated set. In contrast, the a posteriori methods seek to obtain the non-dominated set and one efficient solution in the pre-image of every non-dominated point. The decision maker then selects the most preferred non-dominated point and a corresponding efficient solution out of the set.

2.1.4 Representation and approximation of the non-dominated set

In practice, it might be neither practical nor desirable to generate the non-dominated set. For multi-objective continuous optimisation problems, the non-dominated set consists of infinitely many non-dominated points and it is impractical for a decision maker to examine all of them. Instead, a practical approach is to obtain a discrete representation of the non-dominated set that satisfies some quality requirements (Faulkenberg and Wiecek, 2010, Sayın, 2000).

Definition 2.6 (A representative non-dominated set)**.** A set of discrete non-dominated points $Y' \subset Y_N$ used to represent the non-dominated set is referred to as a *representative* non-dominated set or a *representation* of the non-dominated set.

When the optimisation required to find a non-dominated point is computationally expensive, computing evenly distributed non-dominated points across the non-dominated set with a fine granularity can be time-consuming. To reduce the

computational burden, it is often practical to generate an *approximation* of the non-dominated set or an *approximated* non-dominated set within a tolerable error instead of generating a representative non-dominated set.

Definition 2.7 (An approximated non-dominated set)**.** An approximated non-dominated set or an approximation of the non-dominated set refers to a set that is used as an estimation of the non-dominated set. The points in the approximation are not necessarily non-dominated or feasible.

Methods for generating an approximation of the non-dominated set are reviewed by Ruzika and Wiecek (2005). In particular, a subclass of the approximation methods, the sandwiching methods, are of close relevance to the field of radiotherapy treatment planning optimisation. A sandwiching method is characterised by a procedure that iteratively encloses the space between an upper bounding set $Y^u \subset \mathbb{R}^p$ and a lower bounding set $Y^l \subset \mathbb{R}^p$ of Y_N. A lower (upper) bound set contains points that act as lower (upper) bounds for some points in Y_N (Ehrgott and Gandibleux, 2007). Precisely, let $Y' \subset Y_N$, then $Y' \subseteq Y^l + \mathbb{R}^p_{\geqq}$ and $Y^u \subseteq Y' + \mathbb{R}^p_{\geqq}$. For multi-objective convex optimisation problems, given a set of feasible solutions $\bar{X}$ and let the mapping of $\bar{X}$ in the objective space be $\bar{Y}$, then by convexity, the non-dominated set of the convex hull of $\bar{Y}$ defines an upper bound set of Y_N. A lower bounding set can be obtained from supporting hyperplanes of Y. In a sandwiching method, an upper bound set acts as an approximation of Y_N and a lower bound set provides an error measure of the approximation. Given an approximation of Y_N, the decision maker then wishes to find the most preferred point out of the set and its corresponding solution. In practice, the searching process to find such a point is often conducted by solving interactive multi-objective optimisations in which the feasible set of the original problem is further constrained to the convex hull of $\bar{X}$.

2.2 Introduction to data envelopment analysis

In this section we introduce basics of data envelopment analysis. For a detailed introduction to DEA we refer the interested reader to Chapter 6 of Coelli et al. (2005) and to Cooper et al. (2011).

2.2.1 Efficiency, the production possibility set and the production frontier

DEA is a technique for evaluating the performance of a set of peer entities referred to as decision-making units (DMUs). The performance of each DMU is measured based on how good the DMU is in converting a set of inputs into a set of outputs. In particular, the performance of a DMU is evaluated by comparing its inputs and outputs to a set of inputs and outputs potentially attainable under certain assumptions. Essentially, given a DMU with certain amounts of inputs and outputs, DEA finds a greatest attainable improvement in inputs and outputs for the DMU.

Assume there are N DMUs, each converts a varying amount of I inputs into a varying amount of O outputs. Let $r^j \in \mathbb{R}^I_{\geq}$ and $w^j \in \mathbb{R}^O_{\geq}$ be the amount of inputs consumed and outputs produced by the jth DMU , respectively. The jth DMU is represented by a point $\text{DMU}^j = \left(r^{j^T}, w^{j^T}\right)^T \in \mathbb{R}^{I+O}$. We denote P as a set of potentially attainable inputs and outputs derived by N existing DMUs under certain assumptions (discussed in Subsection 2.2.3). Efficiency in DEA is defined as the following.

Definition 2.8 (DEA efficiency)**.**

- The kth DMU is efficient and (r^{k^T}, w^{k^T}) is non-dominated if there exists no $(r, w) \in P$ such that $r \leq r^k$ and $w \geq w^k$.
- The kth DMU is weakly efficient and (r^{k^T}, w^{k^T}) is weakly non-dominated if there exists no $(r, w) \in P$ such that $r < r^k$ and $w > w^k$.

The concept of efficiency is illustrated in Figure 2.1 in which a set of existing DMUs are represented by points A to J and the set of potentially attainable inputs and outputs is illustrated in grey. In the figure, DMUs represented by A, B, C and D are considered efficient since none of the potentially attainable inputs and outputs show a lower or equal input value and a higher or equal output value simultaneously than these DMUs. Other DMUs are considered inefficient since there are potentially attainable inputs and outputs (e.g., points A, B, C and D) that empirically suggest the inefficient DMUs can be improved in the input/output

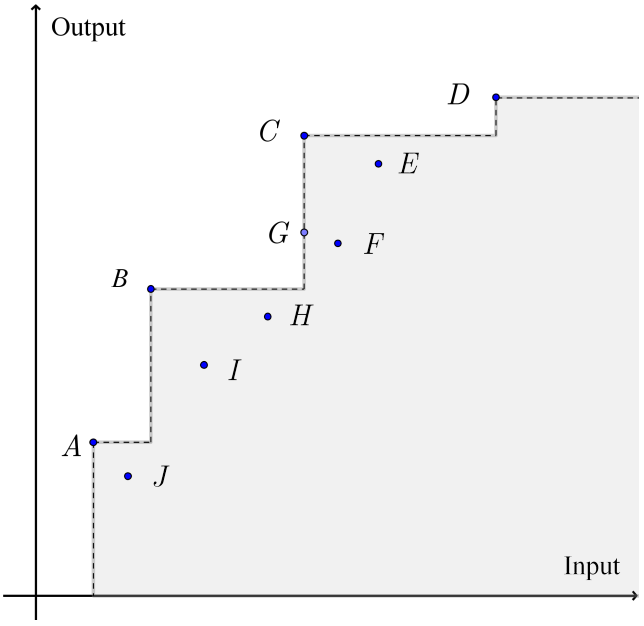

Figure 2.1: Illustrating the concept of efficiency. DMUs represented by points A, B, C and D are considered efficient while other DMUs are inefficient.

without worsening the output/input. Furthermore, the DMU represented by point G is a weakly efficient DMU as it cannot be improved further in input while improvement in output is attainable. The set of potentially attainable inputs and outputs P is referred to as the *production possibility set* in the DEA literature. A production possibility set can be characterised by different assumptions (as discussed in Subsection 2.2.3). The production possibility set shown in Figure 2.1 is referred to as the *free disposal hull* (FDH), which represents a set of attainable DMUs assuming free disposability, i.e., each DMU can consume extra inputs while producing the same levels of outputs or produce less outputs with the same levels of inputs. The frontier that envelops the production possibility set, illustrated by the dashed lines in Figure 2.1, is referred to as the *production frontier*. The production frontier consists of potentially attainable inputs and outputs that cannot be further improved in at least one of the inputs or outputs, i.e., all efficient and weakly efficient DMUs.

2.2.2 An introduction to DEA models

DEA models can have different formulations depending on the intuition behind the comparison of DMUs (the ratio, multiplier or the envelopment form), the orientation of comparison (input-oriented or output-oriented) and the assumptions on the production possibility set (free disposability and returns-to-scale of inputs and outputs). In this subsection, we present input-oriented DEA models with different forms under the assumption of free disposability and constant returns-to-scale (CRS). For a corresponding output-oriented model, readers are referred to Cooper et al. (2011).

Comparing DMUs is difficult when the relative importance or cost of the multiple inputs and outputs are undefined. In the original DEA model proposed by Charnes et al. (1978), as shown in (2.9), comparison of DMUs is facilitated by allowing each DMU to choose weight vectors, one for the inputs and one for the outputs, to maximize its own ratio of weighted output to weighted input. If a DMU is incapable of achieving a superior ratio of weighted output to weighted input than other DMUs using its optimal weights, the DMU is inefficient. The ratio of weighted output to weighted input is referred to as the *efficiency score*.

Let the inputs and outputs of all N DMUs be represented by an input matrix $R \in \mathbb{R}_{\geq}^{I\times N}$ and an output matrix $W \in \mathbb{R}_{\geq}^{O\times N}$ where the jth column vectors of the two matrices are r^j and w^j, respectively. The weights for the inputs and the outputs are denoted by vectors $\rho \in \mathbb{R}^I$ and $\upsilon \in \mathbb{R}^O$, respectively. The optimal ratio of weighted outputs to weighted inputs for the kth DMU can be obtained by solving the following DEA model in the ratio form:

$$\begin{aligned} \max \quad & \frac{\upsilon^T w^k}{\rho^T r^k} \\ \text{s.t.} \quad & \frac{\upsilon^T w^j}{\rho^T r^j} \leqq 1 \qquad \text{for} \quad j = 1, \ldots, N \\ & \upsilon, \rho \geqq 0. \end{aligned} \tag{2.9}$$

Given an optimal solution $\left(\upsilon^{*^T}, \rho^{*^T}\right)$, if the kth DMU is efficient, then the optimal ratio of weighted outputs to weighted inputs for the kth DMU would equal to 1.

Alternatively, if the kth DMU is inefficient, we obtain an optimal ratio of less than 1 due to a binding constraint $\upsilon^T w^{j^*}/\rho^T r^{j^*} = 1$ caused by the j^*th DMU. The j^*th DMU is referred to as the target for the kth DMU due to a superior ratio of weighted outputs to weighted inputs. Note that the target is only guaranteed to be weakly efficient due to the non-negative constraints of the weight vectors υ, ρ. Full development of the model would replace the non-negative constraints with a strict positive constraints, i.e., $\upsilon, \rho \geqq \epsilon e, \epsilon \in \mathbb{R}_>$, to ensure the target is an efficient DMU (Charnes et al., 1979). In the DEA literature, ϵ is referred to as a non-Archimedean element, i.e., a number smaller than any positive real number.

The ratio form has an infinite number of optimal solutions since if $\left(\upsilon^{*^T}, \rho^{*^T}\right)$ is an optimal solution, with any $\alpha \in \mathbb{R}_>$, $\left(\alpha\upsilon^{*^T}, \alpha\rho^{*^T}\right)$ would also be an optimal solution. A particular solution can be selected by setting the denominator of the objective function to one and transforming the linear fractional programme to an equivalent linear programme (Charnes and Cooper, 1962):

$$\begin{aligned} \max \quad & \upsilon^T w^k \\ \text{s.t.} \quad & \upsilon^T w^j - \rho^T r^j \leqq 0 \qquad \text{for} \quad j = 1, \ldots, N \\ & \rho^T r^k = 1 \\ & \upsilon, \rho \geqq \epsilon e. \end{aligned} \tag{2.10}$$

Model (2.10) is referred to as the multiplier form and its dual formulation, as shown in (2.11), is referred to as the envelopment form, with decision variables $\theta \in \mathbb{R}, \lambda \in \mathbb{R}^N, s^+ \in \mathbb{R}^O$ and $s^- \in \mathbb{R}^I$.

$$\begin{aligned} \min \quad & \theta - \epsilon(e^T s^+ + e^T s^-) \\ \text{s.t.} \quad & W\lambda - s^+ = w^k \\ & R\lambda - \theta r^k + s^- = 0 \\ & \lambda, s^+, s^- \geqq 0. \end{aligned} \tag{2.11}$$

Let an optimal solution of (2.11) be $\theta^{k*}, \lambda^{k*}, s^{+*}$ and s^{-*}, where θ^{k*} is the efficiency score of the kth DMU, λ^{k*} is a vector of weights and s^{+*} and s^{-*} are slack variables for outputs and inputs, respectively. Essentially, the model searches for an optimal improvement of DMU^k among the set $\{(r^T, w^T)^T : r = R\lambda, w = W\lambda, \lambda \geqq 0\}$. The

search is facilitated by radial contraction of the input vector through θ and by the slack variables s^+ and s^-. Since the penalty weight, ϵ, for the slack variables is very small, the model seeks for maximal radial contraction (by minimising θ) of the input vector where possible, which projects DMU^k to a point on the production frontier where at least one of the inputs or outputs cannot be improved. Slack variables are then used to search for improvements in the individual inputs/outputs of the projected point.

By solving (2.11), one obtains a non-dominated point with inputs equal to $R\lambda^{k*}$, $R\lambda^{k*} = \theta^* r^k - s^{-*}$ and outputs equal to $W\lambda^{k*}$, $W\lambda^{k*} = w^k + s^{+*}$, representing an optimal input-oriented improvement available for the kth DMU. The point is referred to as the target for the kth DMU and the DMUs associated with a positive weight in forming the target are referred to as the peers of the kth DMU, i.e., the jth DMU is a peer of the kth DMU if $\lambda_j > 0$. The target represents the inputs and outputs that the kth DMU should aim for to make itself efficient. For an efficient DMU, the DMU itself is its own target as well as its only peer and in this case $\theta^* = 1$ and all slack variables equal to 0. Note that by the construction of the model, the value of θ must be at least zero and at most one.

2.2.3 The production possibility set – free disposability and returns-to-scale

In this subsection we review the basic assumptions of a production possibility set. From now on, we focus on DEA models in the envelopment form, which are explained intuitively in terms of inputs and outputs.

A basic assumption of a production possibility set is free disposability of the inputs and the outputs. That is, a DMU can dispose extra inputs (while incurring the cost of the inputs) or dispose the produced outputs freely, if the DMU wishes to do so. Hence, when information of the returns-to-scale of inputs and outputs is not available, the production possibility set of a collection of existing DMUs D_0 is defined by the FDH of D_0, i.e., the set $\{(r^{*T}, w^{*T})^T : \forall (r^T, w^T)^T \in D_0, r^* \leqq r, w^* \geqq w\}$. An illustration of the FDH is shown in Figure 2.1. Graphically, each DMU in

the FDH defines a closed orthant, non-negative in inputs and non-positive in outputs, where any DMU located within the orthant, other than the origin of the orthant, is inefficient. Therefore, it is trivial to see that the target of a DMU must be one of the existing DMUs. Hence, instead of searching the entire FDH for a maximal improvement of a DMU, one can simply restrict the search space to the existing DMUs by adding a binary constraint $\lambda \in \{0,1\}^N$ in the envelopment form:

$$\begin{aligned} \min \quad & \theta - \epsilon(e^T s^+ + e^T s^-) \\ \text{s.t.} \quad & W\lambda - s^+ = w^k \\ & R\lambda - \theta r^k + s^- = 0 \\ & \lambda \in \{0,1\}^N \\ & s^+, s^- \geqq 0. \end{aligned} \tag{2.12}$$

If additional information about the returns-to-scale of inputs and outputs is available, then one can derive a set of potentially attainable inputs and outputs D from D_0 (without assuming free disposability). The production possibility set with the additional information is then defined by the FDH of D, i.e., $\{(r^{*T}, w^{*T})^T : \forall (r^T, w^T)^T \in D, r^* \leqq r, w^* \geqq w\}$. In this chapter we consider two assumptions of returns-to-scale: the CRS and the variable returns-to-scale (VRS) (Banker et al., 1984). The assumption of CRS is appropriate when the inputs and outputs change at a constant proportion. Essentially, CRS assumes that inputs and outputs obtained from linear combinations of existing DMUs with non-negative weights are potentially attainable. The assumption of CRS is facilitated by the constraint $\lambda \geqq 0$ in the multiplier form, as shown in (2.11). An illustration of the production possibility set assuming CRS is shown in Figure 2.2, in which the production frontier is a ray that emanates from the origin and passes through points A and B. These two points represent the best ratio of output to input among all the DMUs.

Alternatively, VRS suggests that as the inputs change, the outputs change at a variable rate and vice versa. The change of scale in VRS is described by lines connecting the points corresponding to the existing DMUs, hence the production possibility set is formed by convex combinations of existing DMUs. Hence, intuitively, in the envelopment form, VRS is facilitated by the convexity constraint

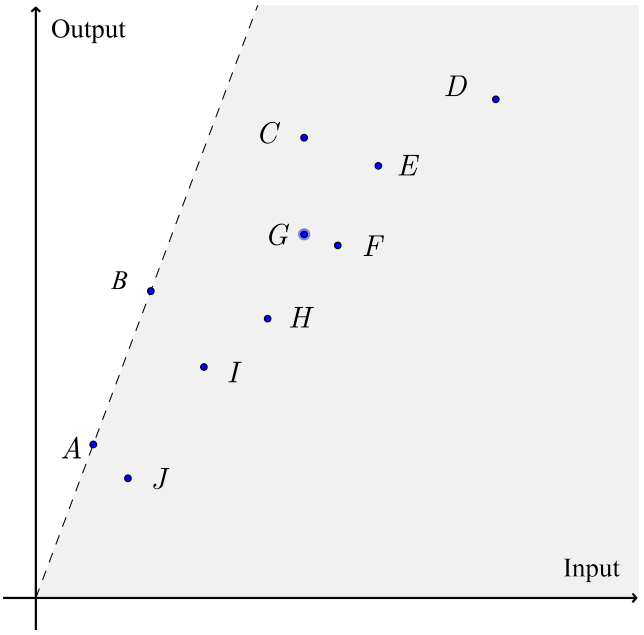

Figure 2.2: An illustration of the production possibility set assuming CRS.

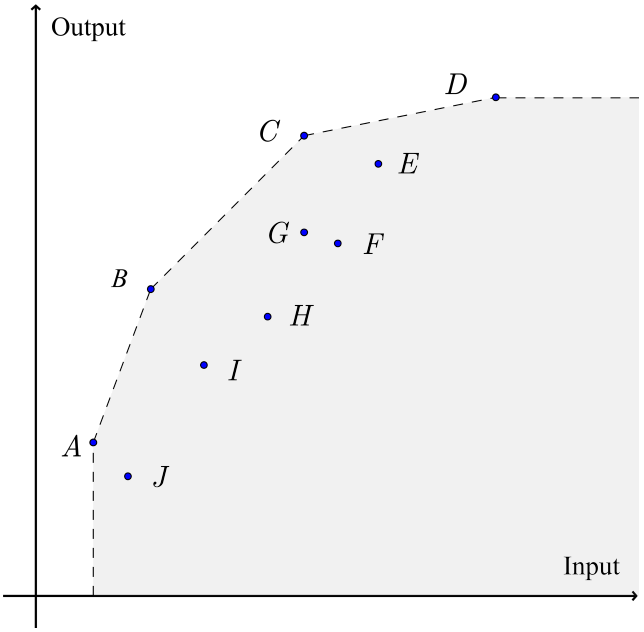

Figure 2.3: An illustration of the production possibility set assuming VRS.

$e^T\lambda = 1$, as shown in (2.13).

$$\begin{aligned} \min \quad & \theta - \epsilon(e^T s^+ + e^T s^-) \\ \text{s.t.} \quad & W\lambda - s^+ = w^k \\ & R\lambda - \theta r^k + s^- = 0 \\ & e^T\lambda = 1 \\ & \lambda, s^+, s^- \geqq 0. \end{aligned} \tag{2.13}$$

An illustration of the production possibility set assuming VRS is shown in Figure 2.3.

2.2.4 An example of an input-oriented DEA model assuming VRS

In this subsection we demonstrate the mechanism of DEA with an input-oriented model assuming VRS in the envelopment form (model (2.13)). The mechanism is illustrated by Figure 2.4, which shows a set of DMUs, each associated with one input and one output, represented by points A to F. The production frontier is shown by the piecewise linear solid line connecting efficient DMUs A, B, and C. These DMUs have a more preferable output to input ratio than the inefficient DMUs D, E and F, for the corresponding input level. The efficiency of each DMU is determined by the amount of radial contraction of inputs required to shift the corresponding point to the production frontier. DMUs corresponding to points A, B and C have an efficiency of 1 since no radial contraction of the input is required to shift the corresponding points to the frontier. On the other hand, DMUs corresponding to points D, E and F are inefficient as radial contractions of the inputs are required to shift these points to points D^*, E^* and F^* on the production frontier. Once a point is on the production frontier, DEA attempts to further improve the point by decreasing the individual inputs and increasing the individual outputs where possible. This is illustrated by point D^* which is shifted to point A by increasing its output value by 2 units (hence $s^+ = 2$). The efficiency scores, input and output values, target values, value of slacks and the

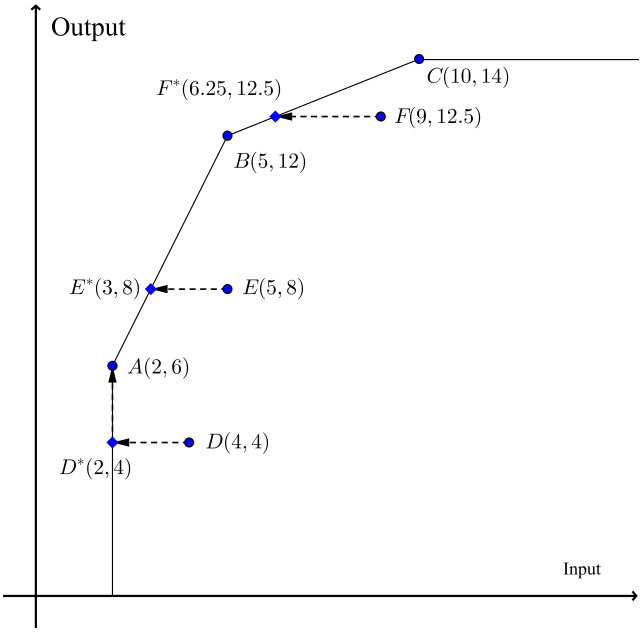

Figure 2.4: A demonstration of the mechanism of an input-oriented DEA model assuming VRS (model (2.13)).

corresponding peers and weights of the DMUs in Figure 2.4 are summarized in Table 2.1.

Table 2.1: The efficiency scores, inputs, outputs, targets, slacks, and the corresponding peers and weights for the six DMUs shown in Figure 2.4.

DMU	θ	(r, w)	Target	(s^-, s^+)	peer(s)	weight(s)
A	1	(2,6)	(2,6)	(0,0)	A	1
B	1	(5,12)	(5,12)	(0,0)	B	1
C	1	(10,14)	(10,14)	(0,0)	C	1
D	0.5	(4,4)	(2,6)	(0,0)	A	1
E	0.6	(5,8)	(3,8)	(0,0)	A,B	0.667,0.333
F	0.695	(9,12.5)	(6,12.5)	(0,2)	B,C	0.750,0.250

Chapter 3

Managing treatment trade-offs in radiotherapy treatment planning

Several criteria need to be considered for a plan to be considered satisfactory: adequate radiation needs to be delivered to the target volume while at the same time radiation delivered to adjacent organs and normal tissues must be tolerable. Due to physical limitations, an ideal dose distribution that simultaneously optimises all the criteria does not exist. Instead, trade-offs between various conflicting criteria need to be made. The planner may rate the criteria differently by considering several factors such as the patient's condition, type of complications expected and/or possible salvage treatment of complications. Ideally, one wishes to find a clinically optimal plan that results in a preferable balance of benefit and risk for the patient.

The multi-criteria nature of the radiotherapy planning problem is conventionally handled by assigning importance scales. Each optimisation criterion, in the form of a mathematical function, is given a weight factor based on the relative importance of that criterion and these criteria are summed up to form an aggregated criterion. The optimisation algorithm then produces a (mathematically) optimised plan according to the composite optimisation criterion. However, it is not clear how a change of the weight factors affects the resulting plan. When an unsatisfactory plan is produced, the treatment planner will have to repeat the optimisation with a different set of weight factors, using a trial-and-error approach, until a satisfactory plan is found. Such a trial-and-error approach is ineffective and time consuming.

In addition, the treatment planner will not know whether the chosen treatment plan is clinically optimal or not because the trade-offs between the treatment goals are unknown and it is unrealistic to explore the trade-offs using a trial-and-error approach.

In this chapter we review planning approaches for managing treatment trade-offs. In Section 3.1, we classify these approaches according to the participation of the planner in the planning process. In Section 3.2, we review the benefits and limitations of a clinically adopted MCO-based planning approach. In Section 3.3, recent advancements in computational efficiency in radiotherapy optimisation are reviewed. The majority of techniques presented in this chapter are in the context of MCO, of which the background materials are provided in Section 2.1.

3.1 Techniques for managing treatment trade-offs

In this section we review techniques for managing treatment trade-offs. We categorise these techniques with MCO classifications as described in Subsection 2.1.3. Note that no planning approach fits the no-preference classification, hence it is not included in this section.

3.1.1 A priori methods

In a priori methods, sufficient information of the decision maker's preferences must be available and be adopted in the solution process. Such information can be specified directly by the decision maker, in the form of prioritised goals or detailed protocols for decision making. Alternatively, one can also extract such information from one or more decisions made previously using methods such as machine learning.

Prioritised goal programming Jee et al. (2007) apply the lexicographic programming method to a prostate case and a head and neck case and demonstrate

that a comprehensive list of planning goals can be optimised in a few priority levels by grouping goals of similar importance into the same level. Wilkens et al. (2007) apply goal programming with slip factor s to slightly relax the constraint of the criteria values obtained from previous steps, i.e., the constraints of (2.8) become

$$f_i(x) \leqq (1+s)\bar{q}_i \quad \text{for} \quad i = 1, \ldots, j-1. \tag{3.1}$$

Application on head and neck cases shows that plans benefit from the use of a slip factor where a small deterioration of higher priority goals allows large improvement of lower priority goals. Similar results are obtained by Clark et al. (2008) where 10 prostate planning problems are considered. Breedveld et al. (2009) use a variant of the lexicographic programming method where the bounds for under-achieving goals are relaxed with a slip factor and over-achieving goals are set to the original goal value. The method is extended to include the beam angle optimisation problem as an automatic plan generation method in Breedveld et al. (2012), in which the beam selection process iteratively finds the most promising beam at the current iteration and adds the beam to the solution, until no significant plan improvement is possible by introducing more beams. This automatic plan generation method is reported superior when compared with plans generated by manual selection of planning parameters (Voet et al., 2013). Goal programming is applied to intensity modulated proton therapy (IMPT) by Falkinger et al. (2012). The study shows that, compared to IMRT cases, the slip factor has little influence in IMPT planning due to the additional freedom in energy modulation and the specific Bragg peak of protons. Kalantzis and Apte (2014) propose a reduced-order prioritised optimisation method. The method consists of three stages. In the first stage, the weighted-sum method is used to generate a representative set of efficient solutions by varying the weight vector. In the second stage, principal component analysis is applied to the representative efficient set to obtain the major eigen modes in terms of bixel intensities. In the third stage, prioritised optimisation is conducted on the reduced search space formed by the major eigen modes. The method is applied to a prostate and a lung case. The results show that plans found in the reduced-order space are comparable to the plans generated without dimension reduction. Moreover, a maximum speedup factor of 49.9 is achieved due to the reduction of the search space.

Optimisation based on a pre-defined planning protocol An alternative way to generate treatment plans is to specify a detailed planning protocol in which the planning parameters and/or the weighting factors are adjusted based on the protocol to achieve planning requirements. Zhang et al. (2011) propose an automated class solution for lung cancers in which the EUD-based planning parameters are found through binary search. Zaghian et al. (2014) propose to satisfy the DV requirements by gradually tightening the hot/cold spot control parameters and increasing the associated weighting factors. Xhaferllari et al. (2013) develop a planning script that iteratively reduces hot or cold spots by automatically incorporating the corresponding volumes as additional criteria with appropriate weights.

Optimisation based on reference plan(s) Plan generation can be conducted based on the information (e.g., the planning parameters or the DVHs) obtained from a set of previously accepted reference plans. This planning approach is most effective if the anatomical geometry of the reference plan(s) closely resembles that of the planning problem. Different approaches have been proposed to quantify geometrical similarity. One way is to consider a so-called overlap volume histogram (OVH) (Wu et al., 2009), which describes the percentage volume of an OAR that is within a specific distance of a PTV. OVH is based on the idea that OARs closer to the target are harder to spare while OARs further away from the target are easier to spare. Wu et al. (2009) use OVHs to identify a group of previous reference plans that are at least as difficult as the planning problem. The minimum achievable DV parameters of the group of reference plans are then used as the DVH objectives for the subsequent plan optimisation. Wu et al. (2011) report that the OVH-based planning approach reduces the average number of optimisation rounds from 27.6 to 1.9. The OVH planing method is extended to VMAT head and neck plans by Wu et al. (2013) in which a database of IMRT plans are used as references. The OVH based planning method is also applied to pancreatic cancer cases (Petit et al., 2012) and prostate cancer cases following hydrogel injection (Yang et al., 2013).

One can also obtain planning information from one of the reference plans that best matches the anatomical geometry of the planning problem. Chanyavanich et al. (2011) find a matching plan based on mutual information of beam's-eye-views, i.e.,

by measuring the mutual dependence of two plans through a statistical comparison of the image histograms of the beam's-eye-views. Planning parameters of the best matching plan are used for the planning problem. Application on prostate cancers shows that the proposed planning approach produces plans with a comparable dose quality as the reference plan. Good et al. (2013) also use mutual information to find the best matching plan. Deformable registration is additionally used to transform the fluence maps of the best matching plan to that of the planning problem. Schreibmann and Fox (2014) rank geometric similarities based on the mean of the closest distances between voxels of a matching plan and the corresponding structure surface of the planning problem. The best matching plan is used as a warm start solution for the subsequent plan optimisation. Application on VMAT prostate cases shows that the planning time reduced from 1-2 hours to 15 minutes due to the reduction of the trial-and-error process of selecting arc directions and lengths as well as the removal of optimisation constraints. Li et al. (2013) propose a DVH-guided planning approach for adaptive radiotherapy in which the voxel weights of the planning problem are iteratively adjusted so that the DVHs of the produced plan match those of the initial plan (which is the reference plan). Zarepisheh et al. (2014) generalise the method of Li et al. to IMRT planning given a set of reference plans. Machine learning is used to build a model for reference plan selection where various geometric features are considered as inputs and DVH similarity as an output.

Machine learning and regression Machine learning and regression can be applied to a set of previously accepted plans to find the correlation between anatomical geometry variations and achievable treatment goals/parameters. Given the correlation, a planner will be able to make an informed evaluation of treatment quality based on the anatomical geometry, hence enables a more consistent planning practice and an improved plan quality.

Moore et al. (2011) build a nonlinear model that predicts the achievable mean dose of the PTV based on the fractional volume of an OAR overlapping the PTV. Application of the model to head and neck cases results in substantial improvement in dose sparing on parotid glands. Zhu et al. (2011) use distance-to-target

histograms (DTH), which have a similar definition as the OVH with minor variations, to account for anatomical geometry variations. In the study, principal component analysis is applied to DVHs and DTHs to identify their salient features. Support vector regression is then applied to establish the correlations between these features. The method successfully predicted the DVHs for 11 out of 14 prostate cancer cases. Appenzoller et al. (2012) develop a model that predicts achievable OAR DVHs based on the minimum distance from an OAR voxel to the PTV surface, in which OAR voxels with the same minimum distances are grouped into subvolumes. Application of the model to 20 prostate cancer cases correctly identified suboptimal plans of which further OAR sparing is achieved after replanning. Yuan et al. (2012) use principal component analysis to identify significant anatomical factors for OAR dose sparing. These factors are used to build a model that predicts OAR DVHs using stepwise multiple regression. The model is applied to the bladder and the rectum for prostate cancer cases and the parotids for head and neck cancer cases. For the bladder and rectum model, the prediction errors of DV parameter values are within 6% for 17 plans and within 10% for 21 plans, out of 24 cases. For the head and neck cases, the prediction error of the median dose values of parotids, are within 6% for 30 plans and within 10% for 40 plans, out of 48 cases.

Other than predicting achievable treatment quality, machine learning methods can also be used to facilitate automatic planning where predicted treatment outcomes are used to set the planning parameters. Yang et al. (2015) use the method of Yuan et al. (2012) to predict OAR DVHs and use the predicted DV parameters as criteria for automatic treatment planning for prostate cancer. The study shows that, compared to conventional planning, the DVH informed planning reduces the average planning time by 5.2 minutes and produces plans with a lower number of MUs while maintaining comparable plan quality. Fogliata et al. (2014) use a combination of principal component analysis and regression to construct a model that predicts the DVHs from various anatomical geometry features, such as OAR volume and overlap volume with targets. Application of the method to lung cancer reduces cases that violate treatment requirements from 11% to 7% for the training set and from 13% to 10% for the validation set. Boutilier et al. (2015) apply machine learning techniques to generate criteria weights for the planning of IMRT prostate treatments. The training dataset consists of 315 treated plans,

each inverse-optimised (i.e., given an intensity pattern, find a corresponding optimal criterion weight vector) using the method of Chan et al. (2014), to obtain a weight vector. Overlap volume ratio and overlap volume histogram slopes of bladder and rectum are used as explanatory variables. The study demonstrates that the generated weight vectors produce plans that closely replicate the clinical plans.

3.1.2 A posteriori methods

In a posteriori methods, a set of potential solutions and the corresponding points in the criterion space are generated and given to the decision maker, who will select a most preferable solution from the set. In radiotherapy treatment design, MCO is primarily applied to a posteriori methods so that only efficient solutions are provided to the decision makers.

Weighted-sum based methods One way to generate a representative non-dominated set is by solving the MOP using the weighted-sum technique repetitively with different weight vectors. Cotrutz et al. (2001) and Lahanas et al. (2003a) use such an approach with entries of the weight vectors normalised and uniformly distributed over zero to one. That is, the weight vectors are chosen from the set

$$\Lambda = \left\{ (\lambda_1, \ldots, \lambda_p) \,|\, \sum_{j=1}^{p} \lambda_j = 1; \lambda_j \in \left\{ \frac{0}{\eta}, \frac{1}{\eta}, \ldots, \frac{\eta - 1}{\eta}, 1 \right\} \right\}, \tag{3.2}$$

where p is the number of criteria and η is the sampling parameter. However, such an approach becomes less practical when the number of criteria is more than 3. In addition, as illustrated in Das and Dennis (1997), this approach is unlikely to produce an even spread of non-dominated points. The examples provided in Das and Dennis suggest that an even spread of non-dominated points often results from a very uneven distribution of weighting factors.

Constraint based methods Alternatively, one can obtain a representative non-dominated set by imposing different constraints. Given a non-dominated point

$f(\hat{x})$, Hamacher and Küfer (2002) propose to find nearby non-dominated points iteratively using relaxed bounds:

$$\begin{aligned} \min \quad & f_i(x) \\ \text{s.t.} \quad & f_i(x) \geq (1-q) f_i(\hat{x}) \\ & f_j(x) \leq f_j(\hat{x}) + q/(p-1) f_i(\hat{x}) \\ & x \in X, \end{aligned} \tag{3.3}$$

where $i \in \{1, \ldots, p\} \setminus \{j\}$. The parameter q is chosen to be between 0.2 and 0.8 and may vary for different criteria. The algorithm starts by finding an efficient solution using a lexicographic programming method and successively expands to other efficient solutions.

Küfer et al. (2003) use the ϵ-constraint method to find non-dominated points where criteria are partitioned into an active set $\mathcal{M}$ with $\mathcal{M} \neq \emptyset$ and a non-active set $\mathcal{N}$. For every possible partition of the criteria, the following problem is solved to obtain a non-dominated point.

$$\begin{aligned} \min \quad & s \\ \text{s.t.} \quad & f_i(x) \leq s && \forall i \in \mathcal{M} \\ & f_j(x) \leq s_j && \forall j \in \mathcal{N} \\ & x \in X, \end{aligned} \tag{3.4}$$

where s_j is an appropriate upper bound chosen by the decision maker. A subsequent optimisation based on the hybrid model (2.6) is then solved to make sure the solution is efficient. The $2^p - 1$ efficient solutions found through each partition define the scope of the approximated non-dominated set. Triangulation and refinement steps are used to add new points between existing non-dominated points to achieve a desired granularity.

The normalised normal constraint method (Messac et al., 2003) is used by Craft et al. (2005) to explore the two-dimensional trade-offs (hence in the study $p = 2$)

between tumour dose homogeneity and OAR sparing. The method first calculates anchor points $y^i, \forall i \in \{1, \ldots, p\}$, which are points found by minimising each criterion individually. The algorithm then derives a set of evenly distributed reference points q (referred to as utopia points) that lie on the line segments connecting the anchor points (referred to as utopia lines). Note that for $p > 2$, equidistant reference points are placed on the utopia plane, which is the convex hull of utopia points. The method finds individual non-dominated points by solving problems of the form:

$$\begin{aligned} \min \quad & f_i(x) \\ \text{s.t.} \quad & N_j^T(f(x) - q) \leq 0 \quad \forall j \in \{1, \ldots, p\} \setminus \{i\} \\ & x \in X, \end{aligned} \tag{3.5}$$

where $N_j = y^i - y^j$. The inequality $N_j^T(f(x) - q) \leq 0$ reduces the feasible criterion space such that a unique non-dominated point can be found for each reference point q by minimising only one criterion.

Sandwiching methods Sandwiching methods create an approximation of the non-dominated set by iteratively enclosing the search area between a lower and an outer bounding set. For multi-objective convex optimisation problems, a common approach to construct a lower and an upper bounding set is by solving scalar subproblems using the weighted-sum method. Let $\bar{x}$ be an optimal solution to a weighted-sum problem with a weight vector $\bar{\lambda}$. A supporting hyperplane of the feasible set in criterion space can be defined by $f(\bar{x})$ and $\bar{\lambda}$ where $\bar{\lambda}$ is the normal of the hyperplane. The supporting hyperplane then forms part of the lower bounding set, along with other supporting hyperplanes found by solving other weighted-sum subproblems each with a different λ. Let $\bar{X}_E$ be a set of efficient solutions found at a given iteration of the sandwiching process and let $\bar{Y}_N$ be the corresponding set of non-dominated points. By convexity, the non-dominated set of the convex hull of $\bar{Y}_N$ defines an upper bounding set, which is also an approximation of Y_N. As more non-dominated points are added to the approximation, the space between the lower bounding set and the upper bounding set shrinks. Given an approximation of Y_N, the decision maker still needs to find the most preferred point and its

corresponding solution. This is achieved by interactively solving a multi-objective optimisation problem in which the feasible set is restricted to the convex hull of $\bar{X}_E$ (see Subsection 3.1.3).

The computational goal of sandwiching methods is to generate an approximation of the non-dominated set of adequate accuracy with as few non-dominated points as possible. Hence, preferably, one wishes to find non-dominated points that maximally reduce the distance between a lower bounding set and an upper bounding set. To achieve this computational goal, Craft et al. (2006) propose a heuristic approach in which the next weight vector is a vector maximally different from the existing weight vectors, chosen from the convex hull of existing weight vectors. Hoffmann et al. (2006) use an ϵ-constraint based sandwiching method where non-dominated points are found by solving scalar subproblems using the ϵ-constraint method while the normal λ of the supporting hyperplane is obtained from the Lagrange multipliers associated with the constraints. The study by Hoffmann et al. investigates different strategies to add new points to the non-dominated set approximation. The results show that strategies that minimise the Hausdorff distance or the total uncertainty area between the upper bounding set and the lower bounding set produce better approximations of the non-dominated set than the strategy that minimises the maximum error. Thieke et al. (2007) propose to successively optimise all combinations of criteria, i.e., first each criterion individually then all combination of 2 criteria and so forth. Bokrantz and Forsgren (2013) propose a vertex enumerative algorithm to find the weight vector of a non-dominated point that maximally reduces the distance between the lower bounding set and the upper bounding set.

Shao and Ehrgott (2008) propose a sandwiching method for MOLPs based on Benson's outer approximation algorithm (Benson, 1998a,b). In Benson's outer approximation algorithm, the lower bounding set is the non-dominated set of a simplex S that encloses the feasible set Y in criterion space. At each iteration, a vertex v of S, with $v \notin Y$, is chosen and the non-dominated point that lies on the line segment between v and the anti-ideal point $y^{AI}, y^{AI}_k := \max\{y_k : y \in Y\}$ is calculated. The non-dominated point is then used to solve a pair of primal and dual linear problems, the solutions of which are used to construct a supporting

hyperplane of Y at the point. S is then reconstructed to incorporate the supporting hyperplane and the vertices are updated. The upper bounding set is formed by the non-dominated set of the convex hull of existing non-dominated points. In Benson's outer approximation algorithm, the procedure repeats with another vertex $v, v \notin Y$, and the procedure ends once all vertices are non-dominated, in which case the non-dominated set of S is the same as Y_N. Shao and Ehrgott (2008) report that the non-dominated surface of some IMRT problems is "curved". "Curved" surfaces lead to the construction of an excessive number of supporting hyperplanes, which inevitably result in a prolonged computational time. Therefore, Shao and Ehrgott propose an approximation version of Benson's algorithm. In the proposed algorithm, a tolerance value ϵ is predefined. If the distance from a vertex to the corresponding non-dominated point is within the tolerance value ϵ, the construction of the supporting hyperplane is omitted. Shao and Ehrgott prove that points of the upper bounding set produced from the algorithm are ε-nondominated where $\varepsilon = \epsilon e$. The authors find this approach significantly reduces the computational time while maintaining the approximation within a tolerable error.

Given a representation or an approximation of the non-dominated set, the decision maker still needs to search for the most preferable point among the set. One way to do so is by the interactive methods, which will be discussed in Subsection 3.1.3.

Evolutionary algorithms Lahanas et al. (2003b) apply the evolutionary algorithm NSGA-II (Deb et al., 2002) and its controlled elitist version NSGA-IIc (Deb and Goel, 2001) to the IMRT planning problem. The solutions generated from NSGA-II and NSGA-IIc are compared to the solutions generated by the gradient based algorithm L-BFGS (Lahanas et al., 2003a), which uses a weighted-sum tactic to find a representative non-dominated set. The results show that the solutions generated from NSGA-IIc cover more of the non-dominated set than the solutions from NSGA-II, which converges prematurely. For both NSGA-ll and NSGA-IIc, some regions of the non-dominated set can only be explored if a number of solutions generated from L-BFGS are used as initial solutions. Holdsworth et al. (2010) propose a hierarchical evolutionary algorithm in which the top level of the algorithm is a multi-criteria evolutionary algorithm and the bottom level is a deterministic algorithm used to solve constrained quadratic problems (Breedveld et al.,

2006). The deterministic algorithm uses a weighted-sum approach to handle the multiple criteria and the top level evolutionary algorithm uses the weight vector as the genes of the individuals. Fiege et al. (2011) apply a genetic algorithm to a combined problem of beam selection and fluence map optimisation. The method is applied to two phantom cases. The results show that conformity-based PTV fitness functions and DVH or EUD based OAR fitness functions produce relatively uniform and conformal PTV doses.

3.1.3 Interactive methods

In interactive methods, the decision maker iteratively adjusts preferences based on existing optimisation results. The conventional planning approach in which a planner iteratively adjusts planning parameters to find a treatment plan can be considered as an interactive multi-criteria method. However, such a planning practice requires a planner to experiment with the planning parameters which can be ineffective. Alternatively, one might want to use an interactive MO method to effectively guide a decision maker to a preferable plan. There are limited studies on the use of interactive MO methods for plan generation. One example is by Ruotsalainen (2009) who uses a classification-based interactive method to solve the planning problem. The method classifies criteria into those that should be improved (optionally with aspiration values), those that are considered satisfactory and those that are allowed to deteriorate (optionally with given bounds).

The lack of interactive methods in radiotherapy treatment design might be due to the long computational time for executing one optimisation round. Using an interactive method, the planning process can be tedious since it can take many optimisation rounds before a satisfactory plan is found. To address this issue, one can conduct interactive methods on the convex hull of a set of efficient solutions generated from a sandwiching method (see Subsection 3.1.2). Let $\left\{x^1, \ldots, x^k\right\}$ be a set of efficient solutions generated by a sandwiching method. Let $\bar{X}$ be a matrix formed by the solution vectors in which $\bar{X}_{\cdot,i} = x^i$. By convexity, solutions formed by convex combinations of the k efficient solutions are also feasible. To obtain the efficient solutions of the set, one solves the following MOP in which the feasible set is determined by the convex hull of the k efficient solutions,

$$\begin{aligned} \min \quad & f(\bar{X}\lambda) \\ \text{s.t.} \quad & e^T\lambda = 1 \\ & \lambda \in \mathbb{R}^k_{\geqq}. \end{aligned} \tag{3.6}$$

As demonstrated in Monz et al. (2008), the optimisation of (3.6) can be solved in real time since the search space is substantially reduced. Furthermore, by convexity, convex interpolation of feasible points in criterion space acts as an upper bound for the criteria values of interpolated feasible solutions because

$$f_j\left(\sum_{i=1}^{k}\lambda_i x^i\right) \leqq \sum_{i=1}^{k}\lambda_i f_j(x^i). \tag{3.7}$$

Hence, the approximated non-dominated set obtained from the sandwiching method is also an upper bounding set for the non-dominated set of (3.6). In other words, if the approximated non-dominated set generated from a sandwiching method is sufficiently accurate, then the point obtained from solving (3.6) will be sufficiently close to be non-dominated.

The interactive process of finding a preferred point from a set of points in the objective space is referred to as *navigation* (Allmendinger et al., 2016). In this paragraph, we review navigation methods proposed in the field of radiotherapy treatment planning in which the set of feasible points are the mapping of the convex hull of a set of efficient solutions $\left\{x^1, \ldots, x^k\right\}$ generated from a sandwiching method. A navigation method should allow a decision maker to specify a "navigation query" that moves a current point to another (non-dominated) point with a desirable improvement. This is achieved by solving a scalarised model of (3.6) incorporating the current point and the navigation query. In Monz et al. (2008), navigation is conducted by solving the following scalarised model of (3.6)

$$\begin{aligned} \min \quad & \max_{i=\{1,\ldots,p\}\setminus\{j\}} && f_i(\bar{X}\lambda) - q_i \\ \text{s.t.} \quad & && f_j(\bar{X}\lambda) = \tau \\ & && f(\bar{X}\lambda) \leqq b \\ & && e^T\lambda = 1 \\ & && \lambda \in \mathbb{R}^k_{\leqq}, \end{aligned} \tag{3.8}$$

where q is the current point, the navigation query is to set the jth objective function value to τ and b represents the minimal requirements for the objective function values set by the planner. To ensure the feasibility of the optimisation problem, the maximal and minimal values for each criterion subject to the constraint $f(\bar{X}\lambda) \leqq b$ are calculated whenever b is modified and search queries τ outside the feasible range are not allowed. In the study, the feasible range of criteria values defined by the individual maxima and minima is referred to as the "planning horizon". Craft and Monz (2010) use a variant of (3.8) to conduct navigation

$$\begin{aligned} \min \quad & t \\ \text{s.t.} \quad & t \geqq f_i(\bar{X}\lambda) - q_i \qquad \text{for} \qquad i = 1, \ldots, p \\ & e^T\lambda = 1 \\ & \lambda \in \mathbb{R}^k_{\leqq}, \end{aligned} \tag{3.9}$$

in which q is a reference point on the convex hull of the individual minima and navigation queries are specified by changing the value of q. The navigation model is applied to a problem with multiple representative non-dominated sets where each set corresponds to a certain beam orientation. Craft and Richter (2013) propose a "2D-cut" navigation which features the 2D trade-off curve of two chosen criteria while other criteria are either relaxed or bounded by certain values. The 2D trade-offs are obtained using the normalised normal constraint method (3.5).

A number of studies apply interactive methods as a fine-tuning tool primarily for plans close to satisfactory. Cotrutz and Xing (2003) and Lougovski et al.

(2010) fine-tune the dose distribution by adjusting the penalty associated with each unsatisfactory voxel. Süss et al. (2013) propose to eliminate cold or hot spots by imposing dose upper/lower bound constraints on the relevant voxels. Otto (2014), Ziegenhein et al. (2014) and Fuss and Salter (2007) propose to fine-tune the dose distribution by adjusting the bixel intensities associated with the voxels that require dose tuning. Real-time interactivity is achieved by Otto (2014) and Ziegenhein et al. (2014) and hence their methods can be used to construct a plan from scratch. Otto (2014) achieves real-time interactivity using an approximated dose calculation, in which the 2D convolution of fluence and dose deposition kernel is performed in the Fourier domain. Ziegenhein et al. (2014) achieves real-time interactivity by exploiting CPU arithmetic capability in which the dose distribution is calculated on-the-fly to avoid the time consuming task of transferring data from the main memory to the arithmetic unit.

3.2 Clinical adoption of a MCO-based planning approach

In recent years, a planning approach based on MCO has been clinically deployed (Craft and Richter, 2013). The method first generates an approximation of the non-dominated set using a sandwiching algorithm (Craft et al., 2006), followed by plan navigation among a set of points mapped from the convex hull of a set of efficient solutions obtained from the sandwiching method (Monz et al., 2008). For convenience, this planning approach will be referred to as MCO-based planning.

Benefits of MCO-based planning have been demonstrated by several studies. Based on a paraspinal case and a prostate case, Thieke et al. (2007) demonstrate that the MCO-based planning approach produces plans of comparable quality to the clinically approved reference plans, while the required planning time for navigation is only on the order of 10 minutes. Hong et al. (2008) apply MCO-based planning to ten locally advanced pancreatic cancer cases and show that navigation can be done within 10 minutes. The study also demonstrates how MCO facilitates active clinical inputs from the planners when planning trade-offs are readily available: 9 out of 10 plans selected from the navigation step have a lower stomach mean dose

and higher kidney dose compared to the previously-treated reference plans. Based on 5 glioblastoma and 5 pancreatic cancer cases, Craft et al. (2012b) show that MCO-based planning reduces planning time from 135 to 12 minutes on average while the active planning time from the planner increases from 4.8 to 8.6 minutes, on average. The plans selected from MCO-based planning are also considered superior to those produced from the conventional planning approach. Wala et al. (2013) find that, based on 9 prostate cancer cases, plans generated from MCO-based planning are considered superior to the clinically approved reference plans generated with the conventional approach, with approximately 10 minutes of navigation time for each case. With 20 head and neck cancer cases, Kierkels et al. (2015) demonstrate that MCO-based planning allows less experienced planners to create plans comparable to those created by experience planners using the conventional approach, with average planning time reduced from 205 minutes to 43 minutes.

The MCO-based planning approach is also shown to benefit other treatment delivery techniques. Craft et al. (2012a) and Bokrantz (2012) apply MCO-based planning to VMAT planning in which the initial step involves fluence map navigation based on static beams, followed by a leaf sequencing step that (approximately) reproduces the selected fluence map. Khan and Craft (2015) apply MCO-based planning with limited segmentation to 10 previously accepted 3D-CRT cases of various disease sites. All plans re-created with MCO are preferred over the original plan, with an average reduction of 17% in MUs. MCO planning can also be used to identify a reference IMRT-plan for VMAT planning, in which the DVHs of the reference IMRT plan are used as criteria and constraints for VMAT planning. Chen et al. (2014) apply this MCO informed VMAT planning approach to prostate and head and neck cancer cases. The study reports that MCO informed VMAT plans are of comparable quality to the reference IMRT plans but with a shorter delivery time and fewer MUs. Chen et al. (2015) apply the MCO informed VMAT planning approach to spinal radiosurgery cases and observe similar plan quality and MUs for both treatment techniques, with a 25% reduced delivery time for VMAT.

Despite the benefits reported in the literature, the MCO-based planning approach is not without limitations. The first limitation is due to plan generation using

convex interpolation of a set of efficient solutions. A plan generated through interpolation is not necessarily efficient and can potentially be improved further. The improvement potential of a plan generated through interpolation depends on how well the non-dominated set is approximated by the finite non-dominated points generated by the sandwiching method. If given a poorly approximated non-dominated set, certain trade-offs in criteria may be neglected or poorly represented, resulting in plans with further improvement potential. To eliminate the approximation error, Bokrantz and Miettinen (2015) propose to project a point, which corresponds to an inefficient plan selected from the navigation process, to the non-dominated set as a post-processing step for MCO-based planning. Application of the method to a prostate case and a head and neck case demonstrates improved dose conformity and better OAR dose sparing.

The second limitation of MCO-based planning is plan degradation from the segmentation process (McGarry et al., 2014). Note that this is also an limitation for the conventional single-objective planning practice in which the FMO problem and the segmentation problem are considered as two separate problems (see, e.g., Rocha et al. (2012)). If a treatment plan becomes unsatisfactory after segmentation, the planner will have to re-optimise and find another plan, which makes the planning process inefficient. To address this issue, Craft and Richter (2013), Salari and Unkelbach (2013) and Fredriksson and Bokrantz (2013) have proposed navigation methods for deliverable plans. These approaches use convex combinations of a set of segmented plans or conical combinations of a set of segments to approximate the non-dominated set. Furthermore, deliverable navigation for VMAT plans with sliding-window MLC approach is proposed by Craft et al. (2014) by averaging the times at which the MLC leafs leave certain control points along the MLC trajectory. However, as stated in the first limitation, these methods also use the interpolation technique to generate treatment plans, hence neglect plan improvement potential.

3.3 Recent advances in computational efficiency in radiotherapy optimisation

Approaches to speed up the optimisation process have been proposed by using specialised algorithms (Chen et al., 2010) and distributed computing (Bokrantz, 2013). Chen et al. (2010) propose a projection based algorithm for multi-criteria IMRT planning. Compared to general purpose algorithms implemented in MOSEK, their algorithm achieves a speed up factor of 2 to 7. However, the algorithm can only be applied to certain types of objective functions with linear constraints. Bokrantz (2013) proposes a scheme for distributed computing for weighted-sum based sandwiching methods. The scheme iteratively computes a set of weight vectors that are likely to minimise the approximation errors based on an artificial approximation of the non-dominated set, followed by solving the corresponding individual weighted-sum problems in parallel. The parallel scheme produces an approximation of the non-dominated set that is comparable in quality to one generated sequentially while the degree of parallelisation is about twice the number of criteria (i.e., a speed up factor of 10 for a typical five criteria problem).

Other studies propose to exploit the parallel arithmetic capabilities of graphical processing units (GPUs) and central processing units (CPUs) to fasten the optimisation. Men et al. (2009) implement a gradient-based FMO algorithm with GPU parallel computing. Based on three test cases, the implementation is capable of generating an optimal fluence map in 0.2 to 2.8 seconds, which is 20-40 times faster than a corresponding implementation using serial CPU computing. The study suggests that the speed up factor depends on the number of multiprocessors available within the GPU. Men et al. (2010a) implement a column generation based planning algorithm (see Men et al. (2007) and Chapter 5 for details) with parallel GPU computing and report a computational time of 0.7 to 3.8 seconds. Performance evaluation of the algorithm shows that the column generation subproblems can be solved efficiently in parallel by a GPU. Men et al. (2010b) apply the GPU-based column generation method to VMAT planning and report an optimisation time of 18-31 seconds, in contrast to a run time of 5 to 8 minutes on a benchmark implementation in serial CPU computing. The relatively small memory size of GPUs cannot be used to store a large dose deposition matrix. To address this

issue, Tian et al. (2015) propose to split the matrix into sub-matrices and use the memory spaces of multiple GPUs to store the sub-matrices. The multi-GPU implementation is capable of optimising VMAT head-and-neck cases within 1 minute and VMAT prostate cases within 24-46 seconds. Note that GPU-based computing has also been applied to imaging-related (e.g., deformable image registration) and therapy-related (e.g., dose calculation) areas of radiotherapy planning, as reviewed by Jia et al. (2014).

Ziegenhein et al. (2013) suggest the parallel architecture of modern multi-core CPUs is not leveraged and implement a performance-optimised approach for solving FMO using a low-cost CPU commonly used for desktop computers. The implementation decomposes the optimisation algorithm (iterative quasi-Newton method) so that the computing associated with each voxel can be executed independently in parallel. The resulting implementation conducts plan optimisation in a few seconds at clinical resolution and quality. The study shows that optimisation time is largely dependent on the transfer rate from the main memory to the arithmetic unit. Ziegenhein et al. (2014) propose an interactive dose shaping method to avoid data transfer of the dose deposition matrix by calculating the resulting dose distribution on-the-fly. The resulting implementation is capable of modifying a dose feature, which involves one modification step and 10 to 30 heuristic recovery steps, within 1 to 2 seconds.

3.4 Summary and remarks

In this chapter we reviewed the literature for managing radiotherapy treatment trade-offs. While a priori methods allow one to (semi-)automate the planning process, the resulting plan is subject to a pre-determined one-for-all planning practice, with limited consideration of different treatment trade-offs. In a posteriori methods, the trade-off information is provided by a representative or an approximated non-dominated set, which can be computationally expensive to generate. The interactive methods require an active planner to iteratively tailor a treatment plan, thus the solution process can be cumbersome if the computational time for one round of optimisation is long.

In clinical practice, computational time for one round of optimisation is usually on the order of minutes (Ziegenhein et al., 2013). This rather long computational time discourages a posteriori and interactive methods as generating a set of solutions and the iterative solution process can be burdensome. To facilitate a posteriori and iterative methods, commercial MCO-based planning restricts the feasible set to the convex hull of some efficient solutions generated by a sandwiching method. By doing so, one can solve the optimisation problems in real-time, thus allowing an efficient interactive navigation process. However, MCO-based planning may generate plans with further improvement potential due to plan degradation after segmentation and the error associated with the approximation of the non-dominated set.

Recent advances in parallel computing have substantially reduced the optimisation time for radiotherapy plan generation. This encourages a posteriori methods that produce finely sampled non-dominated points over the non-dominated set. By doing so, one overcomes the limitations of commercial MCO-based planning, i.e., each point is truly non-dominated and the segmentation process can be applied independently to each plan before being reviewed by the planner. In this thesis, we propose such a method in Chapter 5 and demonstrate its application in radiotherapy treatment planning in Chapter 6. Given a discrete set of solutions, it can be time consuming to examine every solution. In Chapter 7 we propose a navigation method that allows a decision maker to effectively explore a set of points in objective space. In contrast to the navigation methods proposed in radiotherapy planning, the proposed method allows any user-defined navigation criteria, instead of optimisation criteria, hence enhances the relevance of the decision making process to the underlying problem.

Chapter 4

Quality assessment for radiotherapy treatment planning based on data envelopment analysis

Managing treatment trade-offs in conventional radiotherapy treatment planning has been a trial-and-error process in which the planner iteratively adjusts planning parameters (i.e., the objective weights and/or the objectives) in order to find a satisfactory plan. However, because the exact effects of changing the plan parameters cannot be known a priori, it is difficult for the planner to verify if there is further potential to improve a plan. If a plan is deemed to be of inadequate quality it would take further time to produce another plan, without knowing in advance whether the new plan is going to be superior to the previous plan. This trial-and-error aspect of the planning process is time consuming and in order to reduce the planning time, plans tend to be accepted as long as the planning requirements are satisfied, without considering further improvement potential.

This planning dilemma can be addressed by comparing the plan quality against other plans. This comparison will allow the planners to score the plan against an "optimal" plan and therefore allow informed decisions on whether further improvement is possible. The "optimal" plan may not be truly the best one achievable

but rather an indication of what can be achieved. A number of assessment approaches that use past plans as references have been proposed in the literature (see the review in Section 3.1). These quality assessment approaches predict the achievable OAR sparing based on the geometrical relationship of the PTV and the OAR. Thus, given a particular geometrical relationship, if a newly generated plan has a significantly higher dose to the OAR than the predicted dose, the plan is considered of inadequate quality and re-optimisation is required. However, as these approaches do not consider the dose to the PTV, one may unintentionally conclude that more OAR sparing is available without realising that the improvements in OAR sparing would likely deteriorate the PTV dose coverage. This might inadvertently lead to a longer planning process since a high quality plan with good PTV coverage and acceptable OAR sparing may be considered of inadequate quality simply because it does not achieve the maximal OAR sparing.

In this chapter we propose a plan quality assessment method that can take both the dose to the PTV and the OARs into consideration. By doing so, when a plan is generated, a planner can assess the plan quality using its current dose distribution to both the PTV and the OARs rather than assess it against a plan that solely encourages maximal OAR sparing. The method is based on DEA (see Section 2.2). The concept of DEA is directly applicable to the problem of assessing treatment plan quality in radiotherapy in which the doses to OARs are considered as the cost we pay for delivering dose to the PTV. Hence, DEA performs peer evaluation by comparing treatment plans with respect to an ideal defined by historical plans. In healthcare, DEA has been applied in performance assessment of healthcare systems (Chilingerian and Sherman, 2011), including formative evaluation of radiotherapy services (Santos and Amado, 2012) and even to compare prostate cancer treatment options (Ramer et al., 2008). However, to the best of our knowledge, it has never been used for case-based quality assessment for radiotherapy treatment planning. One of the strengths of DEA is its ability to handle multiple inputs and outputs. This strength makes DEA ideal for radiotherapy plan assessment in which several conflicting planning criteria need to be considered.

The DEA based quality control method integrates well with the plan generation methods reviewed in Chapter 3. For planning using a posteriori methods, as the

non-dominated set is determined by the planning parameters, if the planning parameters are poorly chosen, the most preferable plan found from the resulting non-dominated set may be of poor quality. In this case, DEA can be used to reassure the quality of the plan by comparing the plan to past plans. Furthermore, the representative set of plans can consist of a large number of plans and it can be difficult for the planner to explore and compare all the plans thoroughly. A well established DEA model will assist the planner to quickly identify plans of best quality from the representative set and thus enable an even more efficient planning process (as will be demonstrated in Chapter 7). For planning using a priori approaches, since planners are provided with only a single plan that matches the treatment goals/protocols as closely as possible, it is important to have additional measures for quality control, which can be facilitated by DEA. For interactive methods, DEA reduce the need to explore other treatment trade-offs by providing an "optimal" plan for reference.

In this study we apply DEA as a quality assessment tool to a set of prostate VMAT plans and investigate its capability of identifying plans with further improvement potential. In Section 4.1, we present the DEA model used in this study. In Section 4.2 we present a case study in which DEA is applied to assess the quality of prostate radiotherapy plans. The results and discussion are presented in Sections 4.3 and 4.4 respectively.

4.1 A DEA model for quality assessment of radiotherapy treatment plans

In this study we use an input oriented DEA model for the analysis since we are interested in maximal OAR sparing available for a given dose to the PTV. Mathematically, we may assume CRS for the model since we are able to obtain a constant return of dose to the PTV for each unit change of dose to the OARs by scaling the entire dose distribution. However, in practice, treatment planning usually follows strict requirements on the dose of the PTV, thus scaling the entire dose distribution may produce unsatisfactory plans. Thus, in this study we use a DEA model assuming VRS.

When performing DEA, one may want to take environmental factors into account. Environmental factors are factors that may influence the performance of DMUs but are external and out of the control of the DMUs. In our study, the percentage overlap of PTV and rectum is such an environmental variable. This value is not influenced by the planning variables but can adversely affect the attainable quality of the treatment plan. Environmental factors can be incorporated in the DEA formulation as environmental variables. A number of ways to handle environmental variables are introduced in Chapter 7 of Coelli et al. (2005). In this study, we assume that higher values of the environmental variables are likely to impair the efficiency of the DMUs and we incorporate the environmental variables into the DEA model as a so-called non-discretionary output variable (Banker and Morey, 1986). We assume there are L environmental variables represented by a vector $z^i \in \mathbb{R}^L$ for the ith DMU and by a matrix $Z \in \mathbb{R}^{L\times I}$ for the whole dataset. Environmental variables are incorporated into the DEA formulation by adding the constraint $Z\lambda - s^z = z^k$ into formulation (2.13) where $s^z \in \mathbb{R}^L$ is a vector of slack variables. The constraint restricts the feasible set and ensures that DEA searches for the evidence of further improvements among potentially attainable inputs and outputs that are associated with higher or equal values of environmental variables than the kth DMU. Since high values of the environmental variables are considered unfavourable to the efficiency, if a DMU can achieve certain efficiency with higher or equal values of environmental variables than the kth DMU, the kth DMU should be able to achieve at least the same efficiency, otherwise it is considered inefficient. For completeness, the DEA formulation used in this study, which is an input oriented VRS model with a non-discretionary environmental output variable, is shown in formulation (4.1).

$$
\begin{aligned}
\min \quad & \theta - \epsilon(e^T s^+ + e^T s^- + e^T s^z) \\
\text{s.t.} \quad & W\lambda - s^+ = w^k \\
& R\lambda - \theta r^k + s^- = 0 \\
& Z\lambda - s^z = z^k \\
& e^T\lambda = 1 \\
& \lambda, s^+, s^-, s^z \geqq 0.
\end{aligned}
\tag{4.1}
$$

4.2 Application of DEA to prostate radiotherapy treatment planning

One of the most difficult tasks in prostate radiotherapy treatment planning is managing the dose delivered to the PTV and the rectum. The rectum is usually the OAR that most influences the ability to achieve the desired dose to the PTV. In this preliminary study, we only consider the dose delivered to the PTV and the rectum. The goal is to generate a dose distribution that matches the prescription dose in the PTV as closely as possible while maximising rectal sparing. In a loose economic interpretation, the dose delivered to the rectum is considered as the cost for delivering the dose to the PTV. Specifically, we use D_{95} of the PTV as the output and gEUD of the rectum as the input for the DEA model. In addition, we use the percentage volume of the rectum that overlaps the PTV as an environmental variable. The higher the overlap the more difficult it is to achieve good PTV coverage and low OAR dose simultaneously. We handle the environmental variable using the method described in Section 4.2. While many other dose descriptors or anatomical descriptors can be alternatively used for the DEA model, it is out of the scope of this study to investigate the most preferable inputs and outputs. Instead, we empirically select these input and outputs and focus on investigating the validity of using DEA as a quality assessment method in radiotherapy planning.

A series of 37 anonymised clinically intact prostate treatment plans were provided by Auckland Radiation Oncology, following approval and guidelines of New Zealand Health & Disabilities Ethics Committees for observational study. All plans were the actual plans used for the subsequent delivery of treatment and were generated using the same treatment planning system over a 1 year period utilising the same plan acceptability criteria. In all cases, 74 Gy was prescribed to the PTV with requirements that 99% of the PTV received 95% of the prescribed dose and that 99% of the actual prostate receive 99% of the prescribed dose. For rectal criteria, the percentage volume of the rectum that received equal to or more than 40Gy, 60Gy and 70Gy of radiation dose should not exceed 60%, 40% and 10%, respectively, i.e., $V_{40} < 60\%$, $V_{60} < 40\%$ and $V_{70} < 10\%$. Not all rectal criteria were met in all cases, but all plans were nevertheless considered clinically acceptable.

All plans were designed for VMAT delivery with Pinnacle v9 and the SmartArc module (Philips, Netherlands) using a single 360 degree arc. The plan input/output values were extracted using CERR (Deasy et al., 2003). The input/output values of the plans are shown in Table 4.1.

We used an in-house DEA software package, pyDEA, to assess the efficiency of these 37 prostate plans. After obtaining the results from the analysis, five inefficient plans were selected for re-optimisation. We use the term re-optimisation throughout this thesis to describe the process of modifying the treatment planning objectives and re-running the inverse plan optimiser. The input/output values of the re-optimised plans are included in Table 4.1, indicated by the original plan ID with an asterisk. The plans for re-optimisation are selected manually based on two criteria. Firstly, each of the selected plans should have a percentage overlap volume substantially different from the other selected plans. Secondly, each plan should be substantially inefficient for their range of percentage overlap volume. Note that the plans selected for re-optimisation all have an efficiency score less than the average efficiency score (0.985). These plans are selected in order to test the ability of DEA in assessing plans with variations in anatomical structure relationships. A planner was instructed to further improve rectal sparing while maintaining overall clinical acceptance for the selected plans without access to the results of DEA. The original plans were optimised for rectal sparing via the use of a single rectal objective based on gEUD with the volume parameter "a" equal to 1. Plan re-optimisation involved changing the a-value from 1 to 6 thereby increasing the penalty weight of the higher dose components of the rectal DVH. All other aspects of the plan remained the same. After re-optimisation, the plans were included in the dataset and the DEA analysis was repeated.

Table 4.1: The efficiency scores and the input/output values of the plans. Plans 10, 19, 26, 31 and 35 were re-optimised. Those plan IDs with an asterisk indicate re-optimised plans.

Plan ID	Efficiency (original dataset) [1]	Efficiency (re-optimised dataset) [2]	Output: D_{95} PTV (Gy)	Input: Rectal gEUD (Gy)	Fractional overlap
1	0.988	0.975	71.625	61.418	0.048
2	0.982	0.98	71.025	62.478	0.1
3	0.966	0.96	71.275	62.805	0.066
4	0.985	0.977	71.225	61.369	0.053
5	0.994	0.981	71.675	60.744	0.036
6	0.974	0.968	71.825	63.714	0.077
7	0.982	0.975	71.675	62.779	0.074
8	0.978	0.968	71.775	61.745	0.031
9	0.982	0.972	71.775	62.56	0.061
10	0.964	0.955	71.575	63.457	0.065
10*	N/A	0.987	71.525	61.314	0.065
11	0.982	0.97	71.575	61.73	0.049
12	0.979	0.969	71.525	61.61	0.042
13	0.985	0.975	71.675	62.16	0.062
14	0.998	0.987	71.725	60.362	0.03
15	0.98	0.966	71.625	61.934	0.047
16	0.984	0.976	71.475	62.192	0.072
17	0.995	0.986	71.475	61.403	0.069
18	1	0.997	71.825	60.173	0.022
19	0.975	0.963	71.725	62.562	0.052
19*	N/A	0.987	71.325	60.71	0.052
20	0.994	0.985	71.375	60.502	0.037
21	1	1	71.975	63.648	0.111
22	0.978	0.977	72.025	63.87	0.076
23	0.985	0.982	71.875	63.644	0.091
24	0.983	0.983	72.075	63.763	0.077
25	1	1	72.125	61.192	0.033
26	0.98	0.98	71.325	63.545	0.11
26*	N/A	0.991	71.025	62.181	0.11
27	0.976	0.969	71.425	62.746	0.075
28	1	1	72.075	62.019	0.062
29	1	1	70.875	61.572	0.115
30	0.986	0.98	71.625	62.603	0.079
31	0.978	0.968	71.425	61.522	0.038
31*	N/A	1	71.725	59.567	0.038
32	1	1	72.125	64.804	0.119

[1] Efficiency score according to DEA based on original treatment plans 1 to 37 (original dataset).
[2] Efficiency score according to DEA based on original treatment plans 1 to 37 and re-optimised plans 10*, 19*, 26*, 31* and 35* (re-optimised dataset).

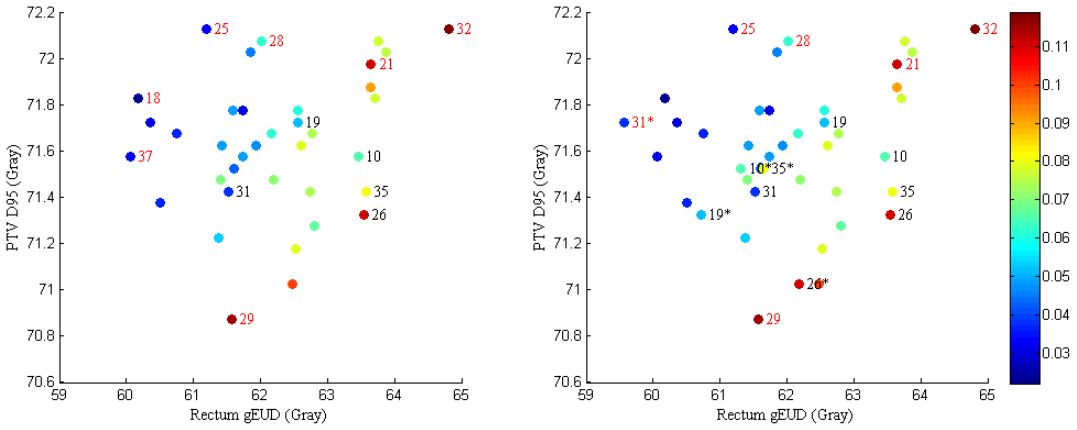

(a) (b)

Figure 4.1: Plot of the data points where colour represents PTV rectum overlap, (a) before and (b) after re-optimisation of a subset of the plans. Red numbers indicate treatment plans considered efficient in each DEA analysis. Black numbers indicate treatment plans selected for re-optimisation.

33	0.993	0.99	72.025	61.854	0.046
34	0.99	0.978	71.775	61.593	0.049
35	0.966	0.96	71.425	63.57	0.08
35*	N/A	0.993	71.525	61.658	0.08
36	0.974	0.97	71.175	62.529	0.079
37	1	0.992	71.575	60.056	0.032
Average	0.985	0.98	71.611	62.098	0.064
Standard deviation	0.01	0.012	0.302	1.17	0.025

The data of Table 4.1 is visualised in Figure 4.1. Notice that due to the presence of the environmental variable, the production frontier is not a single piecewise linear line as in Figure 2.3.

4.3 Results

The efficiency scores before and after including re-optimised plans as well as the input and outputs used in DEA are provided in Table 4.1. Note that in the table we do not show the values of slack variables as the majority of them are zero and the

maximum value of them is only 0.35 Gy. The high efficiency scores (average 0.985, standard deviation 0.01) indicate that the plan qualities of these prostate plans, in terms of PTV D_{95} and rectal gEUD, are quite consistent in general. The PTV D_{95} values among these plans vary slightly from 70.88 to 72.13 Gy while rectal gEUD and percentage overlap volume vary quite considerably from 59.57 to 64.80 Gy and from 2.2% to 11.9%, respectively. We find very strong evidence from the original dataset that the percentage overlap volume is correlated to rectal gEUD ($p = 1.76 \times 10^{-8}$). This suggests that one reason for the large variations in rectal gEUD is due to the large variations in percentage overlap volume. The small variation of PTV D_{95} among the plans is consistent with the fact that meeting the PTV coverage requirement was the highest priority for the planners unless instructed otherwise.

Plans 10, 19, 26, 31 and 35 were deemed sufficiently inefficient for their range of percentage overlap volume and were re-optimised. The efficiency scores of these plans are less than the original average efficiency score (0.985). The re-optimisation of plan 31 produced an additional efficient plan in the dataset. This plan extends the production frontier slightly and results in lower or equal efficiency scores of all other plans compared to those of the original dataset, as shown in Table 4.1. The efficiency score of the re-optimised plans are higher than those of the original plans, with an average improvement of 0.026 units. Note that this improvement is quite substantial since the standard deviation of the efficiency scores is only 0.012 (Table 4.1). All of the re-optimised plans achieved an efficiency score higher than the re-optimised average efficiency score (0.98).

A clinical peer review of the plans was performed using MIM Maestro V5.4 Mac Version (MIM Software Inc. Cleveland, OH). In this review, clinical and DVH parameters not included in the DEA (ie. bladder, femur heads, dose maxima, hot spot percentage and site, etc) were also taken into account. The review confirmed that re-optimised plans 10*, 31* and 35* were deemed superior when comparing with the original plans. Figure 4.2 shows improved conformity for these re-optimised plans. The green shaded areas show iso-dose volume of the original plans and the red shaded areas show the improvement in iso-dose volume after re-optimisation. Re-optimised plan 19* was deemed substantially equivalent to the original plan from a clinical and dosimetric viewpoint. For plan 26, the re-optimised plan was

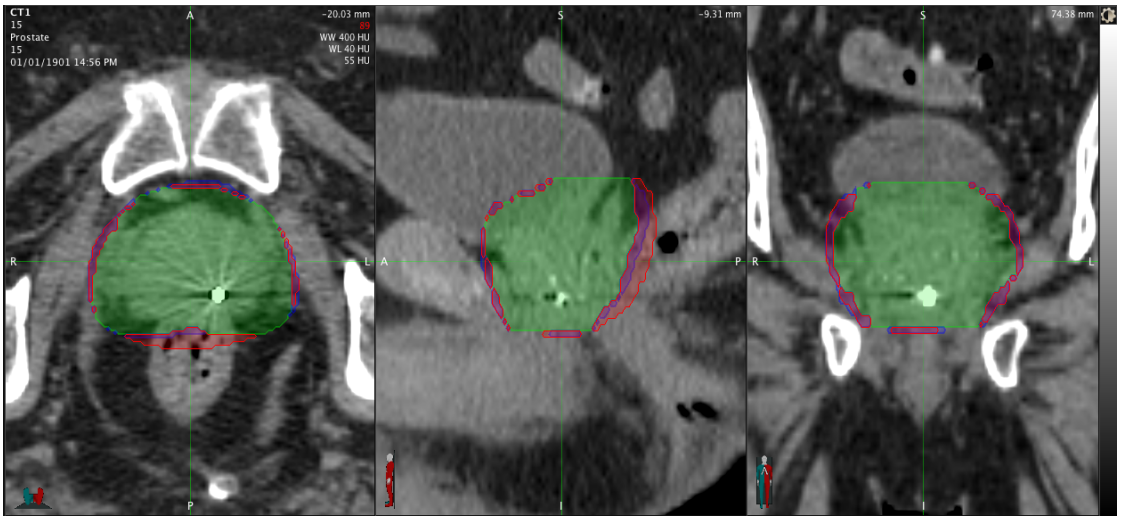

Figure 4.2: Examples of improved conformity for re-optimised plans.

not considered superior to the original plan due to comparatively inferior dose painting and conformity. However, we note that dose painting and conformity index were not taken into account in the DEA analysis and the goal of re-optimisation was to further improve rectal sparing while maintaining overall plan acceptance, which was achieved for all five re-optimised plans. The fact that planners could achieve improvements as instructed is a positive finding with regard to our DEA model. The degradation of the treatment is a product of the DEA information not perfectly in agreement to clinical concerns. The original, re-optimised and the target input values for the selected plans are summarised in Table 4.2. The overlap fractions are not included since they are the same as the corresponding values in Table 4.1. In terms of the chosen input and outputs, re-optimisation of the five plans resulted in an average reduction of 1.84 Gy in rectal gEUD. The reduction in PTV D_{95} is only between -0.3 (an actual increase) and +0.4 Gy with an average of +0.07. Due to this minor change in PTV D_{95} only rectal gEUD (input) is shown in Table 4.2. Among the five plans, the minimum rectal gEUD reduction is 1.364 Gy (Table 4.2). These results suggest that DEA successfully identified plans with potential improvements in terms of the chosen input and outputs. In addition, although planners were not provided with the results of DEA, the values for the re-optimised plans are very close to their targets, with a maximum difference of 0.806 Gy. Note that the target of a re-optimised plan represents DEA's prediction of a best attainable plan. This minor input/output difference between

Table 4.2: The original, re-optimised and the corresponding target input values for the selected plans. The measurement unit for the input is Gy.

Plan	Original	Target (original dataset)	Re-optimised	Target (re-optimised dataset)	Dose reduction [1]	Prediction error [2]
10	63.457	61.198	61.314	60.508	2.143	0.806
19	62.562	61.007	60.71	59.93	1.852	0.78
26	63.545	62.291	62.181	61.632	1.364	0.549
31	61.522	60.167	59.567	59.567	1.955	0
35	63.57	61.409	61.658	61.211	1.912	0.447
Avg	62.931	61.214	61.086	60.57	1.845	0.516

the prediction and the re-optimised plan demonstrates the ability of DEA in predicting potential improvement in terms of the input and outputs. The DVHs for the original plans and the re-optimised plans are shown in Figure 4.3. The improvements in rectal sparing are clearly illustrated in the DVHs while there are no clinically considerable differences between original PTV DVHs and re-optimised PTV DVHs. We also applied the Spearman's rank test to test the significance of the correlation between the change in rectal sparing after re-optimisation and the predicted change in rectal sparing made by DEA. However, possibly due to the small sample size, the result was not statistically significant.

To provide an assessment of the clinical relevance of the plan changes, an analysis of the biological objective function P+ was undertaken. P+ is a scalar quantity that combines the probability of tumour control with the probability of normal tissue complication (Kallman et al., 1992) and is often referred to as the complication free tumour control probability. The average P+ value increased from 0.618 to 0.632, a 2.3% improvement. If, however, we scale the treatment plan to that of a 78Gy prescription, which is now the institutional norm for medium to high risk cases at Auckland Radiation Oncology, then the average P+ increases from 0.661 to 0.685, a 3.6% improvement. As the dose prescription increases, the biological consequences and potential gains also increase. In deriving P+, the endpoint for

[1] Dose reduction is calculated as original input value minus re-optimised input value.

[2] Prediction error is the absolute difference between the re-optimised target input value and the re-optimised input value.

normal tissue complication was necrosis/stenosis according to the values of Ågren-Cronqvist (1995) and tumour control probability from Cheung et al. (2005).

An additional measure that can be extracted from DEA analysis is the peer count, which is the number of times a plan is referred to as a peer. The peer count indicates how often an efficient DMU is used as benchmark for others and thus

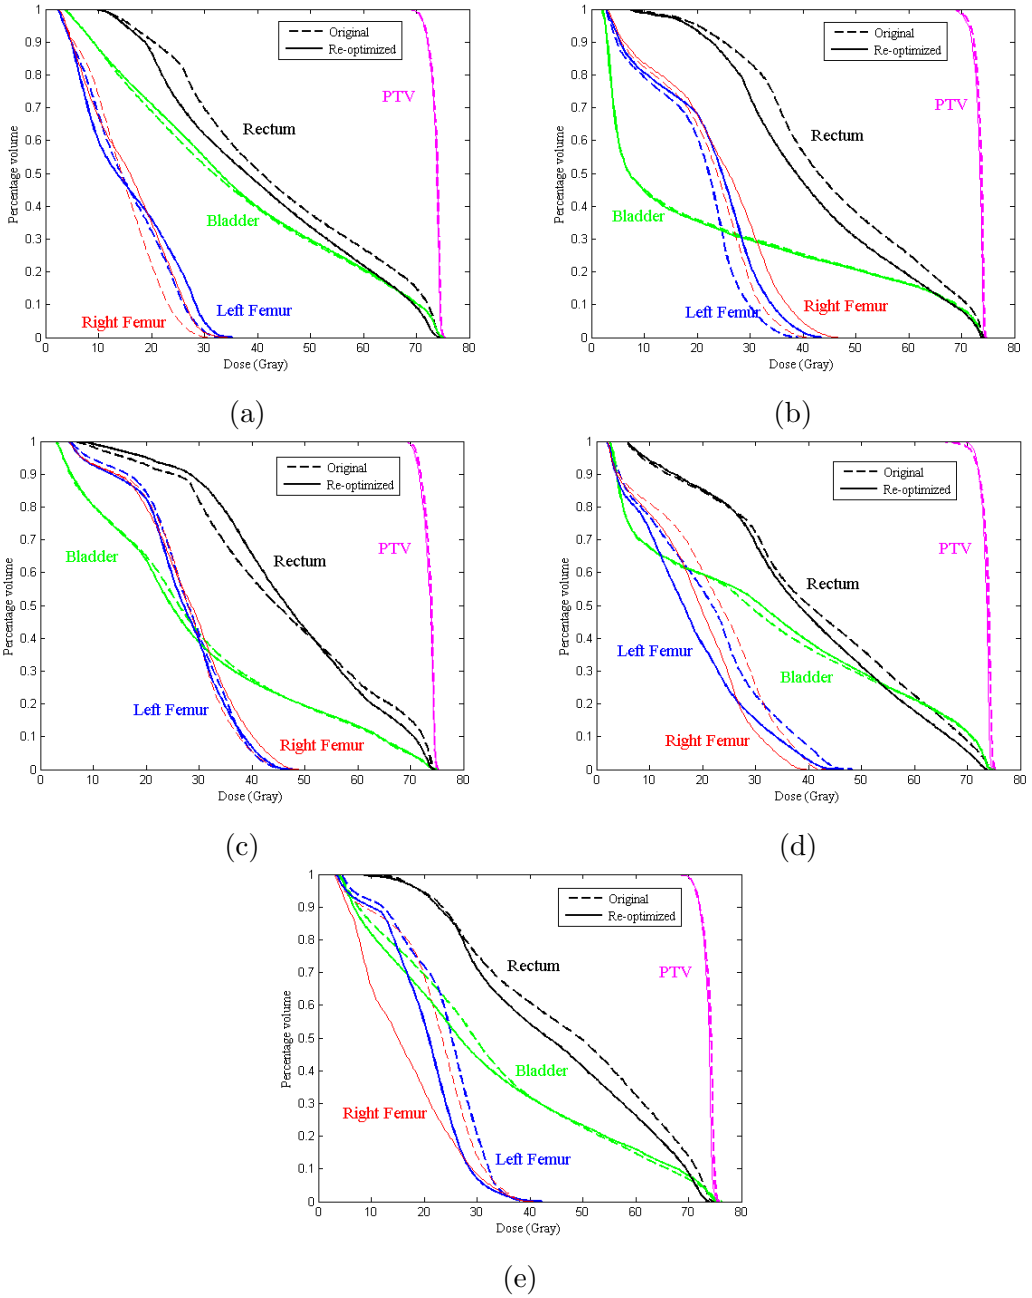

Figure 4.3: The original and re-optimised DVHs of (a) plan 10, (b) plan 19, (c) plan 26, (d) plan 31 and (e) plan 35.

attests to the plan's quality. The peer counts of the efficient plans are shown in Table 4.3. There are 7 efficient plans in the original dataset. After including plan 31* in the dataset, plan 18 and plan 37 are no longer efficient and are not used as a peer for other plans. Plan 25 and plan 32 have relatively low peer counts for both the original and re-optimised dataset, as shown in Table 4.3. One limitation of DEA is that a plan can be considered efficient simply because it has the best value in one of the inputs or outputs. Plan 25 has the highest PTV D_{95} value and plan 32 has the highest overlap fraction. However, the peer counts of these two plans being greater than 1 suggests that these plans have preferable output to input ratios compared to at least one other plan and that they are not considered efficient purely because they have an optimal value in one of the inputs or outputs. Plan 31* has the highest peer count of 37 and is referred to as a peer for all inefficient plans in the re-optimised dataset. A closer look at its input and outputs shows that it has the lowest rectal gEUD value of 59.567 Gy with an above average PTV D_{95} value of 71.725 Gy. This results in the highest output to input ratio of 1.204 within the dataset and explains why the plan is used as a peer for all inefficient plans.

Table 4.3: Peer counts (the number of times a plan is referred to as a peer) of the efficient plans.

Plan	18	21	25	28	29	31*	32	37
Peer count original	22	6	2	24	30	N/A	2	9
Peer count re-optimised	N/A	23	4	5	27	37	2	N/A

To further justify our approach, we carried out re-optimisation for two plans that were characterised as efficient in the initial DEA analysis, plans 28 and 29. The re-optimisation was carried out by the same planner who did the re-optimisation of the 5 inefficient plans with the same instructions. As mentioned before, the planner was unaware of the results of the DEA analysis. The efficiency score of both plans before and after re-optimisation is 1. Once again the values of PTV D_{95} changed by only small amounts (an increase of 0.3 and 0.45 Gy, respectively). Rectal sparing was achieved with a decrease in rectal gEUD of 1.555 and 1.898 Gy, respectively. The DVHs in Figure 4.4 illustrate the plan changes. But in contrast to the re-optimisation of inefficient plans, the improvement in rectal sparing was

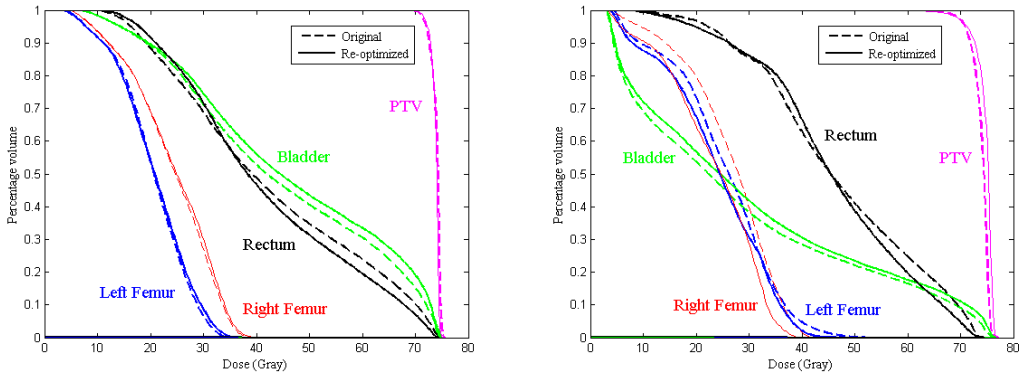

(a) (b)

Figure 4.4: The DVHs of original and re-optimised (a) plans 28 and (b) 29.

accompanied by an increase of dose to the bladder and other areas beyond the rectum, which worsened the conformity index for plan 28 and significantly increased dose to the small bowel beyond what was considered as clinically acceptable for plan 29. Clinical review of these plans suggests that the re-optimised plans are not better than the original plans. These results further support the correctness of our DEA model in identifying plan improvement potential. We also performed the Mann-Whitney U test against the null hypothesis that the improvement in rectal sparing after re-optimisation for the five inefficient plans is not greater than that of the two efficient plans. However, possibly due to the small sample size, the result was not significant.

4.4 Discussion and conclusion

In this study we posited the validity of using DEA as a quality assessment tool for radiotherapy planning. We used DEA to assess the quality of 37 prostate plans using an input-oriented VRS model with rectal gEUD as the input, PTV D_{95} as the output and the percentage volume of rectum overlapping PTV as a non-discretionary output variable. Five plans were considered of lesser quality by DEA were re-optimised with the goal of further improving the dose to the rectum while maintaining clinical acceptance. After re-optimisation, the dose to the rectum for all five plans improved without clinically considerable deterioration

in PTV coverage. In addition, the input and outputs of the re-optimised plans are very close to DEAs prediction of the best achievable plan, with a maximum difference of 0.8 Gy. These results confirm that DEA is capable of identifying plan improvement potential and predicting the best attainable plan in terms of the input and outputs.

There are several advantages of using DEA methodology as a quality assessment tool for radiotherapy treatment planning. Firstly, DEA is non-parametric. For non-parametric approaches, there is no need to assume a functional form for the production frontier whereas such an assumption is required for parametric approaches. While parametric approaches such as regression analysis can also be used to estimate the production frontier, it can be difficult to specify a functional form for the production frontier especially when there are multiple inter-related parameters. In contrast, the non-parametric nature of DEA allows practitioners to select the inputs and outputs that are considered most relevant in assessing the quality of radiotherapy treatment plans without too much concern on the underlying relationship among them. The non-parametric nature of DEA leads to its second advantage: the ability to handle multiple inputs and multiple outputs. This ability is particularly important in assessing radiotherapy treatment plans due to the conflicting nature of treatment objectives. Assessing radiotherapy treatment plans based solely on the maximal OAR sparing might encourage the planners to generate plans with maximal OAR sparing but near minimal acceptable PTV coverage. DEA methodology allows plan assessment based on multiple inputs and outputs and therefore captures the conflicting nature of treatment planning more adequately. Thirdly, DEA constructs a production frontier based on the best attainable results in the dataset of historical treatment plans. This is distinctly different to ordinary least squares regressions that attempt to fit a regression function to the centre of the data spread to estimate the "average" attainable results rather than the best attainable results. In this radiotherapy application, since we are interested in the best attainable results, we consider DEA to be preferable to regression methods. Fourthly, DEA not only provides the efficiency score for the plan being assessed, but also target information, including the peers and the corresponding weights. The target shows the inputs and outputs required to make a radiotherapy treatment plan efficient. A treatment planner can compare the target with the treatment plan being assessed and decide if further planning is required.

If the target is largely composed of a particular peer, a treatment planner can trace back to the peer, assess how the peer is derived and perhaps conduct the re-optimisation using similar treatment objectives and/or objective weights. In addition, DEAs capability to predict the best attainable dose for both the PTV and the OAR allows planners to set achievable plan objectives. Wu et al. (2011) suggested that by setting achievable plan objectives, one can reduce the trial-and-error attempts required to find a satisfactory plan. Last but not least, DEA is readily available in many software packages (Barr, 2004) and can be conducted independently of the treatment planning system with negligible computational effort. This provides clinics with an approach to improve planning efficiency and plan quality without replacing their existing treatment planning system, which is quite an effort from both fiscal and implementation viewpoints.

While DEA offers many advantages, it is not without some potential limitations. One limitation is that the efficiency score for a plan is a relative measure compared to other plans in the dataset. Thus a plan rated fully efficient for a dataset might not be truly the best plan, but simply a superior plan compared to the plans in the dataset. However, as more efficient plans are generated and included in the dataset, DEA will be able to learn from the plans and will be able to approximate the true production frontier more accurately. As a result, this limitation would become less significant over time. Another limitation is that as more inputs and outputs are included in the formulation, DEA starts to lose discrimination power on the performance of the DMUs. Introducing more inputs and outputs imposes more constraints in the formulation and thus further constrains the set of potential inputs and outputs. As a consequence, the DMUs are closer to the production frontier and more DMUs will be deemed efficient or close to be efficient. To address this limitation, we note that in practice at most two or three clinically relevant trade-offs between the objectives need to be considered (Hoffmann et al., 2006). By including only these relevant objectives in the formulation, the number of inputs and outputs can be effectively controlled while maintaining the discrimination power. The last limitation is that a plan can be rated efficient simply because it has an optimal value in one of the DEA inputs or outputs compared to all other plans. For example, the plan with the lowest rectal gEUD and the plan with the highest PTV D_{95} will be rated as efficient regardless the values of other DEA inputs and outputs. This limitation can be alleviated by checking if the plan is

referred to as a peer for other plans. In general, given a database of reasonable size, an efficient plan with preferable output to input ratio is likely to be used as a peer for another plan. In contrast, a plan considered efficient simply because it has an optimal value in one of the input/output values is usually not used as a peer for other plans. Thus by checking the peer counts, we can effectively identify efficient plans that may not be truly desirable.

In this preliminary study, due to limited clinical resources (planners, oncologists and the treatment planning software) available for plan re-optimisation and evaluation, the performance of DEA are evaluated empirically based on a small sample of re-optimised plans. In future studies, the statistical aspects in evaluating the performance of DEA model should be investigated. In particular, statistical significance on the improvements in plan quality after re-optimisation for the inefficient plans and for the efficient plans, as suggested by DEA, should be investigated. Furthermore, the significance on the correlation between the change in plan quality after re-optimisation and the predicted change in plan quality by DEA should be considered. As clinical resources are usually scarce, acquiring clinical resources for plan re-optimisation and evaluation for a decent sample size can be challenging. Alternatively, one may consider collecting all the plans generated from the (trial-and-error) planning phase for the incoming patient cases and subsequently investigate for statistical difference in DEA efficiency scores between plans considered unsatisfactory and plans accepted for actual treatment. Additionally, one can also investigate the significance of correlation between DEA predicted plan improvement and plan improvement realised by the final plan selected for treatment.

Further investigation is required to extend the DEA model to include more and/or other plan assessment criteria. In this study we use PTV D_{95} and rectal gEUD to account for the dose to the PTV and the rectum, respectively. Planning criteria associated with other OARs such as bladder and femur heads are generally unchallenging and, in our opinion, would not require DEA analysis. It may, however, in the case of prostate planning be useful to include second PTV or rectal parameters such as PTV conformity index or rectal V_{70} Gy (percentage volume of a structure that receives at least 70 Gy of radiation). As previously discussed, adding to the number of plan assessment criteria in the DEA model degrades plan quality discrimination. In future research, we will investigate the most effective

plan assessment criteria that should be included in the DEA model, followed by an investigation on how these plan assessment criteria can be incorporated in the DEA model while maintaining sufficient discrimination on the quality of the plans. Other treatment sites likely to require a larger number of inputs and outputs such as head and neck will also be investigated in future work.

Chapter 5

Integrating column generation in a method that generates an evenly distributed representative non-dominated set

MCO deals with optimisation problems involving several conflicting criteria. In MCO, a single solution that simultaneously optimises all objectives generally does not exist. Instead, the decision maker must consider the trade-offs between the efficient solutions. One way to do so is by adopting an a posteriori MCO method, which seeks to obtain the non-dominated set so that the decision maker can select the most preferred non-dominated point and a corresponding efficient solution for the problem at hand. In multi-objective continuous optimisation, the non-dominated set consists of infinitely many non-dominated points. It is therefore impractical for a decision maker to examine all non-dominated points. Instead, a practical approach is to obtain a discrete representation of the non-dominated set satisfying some quality requirements (Faulkenberg and Wiecek, 2010, Sayın, 2000). Many methods that follow this approach have been proposed in the last two decades, as the paper by Faulkenberg and Wiecek (2010) shows. Given this representative non-dominated set, the decision maker can navigate through the non-dominated points and decide on the most preferred point. In this chapter we

consider how column generation can be integrated in an MCO approach to find an approximated representation of the non-dominated set for MOLPs.

Column generation is a technique that solves linear programmes by considering only a subset of the decision variables. The technique is particularly beneficial when the number of variables is much greater than the number of constraints. The idea is based on the fact that, typically, only a subset of variables is required in the basis to reach optimality; other variables are non-basic and have a value of zero. Column generation exploits this fact by only considering variables that have the potential to improve the objective function value, indicated by negative reduced costs. In each iteration of a column generation method, two problems need to be solved successively: the restricted master problem (RMP) and the subproblem (SP). RMP is the original problem with only a subset of variables. By solving the RMP, a vector of dual values associated with the constraints of the RMP is obtained. The dual information is passed on to the SP. The goal of the SP is to identify a new variable and an associated coefficient column with negative reduced cost, which can potentially improve the objective function value of the original problem. If such a variable and column can be identified, then they are added to RMP, which is re-optimised, and the next iteration begins. Otherwise, an optimal solution of RMP is also an optimal solution of the original problem.

Column generation methods in multi-objective optimisation are rare. Moradi et al. (2015) present a column generation approach for the (linear) bi-objective multi-commodity minimum cost flow problem. Their algorithm incorporates column generation within a bi-objective simplex algorithm, which requires a modification of the objective function of the SP to a linear fractional function. The study of Salari and Unkelbach (2013) falls into the domain of non-linear programming, thus the subproblem is based on partial derivatives of individual objective functions. The aim of Salari and Unkelbach (2013) is to approximate the entire non-dominated set using a limited number of variables. The basic idea is to use column generation to identify variables that potentially improve the non-dominated set approximation as a whole. To find such variables, multiple weighted-sum RMPs, where each RMP is associated with a unique non-negative weight vector, are solved. The partial derivatives obtained from solving each RMP are passed to a subproblem, which

aggregates the individual subproblems corresponding to each RMP. A column obtained from solving the aggregated subproblem therefore potentially improves the majority of individual RMPs, thus improving the non-dominated set approximation as a whole. However, due to the use of weight vectors for the RMPs and the use of an aggregated subproblem, their method cannot guarantee that the whole non-dominated set is well approximated.

In this chapter, we propose to use column generation within a procedure that constructs an evenly distributed finite representative set of non-dominated points of an MOLP, i.e., the revised normal boundary intersection (RNBI) method of Shao and Ehrgott (2007). The RNBI method combines aspects of the global shooting method (Benson and Sayin, 1997) and the normal boundary intersection method (Das and Dennis, 1998) and has been proven to generate evenly distributed non-dominated points for MOLPs (Shao and Ehrgott, 2007). Unlike the method of Salari and Unkelbach (2013) in which a subproblem identifies a variable that improves the non-dominated set approximation in general, each of the column generation subproblems in our approach identifies a variable and an associated column to move a point in objective space in a direction that leads to non-dominance. In fact, if the master problems are solved to optimality by column generation, the resulting points will be on the boundary of the feasible set of the MOLP in objective space, thus one obtains a representative non-dominated set identical to the one generated by RNBI. However, we note that column generation is rarely used to solve a master problem to optimality. In this situation, the points generated by the column generation RNBI method will be in the interior of the feasible set in objective space and form an approximated representation of the non-dominated set, i.e., a set of discrete feasible points, not necessarily non-dominated, used to represent the non-dominated set.

In Sections 5.1 and 5.2, we provide background and formulations of single objective column generation and the RNBI method, respectively. In Section 5.3, we introduce the column generation RNBI formulation and discuss implementation issues associated with the method, i.e., the detection of infeasibility through a reference point bounding method and initialisation of the process. The quality of the representative set obtained by our column generation based RNBI method is discussed in Subsection 5.3.3.

5.1 Column generation

Consider a single objective linear programme referred to as the master problem (MP),

$$\begin{aligned} v_{MP} := \min & \sum_{j \in J} c_j x_j \\ \text{s.t.} & \sum_{j \in J} a^j x_j \quad \geqq \quad b, \\ & x_j \quad \geqq \quad 0, \quad j \in J, \end{aligned} \tag{MP}$$

with $|J| = n$ variables and m constraints. Each variable x_j is associated with a cost coefficient c_j and a constraint column $a^j \in \mathbb{R}^m$. The right-hand side constraint coefficients are specified by column $b \in \mathbb{R}^m$. The column generation technique considers a RMP which uses only a subset $J' \subseteq J$ of all variables. Because of this, the optimal solution x^* of RMP is worse than or equal to the optimal solution of MP in terms of the objective function. By solving the RMP, we obtain a dual solution π^* associated with the constraints of the RMP. In the Simplex method, the dual solution is used to calculate the reduced cost of each non-basic variable which indicates the unit change of the objective function value if the variable were to enter the basis. If the reduced cost of all non-basic variables is non-negative, the current basic feasible solution of RMP is an optimal solution to MP. Otherwise, a non-basic variable with negative reduced cost enters the basis, which improves the objective function if the entering variable takes a value greater than zero.

Column generation works in an analogous way. The vector of dual values π^* obtained from solving the RMP is passed into a subproblem,

$$\begin{aligned} \min \bar{c}_j &= c_j - (\pi^*)^T a^j \\ \text{s.t. } & j \in J. \end{aligned} \tag{SP}$$

SP finds a variable x_{j^*} with lowest reduced cost $\bar{c}_{j^*}$. If $\bar{c}_{j^*}$ is negative, the non-basic variable x_{j^*} and the coefficient column $\left(c_{j^*}, a^{j^{*T}}\right)^T$ are added to RMP and RMP is re-solved. Otherwise, an optimal solution of RMP is also an optimal solution of

MP. Note that SP can be solved as an optimisation problem if the set J can be described by the feasible set X_J of an optimisation problem

$$\begin{aligned} &\min\ c(\lambda) - (\pi^*)^T a(\lambda) \\ &\text{s.t. } \lambda \in X_J, \end{aligned} \tag{5.1}$$

in which $c_j = c(\lambda)$ and $a^j = a(\lambda)$ and which has variable vector $\lambda \in X_J$.

The lowest reduced cost $\bar{c}_{j^*}$ can be used to derive a lower bound on the optimal value v_{MP} of MP. Denote the optimal value of the current RMP as v^*_{RMP}. If there exists a constant κ with $\sum_{j \in J} x_j \leq \kappa$ for any optimal solution of MP, then we have a lower bound

$$v^*_{RMP} + \kappa \bar{c}_{j^*} \leqq v_{MP}, \tag{5.2}$$

since we cannot improve the objective function value v^*_{RMP} by more than κ times the lowest reduced cost $\bar{c}_{j^*}$ (Lübbecke, 2010).

5.2 The RNBI method

The RNBI method is designed to solve an MOLP

$$\min\{Cx : x \in X\}, \tag{MOLP}$$

where $C \in \mathbb{R}^{p \times n}$ is the cost coefficient matrix consisting of row vectors $c^k \in \mathbb{R}^n$ for $k = 1, \ldots, p$. For MOLPs we will assume that $X \subseteq \mathbb{R}^n$ is a non-empty compact polyhedral set (a polytope) of feasible solutions. The feasible set in objective space Y defined by

$$Y = \{Cx : x \in X\} \tag{5.3}$$

is also a polytope since it is the image under a linear mapping of the polytope X.

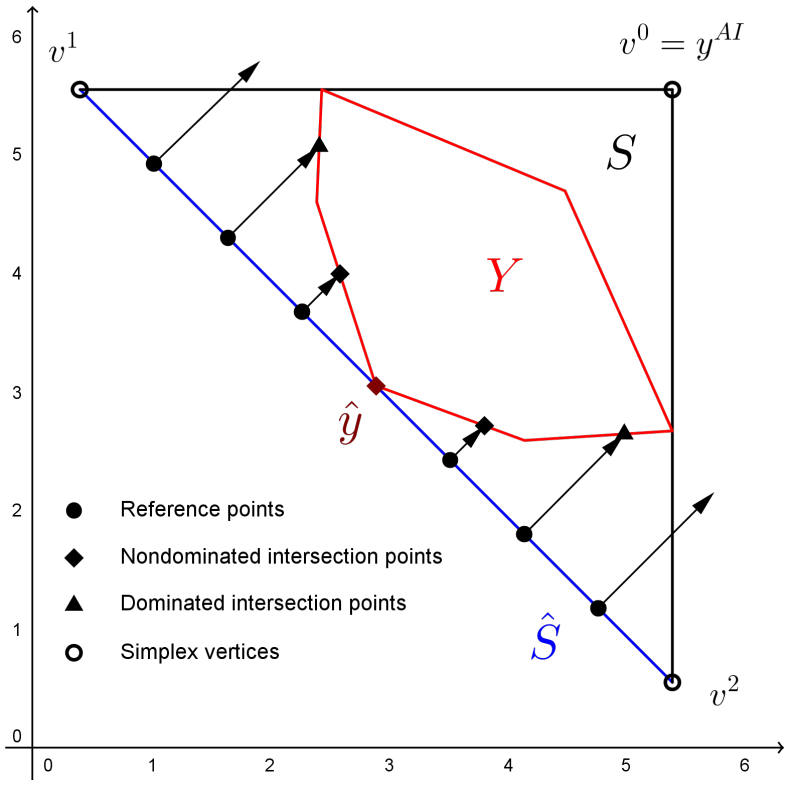

Figure 5.1: The RNBI method: Illustration of the simplex S containing the feasible set in objective space Y, the reference subsimplex $\hat{S}$ and the half-lines emanating from the reference points.

We first explain the general idea of RNBI. A simplex S is constructed such that it contains Y and such that the non-dominated set S_N of S is a subsimplex of S. We denote by $\hat{S} := S_N$ the reference subsimplex. Reference points are positioned on $\hat{S}$ and for each reference point q, a half-line emanating from q in direction e is generated, where e is a vector of all ones. The RNBI subproblem then searches for the intersection point between the half-line and (the boundary of) Y closest to the reference point. As illustrated in Figure 5.1, not all half-lines intersect with Y. In this case the RNBI subproblem will be infeasible. In addition, some intersection points may be dominated. Hence in the last step, the algorithm checks the non-dominance of intersection points by solving one LP for each intersection point. The following subsections outline the RNBI method in more detail.

5.2.1 Constructing the reference subsimplex and choosing reference points

To construct the reference subsimplex $\hat{S}$, we first obtain scalar μ as $\mu := \min\{e^T y : y \in Y\}$. μ is attained at a non-dominated point $\hat{y}$ of Y, as illustrated in Figure 5.1. We then derive the anti-ideal point y^{AI} of the MOLP, where $y_k^{AI} := \max\{y_k : y \in Y\}$ for $k = 1, \ldots, p$ (Ehrgott, 2005). Based on μ and y^{AI}, we can define the $p+1$ vertices $v^k \in \mathbb{R}^p, k = 0, 1, \ldots, p$ of the simplex S that contains Y. Let $v^0 := y^{AI}$. For $k = 1, \ldots, p$ and $l = 1, \ldots, p$ let

$$v_l^k := \begin{cases} y_l^{AI} & \text{if } l \neq k, \\ \mu + y_k^{AI} - e^T v^0 & \text{if } l = k. \end{cases} \tag{5.4}$$

The convex hull of vertices $\{v^k : k = 0, 1, \ldots, p\}$ is a p-dimensional simplex S that contains Y, as shown by Benson and Sayin (1997). The reference subsimplex $\hat{S}$, which is the non-dominated set of S, is defined by the convex hull of vertices $\{v^k : k = 1, 2, \ldots, p\}$. Reference points on $\hat{S}$ can now be chosen as particular convex combinations of the extreme points of $\hat{S}$, i.e. a reference point q is given by

$$q = \sum_{k=1}^{p} \alpha_k v^k, \tag{5.5}$$

where α_k is the weighting of vertex k for $k = 1, \ldots, p$ with $0 \leqq \alpha_k \leqq 1$ and $\sum_{k=1}^{p} \alpha_k = 1$. By varying the weighting for each vertex with a fixed increment η, an evenly distributed discrete set of points on the reference subsimplex $\hat{S}$ can be generated (Benson and Sayin, 1997).

5.2.2 Computing the intersection points and checking non-dominance

Let the set of reference points be Q. For each reference point $q \in Q$, RNBI computes the intersection point y of the half-line $\{q + te : t \geqq 0\}$ and the boundary

of Y by solving the RNBI subproblem

$$\min\{t : q + te \in Y; t \geqq 0\}. \tag{RNBISub}$$

Notice that by construction, the all-ones vector e is the normal of the reference subsimplex $\hat{S}$. As illustrated in Figure 5.1, there are three scenarios for RNBISub:

- RNBISub is infeasible if and only if the half-line $\{q + te : t \geqq 0\}$ does not intersect Y.
- RNBISub has an optimal solution t^*, but the intersection point $q + t^*e$ of the half-line $\{q + te : t \geqq 0\}$ and Y is dominated.
- RNBISub has an optimal solution t^* and $q + t^*e$ is a non-dominated point of Y.

The first case is detected by infeasibility of RNBISub. Because an intersection point may be dominated, it is necessary to check every intersection point for non-dominance. To do so, after obtaining all intersection points, a non-domination filter can be used to exclude some of the dominated points (Messac et al., 2003). This method allows fast elimination of some dominated intersection points but cannot guarantee the remaining points are non-dominated. Hence the non-dominance of the remaining intersection points $\bar{y}$ must be verified, e.g., by solving the linear programme

$$\min\{\lambda^T y : y \leqq \bar{y}; y \in Y\}, \tag{5.6}$$

where $0 < \lambda \in \mathbb{R}^p$ is an arbitrary strictly positive weight vector, for instance $\lambda = e$. Then $\bar{y}$ is non-dominated if and only if the optimal value of (5.6) is equal to $\lambda^T \bar{y}$ (Ehrgott, 2005).

5.3 The RNBI method using column generation

To integrate column generation in the RNBI framework we solve RNBISub using column generation. To do so, we adopt RNBISub as the master problem. Following the background definitions for column generation in Section 5.1, we formulate the

restricted master problem with a subset $J' \cup \{t\}$ of variables and with feasible set defined by constraints $\sum_{j \in J'} a^j x_j \geqq b$, $x_j \geqq 0$ for all $j \in J'$ and $t \geqq 0$. The condition $q + te \in Y$ of RNBISub is rewritten as constraints $q_k + t = \sum_{j \in J'} c_j^k x_j$ for $k = 1, \ldots, p$. In this way, the objective functions of the original MOLP are incorporated in the restricted master problem as constraints of RNBISub, i.e., (5.7b), in addition to the original constraints of MP, i.e., (5.7c). The corresponding RMP is referred to as RMP-RMBISub, and is shown as follows.

$$\min \quad t \tag{5.7a}$$

$$\text{s.t.} \quad q_k + t = \sum_{j \in J'} c_j^k x_j, \quad k = 1, \ldots, p, \tag{5.7b}$$

$$\sum_{j \in J'} a^j x_j \geqq b, \tag{5.7c}$$

$$x_j \geqq 0, \quad j \in J' \tag{5.7d}$$

$$t \geqq 0. \tag{5.7e}$$

Notice that RMP-RNBISub is essentially the same as the RNBI subproblem but with only a subset of variables $j \in J'$. To conduct column generation on this RNBI subproblem, we solve RMP-RNBISub and the corresponding SP sequentially and iteratively. We remark that, in case column generation is terminated early, i.e., an optimal solution of RNBISub is not yet confirmed, the intersection point may be dominated. In contrast to the original RNBI method, non-dominance of the intersection points will not be checked, because only an optimal solution of RNBISub can define a non-dominated point.

As indicated in Section 5.2.2 RNBISub may be infeasible even in the presence of all variables. Hence, if we solve RMP-RNBISub with a subset of variables, it may be infeasible because either the constraints (5.7c) are not satisfied with a subset of variables or because the master problem RNBISub is infeasible, i.e., $\{q+te : t \geqq 0\}$ does not intersect Y. The former case can be dealt with by the use of artificial variables to satisfy constraints (5.7c), see also Section 5.3.1. But in the latter case, many iterations of column generation may be wasted in an attempt to detect

infeasibility. In fact, infeasibility of RNBISub could only be determined once all artificial variables are eliminated from the solution.

It will therefore be beneficial to identify reference points for which this is the case early to avoid attempts to solve RNBISub for such reference points. For convenience, we will from now on refer to reference points for which RNBISub is infeasible as infeasible reference points. In Section 5.3.1 we present a method, which we call reference point bounding, to identify infeasible reference points. To deal with infeasibility due to the restricted number of variables in RMP-RNBISub, we present three methods of initialisation in Section 5.3.2. Finally, we discuss the quality of the discrete approximation generated by column generation RNBI in Section 5.3.3.

5.3.1 Reference point bounding

One issue with the RNBI method, which stems from the use of the anti-ideal point in the definition of the covering simplex S and the reference subsimplex $\hat{S}$, is that there can be infeasible reference points, i.e., reference points for which RNBISub is infeasible such that $\{q + te : t \geqq 0\} \cap Y = \emptyset$. Because the components of y^{AI} may be far larger than the objective values of any non-dominated point, there can potentially be many reference points for which this is also the case, as shown in Figure 5.1. Obviously, any effort invested in solving RNBISub for infeasible reference points is wasted in the sense that it does not contribute to the computation of a representative set of non-dominated points. Therefore, solving RNBISub using column generation if RNBISub is in fact infeasible, can dramatically increase the computational time (see Section 6.2). In order to identify infeasible reference points we provide Theorem 5.2 characterising infeasible reference points and therefore defining the subset of feasible reference points of $\hat{S}$. We first state a lemma concerning the set of all feasible reference points.

Lemma 5.1. *The subset $\hat{Q} \subset \hat{S}$ of points q such that $\{q + te : t \geqq 0\} \cap Y \neq \emptyset$ is a polytope.*

Proof. The result follows, because $\hat{Q}$ is the projection of polytope Y onto $\hat{S}$, which is a simplex on the hyperplane $e^T y = \mu$. □

Theorem 5.2. *Let $q \in \hat{S}$ be a reference point. Then q is infeasible if and only if there is some $d \in \mathbb{R}^p \setminus \{0\}$ such that $d^T q < \min\{d^T z : z + te \in Y, z \in \hat{S}, t \geqq 0\}$.*

Proof. We first observe that the LP $\min\{d^T z : z + te \in Y, z \in \hat{S}, t \geqq 0\}$ is always feasible, because $z^* = \hat{y}, t^* = 0$ with $\hat{y}$ as defined in Section 5.2.1 is a feasible solution. It is also bounded, because Y is a compact set by assumption. Then $d^T q < \min\{d^T z : z + te \in Y, z \in \hat{S}, t \geqq 0\}$ implies that q does not satisfy $q + te \in Y$ for any $t \geqq 0$. Now let q be an infeasible reference point. Then $q \notin \hat{Q}$ as defined in Lemma 5.1. Hence there exists a hyperplane strictly separating q from $\hat{Q}$, i.e., there is $d \in \mathbb{R}^p \setminus \{0\}$ such that $d^T q < \min\{d^T z : z \in \hat{Q}\} = \min\{d^T z : z + ty \in Y, z \in \hat{S}, t \geqq 0\}$. □

Although Theorem 5.2 provides a theoretical characterisation of all feasible reference points, it is clearly impractical for implementation. Hence, we restrict ourselves to finding minimum and maximum values of each individual co-ordinate z_k of points on the reference subsimplex $\hat{S}$ that are feasible reference points, i.e., we use the sufficient condition of Theorem 5.2 and apply it to vectors $d = e^k$ and $d = -e^k$ for $k = 1, \ldots, p$, where e^k is the kth unit vector. We call this method reference point bounding.

The linear programme $\min\{d^T z : z + te \in Y, z \in \hat{S}, t \geqq 0\}$ is solved for $d = e^k$ and $d = -e^k$ for $k = 1, \ldots, p$. Let the optimal values be z_k^{min} and z_k^{max}, respectively. Then according to Theorem 5.2, reference points q with $q_k < z_k^{min}$ or $q_k > z_k^{max}$ for any $k \in \{1, \ldots, p\}$ will be infeasible. Corollary 5.1 summarises the above argument.

Corollary 5.1. If q is a reference point with $q_k < z_k^{min}$ or $q_k > z_k^{max}$ for some $k \in \{1, \ldots, p\}$, then $\{q + te : t \geqq 0\} \cap Y = \emptyset$.

Reference points that satisfy the condition of Corollary 5.1 are eliminated from the set Q of reference points and the corresponding RNBI subproblems are left unsolved.

5.3.2 Initialisation of RMP-RNBISub

Constraints (5.7b) may not be feasible given a limited set of variables. In addition, even after the reference point bounding procedure is applied, infeasible reference points may remain due to $\{q + te : t \geqq 0\} \cap Y = \emptyset$. In this section we discuss how the infeasibility of RMP-RNBISub can be managed.

One way to handle the infeasibility is the Phase-1 approach, see, e.g., Chvátal (1983), which adds non-negative artificial variables to satisfy constraints (5.7b) and (5.7c) while changing the objective function of the problem to minimise the sum of the artificial variables. The Big-M approach assigns large costs M to the artificial variables and minimises the sum of the original objective function plus the sum of the costed artificial variables. Using artificial variables, feasibility of RMP-RNBISub is assured. As soon as any of the artificial variables has a value of zero in a solution, the artificial variable can be removed. If any of the artificial variables remain positive when the optimality condition is satisfied, we can conclude that RMP-RNBISub is infeasible because $\{q + te : t \geqq 0\} \cap Y = \emptyset$.

We notice that in practice, column generation is rarely used to solve a (single objective) linear programme to optimality. In this situation, a possible approach is to perform column generation iterations on RMP-RNBISub until a specified termination condition, such as a pre-specified number of columns, is reached. One can, for example, conclude that a reference point is (approximately) feasible, if the solution satisfies constraints (5.7c) and the remaining total infeasibility in constraints (5.7b) is small enough, i.e., below a certain pre-determined threshold.

An alternative approach to manage infeasibility is to generate coefficient columns that show that the RMP is feasible (Andersen, 2001). The method is based on Farkas' lemma, which states that either $Ax = b, x \geqq 0$ is feasible or there is a vector π with $\pi^T A \geqq 0$ and $\pi^T b < 0$. The vector π corresponds to the dual vector of a linear programme. A linear programme is proved to be infeasible by finding a dual vector such that the condition $\pi^T Ax = \pi^T b$ can never be met due to opposite signs on the right-hand side and the left-hand side of the equation. Thus to prove that the restricted master problem is feasible, we can add a column a to A with $\pi^T a \leqq 0$. Such a column can be found by solving $\min \{\pi^T a(\lambda) : \pi^T a(\lambda) \leqq 0, \lambda \in X_J\}$. If no

such column exists, we can conclude that the corresponding master problem is infeasible. We will refer to this approach as Farkas pricing.

5.3.3 Quality of the representative set computed by the column generation RNBI method

Sayın (2000) defines three measures, coverage, uniformity and cardinality, to quantify the quality of a discrete representation of a set. A good representation of the non-dominated set should not contain an excessive number of points (low cardinality), should have points significantly different from one another (as indicated by high uniformity level) and should not neglect large portions of the non-dominated set (low coverage error).

Let $G \subset Y_N$ be a finite set of non-dominated points generated by the standard RNBI method using reference points $q \in Q$. Let H be the representative set generated by the column generation RNBI method based on the same set of reference points. We shall write $g(q)$ and $h(q)$, respectively, to indicate the dependence of representative points on reference point q. The distance between two adjacent reference points is denoted as dq. Cardinality represents the number of points contained in the representation. It is clear that the number of points contained in H depends on the distance between adjacent reference points. In the rest of this section we discuss the quality of H in terms of uniformity level and coverage error.

The uniformity level δ of a representative set is measured by the distance between a pair of closest points in the set. The uniformity level of H can therefore be expressed by

$$\delta := \min_{h^i, h^j \in H, h^i \neq h^j} d(h^i, h^j) \tag{5.8}$$

with d being a metric. Here we use the Euclidean distance as metric. Assume h^k and h^l are the two closest points in H and let q^k and q^l be the corresponding reference points, as illustrated in Figure 5.2. By definition of the RNBI method, we know that vector $v^N = h^l - q^l$ must be perpendicular to vector $v^q = q^l - q^k$. Hence we have $\cos\theta = \|v^q\|/\|v^h\|$ where vector $v^h = h^l - h^k$ and θ is the angle between vectors v^h and v^q. To satisfy $h^l = q^l + t_l e$ with $t_l \geqq 0$ being the absolute difference between q^l and h^l in all objectives, we must have $0 \leqq \theta < \pi/2$, which

corresponds to $0 < \cos\theta \leqq 1$. Therefore, minimal $\|v^h\|$ occurs when $\cos\theta = 1$ and in that case the distance between h^l and h^k is $\|v^h\| = \|v^q\| = dq$. Therefore the lower bound on the uniformity level of H is dq, which is the same as that of G (Shao and Ehrgott, 2007).

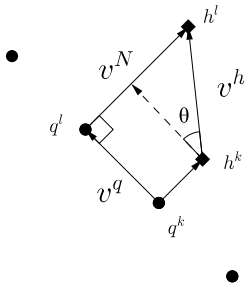

Figure 5.2: An illustration of the uniformity level for a representation produced by the RNBI method using column generation. The two diamonds represent a pair of closest representative points h^l and h^k and the circles represent the corresponding reference points q^l and q^k.

The coverage error ϵ indicates how accurately set H represents Y_N and can be expressed as

$$\epsilon := \max_{y \in Y_N} \min_{h \in H} d(y, h). \tag{5.9}$$

Essentially, the coverage error ϵ is the maximum distance between a point in the non-dominated set and its closest point in the representation H. Notice that if RMP-RNBISub is not solved to optimality by column generation, $h \in H$ can be an intersection point of $\{q + te : t \geqq 0\}$ with Y that is dominated even though $\{q + te : t \geqq 0\}$ intersects Y in a non-dominated point.

Shao and Ehrgott (2007) show that the coverage error of G is at most $(\sqrt{p}dq)/2$. Hence the coverage error of H is bounded by the maximum distance between points $g(q)$ and $h(q)$ of H and G generated for reference points $q \in Q$ plus the coverage error of G which can be expressed as

$$\epsilon \leqq \max_{q \in Q} d\left(g(q), h(q)\right) + \sqrt{p}dq/2. \tag{5.10}$$

The term $d(g(q), h(q))$ can be derived from the difference between the objective function values of RNBISub and RMP-RNBISub for reference point q. If RNBISub

is not solved to optimality, then one can use a lower bound on the optimal value of RNBISub as described in Section 5.1 to estimate the coverage error.

Based on the above discussion, we can see that the quality of a representation generated by column generation RNBI depends on the distance between adjacent reference points. As the distance decreases, cardinality increases, the uniformity level decreases and the coverage error decreases. In addition, the coverage error also depends on the maximum distance between representative points $g(q)$ and $h(q)$ for reference points $q \in Q$.

5.4 Summary

In this chapter, we have proposed the integration of column generation in the RNBI method to compute a discrete approximation of the non-dominated set of an MOLP. The method is based on solving each RNBI subproblem using column generation, which moves the current point in objective space of the MOLP towards the non-dominated set. Since RNBI subproblems may be infeasible, we attempt to detect this infeasibility early. First, a reference point bounding method is proposed to eliminate reference points that lead to infeasible RNBI subproblems. Furthermore, different initialisation approaches are also discussed, including Farkas pricing, which provides a mechanism to conclude the infeasibility of a RNBI subproblem.

In terms of the quality of a representation computed with the column generation RNBI method, we showed that the uniformity level is at least the same as the distance between two closest reference points (dq). The coverage error is bounded by the the distance between column generation RNBI points and the non-dominated set plus $\sqrt{p}dq/2$. This property allows the decision maker to choose an appropriate value of dq and select appropriate termination condition for the column generation process to produce a representative set tailored for the decision making problem at hand.

Chapter 6

Application of the column generation RNBI method in radiotherapy treatment design

The clinical deployment of MCO-based planning has been shown to benefit the planning practice in both the planning time and in managing the treatment quality (see Subsection 3.2). However, one limitation of MCO-based planning is plan degradation from the segmentation process (McGarry et al., 2014, Rocha et al., 2012). The plan selected from the navigation process represents an intensity pattern that may not be efficiently delivered. In order to make the plan deliverable, the segmentation process finds a limited set of deliverable segment intensities that reproduce the intensity pattern. Because the number of segments allowed in a treatment session is limited, the intensity pattern is only approximately reproduced, thus the plan quality deteriorates. If the segmented plan becomes unsatisfactory, the planner will have to find another plan, resulting in an ineffective planning process. In this chapter, we show how this limitation of MCO-based planning can be overcome by applying the column generation RNBI method to the planning problem.

Column generation has been used to generate deliverable plans for single objective radiotherapy plan optimisation (Preciado-Walters et al., 2004, Romeijn et al.,

2005). The physical delivery constraints in the segmentation process are considered in the column generation subproblem. Essentially, each column generated from the subproblem represents a segment that is likely to improve the objective function value. As a result, the plans produced from column generation can be delivered without additional segmentation. By applying column generation within the RNBI framework, one produces a set of plans that are deliverable and each represents a certain trade-off in optimisation criteria. As will be demonstrated by our results, column generation RNBI produces plans that are near-optimal and can be delivered with dramatically lower monitor units than the corresponding optimal plans generated by the standard RNBI method followed by segmentation. Such plans are desirable as they can be delivered with a shorter treatment time, i.e., exposing patients to radiation for a shorter time, and with lower radiation leakage from the MLCs (Broderick et al., 2009). In fact, near-optimal plans that can be delivered efficiently and accurately are often preferable to complex optimal solutions in practice (Carlsson and Forsgren, 2014). The column generation RNBI approach therefore combines the advantage of considering multiple objectives in the FMO problem with the advantage of producing deliverable intensity maps without potential deterioration of treatment quality, which column generation delivers.

In this chapter, we explain the column generation reformulation of a FMO model in Section 6.1. Application of the method to a prostate radiotherapy planning problem is described in Section 6.2 and the results are shown in Section 6.3, followed by discussion and conclusion in Section 6.4.

6.1 Formulation

Consider a FMO problem with m voxels and n bixels with the dose calculated by $d = Ax$ where $d \in \mathbb{R}^m$ is the dose vector, $x \in \mathbb{R}^n$ is the radiation intensity vector and $A \in \mathbb{R}^{m \times n}$ is the dose deposition matrix. We denote $A_T \in \mathbb{R}^{m_T \times n}$, $A_C \in \mathbb{R}^{m_C \times n}$ and $A_N \in \mathbb{R}^{m_N \times n}$ as the dose deposition matrices for the PTV T with m_T voxels, critical organs C with m_C voxels and normal tissue N with m_N voxels, respectively.

The formulation used in this study is based on the model of Holder (2003), as shown in (6.1), but with a slight variation. The parameters include the dose lower bound $LB_T \in \mathbb{R}^{m_T}$ for the tumour and upper bounds $UB_T \in \mathbb{R}^{m_T}$, $UB_C \in \mathbb{R}^{m_C}$ and $UB_N \in \mathbb{R}^{m_N}$ for the tumour, critical organs and normal tissue, respectively. Variables $\alpha \in \mathbb{R}^{m_T}$, $\beta \in \mathbb{R}^{m_C}$ and $\gamma \in \mathbb{R}^{m_N}$ are voxel-wise one-sided dose deviations from tumour lower bound, critical organ upper bound and normal tissue upper bound, respectively.

$$
\begin{aligned}
\min \quad & \left(\frac{1}{m_T}e^T\alpha, \frac{1}{m_C}e^T\beta, \frac{1}{m_N}e^T\gamma\right) \\
\text{s.t.} \quad & LB_T - \alpha \leqq A_T x \leqq UB_T \\
& A_C x \leqq UB_C + \beta \\
& A_N x \leqq UB_N + \gamma \\
& \alpha \leqq UB_\alpha \\
& \beta \leqq UB_\beta \\
& \gamma \leqq UB_\gamma \\
& 0 \leqq x, \alpha, \beta, \gamma.
\end{aligned}
\tag{6.1}
$$

Different from Holder's model, we introduce upper bounds $UB_\alpha \in \mathbb{R}^{m_T}$, $UB_\beta \in \mathbb{R}^{m_C}$ and $UB_\gamma \in \mathbb{R}^{m_N}$ for α, β and γ, respectively. These upper bounds can easily be set so that Y is bounded and thus allows us to compute y^{AI}. Note that the unit of the objective functions is Gy, since we are trying to minimise dose deviations. The RNBI subproblem of (6.1) is simply (5.7) with constraints (5.7c) and (5.7d) replaced by the constraints of (6.1) and with the objective functions of (6.1) incorporated in the form of constraint (5.7b).

An optimal fluence map obtained by solving the RNBI subproblem of (6.1) with variables representing bixel intensities is not directly deliverable. Alternatively, deliverable plans can be generated from a reformulation in which the bixel intensity variables x are replaced by segment intensity variables $\bar{x}$ and the dose deposition matrix A based on bixel columns is replaced by the dose deposition matrix $\bar{A}$ using segment columns. Column s of $\bar{A}$, denoted as $\bar{a}^s \in \mathbb{R}^m$, represents the dose

deposited to the m voxels by segment s at unit intensity. Column $\bar{a}^s$ is derived by

$$\bar{a}^s = Au^s, \tag{6.2}$$

where $u^s \in \{0,1\}^n$ is a vector defining segment s with $u^s_j = 1$ if bixel j is open and $u^s_j = 0$ if bixel j is closed in segment s. Note that the two formulations have the same optimal values since feasible solutions in terms of variables x can be represented by variables $\bar{x}$ and vice versa, through the relationship

$$x = U\bar{x}, \tag{6.3}$$

with U being a matrix containing all feasible segment columns u.

Due to the large number of feasible segments, $\bar{A}$ has a much larger number of columns than A, which makes the reformulation hard to solve. Therefore it is beneficial to use column generation to solve the reformulation in which we only consider a subset of segments in the RMP. By solving the RMP, we obtain a vector of dual values π^*, which is passed to the subproblem to find a non-basic variable (representing the radiation intensity for a segment) with the most negative reduced cost. Here the subproblem is

$$\min\{-\pi^{*T} Au : u \in U\}, \tag{6.4}$$

where, as before, $u \in \{0,1\}^n$ is a vector defining a segment and U is the set of all feasible segment columns satisfying the MLC constraints. Note that the objective function coefficients of the segment intensity variables, with model (6.1) reformulated to the form of RNBISub, are zero, thus they are not included in subproblem (6.4). Let u^* be an optimal solution of (6.4). Given u^*, we can derive the dose deposition column $a^* = Au^*$ to be added to the RMP-RNBISub reformulation of (6.1).

In this study we consider the basic constraints of the MLC leaves, i.e., the leaf collision constraint, which prevent opposing leaves to overlap, and the constraint that ensures the opening for each row of collimator leaves is contiguous. Since we consider only these two constraints, all MLC rows are independent of one another. Therefore, (6.4) can be further decomposed by MLC row. For a given row, the

objective of the decomposed problem is to find the leaf positions that result in the lowest reduced cost for the MLC row. To decompose (6.4) by MLC row, one needs to change the index of each bixel to its corresponding beam index, MLC row position and MLC column position. Let $\tau = -\pi^{*^T} A$ and denote by $\tau_{b,r,c}$ the objective function coefficient of (6.4) corresponding to the bixel at beam b, MLC row position r and MLC column position c. We assume that each MLC row consists of t bixels, with MLC column positions indexed incrementally from left to right. Denote by t_1 the column index of the right-most bixel covered by the left leaf and by t_2 the column index of the left-most bixel covered by the right leaf. Then the decomposed subproblem (6.4) for beam b and MLC row r is

$$\begin{aligned} \min \quad & \sum_{c=t_1+1}^{t_2-1} \tau_{b,r,c} \\ \text{s.t.} \quad & t_1 \leqq t_2 - 1 \\ & t_1 \in \{0, \ldots, t\} \\ & t_2 \in \{1, \ldots, t+1\}. \end{aligned} \tag{6.5}$$

This decomposed problem is then solved by the algorithm described in Section 3.1.1 of Romeijn et al. (2005). It is important to note, that the non-dominated set of the original formulation of (6.1) and its column generation reformulation are identical.

6.2 The test case

We apply both the original RNBI method and the column generation RNBI method to a prostate radiotherapy treatment design problem. The RNBI method solves the RNBISub reformulation of (6.1) with bixel intensity variables x, producing a set of (not necessarily deliverable) intensity patterns that define a representatve set of non-dominated points for MOLP (6.1). The column generation RNBI method solves the RMP-RNBISub reformulation of (6.1) with a subset of segment intensity variables $\bar{x}$. We consider three objective functions: one for the PTV (objective 1), one for the rectum (objective 2) and one for the bladder (objective 3). Other clinically relevant structures such as the prostate, the right and left femur head

and normal tissues, are involved in the formulation as constraints, e.g., voxels of the prostate are given a lower bound and an upper bound on the delivered dose and voxels of femur heads and normal tissues are given structure specific upper bounds. By only involving three objective functions, we are able to illustrate the results graphically.

The dose deposition matrix A consists of 593 columns (corresponding to bixels) for 11 equi-spaced coplanar beam angles and 20000 rows (corresponding to voxels). Both methods use the set of reference points generated by the standard RNBI method. Therefore, we are able to identify feasible reference points, and we apply the column generation RNBI method only to feasible reference points. The column generation process terminates when any of the following termination conditions is satisfied:

- no variable with a negative reduced cost can be found,
- the number of segments assigned with a positive intensity in a solution exceeds 100, or
- the number of column generation iterations (or equivalently the number of segments) exceeds 150.

We implement the Phase-1 approach, the Big-M approach and Farkas pricing to handle infeasibility of RMP-RNBISub. The initialisation phase stops when the feasibility of RMP-RNBISub without artificial variables is guaranteed or when the termination condition is reached. The column generation model starts with only one coefficient column representing fully closed MLC segments (i.e., $u = 0$) in all beam directions. If a segment is assigned a positive intensity, we refer to this segment as a "positive segment". Positive segments are those segments that will be delivered in a treatment plan. This is in contrast to the zero-intensity segments, which will not be part of the treatment plan. We limit the number of positive segments in each plan so that the plans can be delivered in a reasonable treatment time. Additionally, we group plans with similar number of positive segments together to allow fair comparison, since radiotherapy treatment quality is influenced by the number of segments involved in a treatment. Specifically, a solution is separately recorded when the number of positive segments in the

solutions first exceeds 40, 50, 60, 70, 80, 90 and 100. These solutions are grouped into representative sets according to the number of positive segments.

6.3 Results

For convenience, representative points generated using RNBI and column generation RNBI will be referred to as RNBI points and the CG-RNBI points, respectively. The representative sets of the CG-RNBI points, grouped according to the number of positive segments, will be denoted as CG-*number* with *number* being the corresponding number of positive segments.

Using the standard RNBI method with increment $\eta = 0.08$ (see Section 5.2.1) or a distance of 3.2153 Gy between closest reference points, we identify that 17 of 91 reference points are feasible. Figure 6.1 illustrates the reference points and the RNBI points. We then initialise the column generation RNBI subproblems with the Phase-1 approach, the Big-M approach or Farkas pricing, followed by RMP-RNBISub when feasibility of RMP-RNBISub is guaranteed.

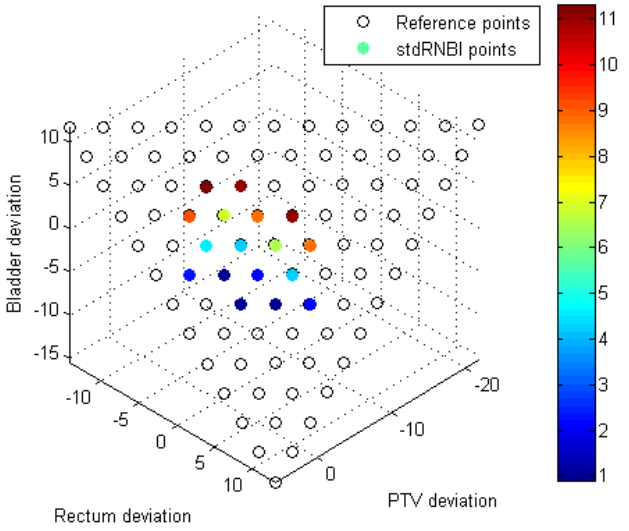

Figure 6.1: Illustration of the reference points and standard RNBI intersection points. The colour of the intersection points indicates the value of bladder deviation.

Figure 6.2 shows RMP-RNBISub objective function values versus column generation iteration for the first 4 of the 17 reference points after the initialisation stage. In each column generation iteration, one newly generated column, which represents a segment, is added to RMP-RNBISub. The red dashed line, blue solid line and green dash-dot line indicate the objective function values of the corresponding RMP-RNBISub problem initialised with the Big-M approach, the Phase-1 approach and Farkas pricing, respectively. The starting point of these lines indicates the number of iterations used to reach RMP-RNBISub feasibility using the three approaches. The dark line parallel to the horizontal axis represents the objective function value obtained using the standard RNBI method for the same reference point. The results show that, for all 17 cases, an initialisation using the Big-M approach reaches feasibility with the same or a smaller number of iterations compared to an initialisation with the other two approaches. In addition, during the early stages of the column generation process, initialisation with the Big-M approach generally produces a lower objective function value compared to initialisations with the other two approaches. The results indicate that the Big-M approach is superior to the other two approaches in terms of identifying columns that contribute to the objective function value. However, the difference in objective function values among different initialisation strategy diminishes as more columns are added to the model.

The next results are based on the solutions obtained from runs that employ an initialisation with the Big-M approach. Figure 6.3 shows the RNBI points (solid circle) and the points of CG-40 (asterisk) and CG-100 (empty circle). The figure illustrates how intersection points gradually move toward the boundary of the feasible set in objective space during the column generation process. Notice that in each of the column generation representative sets, there can be points dominated by other points. Our results show that the number of dominated points in a representative set ranges from one point in CG-60 to five points in CG-90.

Our results show that CG-RNBI points are close to the corresponding optimal RNBI points. The objective function values and the average computation time of the RNBI points, CG-40 and CG-100 are shown in Table 6.1. Recall that these points are obtained by solving the RNBI reformulation (RNBISub) and CG-RNBI reformulation (5.7) of Holder's model (6.1). The objective function values of the

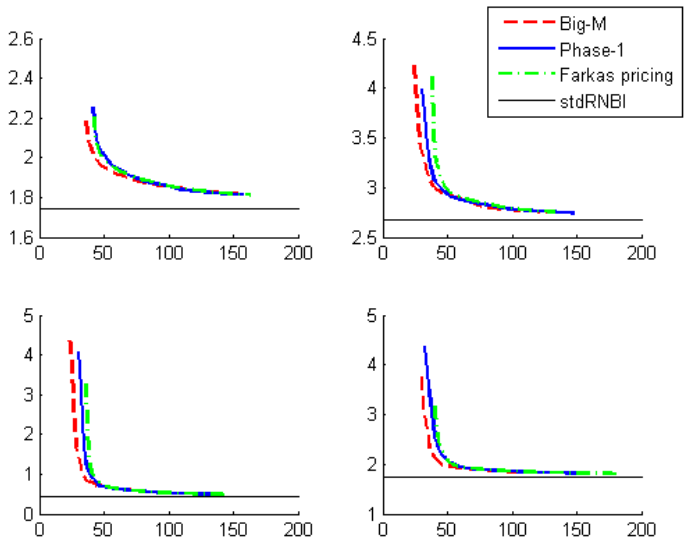

Figure 6.2: RMP-RNBISub objective function values (vertical axis) versus column generation iteration (horizontal axis) for the first four reference points.

points in CG-40 and CG-100 are on average 0.3647 and 0.1264 Gy higher than the objective function values of the RNBI points. In agreement with the remark in Lübbecke (2010) that column generation is in general not a competitive technique in solving linear programmes, we observe that computation times using column generation are longer. The average computation time used to obtain the RNBI points, CG-40 and CG-100 are approximately 17, 43 and 553 seconds, respectively. In addition, we observe that, as the number of generated columns increases, the computation time required for solving RMP-RNBISub increases as well.

Next we compare the plan complexity, in terms of monitor units, of plans generated by the RNBI and the CG-RNBI method. The segmentation algorithm by Engel (2005) is used to derive the segments and the monitor units needed for delivering the optimal intensity patterns generated by the standard RNBI method. The results, as included in Table 6.2, show that the monitor units of the segmented pattern (an average of 408.8 MU) is much higher than that of the intensity patterns generated by column generation (an average of 76.7 MU and 123.8 MU for CG-40 and CG-100, respectively).

We use (5.8) to measure the uniformity of the representative sets. The results show

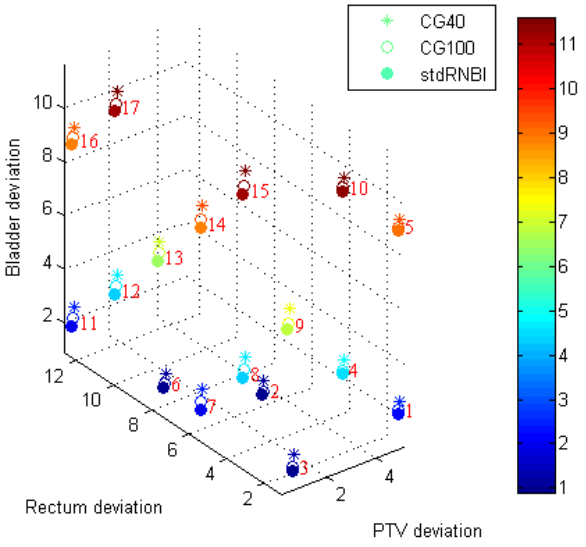

Figure 6.3: The RNBI points (solid circle) and the points in CG-40 (asterisk) and CG-100 (empty circle). Colour indicates the value of bladder deviation.

that the uniformity levels for all representative sets are the same up to 4 decimal places, with a value of 3.2153 Gy, which is the same as the distance between any two closest reference points. However, the two closest intersection points that define the uniformity level are different for the different representative sets.

Next, we apply the reference point bounding method described in Section 5.3.1 to the column generation RNBI method. Column generation is typically used for problems that cannot be practically solved by standard linear optimisation algorithms due to a large number of variables. When applying column generation RNBI to these problems, one also needs to solve the reference point bounding problem with column generation. Therefore, to assess how reference point bounding can be affected by column generation, reference point bounding is firstly solved to optimality (with the bixel intensity based formulation) and then solved using column generation (with the segment intensity based formulation). The column generation process terminates when the number of positive segments exceeds 40. The objective function values for reference point bounding are shown in Table 6.3. We can see that the minima and maxima produced using column generation are close to the corresponding minima and maxima solved to optimality, with a maximum absolute difference of 0.4494. In fact, using either set of minima and

Table 6.1: Objective values and average computation time (rounded to seconds) of RNBISub and CG-RNBISub with 40 positive segments (CG-40) and 100 positive segments (CG-100).

Reference point	Standard	CG-40	CG-100
1	1.7450	1.9880	1.8140
2	2.6806	2.9661	2.7525
3	0.4216	0.7765	0.5019
4	1.7451	2.0407	1.8215
5	4.0186	4.2616	4.0875
6	2.6941	2.9738	2.7695
7	1.5498	1.9892	1.7157
8	1.5417	1.9638	1.7104
9	1.8497	2.2986	1.9810
10	4.0186	4.2959	4.0916
11	3.8230	4.2212	3.9890
12	3.8108	4.1933	3.9770
13	3.8108	4.2176	3.9943
14	3.8112	4.2856	3.9865
15	3.8153	4.3023	3.9846
16	6.0843	6.4450	6.2286
17	6.0843	6.4846	6.2481
Average time	17	43	553

maxima, we are able to eliminate 67 out of 91 (73.63%) reference points. Out of the remaining 24 reference points, only 7 lead to RNBISub infeasibility.

We also test the performance of Farkas pricing in concluding the infeasibility of RNBISub instances. Note that we have 91 reference points in total, with 74 leading to RNBISub infeasibility. Table 6.4 shows that Farkas pricing is capable of concluding the infeasibility of 66 out of 74 RNBISub instances with 10 or fewer iterations. The average computation time for these 66 instances is 0.4 seconds. However, Farkas pricing is not capable of concluding infeasibility within 150 iterations for the remaining 8 instances. The computation time for each of these 8 instances ranges from 1179 to 7084 seconds, with an average of 4396 seconds and a standard deviation of 1676 seconds. For comparison, we apply column generation with the big-M initialisation to 10 reference points leading to RNBISub infeasibility. With a termination condition of 150 column generation iterations, the average computation time for solving each of the 10 reference points is 1213 seconds, with

Table 6.2: Total monitor units of intensity patterns generated by RNBI and CG-RNBI with 40 positive segments (CG-40) and 100 positive segments (CG-100).

Reference point	Standard	CG-40	CG-100
1	494	97.8	155.4
2	474	82.2	146.6
3	444	79.1	138.0
4	493	97.2	152.3
5	494	97.8	155.4
6	469	72.9	136.8
7	341	66.4	119.7
8	354	70.0	116.3
9	446	79.7	135.6
10	493	87.1	144.8
11	343	65.8	103.6
12	350	67.2	94.6
13	350	60.8	106.5
14	351	68.9	95.9
15	354	77.9	108.0
16	350	65.7	98.8
17	350	66.7	96.3
Average MU	408.8	76.7	123.8

Table 6.3: Minimum (min.) and maximum (max.) value for each objective based on the reference point bounding solved to optimality and solved by column generation (CG).

	Optimal min.	Optimal max.	CG min.	CG max.
Objective 1	-7.5352	3.9529	-7.2641	3.5035
Objective 2	-4.4038	9.159	-4.0101	8.9119
Objective 3	-4.5379	9.1226	-4.2902	8.8587

a standard deviation of 71 seconds. The results suggest that Farkas pricing can potentially be quite time consuming. Thus, if Farkas pricing cannot identify the infeasibility of a RNBISub instance in a small number of iterations, it would be beneficial to change the initialisation method to another approach.

Table 6.4: Number of iterations required for Farkas pricing to identify RNBISub infeasibility.

Iterations used to identify infeasibility	2	3	4	5	10	>150
Number of RNBISub instances	48	10	1	6	1	8

6.4 Discussion and conclusion

In this study we apply the column generation RNBI method to an MOLP formulation of a radiotherapy treatment design problem, which can be solved by both the standard RNBI and the column generation RNBI method. Column generation allows us to use variables representing segment intensities, as opposed to bixel intensities which are used in the conventional formulation. Because the number of possible segment shapes is much greater than the number of bixels, there is not yet a practical approach to solve segment-based model to optimality. Instead, one must consider trade-offs between plan quality and plan complexity (i.e. number of segments). In this study, we demonstrate with a test case that the plans generated by column generation are near-optimal and can be delivered with dramatically lower monitor units than the corresponding optimal ones followed by segmentation. Although these plans require longer computational time than those generated with a bixel-based model, the resulting reduced delivery complexity is desirable in practice due to shorter treatment time and lower radiation leakage from the MLCs (hence better delivery accuracy) (Broderick et al., 2009, Carlsson and Forsgren, 2014). To adopt the method clinically, further investigation on the performance of column generation is needed. In particular, one should check the robustness of the method, i.e., if the benefits of using column generation technique (near-optimality with delivery efficiency) are consistently observed for different treatment cases.

Fredriksson and Bokrantz (2013) introduce a concept of non-dominance called the "n-aperture Pareto set", which is a set of efficient plans given that each plan is formed by only n segments. However, to our knowledge, there is no practical method available to generate the n-aperture Pareto set. The concept of the n-aperture Pareto set can be generalised to the n-column non-dominated set for problems solved by column generation. Further research is required to extend the column generation RNBI method to ensure n-column non-dominance. Another

topic for future research is the extension of the column generation RNBI method to nonlinear multi-objective optimisation problems. This will, for example, allow us to consider other formulations of the radiotherapy treatment design problem.

Chapter 7

Multi-objective navigation of external radiotherapy plans based on clinical criteria

In this chapter we consider the process of finding the most preferable plan from a set of plans generated by an a posteriori approach. As the number of treatment plans in the set can be large, it can be time consuming to examine the quality of every plan in the set. Instead, a navigation method that effectively guides the planner toward the most preferable plan is needed (Allmendinger et al., 2016). Essentially, a navigation method should allow the planner to move from one plan to another with more desirable planning trade-offs until the most preferable plan is found. Desirable planning trade-offs can be expressed as a navigation query which specifies desirable improvements in certain criteria while allowing other criteria to deteriorate.

The navigation methods proposed for radiotherapy treatment planning so far are based on solving interactive multi-objective optimisation problems on the convex hull of a set of efficient solutions generated by a sandwiching method (see Subsection 3.1.3). In these methods, navigation queries are specified in terms of optimisation criteria, i.e., the objective functions used in the optimisation. These criteria are convex functions so that optimisation algorithms can be used to find

an optimal solution efficiently. Convex criteria along with a convex feasible set preserve convexity, hence allowing continuous navigation of treatment trade-offs.

However, optimisation criteria are only surrogates used to generate treatment plans and may not necessarily be relevant to plan evaluation (Li et al., 2012). In this study, we also show that the optimisation criterion values may not correctly reflect the plan quality, thus can mislead a planner to select a "locally optimal" plan. Instead, we propose a navigation method that uses planner-specified clinical criteria, which are criteria used to evaluate plan quality in clinical practice, to specify a navigation query. While many clinical criteria are convex or can be reformulated as convex functions (see, e.g., Hoffmann et al. (2008), Romeijn et al. (2004)), some clinical criteria are non-convex and non-continuous, for example, the dose-volume parameters (Deasy, 1997, Llacer et al., 2003) and treatment time (where changes in beam orientation and/or the number of segments result in discrete changes in treatment time, see, e.g., Bortfeld and Schlegel (1993)). For a non-continuous set, continuous navigation, as implemented in existing navigation models, is not applicable.

Instead, in this study we consider navigation on a discrete set of deliverable plans of which the beam orientation, segment shapes and segment weights of each plan are specified. The clinical criterion values are extracted from the plans and form the underlying data for the navigation model. A few studies propose navigation/decision support methods with such a set of deliverable plans. Rosen et al. (2005) propose to fit the dose-volume parameters of a set of plans to linear functions. Navigation is conducted based on the linear functions that estimate how the DVHs change with regards to a navigation query. However, due to the use of fitted linear functions, i.e., by fitting the data with linear functions at the centre of the data spread, the dose-volume parameters are the most preferred "average" values instead of the actual most preferred values. Ehrgott and Winz (2008) propose a decision support system based on filtering out plans with inferior criterion values. To conduct navigation in this method, the planner needs to iteratively impose, relax or tighten filters in a trial-and-error manner, which can be a cumbersome process. Ripsman et al. (2015) propose a ranking based navigation system in which the planner assigns weights to the clinical criteria and rank the plans with the sum of the weighted criterion values. However, the change of weight vector required to

move from one efficient plan to another is unknown and can only be discovered through a trial-and-error process.

Our proposed navigation method is inspired by DEA. To allow intuitive navigation, clinical criteria are categorised into inputs or outputs where decreases in inputs and increases in outputs are considered favourable. This is in contrast to optimisation-criteria based navigation methods where all optimisation criteria are considered favourable with lower values although some of the corresponding clinical criteria may be considered favourable with higher values. In our navigation method, the planner sets aspiration values for the clinical criteria and an optimisation model is solved to identify a plan that best satisfies the aspiration values. These aspiration values are specified as soft constraints for the optimisation model, hence the planner can freely set the ideal clinical criterion values, without the risk of generating an infeasible model. The planner can then inspect the plan identified by the navigation method and iteratively adjust the aspiration values until the most preferable plan is identified.

The contribution of this study is two-fold. Firstly, we propose a DEA based navigation model that allows a decision maker to freely express his or her preferences in criterion values. Secondly, we propose a clinically oriented navigation practice that uses clinical criteria, instead of identifying the most preferable plan from a set using optimisation criteria. In Section 7.1, we introduce and interpret the proposed navigation model. In Section 7.2, we present an implementation of the proposed model which is applied to a prostate planning problem. Discussion and conclusion then follow in Section 7.3.

7.1 Method

Typically, given a set of DMUs, the decision maker would prefer DMUs that are characterised as efficient. A DMU is efficient if and only if none of its inputs or outputs can be improved without deteriorating some of its other inputs or outputs (Definition 2.8). Under the context of a set of DMUs, we consider efficiency as a relative measure. That is, a DMU is considered efficient if no other points in the production possibility set indicate that some of its inputs or outputs can be

improved without worsening some of its other inputs or outputs. In this study, we consider a production possibility set referred to as the FDH (see Section 2.2), which represents a set of DMUs assuming free disposability, i.e., each DMU can consume extra inputs while producing the same levels of outputs or produce less outputs with the same levels of inputs.

Consider a set of N DMUs that are evaluated by the decision maker based on a set of criteria. The criteria are categorised into I inputs and O outputs where inputs are considered as costs to produce outputs. The decision maker wants to identify the most preferable DMU out of the set. Given a large set, it can be impractical to examine each DMU from the set. Instead, the decision maker can propose a set of aspired criterion values (or aspiration values) and identify a DMU that best satisfies the aspiration values. In the following, we propose a model that allows the decision maker to accomplish such a task while at the same time ensuring the DMU identified by the model is efficient.

Let $r \in \mathbb{R}^I_>$ and $w \in \mathbb{R}^O_>$ represent the aspiration values for the inputs and outputs, respectively. In addition, let the criterion values of the N DMUs be grouped into the input matrix $R \in \mathbb{R}^{I\times N}_\geq$ and the output matrix $W \in \mathbb{R}^{O\times N}_\geq$ in which the ith columns of R and W contain the criterion values of the ith DMU. Once the decision maker specifies the aspiration values, the following model is solved to identify a DMU that best satisfies $(r^T, w^T)^T$:

$$\begin{aligned}
\max \quad & \beta + \epsilon \left(e^T s^- + e^T s^+ \right) && \text{(7.1a)}\\
\text{s.t.} \quad & R\lambda \leqq (1-\beta)r - s^- && \text{(7.1b)}\\
& W\lambda \geqq (1+\beta)w + s^+ && \text{(7.1c)}\\
& e^T\lambda = 1 && \text{(7.1d)}\\
& \lambda \in \{0,1\}^N && \text{(7.1e)}\\
& s^+, s^- \geqq 0, && \text{(7.1f)}
\end{aligned}$$

where $\beta \in \mathbb{R}$, $\lambda \in \{0,1\}^N$, $s^- \in \mathbb{R}^I$, $s^+ \in \mathbb{R}^O$ are the decision variables. e represents a vector of ones with an appropriate dimension and $\epsilon > 0$ is a very small number. In practice, there is no need to determine the value of ϵ. Instead, model (7.1) can be implemented using a 2-step lexicographic optimisation, as described in Subsection 7.1.1. Since the production possibility set of model (7.1) is defined by the FDH

of the set of DMUs, the efficient DMU identified by the model must be one of the existing DMUs.

Proposition 7.1. *Model* (7.1) *is always feasible.*

Proof. Let $j^* \in \{1, \ldots, N\}$ and let $\lambda^{j^*} \in \mathbb{R}^N$ be a vector with one in the j^*th entry and zeros in the other entries. Clearly λ^{j^*} satisfies constraints (7.1d) and (7.1e). By setting $\lambda = \lambda^{j^*}$ and re-arranging the variables, constraints (7.1b) and (7.1c) can be equivalently expressed as

$$\begin{aligned} \beta &\leq \frac{R_{i,j^*} - r_i + s_i^-}{-r_i}, \quad i = 1, \ldots, I \\ \beta &\leq \frac{W_{o,j^*} - w_o - s_o^+}{w_o}, \quad o = 1, \ldots, O. \end{aligned} \tag{7.2}$$

Let s^{-*} and s^{-*} be zero vectors of dimension I and dimension O, respectively and let $\beta_{j^*} := \min \left\{ \frac{R_{i,j^*} - r_i}{-r_i}, \frac{W_{o,j^*} - w_o}{w_o} : i = 1, \ldots, I, o = 1, \ldots, O \right\}$. It is obvious that $\beta = \beta_{j^*}, s^- = s^{-*}, s^+ = s^{+*}$ satisfies (7.2) and thus $(\beta_{j^*}, \lambda^{j^*}, s^{-*}, s^{-*})$ is a feasible solution of (7.1). □

The model can be interpreted as a mechanism that adjusts the aspiration values to identify a best matching DMU. The adjusted aspiration values are expressed as the right-hand side of constraint (7.1b) and (7.1c), i.e., $(1-\beta)r - s^-$ and $(1+\beta)w + s^+$. Constraints (7.1d) and (7.1e) specify that only one element in vector λ can have a value of 1 with the rest of the elements having a value of zero. Thus, essentially, the model specifies that there must be one DMU with inputs and outputs better than or equal to the adjusted aspiration values.

Two types of adjustments are used in the model: directional and additive. Directional adjustment is inspired by Chambers et al. (1996, 1998) in which the directional distance function $R\lambda \leq (1-\beta)\delta r$ and $W\lambda \geq (1+\beta)\delta w$ allows the decision maker to specify a direction $(-\delta r^T, \delta w^T)^T$ for the adjustment of the aspiration values. In model (7.1) the direction is specified as $(-r^T, w^T)^T$ and β represents the magnitude of adjustment in the specified direction. As model 7.1 is always feasible

(Proposition 7.1), the decision maker can choose the aspiration values freely even if the values are out of the production possibility set. The model simply deteriorates the aspiration values with a negative β value, which is penalised in the objective function.

The other adjustment mechanism, additive adjustment, adjusts the individual aspiration values through s^- and s^+. In contrast to directional adjustment, which adjusts all criteria simultaneously, additive adjustment applies to each input/output independently and the adjustment only improves the aspiration values. Since the cost coefficient for additive adjustment, ϵ, is a very small value, additive adjustment is applied to the aspiration values with a lower priority than directional adjustment. Hence, the adjustment of aspiration values in model (7.1) can be considered as a two step process where directional adjustment is conducted first, followed by additive adjustment. Directional adjustment moves a point to the production frontier where at least one input or one output cannot be further improved. The point on the boundary is then adjusted by individual input/output through additive adjustment of s^+ and s^-. Additive adjustment ensures that the adjusted aspiration values equal the criterion values of one of the DMUs (since s^+ and s^- are maximised in the objective function, thus making constraints (7.1b) and (7.1c) binding). In addition, the DMU corresponding to the adjusted aspiration values must be efficient (Proposition 7.2).

Proposition 7.2. *The DMU selected by the optimal solution of model* (7.1) *must be efficient.*

Proof. If the selected DMU is not efficient, there must be another DMU from the set with better input/output values. Thus the adjusted aspiration values can be further improved, resulting in a better objective function value, which is contrary to the optimality of the solution. □

Figure 7.1 illustrates the use of model (7.1), in which a two-criteria problem with one input and one output is considered. In the figure, points A, B, C and D represent a set of efficient DMUs and points a, b and c each represents a set of aspiration values specified by the decision maker. Directional adjustment of aspiration values moves points a and b to points a^* and b^*. Notice that directional adjustment is

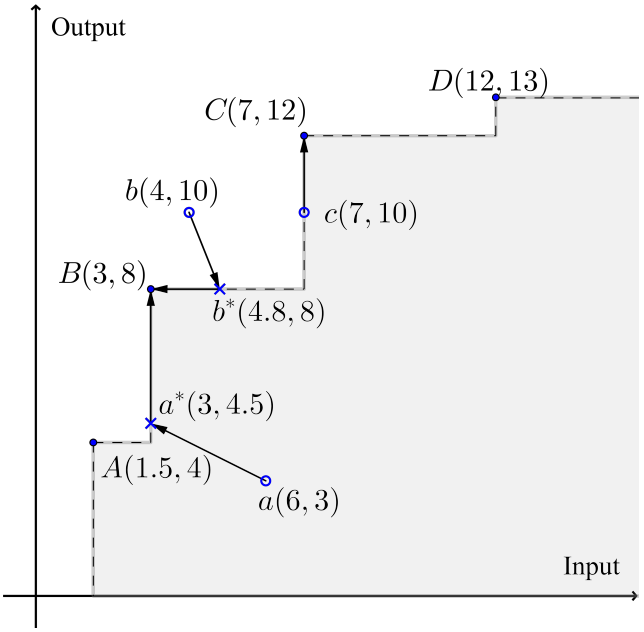

Figure 7.1: Illustration of model (7.1) with points A, B, C and D representing DMUs and points a, b and c representing aspiration values.

determined by the corresponding aspiration values. Directional adjustment cannot be applied to point c as the point is already on the production frontier. Additive adjustment is then applied to points a^*, b^* and c to identify the corresponding efficient DMUs. β, s^- and s^+ values for points a, b and c are shown in Table 7.1.

Table 7.1: β, s^- and s^+ values for the adjustment of points a, b and c in Figure 7.1.

Point	$a(6,3)$	$b(4,10)$	$c(7,10)$
β	0.5	-0.2	0
(s^-, s^+)	(0, 3.5)	(1.8, 0)	(0, 2)
Identified DMU	B	B	C

Directional adjustment, which can be specified as $\beta(-r^T, w^T)^T$, allows the multiple inputs and the multiple outputs of the aspiration values to remain in the same proportion after adjustment. This is illustrated in Figure 7.2a where a two-criteria problem with only inputs is considered. In the figure, directional adjustment only moves aspiration points d and e along the corresponding ray emanating from the origin (shown in the dotted line) where any point on the ray has the same proportion in inputs as the corresponding aspiration point. Hence model (7.1) can be

interpreted as a two step mechanism. The first step finds a point on the production frontier such that the point has criterion values in the same proportion among inputs and among outputs as the aspiration point. The second step then improves the adjusted aspiration point such that it equals one of the efficient DMUs.

With a slight modification, our navigation model can also be applied to a convex set where convex combinations of a set of DMUs are considered attainable. This convexity property is used in other navigation methods where a convex combination of plans (or precisely, bixel intensities) is also a valid plan. To apply (7.1) to a convex set, one simply replaces the binary constraint of λ, (7.1e), to $0 \leq \lambda_i \leq 1 \ \forall i = 1, \ldots, n$. When (7.1) is applied to a convex set, the majority of aspiration values would be adjusted to the efficient set through directional adjustment (and hence no additive adjustment is required) since the production frontier is mainly composed of the non-dominated set. This is illustrated in Figure 7.2b where aspiration values represented by points d and e are adjusted to efficient points d^* and e^* while maintaining the same proportion in inputs. Additive adjustment is needed only when a set of aspiration values is projected to the weakly non-dominated set of the production frontier.

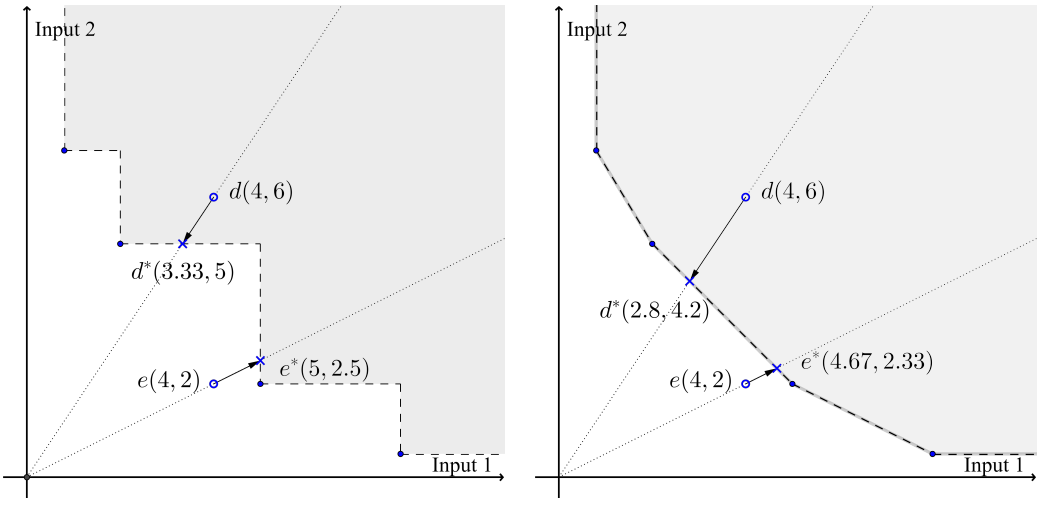

(a) (b)

Figure 7.2: Illustration of proportional adjustment with a two-input problem (a) with the assumption of FDH (b) with the assumption of convexity.

7.1.1 Implementation

In practice, it can be non-trivial to determine the value of ϵ. Instead of solving model (7.1) directly, one can solve models (7.3) and (7.4) sequentially to get the same optimal solution without defining ϵ. Model (7.3) maximises β solely, without considering s^- and s^+, to obtain the optimal β value β^*. Model (7.4) maximises $e^T s^- + e^T s^+$ with the β value in the constraints equal to β^*.

$$\begin{aligned} \beta^* := \max\, & \beta \\ \text{s.t.} \quad & R\lambda \leqq (1-\beta)r \\ & W\lambda \geqq (1+\beta)w \\ & e^T\lambda = 1 \\ & \lambda \in \{0,1\}^n. \end{aligned} \tag{7.3}$$

$$\begin{aligned} \max \quad & e^T s^- + e^T s^+ \\ \text{s.t.} \quad & R\lambda \leqq (1-\beta^*)r - s^- \\ & W\lambda \geqq (1+\beta^*)w + s^+ \\ & e^T\lambda = 1 \\ & \lambda \in \{0,1\}^n \\ & s^+, s^- \geqq 0. \end{aligned} \tag{7.4}$$

Proposition 7.3. *Models (7.3) and (7.4) are always feasible.*

Proof. Let $j^* \in \{1,\ldots,N\}$ and let $\lambda^{j^*} \in \mathbb{R}^N$ be a vector with one in the j^*th entry and zeros in the other entries and let

$$\beta_{j^*} := \min\left\{\frac{R_{i,j^*} - r_i}{-r_i}, \frac{W_{o,j^*} - w_o}{w_o} : i = 1,\ldots,I, o = 1,\ldots,O\right\}.$$

Following the proof of Proposition 7.1, it is easy to see that λ^{j^*} and β_{j^*} is a feasible solution for model (7.3). Consequently, given β^* and λ^* that satisfy constraints in model (7.3), constraints in model (7.4) will always be satisfied by s^+ and s^- equal to zero vectors. □

7.1.2 Navigation mechanism with hard constraints

While the use of soft constraints in model (7.1) allows a decision maker to freely express the aspired criterion values, soft constraints do not allow exploration of alternative DMUs. In practice, changing the aspiration values may result in identifying the same DMU and the changes in (one or more) aspiration values required to identify a neighbouring DMU are not trivial. Instead, alternative DMUs can be accessed through the use of hard constraints. In this subsection, we illustrate the use of hard constraints in navigation. For convenience, we express the hard constraints for inputs only since extending the constraints to outputs is trivial.

Let the ith DMU be the DMU currently identified by the navigation system and let $e_j \in \mathbb{R}^I$ be a vector with one in the jth entry and zeros in the other entries. Constraint (7.5a) or (7.5b) can be added to model (7.1) in order to move from the current DMU to an alternative DMU where the kth input criterion is improved by ϵ_k or worsened by ϵ_k, respectively.

$$e_k^T R\lambda \leqq R_{k,i} - \epsilon_k \tag{7.5a}$$

$$e_k^T R\lambda \geqq R_{k,i} + \epsilon_k. \tag{7.5b}$$

In our implementation (see Subsection 7.2.2), ϵ_k is the step size of the criterion slider bar defined as 1% of the value range of the kth input criterion, i.e.,

$$\epsilon_k = 0.01 \times (\max\{R_{k,j} : j = \{1, \ldots, I\}\} - \min\{R_{k,j} : j = \{1, \ldots, I\}\}).$$

While navigating, the decision maker may only be interested in DMUs with a certain set of input criteria, $\mathcal{L}$, better than or equal to certain values. This can be achieved by constraints of the form

$$e_l^T R\lambda \leqq b_l \quad \forall l \in \mathcal{L}, \tag{7.6}$$

where vector b records the upper bounding values for input criteria $l \in \mathcal{L}$. The decision maker may also wish to combine the two navigation mechanisms together

so that neighbouring DMUs with input criteria $l \in \mathcal{L}$ better than or equal to certain values can be found.

Model (7.1) with constraints (7.5) and (7.6) forms the complete navigation model for this study. Notice that when hard constraints are imposed, soft constraints (7.1b) and (7.1c) still apply. That is, the navigation model still relies on the directional adjustment and the additive adjustment to identify an efficient DMU, although now the production possibility set is restricted by the hard constraints. In other words, the navigation mechanism still finds a DMU that best matches the aspiration values among DMUs that satisfy the hard constraints. Notice that the navigation model can be infeasible when one or more hard constraints exclude all DMUs from the feasible set. When infeasibility occurs, the decision maker will have to accept that the hard constraint(s) cannot be satisfied and should consider weakening the constraint(s) in order to find a solution. In our implementation, the hard constraints can only be imposed one at a time (by pressing a slider bar arrow button or by clicking a constraint check box) and the navigation model is solved when a new hard constraint is added. Hence if a new constraint results in an infeasible model, the decision maker will be immediately notified.

7.2 Application to radiotherapy treatment planning

We apply the proposed navigation method to the set of prostate radiotherapy treatment plans generated by the column generation RNBI method (as described in Chapter 6). The clinical criteria used in this case study are shown in Subsection 7.2.1. Subsection 7.2.2 presents the graphical user interface of the proposed navigation system and demonstrates the functionality of the system.

7.2.1 Navigation criteria as inputs and outputs

The navigation criteria, in terms of inputs and outputs, used for the navigation process are shown in Table 7.2. Left femur head and right femur head are abbreviated as LFH and RFH. For dose conformity, we use the CI proposed by Lomax and Scheib (2003)

$$\text{CI} = \frac{V_{(T,d)}}{V_d}, \tag{7.7}$$

where $V_{(T,d)}$ is the d Gy isodose volume of the target and V_d is the d Gy isodose volume of the total patient volume. The CI value is considered as an output as higher values of CI are preferable (with a maximum value of 1). For target dose homogeneity, we use the HI proposed by Yoon (2007)

$$\text{HI} = \sqrt{\sum_i (D_i - D_{mean})^2 \times \frac{v_i}{V}}, \tag{7.8}$$

where D_i and v_i are the dose and volume of voxel i of the target, respectively. D_{mean} is the mean dose and V is the volume of the target. Essentially, (7.8) measures the standard deviation of dose in the target volume. Hence HI is considered as an input where smaller dose deviation in the target is considered preferable.

Table 7.2: Navigation criteria used in the planning problem in terms of inputs and outputs.

Inputs	Outputs
PTV HI	PTV D_{95}
Rectum gEUD	PTV CI
Rectum D_5	
Bladder D_{50}	
Bladder D_{25}	
LFH D_{10}	
RFH D_{10}	
Number of segments	

It can be seen from Figure 7.3 that optimisation criteria may not correctly reflect the quality of clinical criteria. Ideally, a clinical criterion value, if plotted against the corresponding optimisation criterion value, should be either strictly increasing or strictly decreasing. That is, an improvement in the optimisation criterion value

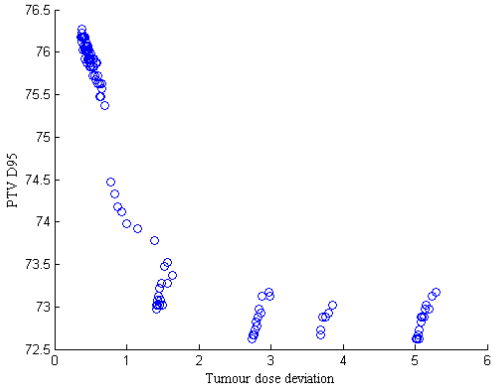

(a) PTV D_{95} against Tumour dose deviation.

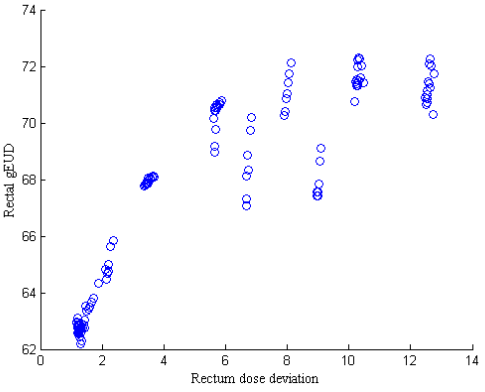

(b) Rectum gEUD against rectum dose deviation.

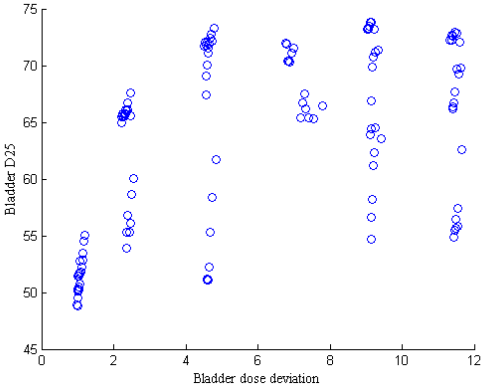

(c) Bladder D_{25} against bladder dose deviation.

Figure 7.3: Plots showing that optimisation criterion values may not be appropriate quality indicator for clinical criteria.

should always lead to an improvement in the clinical criterion value. However, Figures 7.3a to 7.3c show that this is not the case. Thus when conducting navigation with optimisation criteria, it is possible for a planner to conclude optimality of a plan prematurely when a "locally optimal" plan, where further improvement in optimisation criteria does not lead to improvements in clinical criteria, is identified.

7.2.2 Graphical user interface

A screen shot of the navigation system is shown in Figure 7.4. Three main components are used to conduct navigation: criterion slider, aspiration slider and constraint check box. The value range of each slider is limited by the maximum and minimum of the corresponding criterion in the database. The user can use either the slider or the text box below each slider to change the corresponding criterion or aspiration value. When any of the aspiration values change, model (7.1) is updated accordingly and solved to obtain the best plan that satisfies the aspiration values. When the value of a criterion slider is changed, a hard constraint of the form of equation (7.5a) or (7.5b) is imposed on model (7.1) and the model is solved to obtain a plan that best satisfies the change of criterion value with respect to current aspiration values. If a current plan is replaced by an alternative plan, the values of criterion sliders are changed to the values of the alternative plan. When a constraint check box is activated, hard constraints of the form of (7.6) are added into model (7.1). If an aspiration value is/is not satisfied by the corresponding criterion value, the background colour of the criterion text box appears green/red. Similarly, the "Feasibility" text box shown on the top-left corner indicates if a navigation step is feasible (with green background colour) or not (with red background colour).

The navigation system is implemented in MATLAB and the optimisation models are solved with Gurobi. Clinical criterion values are extracted with CERR (Deasy et al., 2003). As the optimisation problem size is small, the optimisation is solved in real time, hence allowing immediate interaction between the planner and the navigation system.

7.2.3 Demonstration of the navigation system

To demonstrate the use of the navigation system, we show a sequence of navigation steps. The navigation steps are composed of the following actions:

a Enter aspiration values

b Click on the right arrow of PTV D_{95} criterion value slider bar

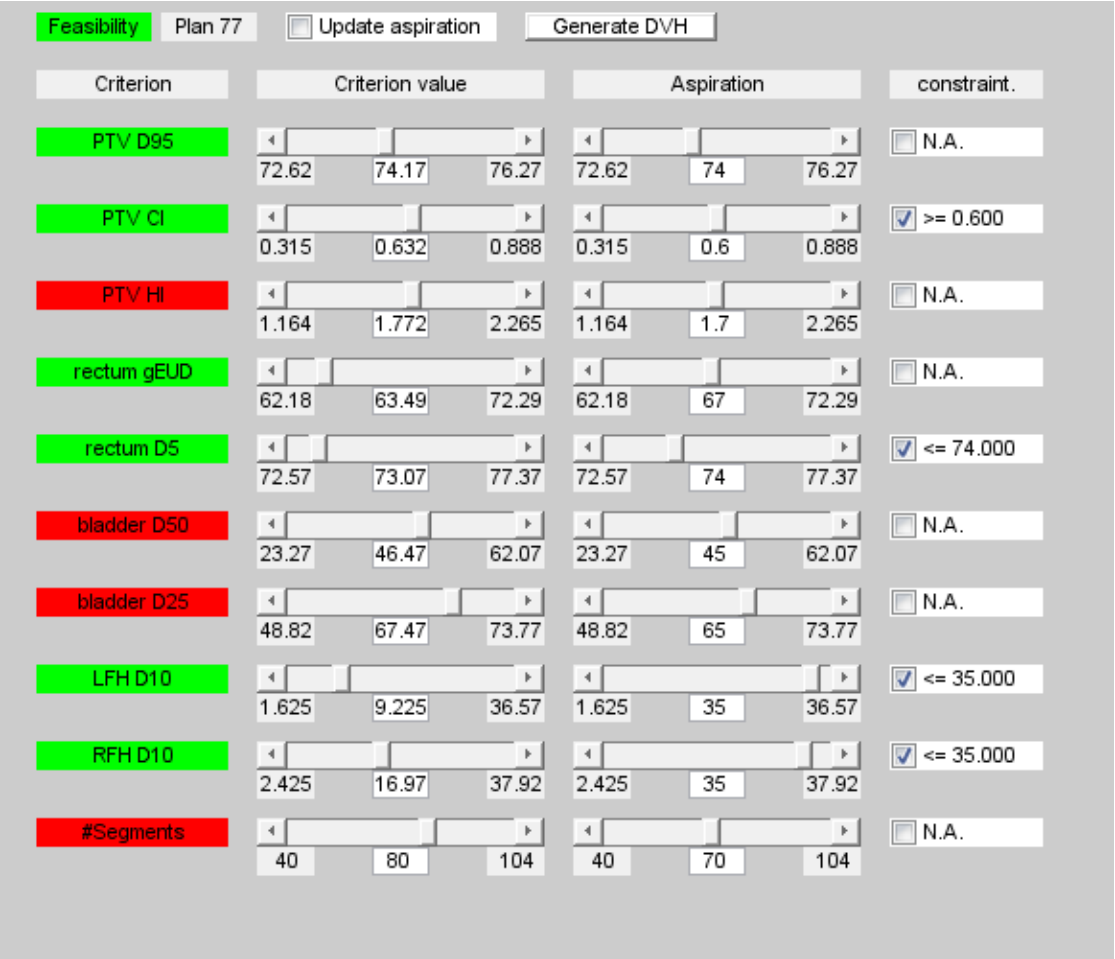

Figure 7.4: Screen shot of the navigation system showing plan 77 with a set of aspiration values and constraints.

c Click on rectal D_5 constraint check box

d Click on bladder D_{25} constraint check box

The action(s) performed and the plan identified in each step are shown in Table 7.3. Action(s) in brackets represent action(s) imposed from previous steps. The aspiration values used and the criterion values of the plans identified by the navigation system from the steps are shown in Table 7.4.

In step 1, aspiration values for each criterion are set and the navigation system identifies plan 5 as the most preferable plan. All the criterion values of plan 5 succeed the corresponding aspiration values except for PTV D_{95}. Thus in step

Table 7.3: Action(s) in each of the navigation step and the identified plan.

Step	Action(s)	Plan ID
1	a	5
2	(a), b	66
3	(a), c, d	5
4	(a, c, d), b	infeasible
5	(a, c, d), d, b	9
6	(a, c, b), b	26
7	(a, c, b, b), b	60

Table 7.4: Aspiration values used and plans identified by the navigation steps in Table 7.3. Criterion values worse than the corresponding aspiration values are shown in bold.

Criteria	Aspiration	Plan 5	Plan 66	Plan 9	Plan 26	Plan 60
PTV D_{95}	74	**73.18**	75.83	**73.78**	**73.93**	74.13
PTV CI	0.6	0.77	**0.59**	0.67	0.67	0.65
PTV HI	1.7	1.62	1.32	**1.80**	**1.80**	**1.81**
rectum gEUD	67	63.54	**68.05**	64.83	64.35	63.68
rectum D_5	74	73.03	**76.18**	73.63	73.28	73.13
bladder D_{50}	45	36.13	35.58	41.93	41.93	44.43
bladder D_{25}	65	63.53	**67.68**	**66.43**	**65.33**	**66.18**
LFH D_{10}	35	15.78	17.58	7.98	8.33	9.23
RFH D_{10}	35	28.63	8.23	15.68	18.68	19.73
#Segments	70	40	**72**	42	50	70

2, action b is used to find a plan with an improved PTV D_{95} value, which results in plan 66 being identified. Although plan 66 achieves the aspired PTV D_{95} value, many of the criterion values become worse than the corresponding aspiration values. In particular, we want to achieve the aspired values for rectum D_5 and bladder D_{25}. This can be specified by action c and d. It is not surprising that plan 5 is identified by step 3 since from step 1 we know plan 5 is the most preferable plan for the given aspiration values while at the same time satisfying the aspired rectum D_5 and bladder D_{25} values. In step 4, an attempt to find a plan with a better PTV D_{95} is conducted through action b, however, the navigation system returned infeasibility as the output. The infeasibility suggests that one cannot further improve the PTV D_{95} value without sacrifice the aspired value for rectum D_5 or bladder D_{25}. In step 5, we relax the constraint on bladder D_{25} and try to

improve PTV D_{95} through actions d and b, respectively. An improvement in PTV D_{95} is achieved, as shown by plan 9, although the aspired PTV D_{95} value is still not achieved. Further improvement in PTV D_{95} value is conducted in step 6 and 7, resulting in plan 26 and 60, respectively. Plan 60 shows that the aspired PTV D_{95} value is achieve at the cost of sacrificing the aspired values for inhomogeneity index and bladder D_{25}.

7.3 Discussion and conclusion

Previously proposed navigation methods are designed for commercial MCO-based planning in which the feasible set is the convex hull of a set of efficient fluence maps. By doing so, one can quickly generate a large number of feasible fluence maps using interpolation with modest computational expense. However, fluence maps generated by linear interpolation are not efficient and may be further improved (Bokrantz and Miettinen, 2015). The potential improvement in plan quality of these plans are sacrificed in order to reduce the computational expense. In addition, linear interpolation of fluence maps does not consider plan delivery. Thus, a plan selected from the approximated efficient set needs to go through the segmentation step, which can significantly deteriorate the plan quality (McGarry et al., 2014, Rocha et al., 2012). Furthermore, linear interpolation cannot be applied to different beam angle configurations, thus plans generated from interpolation are subjected to the same beam angle configuration. However, the best beam configuration to achieve different treatment trade-offs can be different (see, e.g., Cabrera et al. (2016)). Excluding the search of different beam configurations in the plan generation process prevents the identification of the best-quality plans. The overall effect of these limitations may lead to generation of plans with considerable improvement potential.

An alternative MCO practice is to generate a finely-sampled discrete representation of the efficient set in which each plan is generated to achieve a specific treatment trade-off. Using this approach, each plan is given freedom in beam angle configuration, segment shapes and segment intensities and hence the method allows one

to achieve the best-quality plans for different treatment trade-offs. Plan generation with this approach can be computationally expensive, especially for problems with many optimisation criteria. However, recent advancement of high performance computing, such as Jia et al. (2014), Tian et al. (2015) and Ziegenhein et al. (2013), alleviates this potential drawback. To further reduce the computational expense, one can also consider a two-stage planning practice where in the first stage, a coarser sample of the efficient set is generated and navigation is used to identify a close-to-ideal plan, followed by fine-tuning (see, e.g., Otto (2014), Ziegenhein et al. (2014)) of dose distribution in the second stage. The navigation method proposed in this study best facilitates this MCO practice where navigation is conducted on a discrete set of plans.

The proposed navigation method uses soft constraints to set the aspiration values, hence allows the decision maker to freely express his/her preferences without resulting in an infeasible model. This is in contrast to previously proposed navigation methods. In Monz et al. (2008), navigation queries are specified by hard constraints. To avoid specifying an infeasible navigation query, the maxima and minima of individual objective functions are calculated when the upper bound constraints of the individual objective functions change. Navigation queries are restricted to be within the individual maxima and minima. In Craft and Monz (2010) and Craft and Richter (2013), navigation queries are specified by changing the value of a reference point that lies on the convex hull of individual minima. The reference point projects in a direction to the non-dominated set to find a corresponding non-dominated point. However, as the projection of the convex hull of the individual minima does not cover the whole non-dominated set, as shown in Shao and Ehrgott (2007), some portions of the non-dominated set may be neglected during the navigation process. Our proposed model allows the decision maker to specify the desirable aspiration values freely and the model returns a DMU (or a solution in terms of MCO) that best matches the aspiration values. Consequently, this flexibility in specifying aspiration values allows us to apply the model to a discrete set, which previous navigation methods are not capable of.

In this study we apply the navigation model on a set of discrete points each representing a radiotherapy treatment plan with certain beam angle configuration,

segment shapes and segment intensities. As the set for navigation consists of discrete plans, one can extract any planner-defined clinical evaluation criteria from each plan, including the commonly used DV parameters, which are known to be non-convex. This is in contrast to previous navigation methods which use convex optimisation criteria for navigation. It is shown in this study that optimisation criteria can be misleading in determining plan quality and may lead to a "sub-optimal" plan during navigation. Instead, using clinical criteria enables the planner to assess plan quality and conduct navigation (including setting constraints) in a more intuitive manner. The knowledge derived from the navigation steps can then be used to form new preferences and the iterative process of navigation and forming new preference continues until the most preferable plan is identified.

Chapter 8

Conclusion

In this thesis, we first review planning approaches for managing treatment trade-offs between tumour control and normal tissue sparing. The a priori approaches generate plans based on pre-determined planning protocols. The a posteriori approaches first generate a set of plans featuring different treatment trade-offs, followed by a planner who selects the most preferable plan from the set. The interactive approaches involve a planner who iteratively adjusts treatment preferences and re-optimises the current plan.

We propose a quality control method based on DEA, which assesses the quality of a plan by comparing it to other reference plans and searches for evidence of improvement potential. The quality control method integrates well with other planning approaches by providing a means to reassure plan quality and thus improves the planning efficiency. Application of the method to prostate cases demonstrates the capability of DEA in correctly identifying plans with further improvement potential.

We integrate column generation in the RNBI framework to generate an approximated representative set of non-dominated points for an MOLP. Since RNBI subproblems may be infeasible, we attempt to detect this infeasibility early to reduce computational expense. First, a reference point bounding method is proposed to eliminate reference points that lead to infeasible RNBI subproblems. Furthermore,

different initialization approaches for column generation are implemented, including Farkas pricing. The quality of the representation obtained by the method is also investigated.

The column generation RNBI method is applied to an MOLP in radiotherapy treatment planning. In contrast to the commercial MCO-based planning method, which uses convex interpolation of fluence maps to generate plans, the column generation RNBI method produces a discrete set of plans with each plan featuring a certain treatment trade-offs. By doing so, each plan can have a unique beam configuration and segments, thus allowing one to fully utilise the treatment potential. As demonstrated by our results, column generation RNBI produces plans that are near-optimal and can be delivered with dramatically lower total fluence than the corresponding optimal plans generated by the standard RNBI method followed by segmentation. Such plans are desirable as they can be delivered with a shorter treatment time. We also note that reference point bounding dramatically reduces the number of RNBI subproblems that need to be solved.

Given a set of radiotherapy plans, one requires an effective method to search for the most preferable plan in the set. Previous navigation methods are conducted based on optimisation criteria. We show that optimisation criterion values may not accurately reflect plan quality, thus can mislead a planner to overlook other preferable plans. Instead, we propose to conduct navigation with planner-defined clinical criteria. The proposed navigation model specifies the aspiration values as soft constraints, thus enables the planner to freely express his or her preferences without leading to an infeasible model. Consequently, the soft constraints enable navigation among a discrete set, which previous navigation methods are not capable of.

Bibliography

A.-K. Ågren-Cronqvist. *Quantification of the response of heterogeneous tumours and organized normal tissues to fractionated radiotherapy.* PhD thesis, Stockholm University, 1995.

R. Allmendinger, M. Ehrgott, X. Gandibleux, M. J. Geiger, K. Klamroth, and M. Luque. Navigation in multiobjective optimization methods. *Journal of Multi-Criteria Decision Analysis*, 2016. To appear.

E. D. Andersen. Certificates of primal or dual infeasibility in linear programming. *Computational Optimization and Applications*, 20(2):171–183, 2001.

L. M. Appenzoller, J. M. Michalski, W. L. Thorstad, S. Mutic, and K. L. Moore. Predicting dose-volume histograms for organs-at-risk in IMRT planning. *Medical Physics*, 39(12):7446–7461, 2012.

D. Baatar, H. W. Hamacher, M. Ehrgott, and G. J. Woeginger. Decomposition of integer matrices and multileaf collimator sequencing. *Discrete Applied Mathematics*, 152(1):6–34, 2005.

R. D. Banker and R. C. Morey. Efficiency analysis for exogenously fixed inputs and outputs. *Operations Research*, 34(4):513–521, 1986.

R. D. Banker, A. Charnes, and W. W. Cooper. Some models for estimating technical and scale inefficiencies in data envelopment analysis. *Management Science*, 30(9):1078–1092, 1984.

R. S. Barr. DEA software tools and technology. In W. W. Cooper, L. M. Seiford, and J. Zhu, editors, *Handbook on Data Envelopment Analysis*, pages 539–566. Springer, Boston, MA, 2004.

R. Baskar, K. A. Lee, R. Yeo, and K.-W. Yeoh. Cancer and radiation therapy: Current advances and future directions. *International Journal of Medical Sciences*, 9(3):193–199, 2012.

J. L. Bedford. Treatment planning for volumetric modulated arc therapy. *Medical Physics*, 36(11):5128–5138, 2009.

H. Benson. Hybrid approach for solving multiple-objective linear programs in outcome space. *Journal of Optimization Theory and Applications*, 98(1):17–35, 1998a.

H. P. Benson. An outer approximation algorithm for generating all efficient extreme points in the outcome set of a multiple objective linear programming problem. *Journal of Global Optimization*, 13(1):1–24, 1998b.

H. P. Benson and S. Sayin. Towards finding global representations of the efficient set in multiple objective mathematical programming. *Naval Research Logistics*, 44(1):47–67, 1997.

R. Bokrantz. Multicriteria optimization for volumetric-modulated arc therapy by decomposition into a fluence-based relaxation and a segment weight-based restriction. *Medical Physics*, 39(11):6712–6725, 2012.

R. Bokrantz. Distributed approximation of pareto surfaces in multicriteria radiation therapy treatment planning. *Physics in Medicine and Biology*, 58(11): 3501–3516, 2013.

R. Bokrantz and A. Forsgren. An algorithm for approximating convex pareto surfaces based on dual techniques. *INFORMS Journal on Computing*, 25(2): 377–393, 2013.

R. Bokrantz and K. Miettinen. Projections onto the Pareto surface in multicriteria radiation therapy optimization. *Medical Physics*, 42(10):5862–5870, 2015.

T. Bortfeld. Optimized planning using physical objectives and constraints. *Seminars in Radiation Oncology*, 9(1):20–34, 1999.

T. Bortfeld and W. Schlegel. Optimization of beam orientations in radiation therapy: Some theoretical considerations. *Physics in Medicine and Biology*, 38(2): 291–304, 1993.

J. J. Boutilier, T. Lee, T. Craig, M. B. Sharpe, and T. C. Chan. Models for predicting objective function weights in prostate cancer IMRT. *Medical Physics*, 42(4):1586–1595, 2015.

A. Brahme. Optimization of stationary and moving beam radiation therapy techniques. *Radiotherapy and Oncology*, 12(2):129–140, 1988.

S. Breedveld, P. R. Storchi, M. Keijzer, and B. J. Heijmen. Fast, multiple optimizations of quadratic dose objective functions in IMRT. *Physics in Medicine and Biology*, 51(14):3569–3579, 2006.

S. Breedveld, P. R. Storchi, and B. J. Heijmen. The equivalence of multi-criteria methods for radiotherapy plan optimization. *Physics in Medicine and Biology*, 54(23):7199–7209, 2009.

S. Breedveld, P. R. Storchi, P. W. Voet, and B. J. Heijmen. iCycle: Integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans. *Medical Physics*, 39(2):951–963, 2012.

M. Broderick, M. Leech, and M. Coffey. Direct aperture optimization as a means of reducing the complexity of intensity modulated radiation therapy plans. *Radiation Oncology*, 4(8):1–7, 2009.

G. Cabrera, M. Ehrgott, A. J. Mason, and A. Raith. A matheuristic approach to solve the multiobjective beam angle optimization problem in intensity-modulated radiation therapy. *International Transactions in Operational Research*, 2016. doi: 10.1111/itor.12241.

F. Carlsson and A. Forsgren. On column generation approaches for approximate solutions of quadratic programs in intensity-modulated radiation therapy. *Annals of Operations Research*, 223(1):471–481, 2014.

Y. Censor and J. Unkelbach. From analytic inversion to contemporary IMRT optimization: Radiation therapy planning revisited from a mathematical perspective. *Physica Medica*, 28(2):109–118, 2012.

R. G. Chambers, Y. Chung, and R. Färe. Benefit and distance functions. *Journal of Economic Theory*, 70(2):407–419, 1996.

R. G. Chambers, Y. Chung, and R. Färe. Profit, directional distance functions, and nerlovian efficiency. *Journal of Optimization Theory and Applications*, 98 (2):351–364, 1998.

T. C. Chan, T. Craig, T. Lee, and M. B. Sharpe. Generalized inverse multiobjective optimization with application to cancer therapy. *Operations Research*, 62(3): 680–695, 2014.

V. Chanyavanich, S. K. Das, W. R. Lee, and J. Y. Lo. Knowledge-based IMRT treatment planning for prostate cancer. *Medical Physics*, 38(5):2515–2522, 2011.

A. Charnes and W. W. Cooper. Programming with linear fractional functionals. *Naval Research Logistics Quarterly*, 9(3-4):181–186, 1962.

A. Charnes, W. W. Cooper, and E. Rhodes. Measuring the efficiency of decision making units. *European Journal of Operational Research*, 2(6):429–444, 1978.

A. Charnes, W. W. Cooper, and E. Rhodes. Short communication: Measuring the efficiency of decision making units. *European Journal of Operational Research*, 3(4):339, 1979.

H. Chen, D. L. Craft, and D. P. Gierga. Multicriteria optimization informed VMAT planning. *Medical Dosimetry*, 39(1):64–73, 2014.

H. Chen, B. A. Winey, J. Daartz, K. S. Oh, J. H. Shin, and D. P. Gierga. Efficiency gains for spinal radiosurgery using multicriteria optimization intensity modulated radiation therapy guided volumetric modulated arc therapy planning. *Practical Radiation Oncology*, 5(1):49–55, 2015.

W. Chen, D. Craft, T. M. Madden, K. Zhang, H. M. Kooy, and G. T. Herman. A fast optimization algorithm for multicriteria intensity modulated proton therapy planning. *Medical Physics*, 37(9):4938–4945, 2010.

R. Cheung, S. L. Tucker, A. K. Lee, R. de Crevoisier, L. Dong, A. Kamat, L. Pisters, and D. Kuban. Dose–response characteristics of low-and intermediate-risk prostate cancer treated with external beam radiotherapy. *International Journal of Radiation Oncology Biology Physics*, 61(4):993–1002, 2005.

J. A. Chilingerian and H. D. Sherman. Health-care applications: From hospitals to physicians, from productive efficiency to quality frontiers. In W. W. Cooper, L. M. Seiford, and J. Zhu, editors, *Handbook on Data Envelopment Analysis*, pages 445–493. Springer, New York, 2011.

V. Chvátal. *Linear Programming*. W.H. Freeman and Company, New York, 1983.

V. Clark, Y. Chen, J. Wilkens, J. Alaly, K. Zakaryan, and J. Deasy. IMRT treatment planning for prostate cancer using prioritized prescription optimization and mean-tail-dose functions. *Linear Algebra and its Applications*, 428(5):1345–1364, 2008.

T. J. Coelli, D. S. P. Rao, C. J. O'Donnell, and G. E. Battese. *An Introduction to Efficiency and Productivity Analysis*. Springer Science & Business Media, New York, 2005.

W. W. Cooper, L. M. Seiford, and J. Zhu. Data envelopment analysis: History, models, and interpretations. In W. W. Cooper, L. M. Seiford, and J. Zhu, editors, *Handbook on Data Envelopment Analysis*, pages 1–39. Springer, New York, 2011.

C. Cotrutz and L. Xing. IMRT dose shaping with regionally variable penalty scheme. *Medical Physics*, 30(4):544–551, 2003.

C. Cotrutz, M. Lahanas, C. Kappas, and D. Baltas. A multiobjective gradient-based dose optimization algorithm for external beam conformal radiotherapy. *Physics in Medicine and Biology*, 46(8):2161–2175, 2001.

D. Craft and M. Monz. Simultaneous navigation of multiple Pareto surfaces, with an application to multicriteria IMRT planning with multiple beam angle configurations. *Medical Physics*, 37(2):736–741, 2010.

D. Craft and C. Richter. Deliverable navigation for multicriteria step and shoot IMRT treatment planning. *Physics in Medicine and Biology*, 58(1):87–103, 2013.

D. Craft, T. Halabi, and T. Bortfeld. Exploration of tradeoffs in intensity-modulated radiotherapy. *Physics in Medicine and Biology*, 50(24):5857–5868, 2005.

D. Craft, D. McQuaid, J. Wala, W. Chen, E. Salari, and T. Bortfeld. Multicriteria VMAT optimization. *Medical Physics*, 39(2):686–696, 2012a.

D. Craft, D. Papp, and J. Unkelbach. Plan averaging for multicriteria navigation of sliding window IMRT and VMAT. *Medical Physics*, 41(2):021709–1–021709–5, 2014.

D. L. Craft, T. F. Halabi, H. A. Shih, and T. R. Bortfeld. Approximating convex Pareto surfaces in multiobjective radiotherapy planning. *Medical Physics*, 33(9): 3399–3407, 2006.

D. L. Craft, T. S. Hong, H. A. Shih, and T. R. Bortfeld. Improved planning time and plan quality through multicriteria optimization for intensity–modulated radiotherapy. *International Journal of Radiation Oncology Biology Physics*, 82(1): e83–e90, 2012b.

I. Das and J. E. Dennis. A closer look at drawbacks of minimizing weighted sums of objectives for pareto set generation in multicriteria optimization problems. *Structural Optimization*, 14(1):63–69, 1997.

I. Das and J. E. Dennis. Normal-boundary intersection: A new method for generating the Pareto surface in nonlinear multicriteria optimization problems. *SIAM Journal on Optimization*, 8(3):631–657, 1998.

J. Deasy. Multiple local minima in radiotherapy optimization problems with dose–volume constraints. *Medical Physics*, 24(7):1157–1161, 1997.

J. O. Deasy, A. I. Blanco, and V. H. Clark. CERR: a computational environment for radiotherapy research. *Medical Physics*, 30(5):979–985, 2003.

K. Deb and T. Goel. Controlled elitist non-dominated sorting genetic algorithms for better convergence. In E. Zitzler, L. Thiele, K. Deb, C. A. Coello Coello, and D. Corne, editors, *Evolutionary Multi-Criterion Optimization*, pages 67–81. Springer, Berlin, 2001.

K. Deb, A. Pratap, S. Agarwal, and T. Meyarivan. A fast and elitist multiobjective genetic algorithm: NSGA-II. *IEEE Transactions on Evolutionary Computation*, 6(2):182–197, 2002.

G. Delaney, S. Jacob, C. Featherstone, and M. Barton. The role of radiotherapy in cancer treatment. *Cancer*, 104(6):1129–1137, 2005.

Department of Health Cancer Policy Team. Radiotherapy services in England 2012. Technical report, Department of Health, 2012.

M. Ehrgott. *Multicriteria Optimization*. Springer, Berlin, 2nd edition, 2005.

M. Ehrgott and X. Gandibleux. Bound sets for biobjective combinatorial optimization problems. *Computers & Operations Research*, 34(9):2674–2694, 2007.

M. Ehrgott and D. M. Ryan. Constructing robust crew schedules with bicriteria optimization. *Journal of Multi-Criteria Decision Analysis*, 11(3):139–150, 2002.

M. Ehrgott and I. Winz. Interactive decision support in radiation therapy treatment planning. *OR Spectrum*, 30(2):311–329, 2008.

M. Ehrgott, Ç. Güler, H. W. Hamacher, and L. Shao. Mathematical optimization in intensity modulated radiation therapy. *4OR*, 6(3):199–262, 2008a.

M. Ehrgott, H. Hamacher, and M. Nußbaum. Decomposition of matrices and static multileaf collimators: A survey. In C. Alves, P. Pardalos, and L. Vicente, editors, *Optimization in Medicine*, volume 12 of *Springer Optimization and its Applications*, pages 27–48. Springer, Berlin, 2008b.

M. Ehrgott, A. Holder, and J. Reese. Beam selection in radiotherapy design. *Linear Algebra and its Applications*, 428(5):1272–1312, 2008c.

M. Ehrgott, Ç. Güler, H. W. Hamacher, and L. Shao. Mathematical optimization in intensity modulated radiation therapy. *Annals of Operations Research*, 175 (1):309–365, 2010.

K. Engel. A new algorithm for optimal multileaf collimator field segmentation. *Discrete Applied Mathematics*, 152(1):35–51, 2005.

M. Falkinger, S. Schell, J. Müller, and J. J. Wilkens. Prioritized optimization in intensity modulated proton therapy. *Zeitschrift fur Medizinische Physik*, 22(1): 21–28, 2012.

S. L. Faulkenberg and M. M. Wiecek. On the quality of discrete representations in multiple objective programming. *Optimization and Engineering*, 11:423–440, 2010.

L. Feuvret, G. Noël, J.-J. Mazeron, and P. Bey. Conformity index: A review. *International Journal of Radiation Oncology Biology Physics*, 64(2):333–342, 2006.

J. Fiege, B. McCurdy, P. Potrebko, H. Champion, and A. Cull. PARETO: A novel evolutionary optimization approach to multiobjective IMRT planning. *Medical Physics*, 38(9):5217–5229, 2011.

A. Fogliata, F. Belosi, A. Clivio, P. Navarria, G. Nicolini, M. Scorsetti, E. Vanetti, and L. Cozzi. On the pre-clinical validation of a commercial model-based optimisation engine: Application to volumetric modulated arc therapy for patients with lung or prostate cancer. *Radiotherapy and Oncology*, 113(3):385–391, 2014.

A. Fredriksson and R. Bokrantz. Deliverable navigation for multicriteria IMRT treatment planning by combining shared and individual apertures. *Physics in Medicine and Biology*, 58(21):7683–7697, 2013.

M. Fuss and B. J. Salter. Intensity-modulated radiosurgery: Improving dose gradients and maximum dose using post inverse-optimization interactive dose shaping. *Technology in Cancer Research & Treatment*, 6(3):197–203, 2007.

D. Good, J. Lo, W. R. Lee, Q. J. Wu, F.-F. Yin, and S. K. Das. A knowledge-based approach to improving and homogenizing intensity modulated radiation therapy planning quality among treatment centers: An example application to prostate cancer planning. *International Journal of Radiation Oncology Biology Physics*, 87(1):176–181, 2013.

H. W. Hamacher and K.-H. Küfer. Inverse radiation therapy planning–a multiple objective optimization approach. *Discrete Applied Mathematics*, 118(1):145–161, 2002.

A. L. Hoffmann, A. Y. D. Siem, D. Den Hertog, J. H. A. M. Kaanders, and H. Huizenga. Derivative-free generation and interpolation of convex Pareto optimal IMRT plans. *Physics in Medicine and Biology*, 51(24):6349–6369, 2006.

A. L. Hoffmann, D. den Hertog, A. Y. Siem, J. H. Kaanders, and H. Huizenga. Convex reformulation of biologically-based multi-criteria intensity-modulated radiation therapy optimization including fractionation effects. *Physics in medicine and biology*, 53(22):6345–6362, 2008.

A. Holder. Designing radiotherapy plans with elastic constraints and interior point methods. *Health Care Management Science*, 6(1):5–16, 2003.

A. Holder. Radiotherapy treatment design and linear programming. In M. L. Brandeau, F. Sainfort, and W. P. Pierskalla, editors, *Operations Research and Health Care: A Handbook of Methods and Applications*, pages 741–774. Springer, Boston, MA, 2005.

A. Holder. Partitioning multiple objective optimal solutions with applications in radiotherapy design. *Optimization and Engineering*, 7(4):501–526, 2006.

C. Holdsworth, M. Kim, J. Liao, and M. H. Phillips. A hierarchical evolutionary algorithm for multiobjective optimization in IMRT. *Medical Physics*, 37(9): 4986–4997, 2010.

T. S. Hong, D. L. Craft, F. Carlsson, and T. R. Bortfeld. Multicriteria optimization in intensity–modulated radiation therapy treatment planning for locally advanced cancer of the pancreatic head. *International Journal of Radiation Oncology Biology Physics*, 72(4):1208–1214, 2008.

K. W. Jee, D. L. McShan, and B. A. Fraass. Lexicographic ordering: Intuitive multicriteria optimization for IMRT. *Physics in Medicine and Biology*, 52(7): 1845–1861, 2007.

U. Jeleń and M. Alber. A finite size pencil beam algorithm for IMRT dose optimization: Density corrections. *Physics in Medicine and Biology*, 52(3):617–633, 2007.

R. Jeraj, C. Wu, and T. R. Mackie. Optimizer convergence and local minima errors and their clinical importance. *Physics in Medicine and Biology*, 48(17): 2809–2827, 2003.

X. Jia, P. Ziegenhein, and S. B. Jiang. GPU-based high-performance computing for radiation therapy. *Physics in Medicine and Biology*, 59(4):R151–R182, 2014.

G. Kalantzis and A. Apte. A novel reduced-order prioritized optimization method for radiation therapy treatment planning. *IEEE Transactions on Biomedical Engineering*, 61(4):1062–1070, 2014.

P. Kallman, B. Lind, and A. Brahme. An algorithm for maximizing the probability of complication-free tumour control in radiation therapy. *Physics in Medicine and Biology*, 37(4):871–890, 1992.

T. Kataria, K. Sharma, V. Subramani, K. Karrthick, and S. S. Bisht. Homogeneity index: An objective tool for assessment of conformal radiation treatments. *Journal of Medical Physics/Association of Medical Physicists of India*, 37(4): 207–213, 2012.

P. J. Keall and P. W. Hoban. Superposition dose calculation incorporating Monte Carlo generated electron track kernels. *Medical Physics*, 23(4):479–485, 1996.

F. Khan and D. Craft. Three-dimensional conformal planning with low-segment multicriteria intensity modulated radiation therapy optimization. *Practical Radiation Oncology*, 5(2):e103–e111, 2015.

R. G. Kierkels, R. Visser, H. P. Bijl, J. A. Langendijk, A. A. van 't Veld, R. J. Steenbakkers, and E. W. Korevaar. Multicriteria optimization enables less experienced planners to efficiently produce high quality treatment plans in head and neck cancer radiotherapy. *Radiation Oncology*, 10(1):1–9, 2015.

P. Kolmonen, J. Tervo, and T. Lahtinen. Use of the cimmino algorithm and continuous approximation for the dose deposition kernel in the inverse problem of radiation treatment planning. *Physics in Medicine and Biology*, 43(9):2539–2554, 1998.

K.-H. Küfer, A. Scherrer, M. Monz, F. Alonso, H. Trinkaus, T. Bortfeld, and C. Thieke. Intensity-modulated radiotherapy–a large scale multi-criteria programming problem. *OR Spectrum*, 25(2):223–249, 2003.

M. Lahanas, E. Schreibmann, and D. Baltas. Multiobjective inverse planning for intensity modulated radiotherapy with constraint-free gradient-based optimization algorithms. *Physics in Medicine and Biology*, 48(17):2843–2871, 2003a.

M. Lahanas, E. Schreibmann, N. Milickovic, and D. Baltas. Intensity modulated beam radiation therapy dose optimization with multiobjective evolutionary algorithms. In C. M. Fonseca, P. J. Fleming, E. Zitzler, L. Thiele, and K. Deb, editors, *Evolutionary Multi-Criterion Optimization*, pages 648–661. Springer, Berlin, 2003b.

N. Li, M. Zarepisheh, A. Uribe-Sanchez, K. Moore, Z. Tian, X. Zhen, Y. J. Graves, Q. Gautier, L. Mell, L. Zhou, X. Jia, and S. Jiang. Automatic treatment plan re-optimization for adaptive radiotherapy guided with the initial plan DVHs. *Physics in Medicine and Biology*, 58(24):8725–8738, 2013.

X. A. Li, M. Alber, J. O. Deasy, A. Jackson, K.-W. K. Jee, L. B. Marks, M. K. Martel, C. Mayo, V. Moiseenko, A. E. Nahum, A. Niemierko, V. A. Semenenko, and E. D. Yorke. The use and QA of biologically related models for treatment planning: Short report of the TG-166 of the therapy physics committee of the AAPM. *Medical Physics*, 39(3):1386–1409, 2012.

G. Lim, W. Cao, and R. Mohan. Recent advances in intensity modulated proton therapy treatment planning optimization. In J. Lee and H. Jung, editors, *Proceedings of the 15th Asia Pacific Industrial Engineering and Management Systems Conference*, pages 1520–1525, Jeju, Korea, 2014.

G. J. Lim, M. C. Ferris, S. J. Wright, D. M. Shepard, and M. A. Earl. An optimization framework for conformal radiation treatment planning. *INFORMS Journal on Computing*, 19(3):366–380, 2007.

K.-M. Lin, J. Simpson, G. Sasso, A. Raith, and M. Ehrgott. Quality assessment for VMAT prostate radiotherapy planning based on data envelopment analysis. *Physics in Medicine and Biology*, 58(16):5753–5769, 2013.

J. Llacer, J. O. Deasy, T. R. Bortfeld, T. D. Solberg, and C. Promberger. Absence of multiple local minima effects in intensity modulated optimization with dose-volume constraints. *Physics in Medicine and Biology*, 48(2):183–210, 2003.

N. J. Lomax and S. G. Scheib. Quantifying the degree of conformity in radiosurgery treatment planning. *International Journal of Radiation Oncology Biology Physics*, 55(5):1409–1419, 2003.

P. Lougovski, J. LeNoach, L. Zhu, Y. Ma, Y. Censor, and L. Xing. Toward truly optimal IMRT dose distribution: Inverse planning with voxel-specific penalty. *Technology in Cancer Research & Treatment*, 9(6):629–636, 2010.

M. E. Lübbecke. Column generation. *Wiley Encyclopedia of Operations Research and Management Science*, 2010. doi: 10.1002/9780470400531.

L. B. Marks, E. D. Yorke, A. Jackson, R. K. Ten Haken, L. S. Constine, A. Eisbruch, S. M. Bentzen, J. Nam, and J. O. Deasy. Use of normal tissue complication probability models in the clinic. *International Journal of Radiation Oncology Biology Physics*, 76(3):S10–S19, 2010.

C. K. McGarry, R. Bokrantz, J. M. OSullivan, and A. R. Hounsell. Advantages and limitations of navigation-based multicriteria optimization (MCO) for localized prostate cancer IMRT planning. *Medical Dosimetry*, 39(3):205–211, 2014.

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*, 52(24):7333–7352, 2007.

C. Men, X. Gu, D. Choi, A. Majumdar, Z. Zheng, K. Mueller, and S. B. Jiang. GPU-based ultrafast IMRT plan optimization. *Physics in Medicine and Biology*, 54(21):6565–6573, 2009.

C. Men, X. Jia, and S. B. Jiang. GPU-based ultra-fast direct aperture optimization for online adaptive radiation therapy. *Physics in Medicine and Biology*, 55(15): 4309–4319, 2010a.

C. Men, H. E. Romeijn, X. Jia, and S. B. Jiang. Ultrafast treatment plan optimization for volumetric modulated arc therapy (VMAT). *Medical Physics*, 37 (11):5787–5791, 2010b.

A. Messac, A. Ismail-Yahaya, and C. A. Mattson. The normalized normal constraint method for generating the Pareto frontier. *Structural and Multidisciplinary Optimization*, 25(2):86–98, 2003.

K. Miettinen. *Nonlinear multiobjective optimization*, volume 12 of *International Series in Operations Research & Management Science*. Springer Science & Business Media, New York, 1999.

M. Monz, K. Küfer, T. Bortfeld, and C. Thieke. Pareto navigation–algorithmic foundation of interactive multi-criteria IMRT planning. *Physics in Medicine and Biology*, 53(4):985–998, 2008.

K. L. Moore, R. S. Brame, D. A. Low, and S. Mutic. Experience-based quality control of clinical intensity-modulated radiotherapy planning. *International Journal of Radiation Oncology Biology Physics*, 81(2):545–551, 2011.

S. Moradi, A. Raith, and M. Ehrgott. A bi-objective column generation algorithm for the multi-commodity minimum cost flow problem. *European Journal of Operational Research*, 244(2):369–378, 2015.

A. Niemierko. A generalized concept of equivalent uniform dose. *Medical Physics*, 26(6):1100, 1999.

A. Niemierko. Biological optimization. In T. Bortfeld, R. Schmidt-Ullrich, W. De Neve, and D. E. Wazer, editors, *Image-Guided IMRT*, pages 199–216. Springer, Berlin, 2006.

K. Otto. Real-time interactive treatment planning. *Physics in Medicine and Biology*, 59(17):4845–4859, 2014.

S. F. Petit, B. Wu, M. Kazhdan, A. Dekker, P. Simari, R. Kumar, R. Taylor, J. M. Herman, and T. McNutt. Increased organ sparing using shape-based treatment plan optimization for intensity modulated radiation therapy of pancreatic adenocarcinoma. *Radiotherapy and Oncology*, 102(1):38–44, 2012.

F. Preciado-Walters, R. Rardin, M. Langer, and V. Thai. A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer. *Mathematical Programming*, 101(2):319–338, 2004.

R. Ramer, A. Holder, and N. Papanikolaou. SU-GG-T-392: Utilization of data envelopment analysis (DEA) to compare prostate treatment options. *Medical Physics*, 35(6):2815–2815, 2008.

N. Reynaert, S. Van der Marck, D. Schaart, W. Van der Zee, C. Van Vliet-Vroegindeweij, M. Tomsej, J. Jansen, B. Heijmen, M. Coghe, and C. De Wagter. Monte Carlo treatment planning for photon and electron beams. *Radiation Physics and Chemistry*, 76(4):643–686, 2007.

U. Ringborg, D. Bergqvist, B. Brorsson, E. Cavallin-Ståhl, J. Ceberg, N. Einhorn, J.-e. Frödin, J. Järhult, G. Lamnevik, C. Lindholm, et al. The Swedish council on technology assessment in health care (SBU) systematic overview of radiotherapy for cancer including a prospective survey of radiotherapy practice in Sweden 2001–summary and conclusions. *Acta Oncologica*, 42(5–6):357–365, 2003.

D. A. Ripsman, D. M. Aleman, and K. Ghobadi. Interactive visual guidance for automated stereotactic radiosurgery treatment planning. *Expert Systems with Applications*, 42(21):8337–8348, 2015.

H. Rocha, J. M. Dias, B. C. Ferreira, and M. C. Lopes. Discretization of optimal beamlet intensities in IMRT: A binary integer programming approach. *Mathematical and Computer Modelling*, 55(7–8):1969–1980, 2012.

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*, 49 (10):1991–2013, 2004.

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*, 15(3):838–862, 2005.

I. Rosen, H. H. Liu, N. Childress, and Z. Liao. Interactively exploring optimized treatment plans. *International Journal of Radiation Oncology Biology Physics*, 61(2):570–582, 2005.

H. Ruotsalainen. *Interactive multiobjective optimization in model-based decision making with applications*. PhD thesis, University of Kuopio, Kuopio, Finland, 2009.

S. Ruzika and M. M. Wiecek. Approximation methods in multiobjective programming. *Journal of Optimization Theory and Applications*, 126(3):473–501, 2005.

E. Salari and J. Unkelbach. A column-generation-based method for multi-criteria direct aperture optimization. *Physics in Medicine and Biology*, 58(3):621–639, 2013.

S. P. Santos and C. A. Amado. Using data envelopment analysis for formative evaluation of radiotherapy services: An exploratory study. In E. Tanfani and

A. Testi, editors, *Advanced Decision Making Methods Applied to Health Care*, pages 173–190. Springer, Mailand, 2012.

S. Sayın. Measuring the quality of discrete representations of efficient sets in multiple objective mathematical programming. *Mathematical Programming*, 87 (3):543–560, 2000.

W. Schlegel and A. Mahr. *3D Conformal Radiation Therapy: Multimedia Introduction to Methods and Techniques*. Springer, Berlin, 2nd edition, 2007.

E. Schreibmann and T. Fox. Prior-knowledge treatment planning for volumetric arc therapy using feature-based database mining. *Journal of Applied Clinical Medical Physics*, 15(2):19–27, 2014.

L. Shao and M. Ehrgott. Finding representative nondominated points in multiobjective linear programming. In *IEEE Symposium on Computational Intelligence in Multicriteria Decision Making*, pages 245–252. IEEE, 2007.

L. Shao and M. Ehrgott. Approximately solving multiobjective linear programmes in objective space and an application in radiotherapy treatment planning. *Mathematical Methods of Operations Research*, 68(2):257–276, 2008.

D. M. Shepard, M. C. Ferris, G. H. Olivera, and T. R. Mackie. Optimizing the delivery of radiation therapy to cancer patients. *SIAM Review*, 41(4):721–744, 1999.

S. V. Spirou and C.-S. Chui. A gradient inverse planning algorithm with dose-volume constraints. *Medical Physics*, 25(3):321–333, 1998.

R. E. Steuer and E.-U. Choo. An interactive weighted tchebycheff procedure for multiple objective programming. *Mathematical Programming*, 26(3):326–344, 1983.

M. T. Studenski, V. Bar-Ad, J. Siglin, D. Cognetti, J. Curry, M. Tuluc, and A. S. Harrison. Clinical experience transitioning from IMRT to VMAT for head and neck cancer. *Medical Dosimetry*, 38(2):171–175, 2013.

P. Süss, M. Bortz, K.-H. Küfer, and C. Thieke. The critical spot eraser–a method to interactively control the correction of local hot and cold spots in IMRT planning. *Physics in Medicine and Biology*, 58(6):1855–1867, 2013.

C. Thieke, K.-H. Küfer, M. Monz, A. Scherrer, F. Alonso, U. Oelfke, P. E. Huber, J. Debus, and T. Bortfeld. A new concept for interactive radiotherapy planning with multicriteria optimization: First clinical evaluation. *Radiotherapy and Oncology*, 85(2):292–298, 2007.

Z. Tian, F. Peng, M. Folkerts, J. Tan, X. Jia, and S. B. Jiang. Multi-GPU implementation of a VMAT treatment plan optimization algorithm. *Medical Physics*, 42(6):2841–2852, 2015.

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*, 42(3):1367–1377, 2015.

P. W. Voet, M. L. Dirkx, S. Breedveld, D. Fransen, P. C. Levendag, and B. J. Heijmen. Toward fully automated multicriterial plan generation: A prospective clinical study. *International Journal of Radiation Oncology Biology Physics*, 85 (3):866–872, 2013.

J. Wala, D. Craft, J. Paly, A. Zietman, and J. Efstathiou. Maximizing dosimetric benefits of IMRT in the treatment of localized prostate cancer through multicriteria optimization planning. *Medical Dosimetry*, 38(3):298–303, 2013.

S. Webb. *Intensity-Modulated Radiation Therapy*. CRC Press, 2001.

J. J. Wilkens, J. R. Alaly, K. Zakarian, W. L. Thorstad, and J. O. Deasy. IMRT treatment planning based on prioritizing prescription goals. *Physics in Medicine and Biology*, 52(6):1675–1692, 2007.

B. Wu, F. Ricchetti, G. Sanguineti, M. Kazhdan, P. Simari, M. Chuang, R. Taylor, R. Jacques, and T. McNutt. Patient geometry-driven information retrieval for IMRT treatment plan quality control. *Medical Physics*, 36(12):5497–5505, 2009.

B. Wu, F. Ricchetti, G. Sanguineti, M. Kazhdan, P. Simari, R. Jacques, R. Taylor, and T. McNutt. Data-driven approach to generating achievable dose-volume histogram objectives in intensity-modulated radiotherapy planning. *International Journal of Radiation Oncology Biology Physics*, 79(4):1241–1247, 2011.

B. Wu, D. Pang, P. Simari, R. Taylor, G. Sanguineti, and T. McNutt. Using overlap volume histogram and IMRT plan data to guide and automate VMAT planning:

A head-and-neck case study. *Medical Physics*, 40(2):021714–1–021714–7, 2013. doi: 10.1118/1.4788671.

C. Wu, R. Jeraj, and T. R. Mackie. The method of intercepts in parameter space for the analysis of local minima caused by dose-volume constraints. *Physics in Medicine and Biology*, 48(11):N149–N157, 2003a.

Q. Wu, R. Mohan, M. Morris, A. Lauve, and R. Schmidt-Ullrich. Simultaneous integrated boost intensity-modulated radiotherapy for locally advanced head-and-neck squamous cell carcinomas. I: Dosimetric results. *International Journal of Radiation Oncology Biology Physics*, 56(2):573–585, 2003b.

I. Xhaferllari, E. Wong, K. Bzdusek, M. Lock, and J. Chen. Automated IMRT planning with regional optimization using planning scripts. *Journal of Applied Clinical Medical Physics*, 14(1), 2013.

Y. Yang, E. C. Ford, B. Wu, M. Pinkawa, B. van Triest, P. Campbell, D. Y. Song, and T. R. McNutt. An overlap-volume-histogram based method for rectal dose prediction and automated treatment planning in the external beam prostate radiotherapy following hydrogel injection. *Medical Physics*, 40(1):011709–1–011709–10, 2013.

Y. Yang, T. Li, L. Yuan, Y. Ge, F.-F. Yin, W. R. Lee, and Q. J. Wu. Quantitative comparison of automatic and manual IMRT optimization for prostate cancer: The benefits of DVH prediction. *Journal of Applied Clinical Medical Physics*, 16(2):241–250, 2015.

M. Yoon. A new homogeneity index based on the statistical analysis of dose volume histogram. *Journal of Applied Clinical Medical Physics*, 8(2):9–17, 2007.

C. X. Yu and G. Tang. Intensity-modulated arc therapy: Principles, technologies and clinical implementation. *Physics in Medicine and Biology*, 56(5):R31–R54, 2011.

L. Yuan, Y. Ge, W. R. Lee, F. F. Yin, J. P. Kirkpatrick, and Q. J. Wu. Quantitative analysis of the factors which affect the interpatient organ-at-risk dose sparing variation in IMRT plans. *Medical physics*, 39(11):6868–6878, 2012.

M. Zaghian, G. Lim, W. Liu, and R. Mohan. An automatic approach for satisfying dose-volume constraints in linear fluence map optimization for IMPT. *Journal of Cancer Therapy*, 5(2):198–207, 2014.

M. Zarepisheh, T. Long, N. Li, Z. Tian, H. E. Romeijn, X. Jia, and S. B. Jiang. A DVH-guided IMRT optimization algorithm for automatic treatment planning and adaptive radiotherapy replanning. *Medical Physics*, 41(6):061711, 2014.

X. Zhang, X. Li, E. M. Quan, X. Pan, and Y. Li. A methodology for automatic intensity-modulated radiation treatment planning for lung cancer. *Physics in Medicine and Biology*, 56(13):3873–3893, 2011.

X. Zhu, Y. Ge, T. Li, D. Thongphiew, F.-F. Yin, and Q. J. Wu. A planning quality evaluation tool for prostate adaptive IMRT based on machine learning. *Medical Physics*, 38(2):719–726, 2011.

P. Ziegenhein, C. P. Kamerling, M. Bangert, J. Kunkel, and U. Oelfke. Performance-optimized clinical IMRT planning on modern CPUs. *Physics in Medicine and Biology*, 58(11):3705–3715, 2013.

P. Ziegenhein, C. P. Kamerling, and U. Oelfke. Interactive dose shaping - efficient strategies for CPU-based real-time treatment planning. *Journal of Physics: Conference Series*, 489(1):012066–1–012066–6, 2014.