KTH Engineering Sciences

Multicriteria optimization for managing tradeoffs in radiation therapy treatment planning

RASMUS BOKRANTZ

Doctoral Thesis
Stockholm, Sweden 2013

TRITA MAT 13/OS/07
ISSN 1401-2294
ISRN KTH/OPT/DA-13/07-SE
ISBN 978-91-7501-790-7

Optimization and Systems Theory
Department of Mathematics
KTH Royal Institute of Technology
SE-100 44 Stockholm, Sweden

Akademisk avhandling som med tillstånd av Kungl Tekniska Högskolan framläggs till offentlig granskning för avläggande av teknologie doktorsexamen, fredagen den 14 juni 2013 klockan 10.00 i sal F3, Lindstedtsvägen 26, Kungl Tekniska Högskolan, Stockholm.

Print: Universitetsservice US-AB, Stockholm, 2013

Abstract

Treatment planning for radiation therapy inherently involves tradeoffs, such as between tumor control and normal tissue sparing, between time-efficiency and dose quality, and between nominal plan quality and robustness. The purpose of this thesis is to develop methods that can facilitate decision making related to such tradeoffs. The main focus of the thesis is on multicriteria optimization methods where a representative set of treatment plans are first calculated and the most appropriate plan contained in this representation then selected by the treatment planner through continuous interpolation between the precalculated alternatives. These alternatives constitute a subset of the set of Pareto optimal plans, meaning plans such that no criterion can be improved without a sacrifice in another.

Approximation of Pareto optimal sets is first studied with respect to fluence map optimization for intensity-modulated radiation therapy. The approximation error of a discrete representation is minimized by calculation of points one at the time at the location where the distance between an inner and outer approximation of the Pareto set currently attains its maximum. A technique for calculating this distance that is orders of magnitude more efficient than the best previous method is presented. A generalization to distributed computational environments is also proposed.

Approximation of Pareto optimal sets is also considered with respect to direct machine parameter optimization. Optimization of this form is used to calculate representations where any interpolated treatment plan is directly deliverable. The fact that finite representations of Pareto optimal sets have approximation errors with respect to Pareto optimality is addressed by a technique that removes these errors by a projection onto the exact Pareto set. Projections are also studied subject to constraints that prevent the dose-volume histogram from deteriorating.

Multicriteria optimization is extended to treatment planning for volumetric-modulated arc therapy and intensity-modulated proton therapy. Proton therapy plans that are robust against geometric errors are calculated by optimization of the worst case outcome. The theory for multicriteria optimization is extended to accommodate this formulation. Worst case optimization is shown to be preferable to a previous more conservative method that also protects against uncertainties which cannot be realized in practice.

Keywords: Optimization, multicriteria optimization, robust optimization, Pareto optimality, Pareto surface approximation, Pareto surface navigation, intensity-modulated radiation therapy, volumetric-modulated arc therapy, intensity-modulated proton therapy.

Sammanfattning

En viktig aspekt av planering av strålterapibehandlingar är avvägningar mellan behandlingsmål vilka står i konflikt med varandra. Exempel på sådana avvägningar är mellan tumörkontroll och dos till omkringliggande frisk vävnad, mellan behandlingstid och doskvalitet, och mellan nominell plankvalitet och robusthet med avseende på geometriska fel. Denna avhandling syftar till att utveckla metoder som kan underlätta beslutsfattande kring motstridiga behandlingsmål. Primärt studeras en metod för flermålsoptimering där behandlingsplanen väljs genom kontinuerlig interpolation över ett representativt urval av förberäknade alternativ. De förberäknade behandlingsplanerna utgör en delmängd av de Paretooptimala planerna, det vill säga de planer sådana att en förbättring enligt ett kriterium inte kan ske annat än genom en försämring enligt ett annat.

Beräkning av en approximativ representation av mängden av Paretooptimala planer studeras först med avseende på fluensoptimering för intensitetsmodulerad strålterapi. Felet för den approximativa representationen minimeras genom att innesluta mängden av Paretooptimala planer mellan inre och yttre approximationer. Dessa approximationer förfinas iterativt genom att varje ny plan genereras där avståndet mellan approximationerna för tillfället är som störst. En teknik för att beräkna det maximala avståndet mellan approximationerna föreslås vilken är flera storleksordningar snabbare än den bästa tidigare kända metoden. En generalisering till distribuerade beräkningsmiljöer föreslås även.

Approximation av mängden av Paretooptimala planer studeras även för direkt maskinparameteroptimering, som används för att beräkna representationer där varje interpolerad behandlingsplan är direkt levererbar. Det faktum att en ändlig representation av mängden av Paretooptimala lösningar har ett approximationsfel till Paretooptimalitet hanteras via en metod där en interpolerad behandlingsplan projiceras på Paretomängden. Projektioner studeras även under bivillkor som förhindrar att den interpolerade planens dos-volym histogram kan försämras.

Flermålsoptimering utökas till planering av rotationsterapi och intensitetsmodulerad protonterapi. Protonplaner som är robusta mot geometriska fel beräknas genom optimering med avseende på det värsta möjliga utfallet av de föreliggande osäkerheterna. Flermålsoptimering utökas även teoretiskt till att innefatta denna formulering. Nyttan av värsta fallet-optimering jämfört med tidigare mer konservativa metoder som även skyddar mot osäkerheter som inte kan realiseras i praktiken demonstreras experimentellt.

Nyckelord: Optimering, flermålsoptimering, robust optimering, Paretooptimalitet, Paretofrontsapproximation, Paretofrontsnavigering, intensitetsmodulerad strålterapi, rotationsterapi, intensitetsmodulerad protonterapi.

Acknowledgments

This thesis work was carried out as a joint project between the Division of Optimization and Systems Theory at KTH and RaySearch Laboratories. The project was funded by RaySearch Laboratories, with support from the KTH Center for Industrial and Applied Mathematics (CIAM). I am grateful to the following persons for support during my PhD studies:

First of all, many thanks to my academic advisor Anders Forsgren for guidance and expertise that has been invaluable to my research, and for personal encouragement and support. I am grateful to Kaisa Miettinen, who during her stay as a visiting professor at KTH acted as an academic co-advisor, for sharing your exhaustive knowledge on multicriteria optimization and for joint work on one of the appended papers. I thank my academic co-advisors Johan Håstad and Krister Svanberg for feedback during the semi-annual reference group meetings.

I am grateful to my industrial co-advisor Björn Hårdemark for comments and constructive criticism on my work and for sharing your vast knowledge on anything radiation therapy related. I also thank my former industrial co-advisor Henrik Rehbinder for guidance and encouragement during the initial phase of my PhD. I am grateful to Johan Löf, who as the founder and CEO of RaySearch is the original initiator of my PhD project, for constructive comments and suggestions on directions of research.

I thank Thomas Bortfeld and David Craft at Harvard Medical School for feedback on my research trough a collaboration between RaySearch and Massachusetts General Hospital (MGH). The link to clinical utilization of multicriteria optimization through this collaboration has been a personal motivator, and research published by the MGH group an important source of inspiration. I am grateful to Philip Gill at the University of California, San Diego, for discussions on nonlinear programming and for providing the optimization solver SNOPT, which has been an important tool for early prototype work.

I thank Albin Fredriksson, who applied to the same PhD project opportunity as me. Indecisiveness on our (now) employer's part led to a one-person project becoming two parallel projects—a construction that I personally have found very fruitful. I appreciate you taking interest in my research, and I also enjoyed the joint work on two of the appended papers and travel in conjunction with conferences.

I am grateful to my former and current colleagues at RaySearch, all of whom contribute to making RaySearch a great place to work. I am particularly grateful to Fredrik Löfman for discussions on many aspects of radiation therapy optimization (your dissertation defense was also my first contact with this field), to Göran Sporre for discussions on optimization in general and nonlinear constraints in particular, to Kjell Eriksson for discussions on volumetric-modulated arc therapy, to Agnes Lundborg and Erik Traneus for discussions on physics, to Krister Jacobsson, Henrik Strandemar, and Johan Uhrdin for input on multicriteria optimization and joint product development related to this topic, and to Pelle Jansson for the inspiration to study distributed calculations which resulted in one of the appended papers. I thank my roommates Elin Hynning and Minna Wedenberg for contributing to positive work atmosphere in our office. Thanks also to Minna for discussions on biology.

I am also grateful to my colleagues at the Division of Optimization and Systems Theory at KTH for a rewarding academic environment. I particularly enjoyed taking courses together with fellow students Mikael Fallgren, Anders Möller, Hildur Æsa Oddsdóttir, Tove Odland, Henrik Svärd, and Johan Thunberg.

Lastly, I am grateful for the support from my family and friends.

Stockholm, May 2013

Rasmus Bokrantz

Contents

Notation and terminology

The following mathematical notation and concepts are used in the introduction:

Sets

The absolute value of a finite set denotes its cardinality, and the absolute value of a continuous subset of $\mathbb{R}^3$ its volume. The setwise sum $S+S'$ between two sets S, S' denotes $\{x + x' : x \in S,\ x' \in S'\}$. Setwise differences are defined analogously. The sum between a singleton set $\{x\}$ and a set S is denoted by $x + S$. A set S is said to be *convex* if for any $x, x' \in S$ and $\alpha \in [0, 1]$, it holds that

$$\alpha x + (1 - \alpha)x' \in S.$$

A set S is called a *cone* if $\alpha x \in S$ for any $x \in S$ and $\alpha \geq 0$. The *convex hull* of a set of points $\{x_1, \ldots, x_k\}$ is the set

$$\left\{ \sum_{i=1}^{k} \alpha_i x_i : \sum_{i=1}^{k} \alpha_i = 1,\ \alpha_i \geq 0,\ i = 1, \ldots, k \right\}.$$

A closed halfspace $\{x : a^T x \geq b\}$ defined by some nonzero vector a and a scalar b is said to *support* a set S if

$$a^T x \geq b \text{ for all } x \in S \text{ and } a^T x' = b \text{ for some } x' \in S.$$

A set $\{x : Ax \leq b\}$ defined by some matrix A and vector b is called a *polyhedron*, or a *polyhedral set*.

Functions

A semicolon is used to separate variables from parameters in the arguments of a function. The composition $f \circ g$ of two functions f, g is defined as $f(g(x))$. The

image $f(S)$ of a set S under a function f denotes $\{f(x) : x \in S\}$. A function f is said to be *convex* on a convex set S if for any $x, x' \in S$ and $\alpha \in [0, 1]$, it holds that

$$f(\alpha x + (1 - \alpha)x') \leq \alpha f(x) + (1 - \alpha)f(x').$$

A function f is *concave* if $-f$ is convex. A function f is called *increasing* if $x \leq x'$ implies that $f(x) \leq f(x')$, called *strictly increasing* if $x < x'$ implies that $f(x) < f(x')$, and called *strongly decreasing* if $x \leq x'$ and $x \neq x'$ implies that $f(x') < f(x)$. A function f is *decreasing* if $-f$ is increasing, *strictly decreasing* if $-f$ is strictly increasing, and *strongly decreasing* if $-f$ is strongly increasing. The *expectation* of a function f that depends on a random variable ξ which takes values from a set S is denoted by

$$\mathbb{E}_\pi[f(x; \xi)] = \int_S f(x; s)\pi(s)\, ds,$$

where π is the probability distribution of ξ over S.

Optimization problems

Minimization of a scalar-valued function f according to

$$\underset{x \in X}{\text{minimize}} \quad f(x)$$

is called a *convex problem* if f is a convex function and X a convex set. A point x is *feasible* if $x \in X$, and *optimal* if $x \in X$ and there exists no $x' \in X$ such that $f(x') < f(x)$.

Introduction

Cancer is a leading cause of death worldwide. The mortality rates are particularly high in the western world, where cancer has surpassed cardiovascular disease as the most common cause of death for all but the very elderly (e.g., people younger than 85 years in the US [121]). Many cancers are nevertheless curable. In fact, the expected probability of five year survival after diagnosis is two in three for cancer patients in both Sweden [96] and the US [4] if adjusted for the normal life expectancy in the population. About half of the cancer patients in Sweden [80] and nearly two-thirds of the cancer patients in the US [5] receive radiation therapy during their illness.

This thesis focuses on the decision making during treatment planning for radiation therapy. The forms of decisions that are addressed include:

- Whether to escalate dose in order to improve tumor control, or if to reduce dose in order to avoid normal tissue toxicity
- Whether to protect against large geometric errors, or if to disregard unlikely errors in order to benefit in plan quality with respect to a smaller set of uncertainties
- Whether to sacrifice dose quality in order to shorten the delivery time and thereby decrease the exposure to scatter irradiation and leakage that poses a risk for radiation-induced second cancers

The current standard tools for radiation therapy treatment planning offer limited control of these forms of tradeoffs. Clinical decisions are also often taken with incomplete information because an overview of the possible treatment options is in general not available. The purpose of this thesis is to develop methods that can facilitate more structured and informed decision making during radiation therapy

treatment planning. Improved clinical decisions are ultimately aimed to improve patient care and to make better use of clinical resources.

The thesis is structured into an introductory chapter and six appended papers. The introduction provides background to treatment planning for radiation therapy, formulates the search for the best possible treatment plan as a mathematical optimization problem, and discusses numerical methods to find its solution. The latter part of the introduction poses the treatment planning problem as a multicriteria optimization problem and introduces methods that are aimed to facilitate clinical decision making. The introduction also contains a summary of the appended papers and a discussion on the thesis's main contributions.

1 Radiation therapy

Radiation therapy is a collective term for medical treatments where the patient is exposed to ionizing radiation, the primary application of which is to treat malignant disease. The main delivery techniques are *external beam therapy*, where the patient is irradiated by external fields, and *brachytherapy*, where radioactive seeds are placed within or in the immediate vicinity of the tumor. The purpose of the treatment is generally to deliver a precise radiation dose to a confined target volume that encompasses the malignancy. The absorbed dose in surrounding tissues should simultaneously be minimized in order to avoid damage to healthy organs. Radiation therapy is administered both with the intention to cure and for palliative care, where the goal is to reduce suffering caused by cancer. Cancers where curative treatments are common include tumors in the pelvis, head and neck, lung, and central nervous system. Palliative radiation therapy can be administered for indications such as painful bone metastases and tumors that cause pressure on the spinal cord. Radiation therapy is also commonly used as a complementary treatment for patients that undergo chemotherapy or surgery. Advantages of radiation therapy include that the treatment is non-invasive, potentially organ preserving, and that systemic side effects are generally avoided. Short-term adverse effects include skin burn, fatigue, and sometimes nausea. The possible late side effects depend on the irradiated body site; examples are memory loss, infertility, loss of saliva production, joint problems, and secondary cancers.

1.1 Radiobiology

Ionizing radiation exhibit damage to the cellular DNA through two mechanisms of action: by directly causing ionization events within the DNA, or by inducing the

formation of free radicals that react with the DNA. Most of the radiation-induced DNA lesions can be reversed by cellular repair mechanisms. The repair mechanisms, however, fail with a small probability, which leads to permanent lesions that render the cell unable to undergo cell division. The repair mechanisms of cells in fast proliferating tissue such as tumors generally have an increased likelihood of failure. It is therefore advantageous to partition the treatment into multiple fractions. The treatment fractions are typically delivered with daily intervals, which is a time-scale that permits the cells in normal tissue to recover from the effects of the irradiation. Fractionated delivery also increases the probability that each tumor cell at some point during the treatment is exposed to radiation when it is in a radiosensitive state. The fraction dose and the number of fractions are determined based on the estimated number of tumor cells and their radiosensitivity. A typical fractionation schedule for 10^9 tumor cells with an expected cell kill of 50 % per 2 Gy fraction is 2 Gy $\times$ 30 fractions, which ensures that the expected number of surviving tumor cells is less than one after the last fraction. It is important to note that extinction of all tumor cells at the end of the treatment is often not necessary for long-term survival without recurrence of the cancer: it may instead be sufficient to eradicate the metastatic spread or bring the tumor into partial remission [142]. An extensive overview of radiobiology is contained in Hall and Giaccia [64].

1.2 Intensity-modulated external beam therapy

This thesis focuses on external beam radiation therapy with intensity-modulated fields. External beam treatments constitute about nine-tenths of all radiation therapy treatments [5]. The treatments with intensity-modulated fields are the most sophisticated of the external beam treatments, and also of increasingly more widespread utilization. To exemplify, the fraction of external beam treatments for prostate cancer that in the US were delivered with intensity-modulated fields increased from 0.15 % to 95.9 % between 2000 and 2008 [112].

1.2.1 Treatment machines

The most common medical device for external beam radiation therapy is a linear accelerator that accelerates electrons onto a primary target. The secondary photons that are emitted as the electrons impinge on the target are transmitted through a flattening filter, which produces a therapeutic field with (close to) uniform intensity. The field shape is determined by a *multileaf collimator* (MLC). This device is mounted perpendicular to the radiation field and is composed of pairwise opposing

leaves that can move independently in and out of the treatment field in order to block a fraction of the irradiation, see Figure 1. A given configuration of the MLC leaves is called an *aperture*. The accelerator gantry can be rotated around the patient in order adjust the field incidence angle. The angle of the treatment couch can also be adjusted to allow for non-coplanar fields. The accelerator contains an ionization chamber that quantifies the radiation output in *monitor units* (MUs), which are calibrated to a standardized radiation dose in water. Multileaf collimator-based delivery techniques for photon therapy have been reviewed by Williams [133].

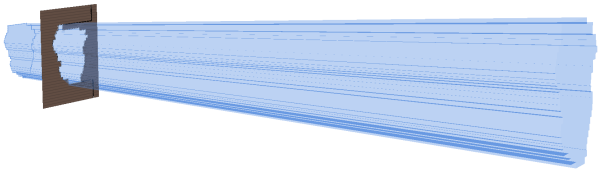

Figure 1. A photon field with field shape determined by tranmission through an MLC.

A small fraction of the external beam treatments are delivered using a narrow beam of accelerated ions that are extracted from a particle accelerator. A therapeutic field is obtained either by *passive scattering*, where the field is broadened through a scattering foil, or by *pencil beam scanning*, where steering magnets are used to scan the particle beam over the target volume. The energy of the incident protons can be adjusted by transmission through a range shifter of variable thickness. A review of beam-delivery techniques for proton therapy is contained in ICRU Report 78 [72].

1.2.2 Physical characteristics

The qualititative differences between proton and photon therapy dose distributions can be understood from the depth-dose curves shown in Figure 2. The photon depth-dose profile shows a short build-up region that is followed by a slow exponential decay. These characteristics make external beam photon therapy best suited for treatment of internal tumors. A relatively large number of overlapping fields ($\sim$5–9) is typically needed to differentiate the absorbed dose in the target volume sufficiently from the absorbed dose in surrounding healthy structures. The depth-dose curve for protons shows a relatively long entrance dose that is followed by a

distinct maximum, which is called the *Bragg peak*. The distal position of the Bragg peak is a function of the proton energy and the density of the traversed medium. The absorbed dose rapidly falls to zero beyond the Bragg peak. A uniform proton dose can be delivered to a spatially extended volume by the superposition of multiple Bragg peaks associated with different energies. The low entrance dose and the lack of exit dose implies that a small number of fields (∼1–3) is often sufficient for proton therapy.

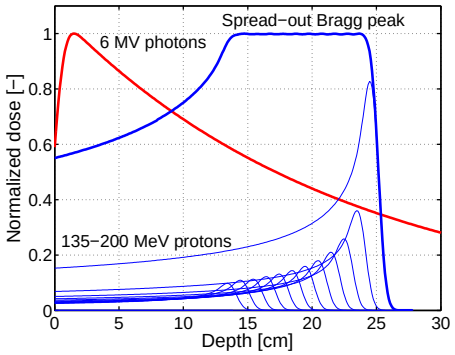

Figure 2. Depth-dose profiles along the central axis for a broad beam of 6 MV photons and 135–200 MeV protons. The superposition of appropriately modulated monoenergetic Bragg peaks produces a spread-out Bragg peak.

1.2.3 Delivery techniques

The term *intensity-modulated radiation therapy* (IMRT) is in this thesis reserved for photon therapy delivered as a set of static beams with modulated intensity[1]. Modulated beam profiles are generated by movements of the MLC, and the accelerator gantry rotated with the field switched off in-between the delivery of one beam to the next. Intensity-modulated radiation therapy is an extension of *three-dimensional conformal radiation therapy* (3DCRT): an earlier delivery technique that uses similar hardware but only a single static aperture per beam. The development of external beam photon therapy from 3DCRT to IMRT has reviewed by Bucci et al. [21].

[1]The term "intensity-modulated" is strictly speaking an abuse of terminology because the field intensity is uniform at any instant in time while the intensity integrated over time (*fluence*) is modulated.

There are two main modes of MLC movements that are used achieve intensity modulation: The first method, called *step-and-shoot* or *segmented MLC* (SMLC), partitions each beam into a set of *segments* that are delivered consecutively. Each segment is defined by a static aperture and a fraction of the beam MU, which is called the *segment weight*. The beam is switched off as the MLC leaves are repositioned before delivery of the next segment. The second method, called *dynamic MLC* (DMLC), uses continuous leaf movement during irradiation. Treatment delivery where the leaves move in unidirectional sweeps back and forth over the beam planes is called *sliding window*. The leaves can either move in synchronized fashion in order to minimize interleaf transmission or in non-synchronized fashion in order to minimize beam-on time. Intensity-modulated radiation therapy has been extensively reviewed, see, e.g., Ahnesjö et al. [1] and Bortfeld [16].

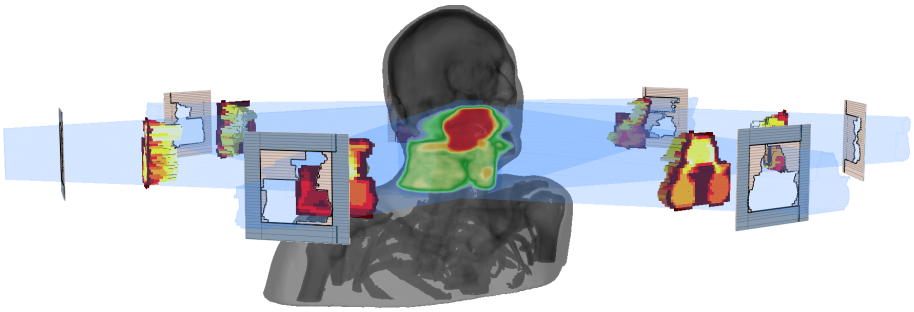

Figure 3. Delivery of an IMRT treatment: The superposition of multiple collimated fields with uniform intensity produces a modulated intensity profile. The depicted treatment plan is designed to deliver a high radiation dose to a tumor located in the nasopharyngeal region and an intermediate dose to surrounding lymphoid tissue.

Volumetric-modulated arc therapy (VMAT) is an extension of IMRT where the gantry rotates continuously during irradiation. Another characteristic property is that the dose rate (the number of MUs delivered per unit of time) and the gantry speed can vary during irradiation in order to allow for modulation in MU as function of gantry angle. These degrees of freedom distinguish VMAT from *intensity-modulated arc therapy* (IMAT): an earlier delivery technique that is limited to constant dose rate and gantry speed[2]. A VMAT treatment can often be de-

[2]The naming conventions for IMAT and VMAT are not consistent in the literature. The definitions given here are in accordance with Webb and McQuaid [131].

livered within a single gantry rotation thanks to its ability to slow down the gantry rotation and increase the dose rate over gantry angle intervals where a high degree of intensity modulation is needed, and the ability to increase the gantry speed and decrease the dose rate over angle intervals where sensitive structures block the field's line of sight. Intensity modulation is mainly achieved using DMLC, but SMLC-like delivery where the irradiation is delivered in high dose-rate burst only when the apertures have been completely formed has also been demonstrated [108]. Volumetric-modulated arc therapy has been reviewed by Yu and Tang [141].

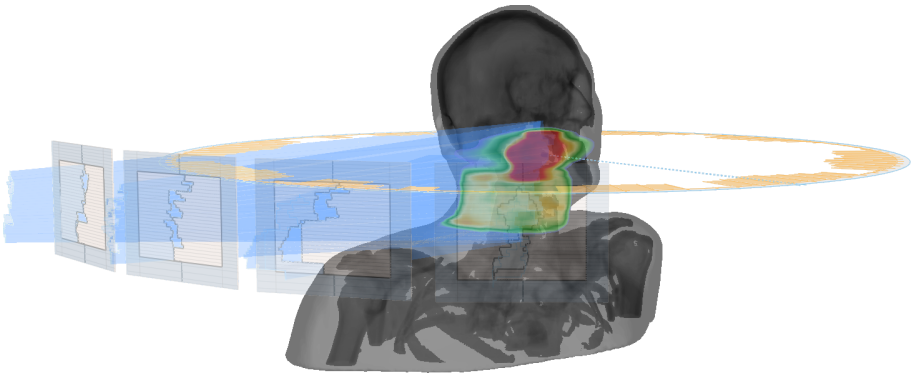

Figure 4. Delivery of a VMAT treatment: The accelerator gantry and the MLC leaves both move continuously during irradiation. The circular histogram depicts the planned number of MUs as a function of gantry angle.

Intensity-modulated proton therapy (IMPT) refers to actively scanned proton therapy where all fields are planned simultaneously. A general IMPT plan is therefore composed of several non-uniform fields that together produce an overall uniform target dose. This delivery technique can be contrasted to single field uniform dose where each beam is planned towards delivery of a uniform dose to the target independent of the other fields. An actively scanned proton beam is represented by a number of *spots*. Each spot is defined by a point in the beam coordinate system and a given particle energy. The fraction of the beam MU that is associated with a given spot is called the spot weight. A therapeutic field with modulated intensity is achieved by varying the spot weights. Intensity-modulated proton therapy has been reviewed by Lomax [83].

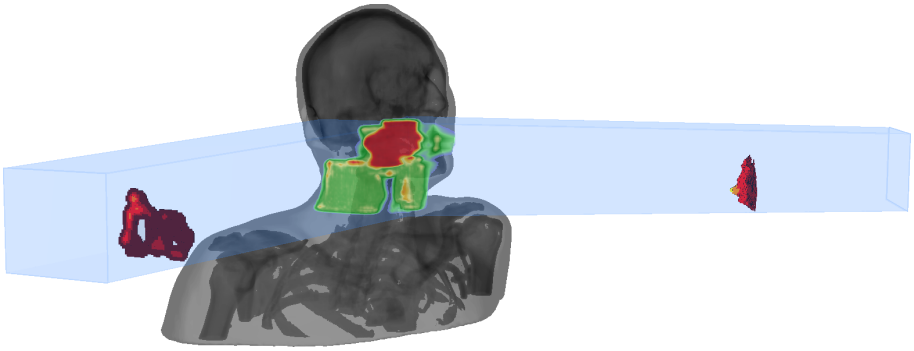

Figure 5. Delivery of an IMPT treatment: The two-dimensional histograms in the beam planes depict the spot weight distribution for a single energy layer.

1.2.4 Clinical benefit

Extensive data show that IMRT is better suited for delivery of concave dose distributions and dose distributions with steep dose gradients than 3DCRT, see, e.g., Purdy [100] and references therein. The improved dose-shaping capabilities of IMRT can be exploited to better spare normal tissue or to escalate dose—and thereby improve local control—in the vicinity of structures that would otherwise be dose-limiting. Comparative trials have shown that IMRT allows safe dose escalation and results in reduced acute and late normal tissue toxicities, see, e.g., Veldeman et al. [126] and Staffurth [115] for reviews and references to the original literature. Despite the benefits of IMRT, there are also some disadvantages. Treatment delivery times for IMRT are generally longer than for 3DCRT, which increases the susceptibility to geometric errors. The absorbed dose outside the fields due to leakage and scatter radiation is also higher for IMRT than for 3DCRT because delivery of IMRT requires two to three times more MUs [65]. Nearly twice as many (1.75 % compared to 1 %) patients are therefore estimated to develop second malignancies within a ten-year period after treatment with IMRT than after treatment with 3DCRT [65]. The advantages and disadvantages of IMRT are both more pronounced for DMLC than for SMLC because DMLC is a comparatively more complex delivery technique that has greater intensity-modulating capabilities, at least in theory. For a discussion on the merits of DMLC and SMLC, see Xia et al. [138].

Volumetric-modulated arc therapy has been extensively compared to IMRT. Most of the conducted planning studies are, however, with respect to dosimetric differences, and the data on clinical outcome is therefore limited. In a review of the current literature, Teoh et al. [119] found that VMAT and IMRT are largely equivalent with respect to target coverage, target homogeneity, and dose conformity with respect to several tumor sites. The significant differences between IMRT and VMAT, according to this review, are that VMAT permits shorter delivery times (about 1–3 minutes compared to 5–15 minutes), that it reduces the total number of MUs, but increases the normal tissue volume that receives low dose radiation.

The physical properties of protons permit proton therapy dose distributions to conform more closely to the target volume than those produced by photon therapy [86]. Proton therapy therefore generally leads to a reduction in dose to healthy structures outside the target volume by a factor two to three compared to photons [61]. This dose reduction makes proton therapy likely to decrease the risk for second cancers [140]; a property that makes proton therapy particularly interesting for treatment of pediatric tumors. An advantage of pencil beam scanning compared to passive scattering is that dose from secondary neutrons that are produced in the scattering foil is avoided [63]. Clinical trials are currently being conducted, but as of April 2013, there is yet very little data available which shows that proton therapy improves clinical outcome compared to photon therapy [93].

2 Treatment planning

The main parameters that need to be determined during treatment planning for intensity-modulated external beam therapy is the number of radiation fields, their orientation, and the intensity modulation of each field. These parameters are most commonly selected both according to the treatment planner's judgment (called *forward planning*) and using computerized automated selection (called *inverse planning*). Forward planning is the most common technique for selection of the number of treatment fields and their orientation, while inverse planning is the only practical method to determine the shape of the intensity profiles.

2.1 Imaging modalities and planning structures

The primary imaging modality for radiation therapy planning is computed tomography (CT), which produces cross-sectional x-ray slices that can be processed into a three-dimensional volume image of the patient volume. Computed tomography can be supplemented by additional functional imaging such as positron emission

tomography or magnetic resonance imaging, which are useful for visualization and staging of regions of pathological tissue.

The CT slices are augmented with contours that specify the location of *regions of interest* (ROIs), such as the target volumes and the *organs at risk* (OARs). The extent of the known macroscopic disease is indicated as the *gross tumor volume*, and the superset of this volume that also contains regions of suspected microscopic disease indicated as the *clinical target volume* (CTV). The planning structures are traditionally delineated by a clinician, but tools that use anatomical atlases for automated segmentation are also becoming available. Details on the delineation of anatomical structures are contained in ICRU Report 62 [71].

2.2 Geometric uncertainties

Radiation therapy is affected by a multitude of errors that can compromise a successful treatment unless they are appropriately accounted for. *Systematic errors* occur during treatment preparation and include errors in the delineation of the ROIs, patient misalignment during image acquisition, and image artifacts due to, e.g., noise, scanner imperfections, and metal implants in the patient's body. *Random errors* occur during the treatment's execution. These errors include daily setup error, inter- and intrafractional organ motion, and anatomic changes induced by tumor shrinkage or weight loss. A systematic error generally results in a shift of the planned dose distribution whereas the accumulated effect of random errors is a blurred total dose (high doses are reduced and lower doses increased). Uncertainties in external beam photon therapy has been reviewed by Van Herk [125], while uncertainties in proton therapy has been reviewed by Lomax [84, 85].

Geometric errors can be mitigated by patient immobilization and positioning according to bony structures or fiducial marker implants (e.g., gold seeds) [12, 76], but not removed entirely. The current recommendation by the International Commission on Radiation Units and Measurements (ICRU) for both photon [73] and proton [72] therapy planning is therefore to expand the CTV into a *planning target volume* (PTV), with the size of the CTV-PTV margin selected so that a certain coverage probability is achieved according to the population distribution of the systematic and random errors. The ICRU also recommends that geometric margins are used for OARs, in particular for structures where a high dose to a small subvolume is sufficient to cause complication.

2.3 Plan evalation criteria

The primary evaluation criteria for assessment of treatment plan quality are dose-volume indices according to:

- *Dose-at-volume* D_x: the highest dose such that at least x % of a given ROI receives this dose or higher
- *Volume-at-dose* V_x: the fraction of the volume of a given ROI that receives a dose of x Gy or higher

Assessment of plan quality with respect to dose-volume statistics is supported by the fact that the current evidence-based knowledge about the outcomes from radiation therapy is mainly with respect to such data [53, 88].

Dose-volume indices constitute points in the (cumulative) *dose-volume histogram* (DVH): a graph that depicts cumulative volume as function of dose for a given ROI, see Figure 6. Such graphs are a standard tool for plan quality assessments because they allow for instantaneous visual inspection of the dose-volume effects in all relevant structures.

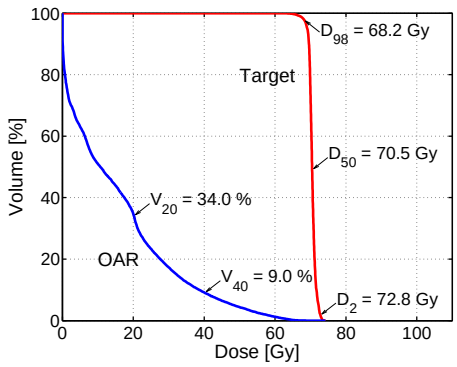

Figure 6. Examples of DVH curves and dose-volume levels for a target structure (red) and an OAR (blue).

Two complementary evaluation criteria are the *homogeneity index* (HI) [73] and the *conformity index* (CI) [71], which are defined according to

$$\mathrm{HI} = \frac{\mathrm{D}_2 - \mathrm{D}_{98}}{\mathrm{D}_{50}} \qquad \text{and} \qquad \mathrm{CI} = \frac{\text{treated volume}}{\text{target volume}},$$

where the treated volume is the volume enclosed by the isodose surface defined at 95 % of the prescription. The homogeneity index measures the dose uniformity within a target volume and has an ideal value of zero. This index is defined with respect to the near minimum and maximum dose (D_{98} and D_2) instead of the exact minimum and maximum because the exact values are susceptible to numerical outliers in the dose distribution data. The conformity index measures how closely the high dose region conforms to the target volume and has an ideal value of one. Appropriate use of this index requires that the treated volume entirely encompasses the target volume.

Another plan evaluation criterion is the *tumor control probability* (TCP), which quantifies the probability that the number of surviving clonogenic tumor cells is zero at the end of the treatment. This probability is calculated from an analytical model of the tumor's response to irradiation. A common assumption is that the survival fraction (SF) of cells exposed to the radiation dose d is described by a linear-quadratic model according to

$$\mathrm{SF} = e^{-\alpha d - \beta d^2}, \tag{1}$$

where α and β biological model parameters that are fitted to empirical data [56]. A related biological evaluation criterion is the *normal tissue complication probability*, which quantifies the probability that some clinical endpoint occurs. The ICRU recommends that biological evaluation criteria are used with caution due to uncertainty in the model parameters [73].

Physical and biological criteria for evaluation of radiation therapy treatment plans have been reviewed by Romeijn and Dempsey [104].

3 Treatment plan optimization

3.1 Problem formulation

Early work on treatment planning for IMRT [18, 33, 34] considered treatment planning as an inverse problem where the desired dose distribution is known (the prescription to all targets and zero dose elsewhere) and the fluence distribution that produces this dose the unknown. The fluence distribution that best realizes the desired dose was then found by analytical inversion. An issue, however, is that the energy fluence can and do become negative at the solution. Direct inversion has therefore been largely abandoned in contemporary treatment planning in favor of numerical optimization techniques. The patient geometry is for optimization

purposes discretized into volume elements called *voxels*, and the beam planes discretized into surface elements called *bixels*.

Inverse planning for radiation therapy is in this thesis posed as a mathematical optimization problem by the introduction of n *objectives* $f_1, \ldots, f_n$ that are to be minimized with respect to some variables x. The objectives are aggregated into a single scalar-valued measure by the introduction of nonnegative weights $w_1, \ldots, w_n$, which are chosen in order to reflect the relative importance of the objectives. Planning aims that must be entirely fulfilled are posed as *constraints* $c_1, \ldots, c_m$ that are required to evaluate to zero or less. The vector x is required to be contained in a set $\{x : Ax \leq b\}$ defined by some matrix A and vector b, which corresponds to the parameter values which can be physically realized. The inverse planning problem is thus formulated

$$\begin{array}{lll} \underset{x}{\text{minimize}} & \displaystyle\sum_{i=1}^{n} w_i f_i(x) & \text{(composite objective function)} \\ \text{subject to} & c_j(x) \leq 0, \quad j = 1, \ldots, m, & \text{(planning constraints)} \\ & Ax \leq b. & \text{(physical constraints)} \end{array} \tag{2}$$

Minimization is considered without loss of generality because maximization of some objective f can be equivalently posed as minimization of $-f$.

The variables x are in this thesis limited to the parameters that determine the intensity modulation of the radiation fields. The treatment planning problem is furthermore assumed to be solved only once and the optimized treatment plan then kept identical during all treatment fractions. Beam orientation optimization and adaptive replanning are both active areas of research, but not within the scope of this thesis.

3.2 Optimization functions

The objectives functions and the planning constraints are in practice chosen in order to reflect the evaluation criteria that are used to judge plan acceptability. Three forms of physical optimization functions are studied in this thesis, namely purely dose-based functions, functions of the equivalent uniform dose (EUD), and functions of the DVH. Biological optimization is also considered with respect to maximization of TCP. More comprehensive reviews of optimization functions for radiation therapy treatment planning are available in Romeijn [105] and Hoffmann et al. [67].

3.2.1 Mathematical formulation

All optimization functions f are for clarity stated with the dose distribution d as the argument and the underlying dependence on x omitted, i.e., $f(d) = f(d(x))$. All criteria are also defined with respect to a single ROI and the subset of the patient volume associated with this ROI denoted by V. The physical criteria are stated on two variants using a linear ramp Θ that is given either by $\Theta(y) = \min\{y, 0\}$ or $\Theta(y) = \max\{y, 0\}$. The choice of ramp function is indicated by the prefix "min" or "max."

Good optimization functions should not only accurately model the clinical goals, but also be differentiable and convex in order to be suitable for optimization. Convexity can often be verified by expressing a function as the composition $h \circ g$ of two functions h, g where either of any of these three conditions hold [17]:

- h is increasing and convex and g convex
- h is decreasing and convex and g concave
- h is convex and g linear

Two properties that will be referenced later are that the function $\min\{y, 0\}^2$ is convex and decreasing and the function $\max\{y, 0\}^2$ convex and increasing.

A *dose function* imposes a penalty on deviation between the dose distribution d and a reference dose level $\hat{d}$ according to

$$f(d) = \int_V \Theta(d(v) - \hat{d})^2 \, dv, \tag{3}$$

where $d(v)$ denotes the dose at a point v in V. The direct sum between a min dose function and a max dose function that have identical reference dose level is called a uniform dose function. Uniform dose functions are infinitely many times continuously differentiable, and min and max dose functions one time continuously differentiable. The integrand of a dose-based function are a composition of a convex function (the squared ramp) and a linear function (its argument) and therefore convex. The integral in (3) preserves convexity, and dose-based functions are hence convex.

An *EUD function* is obtained if a uniform dose with equivalent biological effect is substituted for d in (3). The EUD is in this thesis calculated according the generalized mean

$$\mathrm{EUD}_a(d) = \left(\int_V d(v)^a \, dv \right)^{1/a},$$

for some $a \neq 0$. This quantity, which is originally due to Niemierko [97], permits continuous scaling between the minimum dose ($a \to -\infty$) and the maximum dose ($a \to \infty$). Noteworthy special cases are the harmonic mean ($a = -1$), the geometric mean ($a \to 0$), and the arithmetic mean ($a = 1$). The EUD level is convex if $a \geq 1$ and concave if $0 \neq a \leq 1$ [29]. Min EUD functions are therefore a composition of a convex and decreasing function and a concave function and hence convex if $0 \neq a \leq 1$. Max EUD functions are a composition of a convex and increasing function and a convex function and hence convex if $a \geq 1$. Min EUD and max EUD functions are both one time continuously differentiable.

A *DVH function* imposes a penalty on deviation between the DVH curve associated with a given subvolume V and the reference dose level $\hat{d}$ according to

$$f(d) = \int_I \Theta(D(v;d) - \hat{d})^2 \, dv,$$

where $I = (0, \hat{v}]$ for min DVH and $I = (\hat{v}, 1]$ for max DVH, for some reference volume $\hat{v}$ in $(0, 1]$, and where D is a function that takes cumulative volumes to dose-at-volume levels according to

$$D(v;d) = \max \left\{ d' \in \mathbb{R} : \frac{|\{v' \in V : d(v') \geq d'\}|}{|V|} \geq v \right\}.$$

Dose-volume histograms functions are nonconvex [48] and nonsmooth.

A *TCP function* quantifies the probability that the number of clonogenic tumor cells is zero after the last treatment fraction. The TCP model used in this thesis assumes that the cell kill follows Poisson statistics and does not take repopulation into account. Let the discrete random variable N denote the number of clonogenic tumor cells at the end of treatment. Then, the probability that N equals some integer k is given by the Poisson probability mass function according to

$$\mathbb{P}(N = k) = \frac{\mathbb{E}[N]^k e^{-\mathbb{E}[N]}}{k!},$$

which together with (1) yields an expression for the TCP according to

$$\mathrm{TCP}(d) = \mathbb{P}(N = 0) = e^{-\int_V \rho(v) e^{-\alpha d(v) - \beta \frac{d(v)^2}{n_f}} \, dv}, \tag{4}$$

where $\rho(v)$ is the density of clonogenic tumor cells at a point v in V and n_f the number of treatment fractions. Maximization of TCP according to (4) is convex

under a logarithmic transformation provided that the dose distribution satisfies [67]

$$d(v) > \sqrt{\frac{n_f}{2\beta}} - \alpha\frac{n_f}{2\beta}, \qquad v \in V.$$

Maximization of TCP is studied in Paper F.

3.2.2 Practical use

Optimization with respect to uniform dose functions with reference dose level set to the prescription for targets and to zero otherwise gives the least-squares solution to the inverse problem of radiation therapy. The calculation of this solution has mathematically favorable properties, but generally results in an unacceptable underdosage of the target volume [15]. A common remedy for the underdosage is to use max dose functions for OARs that have the reference dose level increased from zero to a positive threshold that is chosen sufficiently small to avoid complication. Dose-based functions are used in all of the appended papers.

It is also common to augment the formulation of the treatment planning problem with DVH functions, because clinicians have extensive experience with DVH criteria and are aware of how they affect outcome. Max DVH functions are useful for OARs that exhibit a large volume effect, such as the lung or liver, where it is acceptable to deliver a high dose to a subvolume of the organ as long as this subvolume is relatively small. Min DVH functions can be useful if the PTV overlaps with a dose-limiting OAR, where it can permit a controlled underdosage of a subvolume of the PTV. Dose-volume histogram functions are used in Papers B and C.

Structures that exhibit a large volume effect can also be modeled using EUD functions, which have better numerical properties than functions based on the DVH. The parameter a is for EUD functions chosen in order to reflect tissue architecture. Negative values are used for targets, while small positive values are used for OARs where damage to a single functional subunit causes loss of function (*serial organs*, e.g., the spinal cord and esophagus). Larger positive values are used for OARs where loss of function only occurs after damage to a considerable fraction of the functional subunits (*parallel organs*, e.g., the lung and parotid glands). Equivalent uniform dose functions are used in Papers B and C.

A natural extension of (3) is to make $\hat{d}$ a function of the spatial position v, or to introduce a spatially variable weight in the integrand. A variable reference dose is useful for dose-painting: a technique where tumorous regions with increased cell density or higher radioresistance are identified using functional imaging, and then prescribed with a higher dose [11]. Another application of a variable reference dose

is *dose fall-off functions*, where the reference level decreases with distance from the nearest target structure in order to prevent hot spots that are remotely located from targets. Dose fall-off functions are used in Papers B and C. A spatially variable weight has been suggested in order to incorporate the coverage probabilities of the CTV directly in the optimization (in contrast to margins created with respect to such probabilities), see, e.g., Bohoslavsky et al. [13] and references therein. Voxel-specific weights are also used in several methods for fine-tuning of the current dose distribution that are discussed in Section 4.7.

There is considerable debate on whether the nonconvexity of DVH functions leads to local minima or not. Some authors report that if the number of variables are large, then the local minimizers are sufficiently close to the global minimizers for the search after global minimizers to become inconsequential [75, 82, 136]. Other authors report that some local minimizers differ considerably from the global minimizer [135]. Approximate DVH criteria with better numerical properties have nevertheless been suggested: Romeijn et al. [103] observed that a min DVH criteria corresponds to optimization of the minimum dose in the $1 - \hat{v}$ fraction of the volume receiving the highest dose (the value-at-risk), and thereby proposed to optimize the average dose received by this subvolume (the conditional value-at-risk), which is a convex measure. Analogous convex approximations are also possible for max DVH criteria. Zinchenko et al. [147] showed that a DVH curve is determined uniquely by an infinite sequence of EUD criteria [147], and proposed to approximate DVH criteria by minimization of a sequence of EUD functions applied to the difference between the current dose and a reference dose that satisfies the criterion. Halabi et al. [62] formulated DVH criteria using integer variables and then solved the linear relaxation of the resulting mixed-integer program. A large number of heuristics have also been used to account for DVH criteria during optimization, see, e.g., Lan et al. [78] and Zarepisheh et al. [143] for recent summaries of the relevant literature.

3.3 Treatment plan optimization methods

The two main methods for photon therapy optimization is to either consider the energy fluence per bixel as directly controllable variables or to incorporate the underlying dependence on the physical parameters in the optimization. The former approach is called *fluence map optimization* (FMO) and the latter called *direct machine parameter optimization* (DMPO). Treatment plan optimization for IMPT is performed with the spot weights as variables, which has equivalent mathematical

properties with FMO. Fluence map optimization is used in all of the appended papers while DMPO is used in Papers B, C, and E.

3.3.1 Fluence map optimization

A clear advantage of FMO is that the relationship between fluence and dose is linear. This property makes FMO problems convex whenever the objectives and the planning constraints are convex in dose because composition with a linear function preserves convexity. Convex programs can be solved efficiently to global optimality because every local minimizer is also a global minimizer. An FMO problem is generally less computationally expensive to solve than a general linearly constrained problem of the same size because the only physical constraint in FMO is a bound that prevents negative fluence.

A proton therapy plan generated by spot weight optimization is directly deliverable by pencil beam scanning whereas a photon therapy plan optimized by FMO requires conversion by *leaf-sequencing* into deliverable apertures. A survey on leaf-sequencing methods is contained in Ehrgott et al. [50]. Two important results are that leaf-sequencing which minimizes MU can be solved in polynomial time [2, Theorem 1] whereas leaf-sequencing that minimizes the number of apertures is strongly NP-hard [6, Theorem 4.1], meaning there is no algorithm that can find the optimal solution within a polynomial bound on the running time (unless P = NP).

Leaf-sequencing causes a degradation of dose quality, in particular for irradiation of complex target geometries. This degradation is partially due to the fact that an optimal fluence profile often is highly jagged and therefore difficult to decompose into a finite number of apertures. Jaggedness can be counteracted by inclusion of a stabilizing penalty on variation in the fluence planes in the objective function of (2). Several authors have proposed quadratic variational penalties [31, 89, 114], which promote smooth fluence profiles. *Total variation penalties*, which penalize linear variation, have later been shown to promote piecewise constant fluence maps that are better suited for leaf-sequencing [145, 146]. Total variation regularization is utilized in Papers B and D. Variational penalties on fluence are convex and do therefore not interfere with the convexity properties of FMO. Other regularization methods include upper bounds on the admissible fluence [35] and iterative regularization [24], where the optimization is terminated after a relatively small number of iterations.

3.3.2 Direct machine parameter optimization

Direct machine parameter optimization methods consider the leaf positions and the segment weights as variables during optimization (and possibly also other parameters such as the gantry, couch, and collimator angles). The physical constraints are posed on a form that reflects limitations on the apertures, such as interdigitation, connectedness, bounds on the minimum leaf tip gap, bounds on the minimum aperture area, and bounds on the minimum segment weight. Variables that determine the gantry speed and the dose rate are also included during VMAT optimization.

It is considerably more difficult to solve a DMPO problem than its FMO counterpart. This difficulty arises because the relation between leaf positions and fluence is both nonlinear and nonconvex. Background and illustration of this nonconvexity is provided in Figure 7. Direct machine parameter optimization thus amounts to nonconvex optimization regardless of the convexity of its FMO counterpart. A large number of optimization methods have been proposed in order the tackle DMPO, including the following:

- *Simulated annealing methods* [47,98,113], which introduce random changes to the variables and retain feasible changes that improve the objective function value. The algorithm also retains solutions with worse objective values with some probability in order to permit escape from local minima.

- *Column generation methods* [90, 102], which alternate between a subproblem where the new aperture that maximizes the improvement in objective function value is identified and a master problem where the segment weights of the apertures generated so far are optimized.

- *Gradient-based methods* [22,66], which consider the leaf positions and segment weights as variables simultaneously. The referenced methods use a treatment plan generated by FMO and leaf-sequencing as the initial point, and then use first order derivative information and approximate second order derivative information to improve the solution.

- *Genetic algorithms* [37,79], which attempt to mimic the process of natural evolution by maintaining a population of solutions where the fittest individuals are randomly recombined in order to evolve the population.

- *Heuristic methods*, which are based on deterministic rules for the admissible configurations of the leaves [9,47], or similar rules that are defined over the course of the algorithm's execution (tabu-search) [122].

Methods that combine column generation and gradient-based search have also been proposed [23, 26]. Gradient-based optimization methods for DMPO are used in Papers B, C, and E.

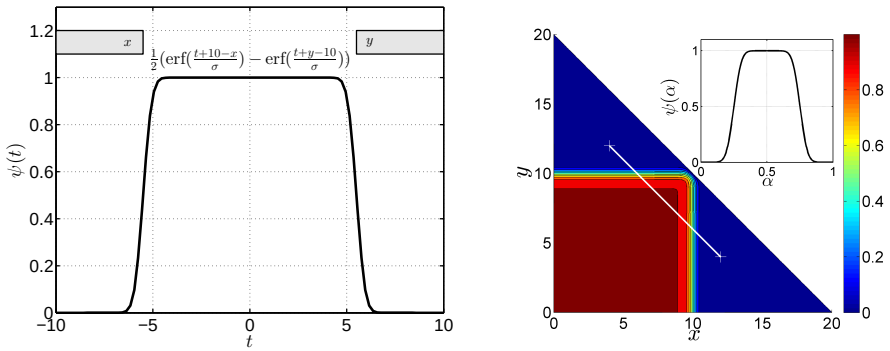

(a) (b)

Figure 7. Relation between leaf positions (x, y) and energy fluence ψ for a two-leaf MLC that blocks a single Gaussian-shaped fluence source with standard deviation σ. (a) The fluence $\psi(t)$ at a point t in the fluence plane is determined by integration over the visible parts of the fluence source, which gives an expression defined by sigmoidal error functions [55]. (b) Fluence at $t = 0$ as a function of x and y. The leaf positions are constrained to $x + y \leq 20$ in order to avoid collision. The inset depicts the trace along the white line segment parameterized by $\alpha \in [0, 1]$, and illustrates that $\psi(0)(x, y)$ is jointly non-convex in x and y.

Direct machine parameter optimization has been shown to both reduce the number of apertures and the number of MUs, and simultaneously provide dose distributions of comparable or improved quality to those generated by FMO and leaf-sequencing, see Broderick et al. [20] for a review and references to the original literature.

4 Multicriteria optimization

4.1 Motivation

The treatment planner's task to find values for the objective weights that condense all clinical requirements into a single number is clearly not trivial. Complicating

issues are that the weights lack a direct clinical interpretation, and that they are chosen without knowledge about how realistic the objectives are to fulfill. It is also in general not known how the objectives are correlated and therefore hard to know how to change the weights in order to introduce a desired modification to the dose distribution. The lack of overview of the possible treatment options also makes it is difficult to know when to terminate the search for better treatment plans. A further difficulty is that the optimized plan often is very sensitive to the choice of weights [70, 130], and tumor-site specific protocols therefore of limited usefulness. It is common to counteract this sensitivity by a relaxation of the planning criteria into requirements that are easier to attain. Such relaxation, however, poses the risk that the requirements become too weak and the optimized treatment plan therefore suboptimal to the initial formulation with sharp criteria.

The problems associated with weights contribute to the fact that treatment planning often is a time-consuming process that involves a considerable amount of manual parameter tuning. Multiple studies have found that treatment plan quality is strongly dependent on both the time-commitment and the experience level of the individual treatment planner [7, 14, 30].

4.2 Multicriteria formulation

The main focus of this thesis is on a generalization of the treatment planning problem to a *multicriteria optimization problem* where the objectives are viewed as components of a vector-valued function and explicit weight factors thereby avoided. The multicriteria counterpart of problem (2) is given by

$$\begin{array}{ll} \underset{x}{\text{minimize}} & \Big[f_1(x) \cdots f_n(x)\Big]^T \\ \text{subject to} & c_j(x) \leq 0, \qquad\qquad j = 1, \ldots, m, \\ & Ax \leq b. \end{array} \tag{5}$$

The notation f is subsequently used to denote the vector of objective function, and the feasible set of (5) denoted by X. The interesting situation occurs when the feasible set is nonempty and no feasible solution exists at which each objective f_i, $i = 1, \ldots, n$, attains its minimum value over X simultaneously. This situation makes (5) a decision problem in the sense that the subjective preferences of some *decision maker* must be taken into account in order for the minimizer to be mathematically well-defined. In practice, the decision maker is the single person that is responsible for the approval of the clinical plan (or a similar group of persons). Decision-making aspects of radiation therapy treatment planning are discussed in the review by Moore et al. [95].

The dependence on a decision maker's preferences is more explicitly shown in an alternative formulation of problem (5) according to

$$\underset{x \in X}{\text{maximize}} \quad u(f(x)) \tag{6}$$

where $u : \mathbb{R}^n \to \mathbb{R}$ is a *utility function* such that the decision maker prefers a feasible x to feasible x' if $u(f(x)) > u(f(x'))$ and is indifferent if $u(f(x)) = u(f(x'))$. A closed-form expression of the decision maker's utility is not available in practice. Multicriteria optimization can therefore equivalently be viewed as optimization under uncertainty in the decision maker's preferences.

Multicriteria optimization methods are in this thesis classified according to the decision maker's participation. The considered classes are summarized below (the classification is adapted from Miettinen [92]):

- *A priori methods*, where the decision maker articulates preferences between the objectives before the optimization
- *A posteriori methods*, where an unbiased approximation of all Pareto optimal solutions is first calculated and the best available alternative then selected by the decision maker
- *No preference methods*, which do not involve an active decision maker
- *Interactive methods*, where the decision maker gradually articulates preferences during the solution process

This thesis mainly focuses on a posteriori methods, which are considered in all of the appended papers. The studied a posteriori methods use an a priori method as a subroutine to generate representations of the Pareto set. An interactive method is studied in Paper B. No preference methods and previous interactive methods are reviewed for completeness. Comprehensive reviews of multicriteria optimization methods are contained in the monographs by Miettinen [92] and Ehrgott [49].

4.3 Pareto optimality

The most common definition of optimality with respect to a vector-valued optimization problem is that a feasible solution is optimal if there exist no other feasible solution that is at least as good in all objectives, and strictly better in at least one objective. Solutions that satisfy this nondominance criterion are called *Pareto*

optimal. A more formal definition is that a feasible solution x^* is Pareto optimal if there exists no feasible x such that

$$f_i(x) \leq f_i(x^*) \text{ for all } i = 1, \ldots, n \text{ and } f_j(x) < f_j(x^*) \text{ for some } j.$$

An equivalent statement is that a feasible x^* is Pareto optimal if there is no feasible x such that

$$f(x) \in f(x^*) - (\mathbb{R}^n_+ \setminus \{0\}). \tag{7}$$

This relation is illustrated in Figure 8(a), which also depicts the *feasible objective space* $Z = f(X)$, the *dominated objective space* $Z_+ = Z + \mathbb{R}^n_+$, the *ideal point* z^{id} and the *nadir point* z^{nad}. The ideal point is the n-vector where the ith component is given by the minimal value of f_i over the Pareto optimal set, and the nadir point the corresponding vector defined with respect to maximization. The ideal and nadir points are in all of the appended papers used for normalization of the objective function to comparable magnitudes according to

$$f_i \;\mapsto\; \frac{f_i - z_i^{\text{id}}}{z_i^{\text{nad}} - z_i^{\text{id}}}, \qquad i = 1, \ldots, n,$$

where it is assumed that $z_i^{\text{id}} < z_i^{\text{nad}}$ holds for $i = 1, \ldots, n$. The depicted situation where the Pareto optimal set forms a convex surface in $\mathbb{R}^n$—meaning a connected surface in the boundary of a convex set—in general only occurs for convex problems [105, Proposition 2.3].

Pareto optimality is a necessary condition for optimality with respect to maximization of utility according to (6) if the decision maker's preferences are *rational* [92, Theorem 2.6.2], meaning that smaller objective values are always preferred to larger ones, or equivalently, that the utility function is strongly decreasing. The connection between Pareto optimality and maximization of utility is illustrated in Figure 8(b): the optimum to problem (6) occurs where the partial derivatives of the Pareto optimal set (the *marginal rates of transformation*) equals the negative of the partial derivatives of the utility function (the *marginal rates of substitution*), if these partial derivatives exist.

4.4 A priori methods

A priori methods attempt to translate the decision maker's preferences into a scalar-valued function that is amenable to optimization. Such methods are therefore often synonymously called *scalarization methods*.

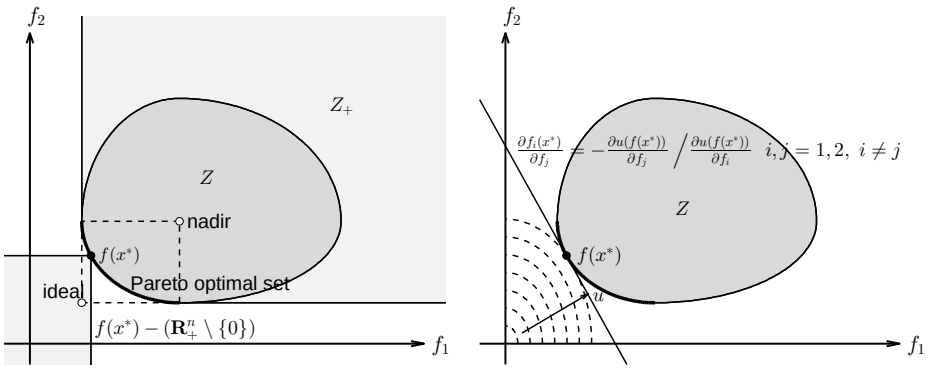

Figure 8. (a) Principle of Pareto optimality: The set Z is indicated in dark gray, the set Z_+ indicated in light gray, and the Pareto optimal set indicated by a thick solid line. The ideal and nadir points are indicated in white. The feasible solution x^* is Pareto optimal because no Pareto optimal point is contained in $f(x^*) - (\mathbb{R}^n_+ \setminus \{0\})$. (b) The Pareto optimal solution x^* is optimal to maximization of a strongly decreasing utility function u. The dashed lines indicate indifference curves with respect to u. The linearization of u equals the linearization of the Pareto optimal set at $f(x^*)$.

The method of weighted sum minimization introduced in Section 3 is an a priori method where the marginal rates of substitution are specified explicitly, and these rates assumed to everywhere constant, see Figure 9(a). The n-vector of objective weights w thus specifies an inwards oriented normal vector to the feasible objective space Z at the optimum. The optimal solution to weighted sum minimization with respect to positive weights is Pareto optimal [92, Theorem 3.1.2]. Optimality to weighted sum minimization for some nonnegative weights is also a necessary condition for Pareto optimality if the optimization problem is convex [92, Theorem 3.1.4], meaning that all Pareto optimal solutions to a convex problem can be obtained by weighted sum minimization.

The *ε-constraint* method is another common a priori method. The decision maker here specifies upper bounds ε_i, $i = 1, \ldots, n$, $i \neq \ell$, for all but one objective, indexed by ℓ, which is minimized. The ε-constraint method constitutes the basic operation for *lexicographic ordering*, a technique where the objectives are first hierarchically ordered and then optimized consecutively subject to the constraint that

all previously optimized objectives must not deteriorate. Lexicographic optimality is a sufficient condition for Pareto optimality [92, Theorem 4.2.1].

The ε-constraint method is a special case of a more general method of *reference point-based optimization* according to

$$\begin{array}{ll} \underset{x,t}{\text{minimize}} & t \\ \text{subject to} & f_i(x) - tp_i \leq \bar{z}_i, \quad i = 1, \ldots, n, \\ & x \in X, \end{array}$$

where $\bar{z}$ is a reference point and p a nonnegative directional vector, both in $\mathbb{R}^n$. This form of optimization corresponds to an assumption on that the marginal rates of substitution are vanishing around a point on the line $\{\bar{z} + tp : \ t \in \mathbb{R}\}$, as illustrated in Figure 9(b). The ε-constraint method is obtained if ε_i is set to $\bar{z}_i$ for $i = 1, \ldots, n$, $i \neq \ell$, and to infinity for $i = \ell$; and p_i set to zero for $i = 1, \ldots, n$, $i \neq \ell$, and to unity for $i = \ell$.

Lexicographic ordering has been studied with respect to IMRT by Jee et al. [74] and Wilkens et al. [132], and studied with respect to IMPT by Falkinger et al. [54]. The latter two studies allowed for a slight relaxation of the constraints in the ε-constraint method using a preselected "slip factor." Wilkens et al. demonstrated that a slip factor is beneficial in treatment planning for IMRT—the solution can otherwise be entirely determined by the first objective—whereas Falkinger et al. found that a slip factor is less crucial for IMPT plans due to the additional degrees of freedom introduced by modulation in depth. A posteriori methods that rely on reference point-based optimization are discussed in the next section, and parallelization of such methods discussed in Paper D. Reference point-based optimization is also utilized in Paper C.

4.5 A posteriori planning

4.5.1 Pareto set approximation

The a posteriori methods studied in this thesis calculate finite representations of Pareto optimal sets using an a priori method that is solved a repeated number of times with respect to different model parameters. Selection of parameters that produce a well-distributed set of Pareto points is, however, nontrivial. Weighted sum minimization imposes no explicit requirements on the location of the Pareto optimal points. Uniformly distributed weights therefore only results in uniformly distributed Pareto points under very special conditions [46]. Reference point-based optimization requires that each Pareto optimal point is contained on the line

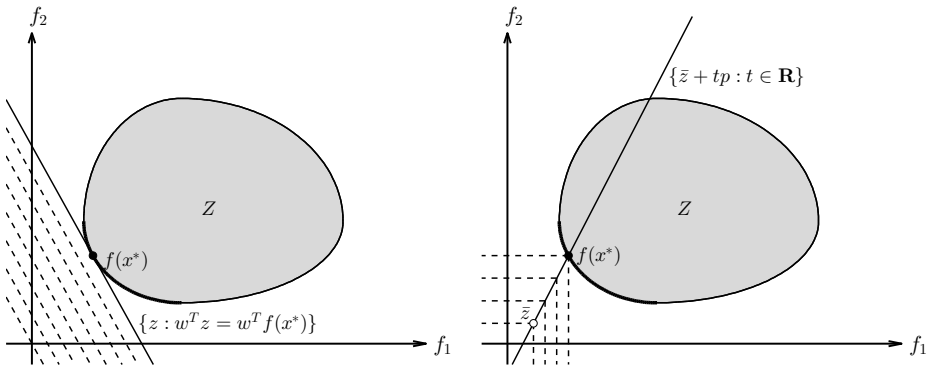

(a) Weighted sum minimization (b) Reference point-based optimization

Figure 9. A priori multicriteria optimization methods. The dashed lines indicate level set curves for the scalarizations.

$\{\bar{z} + tp : \ t \in \mathbb{R}\}$, or that it dominates a point on this line. Uniformly distributed reference points therefore produce uniformly distributed Pareto points [10, 46, 91]. The disadvantage is that it is general not possible to select reference points from a large enough set to obtain the entire Pareto optimal set and simultaneously ensure that all reference points are projected onto this set [91, 110]. Difficulties related to preselected objective weights or reference points are further discussed in Paper D.

Methods that adapt to the shape of the Pareto optimal set have been proposed to account for the difficulties associated with preselected parameters. Küfer et al. [77] proposed to calculate a triangulation of the current set of Pareto points and then iteratively bisect the longest edge of this triangulation. Craft et al. [40] exploited the convexity of Z_+ that occurs for convex problems in order to construct inner and outer approximations of the Pareto optimal set, as illustrated in Figure 10. These approximations are iteratively refined by addition of one point at the time at the location where the distance between the approximations attains its current maximum. Sandwich approximations of this form are if implemented using weighted sum minimization associated with the difficulty that the normal vectors of the bounding faces of the convex hull of the known Pareto points can have negative components in dimension three and higher. These normal vectors are therefore not immediately useful as objective weights. Craft et al. used a heuristic to transform normals with mixed components into nonnegative vectors. Rennen et al. [101] later proposed to augment the inner approximation by setwise summation with $\mathbb{R}^n_+$, which

produces a convex polyhedron that has everywhere nonnegative inwards oriented normal vectors. The algorithm of Rennen et al. is further developed in Paper A, where techniques that improves its computational efficiency are suggested.

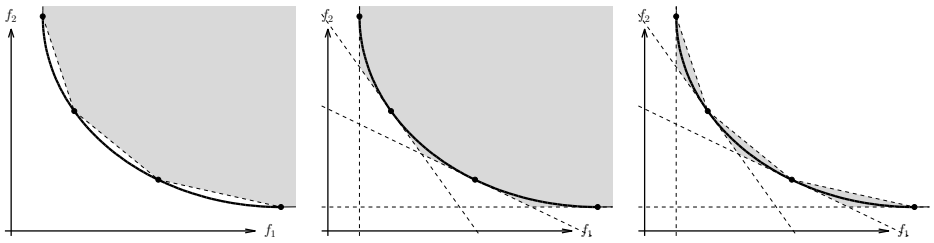

(a) Inner approximation (b) Outer approximation (c) Sandwich approximation

Figure 10. Sandwich approximation of the Pareto surface (thick solid line) using inner and outer polyhedral approximations (dashed lines).

Sandwich approximations that rely on the ε-constraint method have been studied by Hoffmann et al. [68]. These authors also gave a variant of the sandwich method where derivative information is not needed. Sandwich approximations where the vertices of the outer approximation constitute reference points $\bar{z}$ for reference-point based optimization have been suggested by Shao and Ehrgott [111] and Ehrgott et al. [52]. Serna et al. [109] also considered sandwich approximations, and showed that the substitution of a general ordering cone C for the Pareto cone $\mathbb{R}^n_+$ in the definition of Pareto optimality according to (7) is useful to disregard nonrelevant parts of the Pareto optimal set. A smaller set of Pareto optimal solutions is obtained if $\mathbb{R}^n_+ \subseteq C$. Multicriteria optimization with respect to general ordering cones is considered in Paper A.

4.5.2 Pareto set navigation

An equally important aspect of a posteriori planning is the technique used for selection of the clinical plan from the Pareto set representation. Discrete selection of one of the precalculated plans requires a very dense representation that is not feasible to generate in practice. A much coarser representation is sufficient if continuous interpolation between a discrete set of solutions $\{x_1, \ldots, x_k\}$ is allowed

according to

$$\sum_{j=1}^{k} \lambda_j x_j, \quad \text{where} \quad \sum_{j=1}^{k} \lambda_j = 1,\ \lambda_j \geq 0,\ j = 1, \ldots, k. \tag{8}$$

Continuous interpolation of this form is often called *navigation*, see Monz et al. [94]. Navigation is best suited for convex FMO formulations because any convex combination of fluence-based plans is feasible if the feasible set is convex. Convexity in the objectives also enables the objective function values of an interpolated treatment plan to be bounded according to

$$f_i\left(\sum_{j=1}^{k} \lambda_j x_j\right) \leq \sum_{j=1}^{k} \lambda_j f_i(x_j), \qquad i = 1, \ldots, n, \tag{9}$$

i.e., the objective values of an interpolated solution are at least as good as the linear interpolation of the objective values of its constituents. An interpolated solution is therefore near-optimal if the representation of the Pareto set is sufficiently accurate and the coefficients λ in (8) restricted to values such that the right-hand side of (9) defines a point that is not dominated by any other point in the convex hull of $\{f(x_1), \ldots, f(x_k)\}$.

Interfaces for navigation typically contain slider bar controls for each objective function, as exemplified in Figure 11. A movement of a slider results in an update of the current convex combination, and corresponding updates of the displayed dose and DVH. A slider movement also causes the other sliders to move, with the direction of movement being dependent on how the objectives are correlated. The current position of the sliders can be clamped to prevent movements toward increased objective function values. The update of the convex coefficients in (8) can be solved as a reference point-based optimization problem where the reference point is determined by the slider positions [43, 94]. This optimization problem can be solved to approximate optimality in real-time if the exact objective values are approximated by their conservative bound according to (9), which yields a linear programming formulation.

4.5.3 Deliverability

A posteriori planning with respect to fluence-based treatment plans need not sacrifice plan quality compared to an a priori method which utilizes DMPO if the a

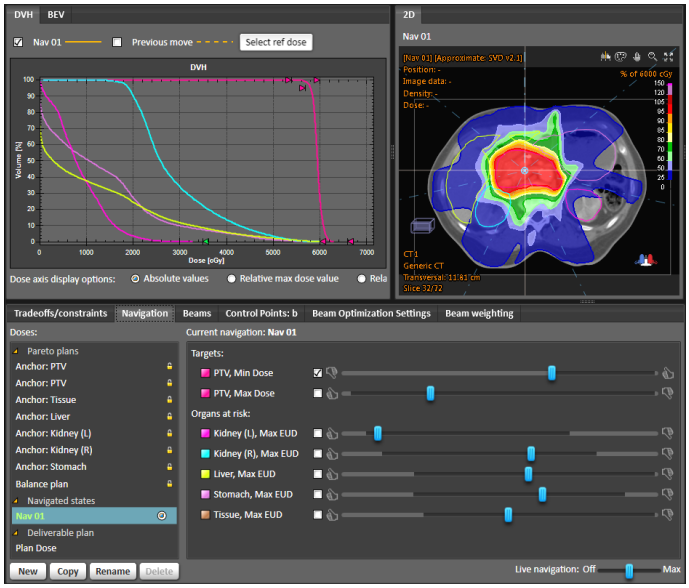

Figure 11. User interface for Pareto surface navigation in the treatment planning system RayStation 2.9 (RaySearch Laboratories, Stockholm, Sweden). The depicted example shows exploration of the tradeoffs between target homogeneity, dose conformity, and sparing of the kidneys, liver and stomach for a pancreas cancer case. The clamp on the min dose objective for the PTV restricts the feasible movements for the remaining sliders.

posteriori method relies on DMPO for conversion of the navigated plan into deliverable apertures. Conversion of fluence-based treatment plans using DMPO has been studied by Craft et al. [39]. These authors proposed a formulation where the error in DVH due to the conversion is minimized using *reference DVH functions* that penalize discrepancy between the current DVH and the DVH associated with a reference dose distribution d^{ref} according to

$$f(d) = \int_0^1 \Theta(D(v;d) - D(v;d^{\text{ref}}))^2 \, dv.$$

Reference DVH functions are used in Papers B and C.

A conversion subsequent to the navigation is nevertheless a complicating factor because the decisions during navigation are then taken with incomplete information about the final deliverable dose distribution. A large discrepancy between the

navigated plan and its best approximation by a deliverable plan can at worst force the treatment planner to make fine-adjustments to the navigated plan and then repeat the conversion, which results in the form of iterative workflow that a posteriori planning is intended to avoid. This shortcoming is addressed by methods for *deliverable navigation* that use precalculated treatment plans on a form such that any convex combination between plans is directly deliverable. Deliverable navigation thus ensures that the treatment plan seen during navigation is exactly the treatment plan that is approved for delivery.

Two methods for deliverable navigation have been proposed for step-and-shoot IMRT. Craft and Richter [44] proposed to generate step-and-shoot plans with s segments, and then restrict the number of positive coefficients in the convex combinations to k, where $1 < k \ll n$, so that a general navigated treatment plan can be delivered within ks segments. Salari and Unkelbach [107] used column generation to calculate a collection of treatment plans that all use a single set of apertures. The subproblems in the column generation method were posed on a form such that apertures that contribute to multiple treatment plans were generated. A navigated plan from this representation is directly deliverable and no limitation exists on the number of positive components in the convex combinations between plans. The methods of Craft and Richter and Salari and Unkelbach are in Paper E unified into a single method where a subset of the apertures are shared across plans and the remaining apertures treated as individual. Deliverable navigation for VMAT is also studied in Paper B.

4.5.4 Clinical benefit

A posteriori planning for IMRT has been investigated in a several retrospective planning studies where treatment plans generated by weighted sum minimization are used as benchmark. Thieke et al. [120] studied a paraspinal and a prostate case and found that a posteriori planning resulted in plans of comparative quality to the benchmark plans, while only requiring in the order of 10 minutes for the navigation step. Hong et al. [69] studied ten pancreatic cancer cases and found that navigation consistently produced acceptable treatment plans within 10 minutes. A posteriori planning also led to different clinical judgment: a lower stomach mean dose was consistently selected, often at the expense of increased kidney dose. Craft et al. [41] studied five glioblastoma and five pancreatic cancer cases and found that a posteriori planning reduced planning time from 135 to 12 minutes on average, while physician involvement time increased from 5 to 9 minutes on average. The plan generated by the a posteriori method was in this study blindly identified as

superior to the benchmark plan for all patient cases. Wala et al. [128] studied nine prostate cases and observed navigation times of approximately 10 minutes per case. These authors also found that physicians ranked the navigated plan as superior to the benchmark plan for all cases.

4.6 No preference methods

No preference methods permit fully automated planning for routine cases, and can provide a starting point for manual planning for more challenging clinical scenarios. Some heuristic is typically used to incorporate previous clinical experience in order to ensure that relevant treatment plans are generated. The dependence on previous data makes the different no preference methods rather disparate from a mathematical point of view.

Several no preference methods are techniques for automated selection of objectives functions and associated weights for weighted sum minimization: Xing et al. [139] proposed an iterative procedure where the objective weights are updated in-between each solve in order to maximize a plan quality score based on DVH criteria. Zhang et al. [144] also solved weighted sum problems in repeated fashion and updated the problem formulation according to a preset list of rules. Xhaferllari et al. [137] similarly updated the formulation by introduction of objectives on the removal of cold and hot spots. Wu et al. [134] used the degree of volume overlap between targets and OARs as a predictor for attainable DVH criteria and queried the current patient geometry against a database of previously treated patients in order to define DVH objectives for weighted sum minimization. Two optimizations were on average found sufficient to generate acceptable treatment plans if the DVH criteria generated by the no preference method were used, whereas 28 optimizations on average were necessary for standard weighted sum minimization.

No preference methods based on lexicographic ordering have been proposed by Clark et al. [32] and Breedveld et al. [19], who used tumor site-specific protocols to define the ordering of the objectives. The no preference method of Breedveld et al. has been compared to weighted sum minimization in a prospective planning study on IMRT for head and neck cancers by Voet et al. [127]. A physician was in this study asked to identify the superior alternative between a plan generated by the no preference method and a plan generated by weighted sum minimization, which resulted in the plan generated by the no preference method being selected on 32 of 33 occasions.

4.7 Interactive methods

There are few fundamentally interactive methods for radiation therapy treatment planning. A rare example is a method for brachytherapy planning proposed by Ruotsalainen et al. [106] where treatment plans are optimized according to a classification of the objectives into those that should be improved (possibly to a given aspiration level), those that should be maintained, and those that are allowed to deteriorate (possibly to a given bound). The decision maker alters the classification in-between each solve after inspection of the current solution.

Several authors have also proposed methods that offer some interactivity in the form of a fine-tuning step, which is useful for the situation when a close to acceptable treatment plan has been generated. Methods that eliminate hot spots using locally increased importance weights have been suggested by Cotrutz and Xing [36] and Lougovski et al. [87]. Süss et al. [117] proposed a similar technique where the removal of hot spots is ensured by constraints, and the error with respect to the current treatment plan then minimized. A suitable balance between elimination of the hot spots and perturbations of the otherwise acceptable treatment plan is here found by continuous navigation between the initial and the corrected treatment plan.

4.8 Extensions

4.8.1 Robustness tradeoffs

Tradeoffs related to robustness against geometric errors can be accommodated by methods that explicitly account for the realization of errors during optimization. The sources of uncertainty are in such methods typically discretized into a set of *scenarios* S. The scenario where no error occurs is called the nominal scenario.

Scenario-based optimization is in this thesis only considered with respect to treatment planning for IMPT. A clear benefit for this treatment modality is that the planned dose distribution can be calculated separately for each scenario and an assumption on the static dose cloud approximation (which predictates that a rigid shift of the patient volume produces a rigid shift of the dose distribution) thereby avoided. The static dose cloud approximation is inaccurate for ion beams due to their strong density dependence. Geometric margins that rely on the static dose cloud approximation have in several studies on treatment planning for IMPT consequently been shown to offer inadequate robustness against geometric errors [3, 28, 58].

Scenario-based optimization is in this thesis considered either by *stochastic programming* or *robust optimization*. Stochastic programming methods minimize the expected value of the objective function conditioned on some estimated probability distribution π, which if applied to problem (2) yields the formulation

$$\underset{x\in X}{\text{minimize}} \quad \mathbb{E}_\pi \left[\sum_{i=1}^{n} w_i f_i(x;\xi) \right] \ , \tag{10}$$

where ξ is a random variable that picks a scenario from S. Robust optimization methods minimize the worst case objective function value, which if applied to problem (2) yields the formulation

$$\underset{x\in X}{\text{minimize}} \quad \max_{s\in\mathcal{S}} \sum_{i=1}^{n} w_i f_i(x;s) \ . \tag{11}$$

It is important to note that robust optimization according to (11) attaches no probability to the scenarios in S. Such optimization is therefore often suitable if no reliable estimate on the probability distribution is available, or against non-repeated uncertainty (systematic errors) where it is sensible to hedge against as large uncertainties as possible in order to ensure a high probability of a satisfactory outcome. Stochastic programming is generally speaking suitable for repeated uncertainties (random errors) where the average outcome tends towards the expected outcome as the number of stochastic events becomes large.

Stochastic programming for IMPT has been studied by Unkelbach et al. [123, 124] and robust optimization for IMPT by Fredriksson et al. [58] and Fredriksson [57]. Treatment plan optimization methods for IMPT that consider the dose to each voxel as independent and protect against the voxel-wise worst case have been suggested by Unkelbach et al. [124], Chan et al. [27], Pflugfelder et al. [99], and Liu et al. [81]. The accuracy of the dose prediction by scenario-wise worst case and voxel-wise worst case has been studied by Casiraghi et al. [25], who found that the less conservative scenario-wise approach provides accurate prediction of the DVH whereas voxel-wise worst case results in overly conservative predictions of the DVH.

Stochastic programming according to (10) can be directly extended to a vector-valued multicriteria program because

$$\mathbb{E}_\pi \left[\Big[f_1(x;\xi) \cdots f_n(x;\xi) \Big]^T \right] = \Big[\mathbb{E}_\pi[f_1(x,\xi)] \cdots \mathbb{E}_\pi[f_n(x;\xi)] \Big]^T, \tag{12}$$

which holds since expectation is a linear operation (that also preserves convexity). The multicriteria counterpart of problem (10) is thus given by

$$\underset{x\in X}{\text{minimize}} \quad \Big[\mathbb{E}_\pi[f_1(x,\xi)]\cdots\mathbb{E}_\pi[f_n(x;\xi)]\Big]^T \ ,$$

which is equivalent to a deterministic multicriteria program under the substitution

$$f_i \ \leftrightarrow \ \mathbb{E}_\pi[f_i(x,\xi)], \qquad i=1,\ldots,n. \tag{13}$$

Robust multicriteria optimization is possible both in analogy with the right-hand side of (12), which yields the vector-valued formulation

$$\underset{x\in X}{\text{minimize}} \quad \Big[\max_{s\in S} f_1(x;s)\cdots\max_{s\in S} f_n(x;s)\Big]^T \ , \tag{14}$$

and in analogy with the left-hand side of (12), which yields the set-valued formulation

$$\underset{x\in X}{\text{minimize}} \quad \max_{s\in S}\Big[f_1(x;s)\cdots f_n(x;s)\Big]^T \ . \tag{15}$$

The set-valuedness occurs because the inner maximization over S is a multicriteria program in itself for each fixed x.

Objective-wise worst case optimization for IMPT according to (14) has been studied by Chen et al. [28]. This formulation is equivalent to a deterministic multicriteria program under a substitution in direct analogy with (13), which preserves convexity because the componentwise maximum of a set of convex functions is convex. *Worst case optimization* with respect to (15) has been studied from a theoretical viewpoint by Ehrgott et al. [51]. The standard definition of Pareto optimality is not directly applicable for characterization of optimality with respect to this formulation due to the set-valuedness of the objective function. Ehrgott et al. therefore proposed an extension of Pareto optimality where a feasible x^* is defined as Pareto optimal if there exists no feasible x such that

$$f(x;S) \subseteq f(x^*;S) - (\mathbb{R}^n_+ \setminus \{0\}). \tag{16}$$

Note that an equivalence holds between this definition and the standard definition of Pareto optimality according to (7) if S contains only the nominal scenario. Worst-case optimization according to (15) is less conservative than objective-wise worst case according to (14) because it only protects against the $|S|$ scenarios that are physically realizable, as opposed to $|S|^n$ combinations between scenarios where the majority are nonphysical.

Tradeoffs in robustness and conservativeness are studied in Paper F. Robustness here reflects the magnitude of errors that are accounted for while conservativeness reflects the degree of protection against variability in the estimated probability distribution.

4.8.2 Time-efficiency tradeoffs

A short treatment delivery time is beneficial not only with respect to patient throughput, but also because it reduces the exposure to scatter irradiation and leakage that poses a risk for second cancers, and because it makes the treatment less susceptible to intrafraction motion. Minimization of treatment delivery time can for sliding-window IMRT be posed as a penalty on the sum of the maximum positive variation in each fluence plane. This penalty is a convex function that is directly proportional to the number of MUs required for delivery [116]. A posteriori planning that includes minimization of the number of MUs as an objective has been studied by Craft et al. [45], who found that considerable MU reductions often are feasible at a very small penalty in dose distribution quality. Penalties on MU in treatment planning for sliding-window IMRT is studied in Paper C. The same fundamental mechanism is also used in the study on treatment planning for VMAT by Craft et al. [42], where fluence profiles at adjacent gantry angles are merged together subsequent to the optimization until the treatment delivery time is sufficiently small. The estimated treatment delivery time of VMAT plans is in Paper B handled by a rigid constraint during DMPO.

4.8.3 Beam orientation and delivery technique tradeoffs

The choice of an appropriate delivery technique, such as VMAT against seven- or nine-field IMRT, is important both in order to make optimal use of clinical resources and in order to maximize the dose quality for the individual patient. A practical method to investigate different delivery techniques is to calculate a Pareto set representation with respect to each relevant alternative and then navigate between these representations in a discrete manner. This form of navigation has been studied by Craft and Monz [43] and Teichert et al. [118] with respect to the closely related problem of choosing an optimal gantry angle configuration. Both of the two references studies found that the optimal gantry angle configuration varies over the Pareto optimal set. Another tradeoff related to beam orientation is the choice of start and stop gantry angles for partial arcs in treatment planning for VMAT. This decision has been studied by Wala et al. [129], who modified the method for mul-

ticriteria VMAT optimization proposed by Craft et al. [42] into a method that finds a partial arc that minimizes treatment delivery time while maintaining dose distribution quality within a user-specified bound.

5 Numerical optimization

Treatment plan optimization is in this thesis primarily considered as general nonlinear program that is solved by *sequential quadratic programming* (SQP). An SQP method finds successively better approximations $\{x_k\}_{k\geq 0}$ of a solution to the *Karush-Kuhn-Tucker* (KKT) conditions. Let F denote the composite objective function of (2) and L the Lagrangian function

$$L(x, \lambda) = F(x) + \sum_{j=1}^{m} \lambda_j c_j(x), \tag{17}$$

where λ is the m-vector of Lagrange multipliers associated with the constraints $c_1, \ldots, c_m$. Then, the KKT conditions of problem (2) are given by

$$\begin{array}{ll} c_j(x) \leq 0, \quad j = 1, \ldots, m, \\ Ax \leq b, & \text{(feasibility)} \\ \nabla_x L(x, \lambda) = 0, & \text{(stationarity)} \\ \lambda \geq 0, & \text{(nonnegativity of the multipliers)} \\ \lambda_i c_j = 0, \quad j = 1, \ldots, m. & \text{(complementarity)} \end{array}$$

The KKT conditions are necessary conditions for global optimality if some constraint qualification holds, such that the gradients of active constraints are linearly independent or the optimization problem convex and strictly feasible (Slater's condition). The KKT conditions are also sufficient conditions for local optimality at a primal-dual feasible solution (x^*, λ^*) if $p^T \nabla^2_{xx} L(x^*, \lambda^*) p > 0$ holds for all $p \neq 0$ in the nullspace of the active constraints. These results are contained in Bazaraa and Shetty [8], which also contains an overview of constraint qualifications for nonlinear programming.

Sequential quadratic programming is for equality constrained optimization problem equivalent to Newton's method applied to the KKT conditions. Newton's method finds the roots to a system $g(x) = 0$ by iterative steps $x_{k+1} = x_k + p_k$, where the Newton step p_k is the solution to $\nabla g(x_k) p_k = -g(x_k)$. The Newton step is also the optimal solution to a second order approximation of the initial problem

around the current iterate, which is a definition that is valid also for inequality constrained problems. A quadratic programming approximation of problem (2) around x_k is given by

$$\begin{array}{ll} \underset{p}{\text{minimize}} & \frac{1}{2}p^T\nabla^2_{xx}L(x_k,\lambda_k)p + \nabla F(x_k)^T p \\ \text{subject to} & c_j(x_k) + \nabla c_j(x_k)^T p \leq 0 \qquad\qquad j = 1,\ldots,m, \\ & A(x_k + p) \leq b. \end{array} \tag{18}$$

The full step p_k defined by the optimal solution to (18) is in general not taken. The iterates x_k are instead updated according to $x_{k+1} = x_k + \alpha_k p_k$, where the step length α_k is calculated by minimization of a merit function M according to

$$\underset{\alpha>0}{\text{minimize}} \quad M(x_k + \alpha p_k, \lambda_k + \alpha d_k). \tag{19}$$

The step d_k in the dual variables is given by $\lambda_{k+1} - \lambda_k$, where λ_{k+1} is the vector of optimal Lagrange multipliers associated with the first m constraints in (18). A common form of M is an augmented Lagrangian defined by the introduction of an additional term to (17) that penalizes violation of the constraints to (2).

Problem (18) is a quadratic program that can be solved efficiently if the Hessian of the Lagrangian $\nabla^2_{xx}L$ is positive semi-definite and (18) thereby convex. The calculation of $\nabla^2_{xx}L$ is, however, costly if the number of variables x is large. A *quasi-Newton* approximation of this matrix is therefore often used for speed. The approximate Hessian is iteratively updated based on first order derivative information, with the update rule designed so that positive-semidefiniteness is retained. It is also common to use approximate line-search, meaning that problem (19) is solved to approximate optimality by evaluation of M with respect to a discrete set of step lengths.

A quasi-Newton SQP method developed by RaySearch Laboratories (Stockholm, Sweden) is used for treatment plan optimization in Papers B–E. The SQP code SNOPT (Stanford Business Software, Inc., Stanford, California), which implements the quasi-Newton SQP method described in Gill et al. [59], is used for treatment plan optimization in Paper F. A general review of SQP methods is provided by Gill and Wong [60].

6 Summary and main contributions

6.1 Summary of the appended papers

Paper A: An algorithm for approximating convex Pareto surfaces based on dual techniques

Paper A is co-authored with Anders Forsgren and has been published in *INFORMS Journal on Computing*, Vol. 25, No. 2, pp. 377-393, 2013.

This paper concerns efficient approximation of Pareto surfaces to convex multicriteria problems. An algorithm is proposed that generate well-distributed sets of Pareto optimal points by calculation of one point from the Pareto surface at the time where the distance between inner and outer approximations of this surface currently attains its maximum. The inner approximation relies on that the convex hull of a set of points is the smallest convex set that contains the points; while the outer approximation relies on that a supporting halfspace exists at any point in the boundary of a convex set. The inner approximation also converges towards the Pareto surface because a closed convex set is the convex hull of the points in its boundary, while the outer approximation converges towards the Pareto surface because a closed convex set is the intersection of its supporting closed halfspaces. It is shown that the calculation of the maximum distance between the inner and outer approximations amounts to maximization of a convex function over a polyhedral set. This nonconvex optimization problem is solved by enumeration of the finitely many vertices of the outer approximation and a solve of a small linear program for each vertex. Vertex enumeration is demonstrated to be orders of magnitude more efficient than a previous technique [101] where the facets of the inner approximation are instead enumerated. The suggested improvements make sandwich approximation of Pareto surfaces practical in up to about ten dimensions.

Paper B: Multicriteria optimization for volumetric-modulated arc therapy by decomposition into a fluence-based relaxation and a segment weight-based restriction

Paper B has been published in *Medical Physics*, Vol. 39, No. 11, pp. 6712-6724, 2012.

This paper considers multicriteria optimization for VMAT that utilizes navigation in two separate stages. The first stage is performed with respect to fluence-based treatment plans, and serves the purpose of defining a coarse tradeoff between objectives. The second stage is performed with respect to deliverable VMAT plans

that share the same set of apertures, and is intended for fine-tuning purposes. The fluence-based plans in stage one are generated by FMO subject to a total variation penalty, which limits the intensity modulation to a degree that is attainable by deliverable VMAT plans. The plan from this representation that is navigated to by the user is converted into a deliverable VMAT plan by DMPO towards minimization of the error in DVH due to the conversion. The VMAT plans that form the basis for the second navigation stage are generated by segment weight optimization with respect to the apertures of the initial converted VMAT plan. The treatment plans generated by the suggested technique are compared to benchmark plans generated by DMPO towards minimization of a weighted sum of the objectives. Numerical results with respect to treatment planning for prostate, pancreas, lung, and head and neck cancer show that the suggested multicriteria method generates VMAT plans that are of comparable quality to the benchmark plans, and simultaneously enables for interactive decision making.

Paper C: Improved plan quality in multicriteria radiation therapy optimization by projections onto the Pareto surface

Paper C is co-authored with Kaisa Miettinen and has been printed as Technical report TRITA-MAT-2012-OS4, Department of Mathematics, Royal Institute of Technology, 2012.

This paper addresses the fact that a general navigated treatment plan has a nonzero approximation error to Pareto optimality because the Pareto set representation is finite. It is shown that the approximation error can be eliminated by reference point-based optimization subsequent to the navigation. This optimization maximizes the minimal improvement upon the objective values of the navigated plan, and also ensures that a Pareto optimal treatment plan is obtained which is at least as good as the navigated plan with respect to all objectives. An augmented formulation is also proposed that ensures that the DVH curves of targets for the final treatment plan are at least as uniform as those of the initial navigated plan, and that the DVH curves of healthy structures never exceed the navigated DVH. The versatility of the suggested technique is demonstrated by application to planning for step-and-shoot IMRT, planning form sliding-window IMRT, and planning for IMPT. Improvements in normal tissue sparing and dose conformity due to the projections are demonstrated for all three delivery techniques.

Paper D: Distributed approximation of Pareto surfaces in multicriteria radiation therapy treatment planning

Paper D has been published in *Physics in Medicine and Biology*, Vol. 58, No. 11, pp. 3501–3516, 2013.

This paper generalizes the method for sandwich approximation of Pareto surfaces presented in Paper A to an algorithm that can take advantage of distributed computational environments. Parallelization is made feasible by application of the previous sequential method to approximation of an inexpensive model of the Pareto surface. The model is used to predict the outcome of each iteration in the sequential algorithm, and the inner and outer approximations then updated accordingly. The objective weights gathered with respect to the model are used for solves of the exact problem that are performed in parallel. The number of plans k that are generated in each batch of solves permits scaling between the previous sequential algorithm ($k = 1$) and a fully parallel algorithm ($k = m$), where m is the total number of plans to be computed. The model of the Pareto surface is given a shape that makes it difficult to approximate by piecewise linear bounds in order to avoid that some parts of the Pareto surface are incorrectly disregarded. A model with a smooth shape also gives the approximation method a behavior that approaches optimization over uniformly distributed weights as $k \to m$. The algorithm's performance as a function of k is studied with respect to a prostate, brain, and pancreas cancer case. Parallelization is demonstrated to yield approximations of comparable quality to those generated by the sequential method for k up to $2n$, where n is the number of objectives. This result translates to speed-up of one order of magnitude in practice because the number of objectives n is typically at least five.

Paper E: Deliverable navigation for multicriteria intensity-modulated radiation therapy planning by combining shared and individual apertures

Paper E is co-authored with Albin Fredriksson and has been printed as Technical report TRITA-MAT-2013-OS4, Department of Mathematics, Royal Institute of Technology, 2013.

This paper concerns deliverable navigation for step-and-shoot IMRT. A deliverable Pareto set representation is calculated by DMPO subject to constraints that some apertures are shared across all treatment plans. Some apertures are also allowed to be individual, and all segment weights treated as individual. Navigation thus produces treatment plans that are deliverable within $s_{\text{sh}} + ks_{\text{ind}}$ apertures,

where s_{sh} and s_{ind} are the number of shared and individual apertures, respectively, and k the number of allowed positive coefficients in the convex combinations between plans. The coupling between the treatment plans imposed by the shared apertures implies that the calculation of the Pareto set representation is not a separable optimization problem in each treatment plan. All treatment plans are therefore optimized simultaneously, so that an optimal pool of shared apertures is obtained. A more elaborate method where apertures are only shared across k-tuples of plans are also discussed, which enables convergence towards the ideal (non-navigable) Pareto set representation where all apertures are individual. Numerical results with respect to a two- and three-objective formulation for a paraspinal case show that a few individual apertures lead to much increased plan quality. The results regarding optimal partitioning into shared and individual apertures are inconclusive.

Paper F: Controlling robustness and conservativeness in multicriteria intensity-modulated proton therapy optimization under uncertainty

Paper F is co-authored with Albin Fredriksson and has been printed as Technical report TRITA-MAT-2013-OS5, Department of Mathematics, Royal Institute of Technology, 2013.

This paper considers multicriteria optimization under uncertainty, posed as objectives and constraints that depend on some random variable. Worst case optimization is used to calculate robust solutions, and mathematical theory for this form of optimization elaborated. Worst case optimization is also contrasted to objective-wise worst case, which is a previous more conservative method [28]. Empirical results are presented for planning of IMPT with respect to a one-dimensional patient geometry subject to systematic setup uncertainty. These results show that worst case optimization better exploits spatial structure and better adapts to a decision maker's preferences than objective-wise worst case. Worst case optimization consequently yields superior dose distributions in the physically realizable scenarios. Tradeoffs in robustness and conservativeness are also studied using a generalization of worst case optimization to *minimax stochastic programming*, which minimizes the worst case expectation subject to upper and lower bounds on the admissible probability distributions. Tradeoffs in robustness are shown to qualitatively correspond to a contraction (or expansion) of the dose distribution, whereas tradeoffs in conservativeness lead to a gradual smoothing (or sharpening) of the lateral dose fall-off. The steep dose fall-off produced by a more conservative method is empirically demonstrated to be preferable with respect to population tumor control.

6.2 Main contributions

The original contributions of this thesis span four topics:

More efficient Pareto surface approximation

- Paper A shows that it is orders of magnitude less costly to calculate the maximum distance between inner and outer approximations of the Pareto surface by vertex enumeration over the outer approximation than by facet enumeration over the inner approximation. This contribution increases the number of dimensions where sandwich approximations can be practically computed from about five to ten, which is sufficient to cover the majority of the formulations that occur in clinical practice [38].

- Paper D introduces the first parallelizable algorithm for sandwich approximation of Pareto surfaces. Batchwise calculation of $2n$ treatment plans at the time, where n is the number of objectives, is demonstrated to be feasible while simultaneously retaining the approximation quality of the finite Pareto set representation. Treatment plan generation for a challenging treatment case with ten conflicting objectives that takes 4 hours with a sequential technique thus becomes feasible within about 12 minutes.

Extensions to new delivery techniques

- Paper B introduces the first method for multicriteria VMAT optimization which utilizes DMPO, and also the first method for deliverable navigation. Direct machine parameter optimization methods are well-established to produce plans that require fewer MUs than those generated by FMO, and simultaneously provide comparable or improved dose quality.

- Paper F provides the first method for multicriteria IMPT optimization where robustness is ensured by worst case optimization. Worst case optimization is preferable to objective-wise worst case because it does not sacrifice plan quality in the physically realizable scenarios in order to gain in performance in nonphysical combinations between the scenarios. Worst case optimization is also shown to be favorable to stochastic programming with respect to population tumor control.

Improvement of treatment plan quality

- Paper C shows that the error to Pareto optimality due to a finite Pareto set representation can be eliminated by a projection onto the Pareto optimal set subsequent to the navigation. This projection can also be posed on a form such that no clinical goals that are initially satisfied become violated due to the projection.

- Paper E extends deliverable navigation to the use of some apertures that are from a collective pool and some apertures that are individual. A small number of individual apertures is shown to greatly improve plan quality. This contribution constitutes a step towards the long standing goal of directly deliverable navigation with respect to the unrestricted Pareto set where all apertures are individual.

Theoretical contributions

- Paper F introduces the concept of *convex hull efficiency*. Novel results associated with this concept are that the convex hull efficient solutions can be found by convex scalarizations, and that this set of solutions is a proper subset of the Pareto optimal solutions in the general situation. Optimality with respect to a strongly increasing convex scalarization is shown to be a sufficient condition for convex hull efficiency, and optimality with respect to a strictly increasing convex scalarization shown to be a necessary condition for convex hull efficiency. The set of convex hull efficient solutions are shown to equivalent to the Pareto optimal solutions with respect to worst case expectation minimization.

6.3 Contributions by co-authors

The study reported in Paper A was performed in collaboration with Anders Forsgren, who supervised the study and suggested directions of research. Kaisa Miettinen had a similar role during the work reported in Paper C. The studies reported in Papers E and F were performed jointly with Albin Fredriksson. The theoretical work reported in these studied were performed in close collaboration, and the computational experiments were also jointly designed. The first author of each paper, Albin Fredriksson for Paper E and I for Paper F, performed the majority of the numerical experiments.

7 Bibliography

[1] A. Ahnesjö, B. Hårdemark, U. Isacsson, and A. Montelius. The IMRT information process—mastering the degrees of freedom in external beam therapy. *Phys. Med. Biol.*, 51(13):R381–R402, 2006.

[2] R. Ahuja and H. Hamacher. A network flow algorithm to minimize beam-on time for unconstrained multileaf collimator problems in cancer radiation therapy. *Networks*, 45(1):36–41, 2005.

[3] F. Albertini, E. Hug, and A. Lomax. Is it necessary to plan with safety margins for actively scanned proton therapy? *Phys. Med. Biol.*, 56(14):4399–4413, 2011.

[4] American Cancer Society. *Cancer Facts & Figures 2012*. American Cancer Society, Atlanta, GA, 2012.

[5] American Medical Association. *Physician Characteristics and Distribution in the United States*. American Medical Association, Atlanta, GA, 2008.

[6] D. Baatar, H. Hamacher, M. Ehrgott, and G. Woeginger. Decomposition of integer matrices and multileaf collimator sequencing. *Discrete Appl. Math.*, 152(1):6–34, 2005.

[7] V. Batumalai, M. Jameson, D. Forstner, P. Vial, and L. Holloway. How important is dosimetrist experience for intensity modulated radiation therapy? A comparative analysis of a head and neck case. *Pract. Radiat. Oncol.*, Published online before print on August 2, 2012.

[8] M. Bazaara and C. Shetty. *Nonlinear Programming: Theory and Algorithms*. John Wiley & Sons, New York, NY, 1979.

[9] J. Bedford and S. Webb. Constrained segment shapes in direct-aperture optimization for step-and-shoot IMRT. *Med. Phys.*, 33(4):944–958, 2006.

[10] H. Benson and S. Sayin. Towards finding global representations of the efficient set in multiple objective mathematical programming. *Nav. Res. Log.*, 44(1):47–67, 1997.

[11] S. Bentzen. Theragnostic imaging for radiation oncology: dose-painting by numbers. *Lancet Oncol.*, 6(2):112–117, 2005.

[12] C. Bert and M. Durante. Motion in radiotherapy: particle therapy. *Phys. Med. Biol.*, 56(16):R113–R144, 2011.

[13] R. Bohoslavsky, M. Witte, T. Janssen, and M. van Herk. Probabilistic objective functions for margin-less IMRT planning. *Phys. Med. Biol.*, 58(11):3563–3580, 2013.

[14] J. Bohsung, S. Gillis, R. Arrans, A. Bakai, C. de Wagter, T. Knöös, B. Mijnheer, M. Paiusco, B. Perrin, H. Welleweerd, and P. Williams. IMRT treatment planning—a comparative inter-system and inter-centre planning exercise of the ESTRO QUASIMODO group. *Radiother. Oncol.*, 76(3):354–361, 2005.

[15] T. Bortfeld. Optimized planning using physical objectives and constraints. *Semin. Radiat. Oncol.*, 9(1):20–34, 1999.

[16] T. Bortfeld. IMRT: a review and preview. *Phys. Med. Biol.*, 51(13):R363–R379, 2006.

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

[18] A. Brahme, J.-E. Roos, and I. Lax. Solution of an integral equation encountered in rotation therapy. *Phys. Med. Biol.*, 27(10):1221–1229, 1982.

[19] S. Breedveld, P. Storchi, P. Voet, and B. Heijmen. iCycle: Integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans. *Med. Phys.*, 39(2):951–963, 2012.

[20] M. Broderick, M. Leech, and M. Coffey. Direct aperture optimization as a means of reducing the complexity of intensity modulated radiation therapy plans. *Radiat. Oncol.*, 4(1):1–8, 2009.

[21] K. Bucci, A. Bevan, and M. Roach 3rd. Advances in radiation therapy: conventional to 3D, to IMRT, to 4D, and beyond. *CA-Cancer J. Clin.*, 55(2):117–134, 2005.

[22] K. Bzdusek, H. Friberger, K. Eriksson, B. Hårdemark, D. Robinson, and M. Kaus. Development and evaluation of an efficient approach to volumetric arc therapy planning. *Med. Phys.*, 36(6):2328–2339, 2009.

[23] F. Carlsson. Combining segment generation with direct step-and-shoot optimization in intensity-modulated radiation therapy. *Med. Phys.*, 35(9):3828–3838, 2008.

[24] F. Carlsson and A. Forsgren. Iterative regularization in intensity-modulated radiation therapy optimization. *Med. Phys.*, 33(1):225–234, 2006.

[25] M. Casiraghi, F. Albertini, and A. Lomax. Advantages and limitations of the 'worst case scenario' approach in IMPT treatment planning. *Phys. Med. Biol.*, 58(5):1323–1339, 2013.

[26] A. Cassioli and J. Unkelbach. Aperture shape optimization for IMRT treatment planning. *Phys. Med. Biol.*, 58(2):301–318, 2013.

[27] T. Chan. *Optimization under Uncertainty in Radiation Therapy*. PhD thesis, Massachusetts Institute of Technology, Boston, MA, 2007.

[28] W. Chen, J. Unkelbach, A. Trofimov, T. Madden, H. Kooy, T. Bortfeld, and D. Craft. Including robustness in multi-criteria optimization for intensity-modulated proton therapy. *Phys. Med. Biol.*, 57(3):591–608, 2012.

[29] B. Choi and J. Deasy. The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning. *Phys. Med. Biol.*, 47(20):3579–3589, 2002.

[30] H. Chung, B. Lee, E. Park, J. Lu, and P. Xia. Can all centers plan intensity-modulated radiotherapy (IMRT) effectively? An external audit of dosimetric comparisons between three-dimensional conformal radiotherapy and IMRT for adjuvant chemoradiation for gastric cancer. *Int. J. Radiat. Oncol.*, 71(4):1167–1174, 2008.

[31] A. Chvetsov, D. Calvetti, J. Sohn, and T. Kinsella. Regularization of inverse planning for intensity-modulated radiotherapy. *Med. Phys.*, 32(2):501–514, 2005.

[32] 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 Appl.*, 428(5):1345–1364, 2008.

[33] A. Cormack. A problem in rotation therapy with X rays. *Int. J. Radiat. Oncol.*, 13(4):623–630, 1987.

[34] A. Cormack and R. Cormack. A problem in rotation therapy with X-rays: dose distributions with an axis of symmetry. *Int. J. Radiat. Oncol.*, 13(12):1921–1925, 1987.

[35] M. Coselmon, J. Moran, J. Radawski, and B. Fraass. Improving IMRT delivery efficiency using intensity limits during inverse planning. *Med. Phys.*, 32(5):1234–1245, 2005.

[36] C. Cotrutz and L. Xing. IMRT dose shaping with regionally variable penalty scheme. *Med. Phys.*, 30(4):544–551, 2003.

[37] C. Cotrutz and L. Xing. Segment-based dose optimization using a genetic algorithm. *Phys. Med. Biol.*, 48(18):2987–2998, 2003.

[38] D. Craft and T. Bortfeld. How many plans are needed in an IMRT multi-objective plan database? *Phys. Med. Biol.*, 53(11):2785–2796, 2008.

[39] D. Craft, F. Carlsson, T. Bortfeld, and H. Rehbinder. Multi-objective IMRT planning which produces deliverable plans. *Med. Phys.*, 35(6):2846, 2008.

[40] D. Craft, T. Halabi, H. Shih, and T. Bortfeld. Approximating convex Pareto surfaces in multiobjective radiotherapy planning. *Med. Phys.*, 33(9):3399–3407, 2006.

[41] D. Craft, T. Hong, H. Shih, and T. Bortfeld. Improved planning time and plan quality through multi-criteria optimization for intensity modulated radiation therapy. *Int. J. Radiat. Oncol.*, 82(1):e83–e90, 2012.

[42] D. Craft, D. McQuaid, J. Wala, W. Chen, E. Salari, and T. Bortfeld. Multicriteria VMAT optimization. *Med. Phys.*, 39(2):686–696, 2012.

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

[44] D. Craft and C. Richter. Deliverable navigation for multicriteria step and shoot IMRT treatment planning. *Phys. Med. Biol.*, 58(1):87–103, 2013.

[45] D. Craft, P. Süss, and T. Bortfeld. The tradeoff between treatment plan quality and required number of monitor units in intensity-modulated radiotherapy. *Int. J. Radiat. Oncol.*, 67(5):1596–1605, 2007.

[46] I. Das and E. Dennis. A closer look at drawbacks of minimizing weighted sums of objectives for Pareto set generation in multicriteria optimization problems. *Struct. Multidiscip. O.*, 14(1):63–69, 1997.

[47] W. De Gersem, F. Claus, C. De Wagter, B. Van Duyse, and W. De Neve. Leaf position optimization for step-and-shoot IMRT. *Int. J. Radiat. Oncol.*, 51(5):1371–1388, 2001.

[48] J. Deasy. Multiple local minima in radiotherapy optimization problems with dose-volume constraints. *Med. Phys.*, 24(7):1157–1161, 1997.

[49] M. Ehrgott. *Multicriteria Optimization*, volume 491 of *Lecture Notes in Economics and Mathematical Systems*. Springer-Verlag, Berlin, Germany, 2nd edition, 2005.

[50] M. Ehrgott, Ç. Güler, H. Hamacher, and L. Shao. Mathematical optimization in intensity modulated radiation therapy. *4OR-Q. J. Oper. Res.*, 6(3):199–262, 2008.

[51] M. Ehrgott, J. Ide, and A. Schöbel. Minmax robustness for multi-objective optimization problems. Technical Report Nr. 1, Institut für Numerische und Angewandte Mathematik, Universität Göttingen, Göttingen, Germany, 2013.

[52] M. Ehrgott, L. Shao, and A. Schöbel. An approximation algorithm for convex multi-objective programming problems. *J. Global Optim.*, 50(3):397–416, 2011.

[53] B. Emami, J. Lyman, A. Brown, L. Cola, M. Goitein, J. Munzenrider, B. Shank, L. Solin, and M. Wesson. Tolerance of normal tissue to therapeutic irradiation. *Int. J. Radiat. Oncol.*, 21(1):109–122, 1991.

[54] M. Falkinger, S. Schell, J. Müller, and J. Wilkens. Prioritized optimization in intensity modulated proton therapy. *Z. Med. Phys.*, 22(1):21–28, 2012.

[55] M. Fippel, F. Haryanto, O. Dohm, F. Nüsslin, and S. Kriesen. A virtual photon energy fluence model for Monte Carlo dose calculation. *Med. Phys.*, 30(3):301–311, 2003.

[56] J. Fowler. The linear-quadratic formula and progress in fractionated radiotherapy. *Brit. J. Radiol.*, 62(740):679–694, 1989.

[57] A. Fredriksson. A characterization of robust radiation therapy treatment planning methods—from expected value to worst case optimization. *Med. Phys.*, 39(8):5169–5181, 2012.

[58] A. Fredriksson, A. Forsgren, and B. Hårdemark. Minimax optimization for handling range and setup uncertainties in proton therapy. *Med. Phys.*, 38(3):1672–1684, 2011.

[59] P. Gill, W. Murray, and M. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Rev.*, 47(1):99–131, 2005.

[60] P. Gill and E. Wong. Sequential quadratic programming methods. In *Mixed Integer Nonlinear Programming*, pages 147–224. Springer, New York, NY, 2012.

[61] M. Goitein and J. Cox. Should randomized clinical trials be required for proton radiotherapy? *J. Clin. Oncol.*, 26(2):175–176, 2008.

[62] T. Halabi, D. Craft, and T. Bortfeld. Dose-volume objectives in multi-criteria optimization. *Phys. Med. Biol.*, 51(15):3809–3818, 2006.

[63] E. Hall. Intensity-modulated radiation therapy, protons, and the risk of second cancers. *Int. J. Radiat. Oncol.*, 65(1):1–7, 2006.

[64] E. Hall and A. Giaccia. *Radiobiology For The Radiologist*. Lippincott Williams & Wilkins, 6th edition, 2006.

[65] E. Hall and C.-S. Wuu. Radiation-induced second cancers: the impact of 3D-CRT and IMRT. *Int. J. Radiat. Oncol.*, 56(1):83–88, 2003.

[66] B. Hårdemark, A. Liander, H. Rehbinder, and J. Löf. Direct machine parameter optimization with RayMachine® in Pinnacle^3®. White paper, RaySearch Laboratories, Stockholm, Sweden, 2003.

[67] A. Hoffmann, D. den Hertog, A. Siem, J. Kaanders, and H. Huizenga. Convex reformulation of biologically-based multi-criteria IMRT optimization including fractionation effects. *Phys. Med. Biol.*, 53(22):6345–6362, 2008.

[68] A. Hoffmann, A. Siem, D. den Hertog, J. Kaanders, and H. Huizenga. Derivative-free generation and interpolation of convex Pareto optimal IMRT plans. *Phys. Med. Biol.*, 51(24):6349–6369, 2006.

[69] T. Hong, D. Craft, F. Carlsson, and T. Bortfeld. Multicriteria optimization in intensity-modulated radiation therapy treatment planning for locally advanced cancer of the pancreatic head. *Int. J. Radiat. Oncol.*, 72(4):1208–1214, 2008.

[70] M. Hunt, C. Hsiung, S. Spirou, C. Chui, H. Amols, and C. Ling. Evaluation of concave dose distributions created using an inverse planning system. *Int. J. Radiat. Oncol.*, 54(3):953–962, 2002.

[71] International Commission on Radiation Units and Measurements. *ICRU Report 62: Prescribing, Recording and Reporting Photon Beam Therapy (Supplement to ICRU Report 50)*. International Commission on Radiation Units and Measurements, Bethesda, MD, 1999.

[72] International Commission on Radiation Units and Measurements. ICRU report 78: Prescribing, recording, and reporting proton-beam therapy. *J. ICRU*, 7(2), 2007.

[73] International Commission on Radiation Units and Measurements. ICRU report 83: Prescribing, recording, and reporting photon-beam intensity-modulated radiation therapy (IMRT). *J. ICRU*, 10(1), 2010.

[74] K.-W. Jee, D. McShan, and B. Fraass. Lexicographic ordering: intuitive multicriteria optimization for IMRT. *Phys. Med. Biol.*, 52(7):1845–1861, 2007.

[75] R. Jeraj, C. Wu, and T. Mackie. Optimizer convergence and local minima errors and their clinical importance. *Phys. Med. Biol.*, 48(17):2809–2827, 2003.

[76] S. Korreman. Motion in radiotherapy: photon therapy. *Phys. Med. Biol.*, 57(23):R161–R191, 2012.

[77] 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.

[78] Y. Lan, C. Li, H. Ren, Y. Zhang, and Z. Min. Fluence map optimization (FMO) with dose–volume constraints in IMRT using the geometric distance sorting method. *Phys. Med. Biol.*, 57(20):6407–6428, 2012.

[79] Y. Li, J. Yao, and D. Yao. Genetic algorithm based deliverable segments optimization for static intensity-modulated radiotherapy. *Phys. Med. Biol.*, 48(20):3353–3374, 2003.

[80] C. Lindholm, E. Cavallin-Ståhl, J. Ceberg, J.-E. Frödin, B. Littbrand, and T. Möller. Radiotherapy practices in Sweden compared to the scientific evidence. *Acta Oncol.*, 42(5&6):416–429, 2003.

[81] W. Liu, X. Zhang, Y. Li, and R. Mohan. Robust optimization of intensity modulated proton therapy. *Med. Phys.*, 39(2):1079–1091, 2012.

[82] J. Llacer, J. Deasy, T. Bortfeld, T. Solberg, and C. Promberger. Absence of multiple local minima effects in intensity modulated optimization with dose volume constraints. *Phys. Med. Biol.*, 48(2):183–210, 2003.

[83] A. Lomax. Intensity modulation methods for proton radiotherapy. *Phys. Med. Biol.*, 44(1):185–205, 1999.

[84] A. Lomax. Intensity modulated proton therapy and its sensitivity to treatment uncertainties 1: the potential effects of calculational uncertainties. *Phys. Med. Biol.*, 53(4):1027–1042, 2008.

[85] A. Lomax. Intensity modulated proton therapy and its sensitivity to treatment uncertainties 2: the potential effects of inter-fraction and inter-field motions. *Phys. Med. Biol.*, 53(4):1043–1056, 2008.

[86] A. Lomax, E. Pedroni, H. Rutz, and G. Goitein. The clinical potential of intensity modulated proton therapy. *Z. Med. Phys.*, 14(3):147–152, 2004.

[87] 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. *Technol. Cancer Res. T.*, 9(6):629–636, 2010.

[88] L. Marks, E. Yorke, A. Jackson, R. Ten Haken, L. Constine, A. Eisbruch, S. Bentzen, J. Nam, and J. Deasy. Use of normal tissue complication probability models in the clinic. *Int. J. Radiat. Oncol.*, 76(3):S10–S19, 2010.

[89] M. Matuszak, E. Larsen, and B. Fraass. Reduction of IMRT beam complexity through the use of beam modulation penalties in the objective function. *Med. Phys.*, 34(2):507–520, 2007.

[90] C. Men, E. Romeijn, X. Jia, and S. Jiang. Ultrafast treatment plan optimization for volumetric modulated arc therapy (VMAT). *Med. Phys.*, 37(2):5787–5791, 2010.

[91] A. Messac and C. Mattson. Normal constraint method with guarantee of even representation of complete Pareto frontier. *AIAA J.*, 42(10):2101–2111, 2004.

[92] K. Miettinen. *Nonlinear Multiobjective Optimization*. Kluwer Academic Publishers, Boston, MA, 1999.

[93] R. Miller, M. Lodge, M. Murad, and B. Jones. Controversies in clinical trials in proton radiotherapy: the present and the future. *Semin. Radiat. Oncol.*, 23(2):127–133, 2013.

[94] M. Monz, K.-H. Küfer, T.Bortfeld, and C. Thieke. Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning. *Phys. Med. Biol.*, 53(4):985–998, 2008.

[95] K. Moore, R. Brame, D. Low, and S. Mutic. Quantitative metrics for assessing plan quality. *Semin. Radiat. Oncol.*, 22(1):62–69, 2012.

[96] National Board of Health and Welfare (Socialstyrelsen). *Cancer incidence in Sweden 2011*. National Board of Health and Welfare, Stockholm, Sweden, 2012.

[97] A. Niemierko. A generalized concept of equivalent uniform dose (EUD). *Med. Phys.*, 26(6):1100, 1999.

[98] K. Otto. Volumetric modulated arc therapy: IMRT in a single gantry arc. *Med. Phys.*, 35(1):310–317, 2008.

[99] D. Pflugfelder, J. Wilkens, and U. Oelfke. Worst case optimization: a method to account for uncertainties in the optimization of intensity modulated proton therapy. *Phys. Med. Biol.*, 53(6):1689–1700, 2008.

[100] J. Purdy. Dose to normal tissues outside the radiation therapy patient's treated volume: a review of different radiation therapy techniques. *Health Phys.*, 95(5):666–676, 2008.

[101] G. Rennen, E. van Dam, and D. den Hertog. Enhancement of sandwich algorithms for approximating higher-dimensional convex Pareto sets. *INFORMS J. Comput.*, 23(4):493–517, 2011.

[102] E. Romeijn, R. Ahuja, J. Dempsey, and A. Kumar. A column generation approach to radiation therapy treatment planning using aperture modulation. *SIAM J. Optimiz.*, 15(3):838–862, 2005.

[103] E. Romeijn, R. Ahuja, J. Dempsey, and A. Kumar. A new linear programming approach to radiation therapy treatment planning problems. *Operations Res.*, 54(2):201–216, 2006.

[104] E. Romeijn and J. Dempsey. Intensity modulated radiation therapy treatment plan optimization. *TOP*, 16(2):215–243, 2008.

[105] E. Romeijn, J. Dempsey, and J. Li. A unifying framework for multi-criteria fluence map optimization models. *Phys. Med. Biol.*, 49(10):1991–2013, 2004.

[106] H. Ruotsalainen, K. Miettinen, J.-E. Palmgren, and T. Lahtinen. Interactive multiobjective optimization for anatomy-based three-dimensional HDR brachytherapy. *Phys. Med. Biol.*, 55(16):4703–4719, 2010.

[107] E. Salari and J. Unkelbach. A column-generation-based method for multi-criteria direct aperture optimization. *Phys. Med. Biol.*, 58(3):621–639, 2013.

[108] B. Salter, V. Sarkar, B. Wang, H. Shukla, M. Szegedi, and P. Rassiah-Szegedi. Rotational IMRT delivery using a digital linear accelerator in very high dose rate 'burst mode'. *Phys. Med. Biol.*, 56(7):1931–1946, 2011.

[109] J. Serna, M. Monz, K.-H. Küfer, and C. Thieke. Trade-off bounds for the Pareto surface approximation in multi-criteria IMRT planning. *Phys. Med. Biol.*, 54(20):6299–6311, 2009.

[110] L. Shao. *Multiple Objective Linear Programming in Radiotherapy Treatment Planning*. PhD thesis, The University of Auckland, Auckland, New Zealand, 2008.

[111] L. Shao and M. Ehrgott. Approximately solving multiobjective linear programmes in objective space and an application in radiotherapy treatment planning. *Math. Method Oper. Res.*, 68(2):257–276, 2008.

[112] N. Sheets, G. Goldin, A.-M. Meyer, Y. Wu, Y. Chang, T. Stürmer, J. Holmes, B. Reeve, P. Godley, W. Carpenter, and R. Chen. Intensity-modulated radiation therapy, proton therapy, or conformal radiation therapy and morbidity and disease control in localized prostate cancer. *JAMA-J. Am. Med. Assoc.*, 307(15):1611–1620, 2012.

[113] D. Shepard, M. Earl, X. Li, S. Naqvi, and C. Yu. Direct aperture optimization: a turnkey solution for step-and-shoot IMRT. *Med. Phys.*, 29(6):1007–1018, 2002.

[114] S. Spirou, N. Fournier-Bidoz, J. Yang, C.-S. Chui, and C. Ling. Smoothing intensity-modulated beam profiles to improve the efficiency of delivery. *Med. Phys.*, 28(10):2105–2112, 2001.

[115] J. Staffurth. A review of the clinical evidence for intensity-modulated radiotherapy. *Clin. Oncol.*, 22(8):643–657, 2010.

[116] J. Stein, T. Bortfeld, B. Dörschel, and W. Schlegel. Dynamic X-ray compensation for conformal radiotherapy by means of multi-leaf collimation. *Radiother. Oncol.*, 32(2):163–173, 1994.

[117] 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. *Phys. Med. Biol.*, 58(6):1855–1867, 2013.

[118] K. Teichert, P. Süss, J. Serna, M. Monz, K.-H. Küfer, and C. Thieke. Comparative analysis of Pareto surfaces in multi-criteria IMRT planning. *Phys. Med. Biol.*, 56(12):3669–3684, 2011.

[119] M. Teoh, C. Clark, K. Wood, S. Whitaker, and A. Nisbet. Volumetric modulated arc therapy: a review of current literature and clinical use in practice. *Brit. J. Radiol.*, 84(1007):967–996, 2011.

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

[121] R. Twombly. Cancer surpasses heart disease as leading cause of death for all but the very elderly. *J. Natl. Cancer I.*, 97(5):330–331, 2005.

[122] S. Ulrich, S. Nill, and U. Oelfke. Development of an optimization concept for arc-modulated cone beam therapy. *Phys. Med. Biol.*, 52(14):4099–4119, 2007.

[123] J. Unkelbach, T. Bortfeld, B. Martin, and M. Soukup. Reducing the sensitivity of IMPT treatment plans to setup errors and range uncertainties via probabilistic treatment planning. *Med. Phys.*, 36(1):149–163, 2009.

[124] J. Unkelbach, T. C. Chan, and T. Bortfeld. Accounting for range uncertainties in the optimization of intensity modulated proton therapy. *Phys. Med. Biol.*, 52(10):2755–2773, 2007.

[125] M. van Herk. Errors and margins in radiotherapy. *Semin. Radiat. Oncol.*, 14(1):52–64, 2004.

[126] L. Veldeman, I. Madani, F. Hulstaert, G. De Meerleer, M. Mareel, and W. De Neve. Evidence behind use of intensity-modulated radiotherapy: a systematic review of comparative clinical studies. *Lancet Oncol.*, 9(4):367–375, 2008.

[127] P. Voet, M. Dirkx, S. Breedveld, D. Fransen, P. Levendag, and B. Heijmen. Toward fully automated multicriterial plan generation: a prospective clinical study. *Int. J. Radiat. Oncol.*, 85(3):866–872, 2013.

[128] 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. *Med. Dosim.*, Published online before print on March 25, 2013.

[129] J. Wala, E. Salari, W. Chen, and D. Craft. Optimal partial-arcs in VMAT treatment planning. *Phys. Med. Biol.*, 57(18):5861–5874, 2012.

[130] S. Webb. Optimizing the planning of intensity-modulated radiotherapy. *Phys. Med. Biol.*, 39(12):2229–2246, 1994.

[131] S. Webb and D. McQuaid. Some considerations concerning volume-modulated arc therapy: a stepping stone towards a general theory. *Phys. Med. Biol.*, 54(14):4345–4360, 2009.

[132] J. Wilkens, J. Alaly, K. Zakarian, W. Thorstad, and J. Deasy. IMRT treatment planning based on prioritizing prescription goals. *Phys. Med. Biol.*, 52(6):1675–1692, 2007.

[133] P. Williams. IMRT: delivery techniques and quality assurance. *Brit. J. Radiol.*, 76(911):766–776, 2003.

[134] 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. *Int. J. Radiat. Oncol.*, 79(4):1241–1247, 2011.

[135] C. Wu, R. Jeraj, and T. Mackie. The method of intercepts in parameter space for the analysis of local minima caused by dose-volume constraints. *Phys. Med. Biol.*, 48(11):N149–N157, 2003.

[136] Q. Wu and R. Mohan. Multiple local minima in IMRT optimization based on dose-volume criteria. *Med. Phys.*, 29(7):1514–1527, 2002.

[137] I. Xhaferllari, E. Wong, K. Bzdusek, M. Lock, and J. Chen. Automated IMRT planning with regional optimization using planning scripts. *J. Appl. Clin. Med. Phys.*, 14(1):176–191, 2013.

[138] P. Xia, J. Ting, and C. Orton. Segmental MLC is superior to dynamic MLC for IMRT delivery. *Med. Phys.*, 34(7):2673–2675, 2007.

[139] L. Xing, J. Li, S. Donaldson, Q. Le, and A. Boyer. Optimization of importance factors in inverse planning. *Phys. Med. Biol.*, 44(10):2525–2536, 1999.

[140] T. Yock and P. Caruso. Risk of second cancers after photon and proton radiotherapy: A review of the data. *Health Phys.*, 103(5):577–585, 2012.

[141] C. Yu and G. Tang. Intensity-modulated arc therapy: principles, technologies and clinical implementation. *Phys. Med. Biol.*, 56(5):R31–R54, 2011.

[142] M. Zaider and L. Hanin. Tumor control probability in radiation treatment. *Med. Phys.*, 38(2):574–583, 2011.

[143] M. Zarepisheh, M. Shakourifar, G. Trigila, P. Ghomi, S. Couzens, A. Abebe, L. Noreña, W. Shang, S. Jiang, and Y. Zinchenko. A moment-based approach for DVH-guided radiotherapy treatment plan optimization. *Phys. Med. Biol.*, 58(6):1869–1887, 2013.

[144] X. Zhang, X. Li, E. Quan, X. Pan, and Y. Li. A methodology for automatic intensity-modulated radiation treatment planning for lung cancer. *Phys. Med. Biol.*, 56(13):3873–3893, 2011.

[145] L. Zhu, L. Lee, Y. Ma, Y. Ye, R. Mazzeo, and L. Xing. Using total-variation regularization for intensity modulated radiation therapy inverse planning with field-specific numbers of segments. *Phys. Med. Biol.*, 53(23):6653–6672, 2008.

[146] L. Zhu and L. Xing. Search for IMRT inverse plans with piecewise constant fluence maps using compressed sensing techniques. *Med. Phys.*, 36(5):1895–1905, 2009.

[147] Y. Zinchenko, T. Craig, H. Keller, T. Terlaky, and M. Sharpe. Controlling the dose distribution with gEUD-type constraints within the convex radiotherapy optimization framework. *Phys. Med. Biol.*, 53(12):3231–3250, 2008.