Robust optimization of radiation therapy accounting for geometric uncertainty
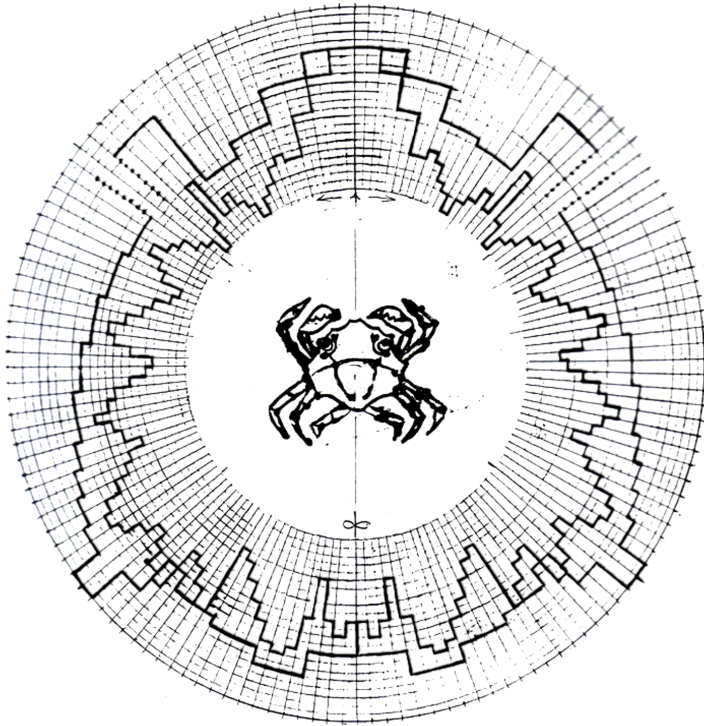

Albin Fredriksson

KTH Engineering Sciences

Robust optimization of radiation therapy accounting for geometric uncertainty

ALBIN FREDRIKSSON

Doctoral Thesis
Stockholm, Sweden 2013

Cover illustration copyright © 1980 by Bob Marshall. Used with permission. It illustrates the first canon of *Das Musikalische Opfer* by Johann Sebastian Bach. This movement is a *canon cancrizans*.

Till min familj

De flesta menar att skinnet skiljer dem från det som omger dem. Men människans skinn är tunt och genomsläppligt, fullt av hål och öppningar likt en trasig rock. Det omänskliga far in och ut genom revorna; jord och vind blåser tvärs igenom oss. Vår hjälplöshet är höggradig.

\- Willy Kyrklund, *Mästaren Ma*

Abstract

Geometric errors may compromise the quality of radiation therapy treatments. Optimization methods that account for errors can reduce their effects.

The first paper of this thesis introduces minimax optimization to account for systematic range and setup errors in intensity-modulated proton therapy. The minimax method optimizes the worst case outcome of the errors within a given set. It is applied to three patient cases and shown to yield improved target coverage robustness and healthy structure sparing compared to conventional methods using margins, uniform beam doses, and density override. Information about the uncertainties enables the optimization to counterbalance the effects of errors.

In the second paper, random setup errors of uncertain distribution—in addition to the systematic range and setup errors—are considered in a framework that enables scaling between expected value and minimax optimization. Experiments on a phantom show that the best and mean case tradeoffs between target coverage and critical structure sparing are similar between the methods of the framework, but that the worst case tradeoff improves with conservativeness.

Minimax optimization only considers the worst case errors. When the planning criteria cannot be fulfilled for all errors, this may have an adverse effect on the plan quality. The third paper introduces a method for such cases that modifies the set of considered errors to maximize the probability of satisfying the planning criteria. For two cases treated with intensity-modulated photon and proton therapy, the method increased the number of satisfied criteria substantially. Grasping for a little less sometimes yields better plans.

In the fourth paper, the theory for multicriteria optimization is extended to incorporate minimax optimization. Minimax optimization is shown to better exploit spatial information than objective-wise worst case optimization, which has previously been used for robust multicriteria optimization.

The fifth and sixth papers introduce methods for improving treatment plans: one for deliverable Pareto surface navigation, which improves upon the Pareto set representations of previous methods; and one that minimizes healthy structure doses while constraining the doses of all structures not to deteriorate compared to a reference plan, thereby improving upon plans that have been reached with too weak planning goals.

Keywords: Optimization, intensity-modulated proton therapy, uncertainty, robust planning, setup error, range error, intensity-modulated radiation therapy, multicriteria optimization.

Sammanfattning

Geometriska fel kan försämra kvaliteten på strålbehandlingar, men optimeringsmetoder som tar hänsyn till felen kan minska deras effekt.

I denna avhandlings första artikel introduceras minimaxoptimering för att ta hänsyn till systematiska fel på räckvidd och positionering i intensitesmodulerad protonterapi. Minimaxmetoden optimerar det värsta utfallet av felen från en given mängd. Metoden prövas på tre patientfall. För dessa leder den till mer robust måltäckning och ökat riskorgansskydd jämfört med konventionella metoder som använder marginaler, likformiga stråldoser och ersatta densiteter. Information om osäkerheterna gör att optimeringen kan motverka effekterna av fel.

I den andra artikeln betraktas slumpmässiga positioneringsfel av osäker sannolikhetsfördelning – utöver de systematiska räckvidds- och positioneringsfelen – i ett ramverk som möjliggör skalning mellan väntevärdes- och minimaxoptimering. Experiment på ett fantom visar att avvägningen mellan måltäckning och riskorgansskydd i det bästa fallet och i medelfallet är likartad mellan metoderna från ramverket, men att avvägningen i värsta fallet förbättras med graden av försiktighet.

Minimaxoptimering tar bara hänsyn till de värsta felen. Detta kan leda till att plankvaliteten blir lidande i fall där planeringsmålen inte går att uppfylla för alla fel. I den tredje artikeln introduceras en metod för sådana fall. Denna metod modifierar mängden av beaktade fel i syfte att maximera sannolikheten att uppfylla planeringsmålen. För två fall behandlade med intensitetsmodulerad foton- och protonterapi ledde metoden till en avsevärd ökning av antalet uppfyllda mål. Sänkta krav på robustheten kan ibland leda till bättre planer.

I den fjärde artikeln utökas teorin för flermålsoptimering till att innefatta minimaxoptimering. Minimaxoptimering visas vara bättre på att utnyttja spatiell information än målvis värsta fallet-optimering, vilket tidigare använts för robust flermålsoptimering.

Artikel fem och sex introducerar metoder för att förbättra strålbehandlingsplaner: en för levererbar navigering av Pareto-ytor, vilken förbättrar tidigare metoders representationer av Pareto-mängder; och en som minimerar doserna till friska strukturer under bivillkor att doserna till alla strukturer inte försämras jämfört med en referensplan, för att på så sätt förbättra planer som har tagits fram med för lågt satta mål.

Nyckelord: Optimering, intensitetsmodulerad protonterapi, osäkerhet, robustplanering, positioneringsfel, räckviddsfel, intensitetsmodulerad strålterapi, flermålsoptimering.

Acknowledgments

First of all, many thanks to my advisor Anders Forsgren, whose encouraging and helpful attitude and mathematical expertise have been invaluable to me during these years. It has been a privilege to work with you.

My industrial co-advisors from RaySearch Laboratories have also been of significant importance: Thanks to Henrik Rehbinder for guiding me during the initial phase of this project. And thanks to Björn Hårdemark, my current co-advisor, for sharing your abounding knowledge of radiation therapy and just about all other topics.

I thank my academic co-advisors Johan Håstad and Krister Svanberg for their helpful comments during reference group meetings.

The work presented in this thesis was co-funded by the Swedish Research Council (VR) and RaySearch Laboratories. I am grateful to both and especially thank Johan Löf, the founder and CEO of RaySearch, for hosting this project and for creating the inspiring workplace that RaySearch is.

Thanks to Rasmus Bokrantz, my academic twin, for all discussions concerning research, typography, line widths; and for co-authorship, traveling companionship, and friendship.

I have deeply appreciated the conversations with previous and present colleagues at RaySearch: with Kjell Eriksson and Fredrik Löfman about optimization of radiation therapy, and with Göran Sporre about optimization in general. Thanks also to Tore Ersmark, Lars Glimelius, and Martin Janson, who have generously shared their knowledge of proton physics.

Further, I am grateful to the faculty members and staff of the Division of Optimization and Systems Theory at KTH. Special thanks to my fellow students for course collaborations and for all laughs during fika: Mikael Fallgren, Johan Markdahl, Anders Möller, Hildur Æsa Oddsdóttir, Tove Odland, Göran Svensson, Henrik Svärd, Johan Thunberg, and Yuecheng Yang.

I am grateful to my other friends for sharing my interests outside of research. A special thanks to Johan Hallgren for drawing my attention to this PhD student position.

Thanks to my family for your unconditional love and support. I hope that I can express how much you mean to me by stating that I cannot. Finally, my deepest love to Sanna, who is my best friend.

Stockholm, April 2013

Albin Fredriksson

Contents

Introduction

It is estimated that every third person in Sweden will get cancer [68]. Half of Swedish cancer patients receive radiation therapy during their illness [54]. The quality of radiation therapy treatment is of high consequence.

This thesis concerns optimization approaches for radiation therapy in the presence of geometric uncertainty. It consists of an introduction and six appended papers. The introduction first presents radiation therapy and relevant optimization concepts. It then introduces optimization approaches to account for uncertainties in radiation therapy. The considered uncertainties are mainly with respect to the patient densities and the alignment of the patient relative to the beams. The topic of facilitating decision making by multicriteria optimization accounting for geometric uncertainty is also discussed, as is the problem of deliverability of plans obtained by Pareto surface navigation. The introduction is concluded with a short summary of the appended papers.

1   Radiation therapy

X-rays were discovered in 1895 by Wilhelm Röntgen. A few months after the announcement of the discovery, x-rays were used to treat skin disease, and within a year, the first accounts of using x-rays to treat cancer were reported. Thus was the birth of radiation therapy.

Radiation therapy is the medical use of ionizing radiation. It is primarily used to treat cancer. In most cases, radiation therapy is given with curative intent. It may also be used in palliative care in cases where the cancer is too advanced for a curative treatment to be possible, but for which symptoms such as pain may be relieved. Radiation therapy is used both as a stand-alone treatment and in combination with other cancer treatments such as surgery and chemotherapy.

Radiation therapy is delivered either by internal radiation sources (*brachytherapy*), which are placed inside or close to the region to be treated, or by external sources (*external beam radiation therapy*). The latter is the most common form of radiation therapy and is the form that this thesis concerns. In external beam radiation therapy, the patient is irradiated by an external radiation source that directs the radiation towards the region requiring treatment. The radiation is commonly delivered in the form of high-energy photon (x-ray), electron, or proton beams, but other particles may also be used. In this thesis, photon and proton beams are considered. Blocking material is used to shape the beams in order to conform the radiation to the target. The superposition of radiation of several beams from different directions enables high doses of radiation to the target while the doses to surrounding healthy tissues can be kept low. Greater amounts of radiation delivered to the target than to healthy tissues increases the probability of eradicating the tumor while sparing critical organs and avoiding radiation-induced second cancers.

1.1 Radiobiology

For curative radiation therapy, the clonogenic cancer cells must be killed to an extent that results in permanent tumor control. Radiation kills cells by damaging the cellular DNA. Sufficient damage to the DNA of a cell disables the ability of the cell to proliferate, ultimately leading to its death.

The cellular DNA is damaged by interaction with ionizing particles. As x-ray photons pass through tissue, they interact with free electrons or electrons with negligible binding energy compared to the photon energy. In the interaction between a photon and an electron, part of the photon energy is given to the electron in the form of kinetic energy. The resulting fast-moving electron may damage the DNA directly or indirectly. In direct action, the electron interacts with the DNA to produce damage. In indirect action, the electron interacts with other molecules, such as water, to produce free radicals, which in turn damage the DNA. Since photon radiation ionizes the absorber (in this case the DNA) via recoil electrons, it is said to be indirectly ionizing. Proton radiation is directly ionizing; it has sufficient energy to ionize the absorber directly.

Cancer cells generally have reduced ability to repair DNA damages compared to healthy cells, and sublethal DNA damages that accumulate over time may eventually lead to lethal damages. Radiation therapy treatment is therefore typically divided into a number ($\sim$30) of treatment fractions that are delivered with daily intervals (with breaks for the weekends). Between fractions, the DNA molecules of healthy cells are repaired to a higher degree than those of cancer cells, and the

healthy tissues often repopulate faster than the tumor tissue. Moreover, fractionation allows the tumor cells to reassort into more radiosensitive phases of the cell cycle, and allows for the reoxygenization of hypoxic tumor regions, which are resistant to indirect action of radiation.

For more details concerning radiobiology, see, e.g., Hall and Giaccia [41].

1.2 Photon therapy

A photon therapy treatment is typically delivered by a gantry-mounted linear accelerator, which accelerates electrons onto a high-density bremsstrahlung target. This results in the scattering of high-energy photons. The photons are filtered to produce a uniform intensity distribution and leave the gantry through a gantry head. The output of the accelerator is measured in *monitor units* (MUs), which are calibrated such that 1 MU yields an absorbed dose of 1 cGy at a specific depth in water for a standardized field. The gantry can be rotated about the patient, which enables the delivery of photon fields from different directions. Photon fields from several directions are combined to yield higher dose to the target than to the surrounding tissues.

Collimating blocks made out of a shielding material such as tungsten are used to shape the beam. Mounted on the gantry head are one or two pairs of opposing blocks called *jaws*, which can create rectangular beam shapes. The gantry head may further be equipped with a *multileaf collimator* (MLC), a device consisting of two opposing rows of shielding leaves, which may individually move in and out of the field to shape the beam. Figure 1 includes illustrations of MLCs. The jaws and the MLC are used to conform the beam to the projection of the target volume onto the beam plane. An arrangement of the jaws and the MLC leaves is called an *aperture*.

Before the invention of three-dimensional (3D) imaging techniques such as computed tomography (CT), two-dimensional (2D) x-ray images were used to plan radiation therapy treatments. The beam setups were typically simple, consisting of one to four beams. With 3D information, it has become practicable to use beams from multiple angles, each shaped as its corresponding target projection. This type of treatment is called *3D conformal radiation therapy* (3DCRT), and enables more conformal dose distributions than 2D planning.

An improvement over 3DCRT came with the introduction of varying fluence over the cross-section of each beam. Such treatment is called *intensity-modulated radiation therapy* (IMRT). In IMRT, the superposition of the fluences transmitted through a succession of apertures forms a field of modulated fluence. Delivery

where the beam is switched off while the MLC leaves move is called *step-and-shoot* delivery. The combination of an aperture and a weight specifying the fraction of the beam MU delivered through the aperture is called a *segment*, and the weight is called the *segment weight*. Delivery where the leaves move during irradiation is called *dynamic MLC* delivery. In this thesis, step-and-shoot delivery is considered. Figure 1 illustrates an IMRT plan for a head and neck case treated with seven equispaced beams.
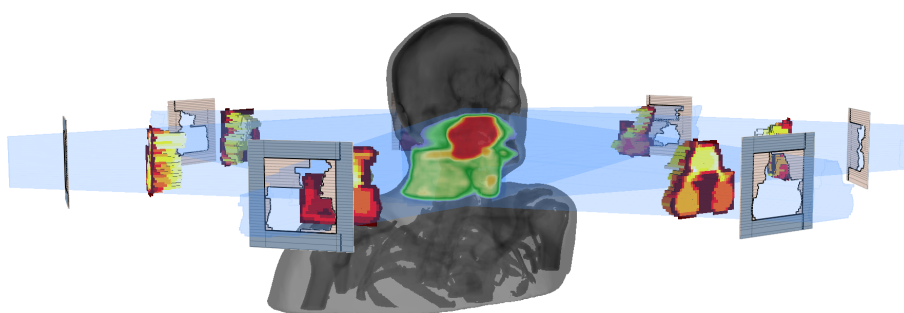
Figure 1. An IMRT plan for a head and neck case. The MLCs shape the beams; a series of MLC apertures for each beam yield the beam fluence distributions; and the fluences from all beams result in the indicated dose distribution in the patient.

It was shown by Brahme et al. [15] and Brahme [14] that modulation of the fluence within each field can yield dose distributions that conform closer to the target than when only uniform beam fluences are used. This enables lower doses to sensitive structures adjacent to the target. Clinical trials show that IMRT reduces acute and late toxicity of healthy structures compared to 3DCRT; for a review, see Staffurth et al. [88]. A drawback of IMRT is that larger volumes are exposed to low doses, which may increase the risk of radiation-induced second malignancies [42, 64]. It also leads to prolonged treatment times and higher beam dose gradients, which increase respectively the risk and the impact of geometric errors [64].

The evolution of photon therapy has been reviewed by Bucci et al. [18]. For reviews of IMRT, see Bortfeld [12] and Ahnesjö et al. [2].

1.3 Proton therapy

A proton therapy treatment is delivered by a narrow proton beam extracted from a particle accelerator. In *pencil beam scanning* of protons, steering magnets are used to scan the narrow proton beam over the treatment volume. Pencil beam scanning

may be contrasted to passive scattering techniques, in which the proton beam is broadened by a scattering foil. In this thesis, pencil beam scanning, or *intensity-modulated proton therapy* (IMPT), is considered.

Protons exhibit two key advantages over photons with respect to therapeutic properties. First, a broad proton beam shows a significant increase in dose deposition at the end of the proton range. The region of increased dose is called the *Bragg peak*. Beyond the Bragg peak, the dose deposition is negligible, which enables improved sparing of healthy tissues behind the target compared to when photon beams are used. Second, the depth of the Bragg peak can be controlled by alteration of the energy of the incident protons. This amounts to an additional degree of freedom as compared to photon therapy. The superposition of pencil beams of different energies allows for spread-out Bragg peaks that cover the full target volume in depth. Depth-dose curves of proton pencil beams, a spread-out Bragg peak, and a photon beam are illustrated in Figure 2. That the doses increase with depth until the end of
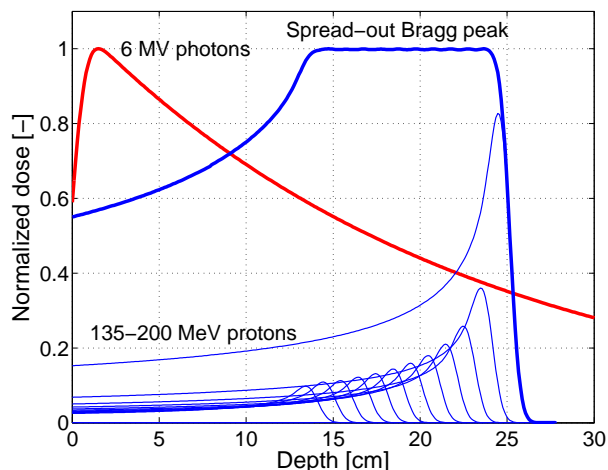
Figure 2. Depth-dose curves of a 6 MV photon beam (red), a proton spread-out Bragg peak (blue, thick), and the 135–200 MeV proton pencil beams constituting the spread-out Bragg peak (blue, thin).

the proton range makes proton treatments feasible within fewer beams than photon treatments. Figure 3 illustrates an IMPT plan for a head and neck case treated with two beams.

The location of the Bragg peak is highly affected by the proton *stopping power* of the traversed medium, which is the average energy loss of the protons per unit
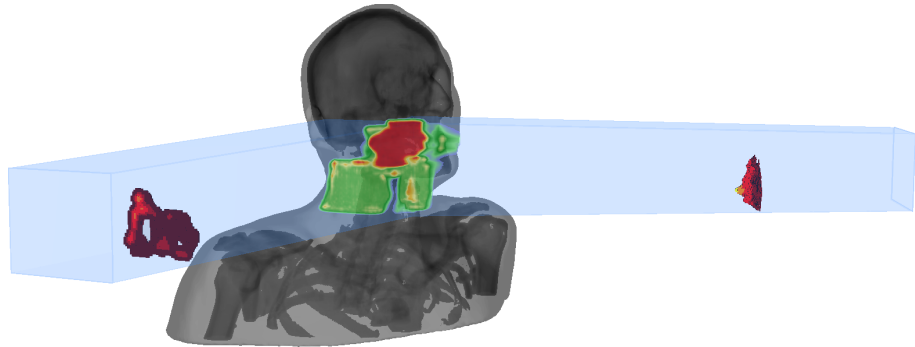
Figure 3. An IMPT plan for a head and neck case. For each beam, the fluence distribution resulting from a specific energy is illustrated. The fluences from all energies result in the indicated dose distribution in the patient.

path length. This amounts to a disadvantage of scanned proton treatments because it makes them much more affected by geometric errors than photon treatments.

A scanned proton beam is represented by a number of *spots*. A spot is defined by a lateral position in the fluence plane through which the narrow proton beam should pass, i.e., a point determining how the scanning magnets should direct the beam, and an energy level determining the depth of the Bragg peak. The fraction of the beam MU delivered by a given spot is controlled via the *spot weight*. Individual spot weights allow for modulated dose distributions in three dimensions from a single beam direction.

Pencil beam scanning results in 2–3 times less dose to uninvolved normal tissues as compared to IMRT [40]. Although there is yet little clinical evidence that proton therapy leads to improved outcomes [66], it has been argued that the high rates of local tumor control after 15 years that proton treatment has yielded would be unlikely to achieve with any other treatment technique [39].

For a historical review of proton therapy, see Smith [85], and for a review of proton therapy treatment planning, see Schwarz [82].

2   Treatment planning

A radiation therapy treatment plan is a specification of the number of beams and the settings that determine the manner in which the beams are delivered to the patient. The goal of treatment planning is to find a plan that yields a high probability of a curative treatment without complications. Since this probability cannot be de-

termined precisely without further assumptions, a plan with high such probability is often approximated as one with an appropriate balance between high target dose and low doses to healthy structures.

Planning based on 2D images, as well as 3DCRT treatment planning, is often performed by *forward planning*, which means that the treatment planner specifies the directions, shapes, and intensities of the beams, calculates and evaluates the resulting dose, and—if the dose is unsatisfactory—determines desirable parameter changes. The process is repeated until a satisfactory dose distribution is obtained.

The large number of parameters (e.g., aperture shapes, segment weights, spot weights) of IMRT and IMPT makes forward planning of all parameters practically impossible. Instead, computerized automated search methods are required. To this end, the treatment planner specifies desired qualities of the treatment plan, such as high target dose and low doses to healthy structures, and an optimization algorithm determines parameters with the aim to achieve these qualities as well as possible. This type of treatment planning is called *inverse planning*.

In this thesis, it is assumed that the treatment plan is identical over the course of the treatment. This is the standard practice of treatment planning today. In *adaptive radiation therapy*, the treatment is modified as new information becomes available [57, 99]. This has the possibility of increasing the probability of tumor control and reducing the doses to healthy structures.

2.1 Patient geometry

Images of the patient geometry guide the treatment planning process and help the treatment planner determine where the tumor and the healthy organs are located and hence which regions to treat and which to avoid.

Tomographic imaging techniques are used to generate 3D representations of the patient geometry. These techniques use 2D projections of the patient from multiple directions to compute cross-sectional images (or "slices") of the patient. The slices can be stacked to reconstruct a 3D representation of the patient. The most common imaging technique in treatment planning is CT, which provides a 3D representation of the patient tissue densities. The densities not only show where the organs are located, but are also required for accurate dose calculation. Other tomographic imaging techniques used in treatment planning are magnetic resonance imaging and positron emission tomography.

The image data is commonly visualized as 2D slices normal to the anteroposterior, superoinferior, or sinistrodextral axis of the patient. The data is used to specify the *regions of interest* (ROIs) of the treatment volume. The ROIs are regions of im-

portance in the planning process, such as the targets and healthy organs. Those representing healthy structures are called *organs at risk* (OARs). The ROIs may either be delineated manually, one 2D slice at a time, or with the help of segmentation software that, e.g., adapts organ models to the image data.

The delineated region of macroscopic disease, which can be determined from the patient images, is called the *gross tumor volume* (GTV). There may also be a microscopic spread of the cancer cells. To account for this, a margin that encompasses the region of suspected microscopic disease is added to the GTV [44]. The GTV plus the margin is called the *clinical target volume* (CTV). The CTV is generally further expanded into a *planning target volume* (PTV), which takes into account uncertainty of positioning, motion, and anatomical changes during the treatment [46, 47]. The PTV is defined to ensure a high probability of delivering sufficient dose to the tumor [94, 95].The different target volumes are described in more detail in ICRU Report 62 [45].

2.2   Evaluation of plan quality

The quality of a given treatment plan is primarily determined by the quality of the resulting dose distribution. The dose distribution is evaluated spatially, i.e., each 2D slice of the dose distribution is inspected, and on the basis of measures of the ROI dose distributions, such as the mean dose of the ROI. Many important physical measures of an ROI dose distribution can be evaluated by inspection of its (cumulative) *dose-volume histogram* (DVH). For a given ROI, its DVH shows how large fraction of the ROI that receives dose at or above each dose level. Examples of DVHs are shown in Figure 4. Some aspects that can be extracted from the DVHs and that are used to evaluate plan quality are the following:

- Dose-at-volume: D_v, the highest dose level d such that at least $v \%$ of a given ROI receives the dose d Gy or higher

- Volume-at-dose: V_d, the fraction of the volume of a given ROI that receives the dose d Gy or higher

- Minimum and maximum dose: D_{100} and D_0

- Near minimum and maximum dose: D_{98} and D_2

- Median dose: D_{50}
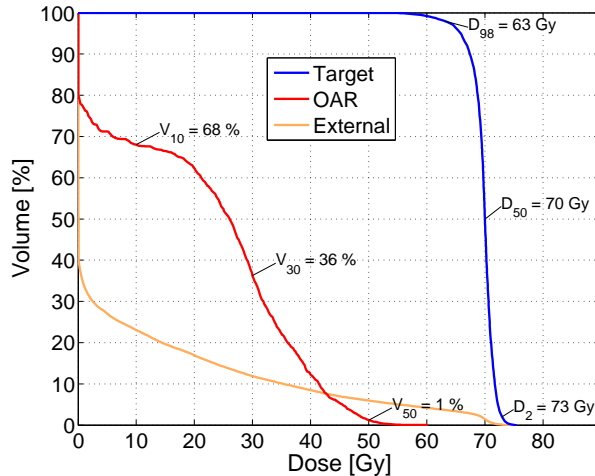
Figure 4. Example of DVHs of three ROIs. The external ROI is the full treatment volume. The 98 % volume of the target DVH shows that 98 % of the target receives 63 Gy or more, and the 50 Gy dose level of the OAR DVH shows that only 1 % of the OAR receives 50 Gy or more.

From these values, other quality measures can be determined, such as:

- Homogeneity index [47]: $(D_2 - D_{98})/D_{50}$

- Conformity index [45]: the ratio between the patient volume that receives 95 % of the prescription and the target volume

Biological measures of dose are also used in the evaluation of the treatment plan quality. For tumors, the probability of achieving tumor control by killing all clonogenic cells is predicted. This probability is called the *tumor control probability* (TCP) and is calculated as the probability that less than one tumor cell survives after the last treatment fraction, under assumption on some cell dose-response model [13, 97]. For healthy structures, the probability of different biological endpoints are predicted and included in the evaluation. An example of an endpoint for a head and neck patient is grade 2 xerostomia (dry mouth of a certain degree). The probability of complications is called the *normal tissue complication probability* (NTCP) [13, 52, 63]. A combination of TCPs and NTCPs can be used to form the measure P_+, which predicts the probability of a complication-free curative treatment [1].

There are also measures of physical dose with biological basis. The *equivalent uniform dose* (EUD) of an ROI is the dose level such that a uniform dose at that level would have an equivalent biological effect as the ROI dose distribution, under some biological model or fit to measured data [69, 70]. For a highly serial organ, which loses its function if one of its subvolumes is damaged, the EUD is mostly related to the maximum dose to the organ. For a parallel organ, the function of which is proportional to the fraction of its volume that is undamaged, the EUD is closer to the mean dose. For a tumor, which may survive unless all its cells are killed, the EUD relates mostly to the minimum dose.

2.3   Optimization functions

Closely related to the evaluation criteria are the optimization functions. These functions steer the optimization towards plans that perform well with respect to the evaluation criteria.

The dose distribution is denoted by d and is a mapping from $\mathbb{R}^3$ to $\mathbb{R}$ that takes each point p in the patient volume to the dose d_p in $\mathbb{R}_+$ deposited at p. Typically, each optimization function related to a physical dose criterion is associated with an ROI and penalizes deviations of the ROI dose distribution from some reference. For a survey of convex functions used in treatment planning, see Romeijn et al. [78]. Commonly used convex optimization functions penalize doses below, above, or other than some reference dose level. There are also nonconvex dose-based functions in use in treatment planning. A notable example is the DVH function [58], which penalizes deviations from a dose-volume criterion that specifies how large subvolume of a given ROI that is allowed to receive a dose above or below the reference level.

In this thesis, physical optimization functions that penalize dose deviations quadratically are primarily considered. The treatment plan is thus fitted to the reference levels in a sense similar to least squares. The functions come in two variants: one that penalizes doses below the reference and one that penalizes doses above the reference. These variants have names prefixed by respectively "minimum" and "maximum." Many of the functions f can be formulated as

$$f(d) = \int_0^1 \rho \left(y(v; d) - \hat{y}(v) \right)^2 dv, \tag{1}$$

where y is the dose-based quantity to be optimized, $\hat{y}$ is a reference that y should go above or below, and ρ is a ramp given by $\rho(z) = \min\{z, 0\}$ for the minimum

variant of the function and $\rho(z) = \max\{z, 0\}$ for the maximum variant, where z is a real number.

For an ROI consisting of the subset $\mathcal{V}$ of the patient volume, let $D(v; d)$ denote the dose level such that the fraction v in $(0, 1]$ of the ROI receives the dose $D(v; d)$ or higher, given the dose distribution d. Thus, $D(v; d)$ parameterizes the DVH of the ROI, and can be defined according to

$$D(v; d) = \max\left\{ d' \in \mathbb{R}_+ : \frac{|\{p \in \mathcal{V} : d_p \geq d'\}|}{|\mathcal{V}|} \geq v \right\},$$

where $|A|$ denotes the volume of the set A. Let $D^{\text{ref}}(v) = D(v; d^{\text{ref}})$ be a reference dose-volume function based on some fixed reference dose distribution d^{ref}.

Now, *minimum reference DVH* and *maximum reference DVH* optimization functions are obtained when y and $\hat{y}$ in (1) are defined according to

$$y(v; d) = D(v; d) \qquad \text{and} \qquad \hat{y}(v) = D^{\text{ref}}(v).$$

Reference DVH functions can be used to replicate dose distributions of reference plans or to ensure that the dose distribution of a given plan does not deteriorate. The latter is the topic of Paper F. *Minimum DVH* and *maximum DVH* functions are obtained when in the corresponding reference DVH functions the reference DVH curve D^{ref}, which defines $\hat{y}$, is given by respectively

$$D^{\text{ref}}(v) = \begin{cases} \hat{d} & \text{if } v \leq \hat{v} \\ 0 & \text{otherwise} \end{cases} \qquad \text{and} \qquad D^{\text{ref}}(v) = \begin{cases} \hat{d} & \text{if } v \geq \hat{v} \\ \infty & \text{otherwise} \end{cases},$$

where $\hat{d}$ is a reference dose level and $\hat{v}$ is a reference volume parameter. A minimum DVH function thus aims to yield a DVH curve that exceeds $\hat{d}$ at the volume $\hat{v}$ and a maximum DVH function aims to yield one that falls below $\hat{d}$ at $\hat{v}$. *Minimum dose* and *maximum dose* functions are obtained when the reference DVH curve D^{ref}, and thus also $\hat{y}$, is given by

$$D^{\text{ref}}(v) = \hat{d}$$

for all values of v. A *uniform dose* function is the sum of a minimum and a maximum dose function with the same reference dose level $\hat{d}$. The minimum, maximum, and uniform dose functions are convex in dose, whereas the DVH functions and reference DVH functions are not.

Functions including the EUDs allow biological dose-response to be taken into account in physical planning. The formula for EUD suggested by Niemierko [70] is based on the generalized mean and is given by

$$\text{EUD}_a(d) = \left(\int_{\mathcal{V}} d_p^a \, dp \right)^{1/a},$$ (2)

where $a \neq 0$. The parameter a depends on the seriality of the considered ROI. For a serial organ, large values of a are used, whereas for a parallel organ, small positive values of a are used. For target structures, negative values of a are used. *Minimum EUD* and *maximum EUD* functions are achieved when y and $\hat{y}$ in (1) are given by

$$y(v; d) = \text{EUD}_a(d) \qquad \text{and} \qquad \hat{y}(v) = \hat{d},$$

independent of the value of v, where $\hat{d}$ is the reference EUD level. The EUD (2) is concave in dose when $0 \neq a \leq 1$ and convex when $a \geq 1$ [23]. Thus, the minimum EUD function is convex when $0 \neq a \leq 1$ and the maximum EUD function is convex when $a \geq 1$.

The functions defined above are all independent from the spatial dose distribution. Two functions that depend on the exact location of the dose are the *reference minimum dose* and *reference maximum dose* functions, which are defined according to

$$f(d) = \int_{\mathcal{V}} \rho \left(d_p - d_p^{\text{ref}} \right)^2 \, dp,$$

where d^{ref} is a fixed reference dose distribution and ρ is as above. These are convex functions of dose. When d_p^{ref} is set to a constant reference level $\hat{d}$ for all p in the considered ROI volume, minimum and maximum dose functions are obtained.

2.4   Optimization problem

Among the parameters that are available to control the treatment plan are, for IMRT, the jaw and MLC leaf positions and the segment weights and, for IMPT, the spot positions and spot weights. For both modalities, the number of beams and the couch and gantry angles can additionally be controlled. Some of these parameters are determined before the optimization or by forward planning, whereas others correspond to the optimization variables that are determined by inverse planning. The optimization variables are denoted by x and the set of feasible optimization variables is denoted by $\mathcal{X}$. Here, the variables determined by the optimization are

assumed to be the MLC leaf positions and segment weights for IMRT and the spot weights for IMPT.

The desired qualities of the treatment plan are specified in the form of an objective function and, possibly, constraint functions. The objective function f penalizes deviation from the desirable qualities, and is thus to be minimized. In this thesis, the objective function f is typically a composite function, consisting of constituents $f_1, \ldots, f_n$ that reflect different desired qualities of the treatment plan. In most cases, these functions penalize conflicting qualities, such as the deviation from a high target dose and the deviations from low doses to critical organs, and do therefore not have a common minimizer. Hence, a tradeoff between the constituents is necessary. Such a tradeoff is commonly specified by the introduction of nonnegative importance weights $w_1, \ldots, w_n$ and the definition

$$f = \sum_{i=1}^{n} w_i f_i.$$

The constraints can be partitioned into physical constraints and planning constraints. Physical constraints are those that are imposed by hardware limitations and the laws of nature, e.g., MLC leaf position limitations and nonnegativity of the fluence. The physical constraints used in this thesis are all linear and are represented by the set $\mathcal{X}$ of feasible variables, which thus takes the form of the set $\{x : Ax \leq b\}$ for some matrix A and vector b. Planning constraints are specified by the treatment planner and reflect the requirements on the treatment plan that must not be compromised. These are defined as optimization functions $c_1, \ldots, c_m$ and are assumed to be satisfied whenever these functions evaluate to zero or less. Examples of objective function constituents and planning constraint functions are the functions defined in Section 2.3. With the dose distribution d as a function of the optimization variables x, the inverse planning problem can now be formulated as

$$
\begin{array}{lll}
\underset{x}{\text{minimize}} & f(d(x)) & \text{(objective function)} \\
\text{subject to} & c_i(d(x)) \leq 0, \quad i = 1, \ldots, m, & \text{(planning constraints)} \\
& x \in \mathcal{X}. & \text{(physical constraints)}
\end{array}
\tag{3}
$$

Numerical optimization is enabled by discretization of the patient geometry into cuboid *voxels*. Photon beams are discretized into rectangular *bixels*, whereas actively scanned proton beams are represented by discrete spots. An incidence matrix P, with one column for each bixel or spot and one row for each voxel, maps

fluence to dose. For photon beams, the dose distribution $d(x)$ is computed via the fluence:

$$d(x) = P\varphi(x),$$

where $\varphi(x)$ is the fluence resulting from the optimization variables x. For pencil beam scanning, the optimization variables x are the spot weights, and the dose distribution is linear in these:

$$d(x) = Px.$$

For IMRT, the mapping from machine parameters to fluence is nonconvex, which makes the IMRT optimization problem nonconvex also when the optimization functions f and $c_1, \ldots, c_m$ are convex functions of dose. For IMPT, the linear mapping from spot weights to dose make the IMPT optimization problem convex whenever the optimization functions are convex functions of dose.

2.5 Optimization method

In all of the appended papers, *sequential quadratic programming* (SQP) is used to solve optimization problems that arise in radiation therapy. An SQP method solves general nonlinear optimization problems, such as the treatment planning problem (3), as a sequence of subproblems. Each subproblem minimizes a quadratic model of the objective and the active constraints, subject to a linearization of the constraints. For a clean exposition, consider the mapping from optimization variables to dose implicit in the optimization functions, i.e., let $f(x) = f(d(x))$ and $c_i(x) = c_i(d(x))$ for $i = 1, \ldots, m$. Starting from point x_k in iteration k, the *search direction* denoted by p_k is determined as the solution to the quadratic programming problem

$$
\begin{aligned}
\underset{p}{\text{minimize}} \quad & \tfrac{1}{2} p^T \nabla^2_{xx} \mathcal{L}(x_k, \lambda_k) p + \nabla f(x_k)^T p \\
\text{subject to} \quad & c_i(x_k) + \nabla c_i(x_k)^T p \le 0, \qquad i = 1, \ldots, m, \qquad (4) \\
& x_k + p \in \mathcal{X},
\end{aligned}
$$

where $\mathcal{L}$ is the Lagrangian function given by $\mathcal{L}(x, \lambda) = f(x) + \lambda^T c(x)$, in which c is the vector of the constraint functions $c_1, \ldots, c_m$, and λ_k is the Lagrange multiplier vector (the dual variables), which is required to be nonnegative. The dual search direction d_k is the difference between the optimal Lagrange multipliers to (4) and the previous estimate λ_k. Note that x_k and λ_k in (4) are constant and that the constraint $x_k + p \in \mathcal{X}$ is equivalent to $A(x_k + p) \le b$ for some matrix A and vector b. For many optimization problems, including those considered in

this thesis, it would be computationally expensive to compute the Hessian of the Lagrangian, $\nabla^2_{xx}\mathcal{L}$, exactly. Therefore, $\nabla^2_{xx}\mathcal{L}$ is approximated by *quasi-Newton* approximations that are updated with *Broyden-Fletcher-Goldfarb-Shanno* updates in each iteration [71]. This makes it possible to represent the approximate Hessian of the Lagrangian as a sum of a small number of vector outer products. If necessary, the approximation of the Hessian of the Lagrangian is modified to maintain positive definiteness throughout the optimization, which ensures that the problem (4) is convex and can be solved to optimality [38]. Once the search directions p_k and d_k have been found, the SQP algorithm proceeds by determining the *step length* α_k by approximately solving the one-dimensional problem

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

where $M(x, \lambda)$ is a merit function that measures the quality of a given primal-dual pair (x, λ) with respect to the objective f and the constraints $c_1, \ldots, c_m$ and $\mathcal{X}$. Given the search directions and the step length, the new primal-dual pair (x_{k+1}, λ_{k+1}) is defined by $(x_{k+1}, \lambda_{k+1}) = (x_k, \lambda_k) + \alpha_k(p_k, d_k)$.

More details concerning the SQP algorithm used in Papers A, B, and D can be found in Gill et al. [38]. In Papers C, E, and F, an SQP algorithm developed by RaySearch Laboratories (Stockholm, Sweden) is used.

3 Uncertainties in radiation therapy

Several sources of uncertainty may affect a radiation therapy treatment. Unless uncertainties are accounted for in the planning process, they may lead to severe degradation of the quality of the delivered treatment as compared to the planned. Major sources of uncertainty include uncertainty in the position of the target relative to the beams, uncertainty regarding the location of the cancer cells, and uncertainty in the patient density data.

The errors in radiation therapy are typically partitioned into systematic (treatment preparation) errors and random (treatment execution) errors. Among the systematic errors are imaging inaccuracies [75, 89], image artifacts [6, 48], errors in the conversion from CT densities to stopping power [80, 81], dose calculation errors [50, 72], target misalignment during image acquisition [62, 94], and errors in the delineation of the ROIs [31, 36]. Among the random errors are daily setup errors [5, 43], organ motion [17, 53], and organ deformation [98]. A review of motion effects in IMRT is given by Webb [96]. Radiation therapy errors and margin recipes to account for these errors have been reviewed by van Herk [94]. Lomax

has reviewed the effects of calculation errors [59] and of motion [60] on IMPT treatments.

In Papers A–D, patient misalignments and density errors are considered and accounted for by robust optimization methods. Uncertainty regarding the alignment of the patient relative to the beams is called *setup uncertainty*. Uncertainty with respect to the densities and the conversion from densities to proton stopping power leads to *range uncertainty* of proton beams: the risk that the beams over- or undershoot. Range errors can result in that a Bragg peak planned to be located in front of an OAR is delivered to the OAR.

Uncertainty can be interpreted as a wider concept. In Paper E, the uncertainty accounted for during the optimization is with respect to the treatment planner's preferences regarding the treatment goals. This means that the importance weights of the optimization functions, as introduced in Section 2.4, are subject to uncertainty during the optimization.

3.1   Optimization under uncertainty

Optimization problems are typically considered with precisely specified parameters. The solutions to such problems may be highly sensitive to errors in the parameter values: if the true values differ from those used during the optimization, the solution may in fact be far from optimal. There are numerous methods that can be used to reach solutions that are robust against uncertainties.

Depending on the nature of the uncertainties, different approaches for taking them into account are preferable. For repeated processes where the sum of the results of the repetitions is an important quantity, it is desirable that the expectation be as good as possible, since by the law of large numbers, the average of the results of a repeated experiment tends towards the expected value. For such processes, the expected value of the objective function is often optimized. This requires an estimate of the probability distribution of the uncertainty. Optimization of the expectation of a function is called *stochastic programming*. Sometimes it is possible to compensate for the effects of the realized uncertainty at some cost, in which case stochastic programming with recourse can be used. For more details on stochastic programming, see, e.g., Shapiro et al. [83].

For a process that is not repeated, a good expected objective value may be insufficient to ensure a high probability of a satisfactory outcome. For such cases, it may be beneficial to hedge against uncertainty in a more conservative manner. To this end, one may optimize risk measures such as the *value at risk* (VaR), which is the best possible outcome conditioned on that one of the α worst possible out-

comes will occur, where $0 < \alpha \leq 1$. This measure is however in general neither subadditive nor convex and is often difficult to optimize. A related measure is the *conditional value at risk* (CVaR) [4], which measures the expected value of the outcome conditioned on that one of the α worst possible outcomes will occur, where $0 < \alpha \leq 1$. The CVaR is convex, has better numerical properties than the VaR, and can be readily optimized [76].

When the probability distributions of the uncertain parameters are unknown, *robust optimization* may provide a suitable means of accounting for uncertainty. In robust optimization, only bounds on the uncertain parameter values are assumed. The optimization then hedges against the worst case outcome of the uncertainty for parameters within the bounds and requires the solution to remain feasible for all realizations of uncertainty within the specified set. For many special classes of robust programs, such as robust linear programs, equivalent deterministic programs can be derived, which can be solved by standard optimization methods [8,9,35,87].

Stochastic programming, CVaR optimization, and robust optimization are special cases of *minimax stochastic programming* [32, 84], in which the probability distribution of the errors is modeled as known only within some bounds, and an optimal solution is sought after for the expectation of the objective function under the worst case probability distribution. This leads to distributionally robust solutions. If the bounds on the probability distribution are tight, the minimax stochastic programming becomes equivalent to stochastic programming. If point distributions, which assign the probability 1 to single realizations of the uncertainty, are included in the set of possible probability distributions, the resulting solution can be made robust in the same sense as by robust optimization.

3.2   Treatment plan optimization under uncertainty

Many optimization approaches that account for uncertainty require some discretization of the possible realizations of the uncertainty in order to yield computational optimization problems. This is often the case in radiation therapy, because errors many times result in nonlinear changes of dose for which there are no known analytical expressions. The realizations are therefore discretized into *scenarios*, where each scenario corresponds to a specific realization of the uncertainty. The scenarios are enumerated by the set $\mathcal{S}$. The dose deposited to the point p in the patient volume under scenario s, given the optimization variables x, is denoted by $d_p(x; s)$. The *nominal scenario* corresponds to no error.

3.2.1 Robust IMRT

For cases of relatively homogeneous density treated with photon-mediated IMRT, the *static dose cloud approximation* [90] is often viable. Under this approximation, the patient is assumed to be moving within a static dose distribution. The effect of a patient setup error thus corresponds to a rigid translation of the dose distribution. This enables easy computation of scenario doses and of moments of the doses to specific points in the patient geometry, such as the expected doses or the expected squared doses, which together yield the dose variances. The expected doses and the dose variances can be used to compute measures such as the expectation of the uniform dose optimization functions defined in Section 2.3.

In the IMRT literature, stochastic programming is often used to account for uncertainties in IMRT. Löf et al. [56] accounted for systematic and random errors by optimizing the expectation of the biological objective function P_+ measuring the probability of complication-free tumor control. Unkelbach and Oelfke [92, 93] also used stochastic programming accounting for systematic and random errors, but applied to uniform dose optimization functions. The stochastic programming problem is formulated

$$\underset{x \in \mathcal{X}}{\text{minimize}} \quad \mathbb{E}_\pi \left[f(d(x; S)) \right], \tag{5}$$

where π is the assumed probability distribution of the uncertainties and S is the random variable picking the scenario s from the set $\mathcal{S}$ with probability π_s. Cromvik and Patriksson [29] accounted for systematic errors by CVaR optimization applied to physical optimization functions. The CVaR optimization problem is formulated

$$\begin{aligned} \underset{\lambda, x}{\text{minimize}} \quad & \lambda + \tfrac{1}{\alpha} \mathbb{E}_\pi \left[\max\{ f(d(x; S)) - \lambda, 0 \} \right] \\ \text{subject to} \quad & x \in \mathcal{X}, \end{aligned} \tag{6}$$

in which $0 < \alpha \leq 1$.

Other authors have used methods from robust optimization to account for uncertainties in IMRT. Chu et al. [24] and Ólafsson and Wright [73] considered the dose to each voxel as a random variable. Using robust linear programming methods from El-Ghaoui and Lebret [35] and Ben-tal and Nemirovski [8], they optimized IMRT plans to ensure a high probability of delivering an adequate dose to each voxel considered independently. Chan et al. [21] optimized IMRT plans accounting for organ motion following an uncertain probability distribution. They showed how a robust linear programming formalism similar to that of Bertsimas and Sim [9] can be used to create treatment plans that deliver a sufficient expected

dose to the tumor under any of the organ motion probability distributions from a given set.

Under the static dose cloud assumption, the effect of robust optimization is similar to margins. Optimization with respect to an uncertain probability distribution [21] and the use of CVaR [29] thus provide continuous scaling between stochastic programming and margin-like planning. These methods for robustness are thus intended to handle uncertainties in a less conservative manner than margins. The benefit of doing so is that the risk of complications may be reduced.

Rather than optimizing the penalty functions of dose directly, Sobotta et al. [86] maximized the probability that the functions evaluated to values within specified intervals for treatments subject to setup uncertainty. They could thereby maximize the probability of achieving a satisfactory treatment. Maximization of the probability of satisfying the planning criteria is the topic of Paper C.

3.2.2 Robust IMPT

The current standard of accounting for uncertainty in proton therapy is to use margins [46]. However, the static dose cloud approximation is not suitable for IMPT: The Bragg peak positions are highly dependent on the stopping powers, and hence the densities, of the traversed tissues. Combined with modulated beam doses with sharp dose gradients, this makes errors lead to deformations of the proton dose distributions. Figure 5 illustrates how the dose distribution of a spot delivered to a lung tumor is distorted as the result of a setup error. Margins may therefore have

(a) Nominal setup   (b) Shifted setup
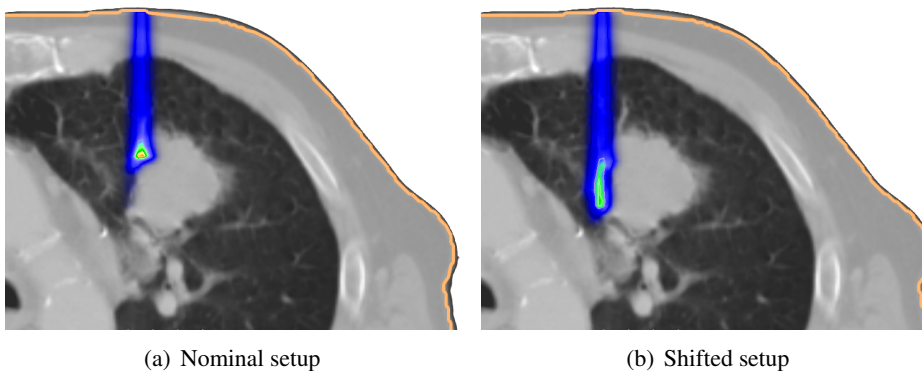
Figure 5. (a) Planned dose distribution of a spot delivered to a lung tumor. (b) Dose distribution of the same spot after a setup shift.

only little of the intended effect for IMPT plans, especially for highly modulated treatments and heterogeneous tissues such as the lung. This has also been found to be the case in empirical studies [3, 22, 37, 55].

Although optimization of the expectation of physical functions of dose has been used to account for systematic range and setup uncertainties in IMPT [90, 91], the lack of protection provided by margins may be the reason behind the higher degree of conservativeness that has often been employed to account for uncertainties in IMPT than in IMRT. In the first approaches to robust IMPT, the dose to each voxel was considered independent from the doses to other voxels: Unkelbach et al. [91] and Chan [20] accounted for systematic range errors in IMPT by optimizing the worst case dose to each voxel considered independently. Pflugfelder et al. [74] used a similar measure to account for systematic range and setup errors in IMPT. They optimized the sum of the objective function applied to the nominal scenario dose distribution and to the worst case dose distribution. The worst case dose distribution was introduced by Lomax et al. [61] and is defined as either $d^{\min}$ or $d^{\max}$, according to

$$d_p^{\min}(x) = \min_{s \in \mathcal{S}} d_p(x; s) \quad \text{or} \quad d_p^{\max}(x) = \max_{s \in \mathcal{S}} d_p(x; s) \tag{7}$$

for each point p in the patient volume, depending on whether the structure considered is a target, for which lower doses are often worse, or an OAR, for which higher doses are worse. Note that the minimum or maximum is taken for each point p considered independently. Liu et al. [55] also used the worst case dose distribution, but accounted for $d^{\max}$, in addition to $d^{\min}$, for the targets in order to avoid overdosage. An optimization problem using worst case dose distributions can be formulated as

$$\underset{x \in \mathcal{X}}{\text{minimize}} \quad \sum_{i \in \mathcal{T}} w_i f_i(d^{\min}(x)) + \sum_{i \in \mathcal{O}} w_i f_i(d^{\max}(x)), \tag{8}$$

possibly with the addition of terms $f_i(d^{\max}(x))$ for i in $\mathcal{T}$, where the optimization functions $i = 1, \ldots, n$ are partitioned into those applied to target structures, indexed by $\mathcal{T}$, and those applied to OARs, indexed by $\mathcal{O}$.

Chan [20] discussed the possibility of minimizing the objective function in the worst case scenario—with correlation between the voxels preserved—as an alternative to optimizing the worst case dose distribution. Formulated as a linear program, this yields too large problems to be computationally tractable. In Papers A and B, nonlinear programs optimizing the worst case scenario objective function value with respect to systematic range and setup errors are considered, which yield

computationally tractable problems. In such formulations, the constituent optimization functions for all ROIs are always considered in the same scenario, so that only physically realizable scenarios are taken into account. This can be formulated as the worst case scenario (or "minimax") optimization problem

$$\underset{x \in \mathcal{X}}{\text{minimize}} \quad \max_{s \in \mathcal{S}} \sum_{i=1}^{n} w_i f_i(d(x; s)). \tag{9}$$

Motivated by an application to multicriteria optimization, Chen et al. [22] used an approach that is intermediate to the approach considering the worst case dose distribution and that considering the worst case scenario. They optimized the worst case scenario for each optimization function considered independently, which can be formulated as the optimization problem

$$\underset{x \in \mathcal{X}}{\text{minimize}} \quad \sum_{i=1}^{n} w_i \max_{s \in \mathcal{S}} f_i(d(x; s)). \tag{10}$$

The formulations (8)–(10) show that the main difference between the approaches for robust IMPT is the level at which the maximum (or minimum) operator is applied. The deeper the level of the operator, the more conservative the method becomes. A more conservative method accounts for a larger number of scenarios. Not all of these scenarios are physically realizable: there is no physical scenario in which each point, considered independently, receives its worst case dose. Casiraghi et al. [19] studied the DVHs defined by the worst case dose distributions and those defined by the worst case scenario, taken as the upper and lower envelopes of the scenario DVHs, which maintains correlation between the voxels. They concluded that the worst case scenario provides accurate predictions of the DVHs, whereas the worst case dose distribution provides overly conservative DVH predictions.

4 Multicriteria optimization of radiation therapy

An unavoidable tradeoff in radiation therapy treatment planning is that between target coverage and sparing of healthy tissues. In conventional planning, this tradeoff is resolved by the introduction of importance weights for the optimization functions reflecting the different criteria, as in the optimization problem (3). However, the effect of the weighting is not fully known before the optimization problem is solved. A manual process of parameter tuning and reoptimization is thus often required before an acceptable plan is found. Multicriteria optimization (MCO)

methods have been introduced to aid in the decision making process. The goal of MCO methods is to avoid the manual tuning loop. Many aspects of general multicriteria optimization are detailed in Miettinen [65].

4.1   Multicriteria optimization

The problem of determining a tradeoff between n conflicting objectives $f_1, \ldots, f_n$, with $n \geq 2$, can be formulated as a multicriteria optimization problem:

$$
\begin{aligned}
\underset{x}{\text{minimize}} \quad & f(d(x)) = \begin{bmatrix} f_1(d(x)) & \cdots & f_n(d(x)) \end{bmatrix}^T \\
\text{subject to} \quad & c_i(d(x)) \leq 0, \qquad\qquad\qquad\qquad i = 1, \ldots, m, \\
& x \in \mathcal{X},
\end{aligned} \tag{11}
$$

where the constraints are as in problem (3). The set of optimal solutions to this vector-valued problem is usually considered to be the set of *Pareto optimal* points. A feasible solution x^* is Pareto optimal if there exists no feasible x such that

$$
f(d(x)) \leq f(d(x^*)) \text{ and } f(d(x)) \text{ is distinct from } f(d(x^*)), \tag{12}
$$

where the inequality is componentwise. The set of Pareto optimal solutions typically consists of infinitely many points and cannot be easily determined. Instead, it may be approximated by a finite set of points, each obtained by minimization of a scalar-valued substitute for the vector-valued objective of problem (11). Scalar-valued substitutes can be achieved by a number of scalarization methods. In this thesis, the common technique is used in which a weighted sum of the optimization functions is minimized, which amounts to solving problem (3) for different importance weights $w_1, \ldots, w_n$. A collection of plans corresponding to optimal solutions to a number of instances of problem (3) with different weights then forms a database of plans that is used to approximate the Pareto set.

The database that approximates the Pareto set may be used in an interactive planning process, in which convex combinations in dose are formed between the database plans. Such convex combination can be formed in real-time, and allow the treatment planner to continuously explore the tradeoffs of the considered patient case. There exist different navigation methods for exploration of the convex combinations of plans [27, 67].

4.2   Robust multicriteria optimization

As in the standard singlecriterion treatment planning problem (3), the parameters of a multicriteria problem like (11) are affected by uncertainty. It is therefore equally

important to account for uncertainty in multicriteria optimization as in singlecriterion optimization. However, not all methods used to account for uncertainty in the singlecriterion case can be directly transferred to a multicriteria setting.

A method that can be transferred to the multicriteria case is stochastic programming on the form (5). It results in the vector-valued objective

$$\mathbb{E}_\pi \left[\left[f_1(d(x; S)) \; \cdots \; f_n(d(x; S)) \right]^T \right],$$

which is equivalent to

$$\left[\mathbb{E}_\pi[f_1(d(x; S))] \; \cdots \; \mathbb{E}_\pi[f_n(d(x; S))] \right]^T \tag{13}$$

because the expectation operation is linear. Deb and Gupta [30] considered uncertain multicriteria optimization in which the expected values of the objective functions were taken as averages over a neighborhood of the current optimization variable values. They also optimized the nominal objective function under constraints on the variability of the functions over the uncertainty set.

Other robust multicriteria methods that account for uncertainty in an objective-wise manner have also been considered. An example is the objective-wise worst case method (10) used by Chen et al. [22], which accounts for the worst case scenario for each objective considered independently. It results in a vector-valued objective on the form

$$\left[\max_{s \in \mathcal{S}} f_1(d(x; s)) \; \cdots \; \max_{s \in \mathcal{S}} f_n(d(x; s)) \right]^T. \tag{14}$$

This form of robustness was considered theoretically by Kuroiwa and Lee [51], who provided scalarization methods and robust optimality conditions.

Since the objectives (13) and (14) are vector-valued, the standard definition (12) of Pareto optimality can be used to define their solution sets. Worst case optimization on the form (9) applied in a multicriteria setting results in the objective

$$\max_{s \in \mathcal{S}} \left[f_1(d(x; s)) \; \cdots \; f_n(d(x; s)) \right]^T. \tag{15}$$

The effect of the maximum operation is not well-defined unless the objective vector is scalarized. It moreover depends on the type of scalarization. Therefore, the objective vectors must generally be considered for all scenarios in $\mathcal{S}$, which results in a set of vectors. The objective (15) is thus set-valued, so Pareto optimality

according to (12) does not apply. This form of robust multicriteria optimization was proposed by Ehrgott et al. [34]. They extended the definition of Pareto optimality to set-valued objectives such as (15): a feasible solution x^* is robust Pareto optimal if there is no feasible x such that

$$\{f(d(x;s)) : s \in \mathcal{S}\} \subseteq \{f(d(x^*;s)) - r : s \in \mathcal{S}, r \in \mathbb{R}^n_+, r \neq 0\},$$

in which f is the vector of the optimization functions $f_1, \ldots, f_n$. Ehrgott et al. [34] showed that neither weighted sum scalarization nor the ε-constraint method [33] can generate all optimal solutions to minimax robust multicriteria programs in the general case. They also considered objective-wise worst case optimization on the form (14), and showed that this method can generate solutions that are optimal to minimax robust multicriteria programs. Extensions of their results are provided in Paper D.

4.3 Deliverability of navigated plans

For IMPT, convex combinations of plans correspond to convex combinations of nonnegative spot weights. Because such combinations remain nonnegative, convex combinations of IMPT plans remain deliverable (although a cutoff for low-weighted spots may be required). For IMRT, convex combinations of deliverable step-and-shoot plans are not directly deliverable within the number of segments of the database plans: the number of segments of the convex combination is the sum of those of the database plans with positive coefficients. Therefore, a conversion step may be beneficial, which creates a new plan replicating the navigated convex combination of dose with a restricted number of segments. The replication may be performed by optimization of the new plan with respect to reference DVH functions, as defined in 2.3, with the reference dose-volume function D^{ref} defined by the navigated dose distribution [11, 26].

The replication that is performed after the navigation results in a deliverable plan, but has the disadvantage that the resulting plan may be clinically unacceptable or correspond to a tradeoff other than the navigated one. It may therefore lead to a manual loop of parameter tuning, thereby defeating the purpose of MCO. To avoid this problem, deliverable navigation for step-and-shoot IMRT has been devised. Craft and Richter [28] proposed using a database of deliverable step-and-shoot plans and restrict the number of positive coefficients in the convex combinations of plans. The resulting convex combinations are thus deliverable with the number of segments increased to the sum of those of a small number of plans. Salari and Unkelbach [79] used column generation [77] to generate a single set of apertures

well-suited for all plans in the database representing the Pareto set. Navigation of such a representation corresponds to modification of the segment weights and thus results in deliverable plans without increasing the number of segments. In Paper E, these two methods for deliverable navigation are generalized by a method that allows for some shared and some individual apertures for the database plans.

5 Summary and main contributions

5.1 Summary of the appended papers

The appended papers are organized thematically: Papers A–D concern radiation therapy treatment plan optimization methods accounting for geometric uncertainty. In Paper D, the geometric uncertainty is considered in a multicriteria setting. The topic of MCO is continued in Paper E, where suboptimal deliverable databases of plans are improved upon. Finally, the method introduced in Paper F improves upon treatment plans that have been obtained with too weak planning criteria.

Paper A: Minimax optimization for handling range and setup uncertainties in proton therapy

Paper A is co-authored with Anders Forsgren and Björn Hårdemark, and has been published in *Medical Physics*, Vol. 38, No. 3, pp. 1672–1684, 2011.

In this paper, IMPT subject to systematic range and setup uncertainties is considered. Minimax optimization is proposed as an alternative to margins to account for the uncertainties. In the minimax method, the possible errors (of the magnitude that margins are intended to protect against) are discretized into scenarios. The optimization then aims to minimize the optimization function penalties in the worst case scenario.

The minimax method is evaluated on three patient cases that represent different treatment planning conditions: a lung case, in which the treatment region is of heterogeneous density; a paraspinal case, in which the tumor surrounds the spinal cord and contains titanium bolts; and a prostate case, which is of homogeneous density. The resulting plans are compared to benchmark plans obtained by planning with conventional margins; margins and uniform beam doses (*single field, uniform dose*); and margins, uniform beam doses, and the densities in the low-density regions overridden during optimization (*material override*). For all cases,

the minimax method resulted in more robust target coverage and lower doses to the OARs than the conventional methods.

Paper B: A characterization of robust radiation therapy treatment planning methods—from expected value to worst case optimization

Paper B has been published in *Medical Physics*, Vol. 39, No. 8, pp. 5169–5181, 2012.

In this paper, random errors of uncertain probability distributions are considered in addition to the systematic errors accounted for in Paper A. To allow for different amounts of conservativeness in the optimization, a minimax stochastic framework is formulated that generalizes many of the previous methods used in the literature and allows for scaling between stochastic programming and minimax optimization. Methods from this framework are characterized empirically by application to a phantom case subject to a variety of uncertainties.

Three special cases of methods from this framework are considered: (i) expected value, (ii) CVaR, and (iii) worst case optimization. These methods are applied to a phantom case with a C-shaped target partly surrounding an OAR. The case is considered when subject to systematic errors, random errors, or simultaneous systematic and random errors. The random errors are assumed to follow known as well as uncertain probability distributions. Systematic errors and the uncertain probability distributions are handled by means of methods (i)–(iii), whereas the random errors are handled by expected value optimization. Tradeoff curves with respect to target coverage and OAR sparing for methods (i)–(iii) show that all methods perform similarly in the best case scenario and in the mean case, but that the tradeoff in the worst case scenario improves with the conservativeness of the method. It is moreover shown that optimization with respect to random errors of known probability distribution can lead to highly heterogeneous dose distributions.

Paper C: Maximizing the probability of satisfying the planning criteria in radiation therapy under setup uncertainty

Paper C is co-authored with Anders Forsgren and Björn Hårdemark., and has been submitted to *Physics in Medicine and Biology*.

In worst case optimization, there is a risk that the planning criteria cannot be fulfilled in all scenarios. This paper introduces a method that determines how large

setup errors that can be accounted for while all planning criteria are satisfied. To this end, the magnitudes of the setup errors that are accounted for are included as variables in the optimization together with the standard variables for IMPT or photon-mediated IMRT. These magnitudes are then maximized (within specified bounds) subject to constraints that enforce the planning criteria under the considered errors. This results in a maximization of the probability of satisfying the planning criteria.

The method is applied to two patient cases subject to IMRT or IMPT treatment and compared to worst case optimization accounting for a priori determined setup errors. For both cases and modalities, the proposed method reduced the size of the region within which the optimization aimed to satisfy the planning criteria, and thereby generated plans that satisfied a larger number of planning criteria under the retracted setup shifts than the method accounting for a priori errors. The proposed method moreover satisfied a larger number of the planning criteria under the a priori setup errors. It thereby enabled better plans than robust planning with respect to a priori determined setup errors.

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

Paper D is co-authored with Rasmus Bokrantz, and has been printed as Technical Report TRITA-MAT-2013-OS5, Department of Mathematics, Royal Institute of Technology, 2013.

In this paper, MCO for IMPT in the presence of systematic setup uncertainty is considered. The uncertainty is accounted for by application of worst case optimization. Mathematical theory for robust Pareto optimality is introduced and a subset of the robust Pareto optimal solutions that are optimal under risk-averse preferences is defined and characterized.

The worst case optimization is contrasted to objective-wise worst case optimization. For a one-dimensional phantom geometry, it is shown that the worst case method better exploits spatial structure and provides the treatment planner with more control over the tradeoffs than does the objective-wise worst case method. The parameter changes in the minimax stochastic formulation of Paper B for systematic errors that allow for control over robustness and conservativeness are detailed. Here, robustness is defined as the magnitude of the uncertainties that are accounted for and conservativeness is defined as the amount of variability in the estimated probability distributions that is protected against. It is shown that by

reduction of the robustness, dose escalation can be made feasible while a sharp lateral dose fall-off is maintained. A decrease in conservativeness is shown to produce a gentle dose fall-off that contributes little to tumor control. The sharp lateral fall-off is shown to be motivated because it minimizes the integral dose under the constraint that the expected TCP must be at least 95 %.

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

Paper E is co-authored with Rasmus Bokrantz, and has been submitted to *Physics in Medicine and Biology*.

The problem of deliverable Pareto surface navigation for step-and-shoot IMRT is considered. This problem amounts to generating a representation of the Pareto set such that convex combinations of plans from the representation remain directly deliverable. In this paper, the Pareto set is represented by plans that have some apertures from a collective pool and some apertures that are individual to the plans. All segment weights are individual. Since some apertures are shared, the number of segments required to deliver convex combinations of plans is reduced compared to when all apertures are individual: combinations of k plans with n_{sh} shared and n_{ind} individual apertures result in plans deliverable within $n_{sh} + k n_{ind}$ apertures. The shared apertures constitute a coupling between the plans representing the Pareto set. Changes to one such plan may thus affect the other plans. All plans representing the Pareto set are therefore optimized simultaneously by direct step-and-shoot optimization with constraints that enforce some of the apertures to be identical across the plans. This method generalizes previous methods for deliverable navigation to allow for some shared and some individual apertures. The introduced method can also be used as a post-processing step to previous methods for deliverable navigation in order to improve upon their plans. Application of the method to subsets of plans of the Pareto set representation enables deliverable Pareto surface navigation between plans of similar quality as those of the unrestricted (non-navigable) Pareto set of plans for which all apertures are individual. The method is applied to a paraspinal case with two or three objectives. The results show that the use of a few individual apertures leads to much increased plan quality compared to plans with all apertures shared.

Paper F: Automated improvement of radiation therapy treatment plans by optimization under reference dose constraints

Paper F has been published in *Physics in Medicine and Biology*, Vol. 57, No. 23, pp. 7799–7811, 2012.

This paper deals with the fact that the level of experience of the treatment planner has large impact on the quality of treatment plans [7, 10, 25], which indicates that many plans are suboptimal. It may thus be possible to improve upon some criterion of a given treatment plan while all other criteria are maintained. In this paper, a method is introduced that, starting from a given plan, improves thereupon by minimizing the doses to the OARs while the doses of all structures are constrained to be at least as good as in the given plan. The constraints that enforce this are based on reference DVH functions and reference dose functions, as defined in Section 2.3. The minimum and maximum operations in these constraints can lead to convergence difficulties when gradient-based optimization methods are warm-started from points where the operations evaluate to zero. The difficulties are counteracted by the introduction of log-sum-exp regularization of these operations.

5.2 Main contributions

The main contributions of the appended papers are within three fields:

Robustness of treatment plans

Paper A introduces minimax optimization into the field of IMPT as a substitute for margins. It is the first method shown to lead to improved target coverage robustness and reduced OAR doses as compared to conventional heuristics for robustness.

Paper B provides a generalized framework that shows how many previous methods for robust treatment plan optimization are related. It extends the method of Paper A to account for random errors of uncertain probability distribution in addition to the systematic errors. The application of this framework to a phantom case subject to IMPT treatment provides the first indication that the more conservative methods to account for uncertainty may yield more attractive tradeoffs.

Paper C introduces a new method for maximizing the probability of satisfying the planning criteria when the treatment is subject to setup uncertainty. This method requires fewer scenarios than previous methods with the same goal [86]. The results show that by asking for a little less, the treatment planner can sometimes reach better plans.

Paper D provides the first MCO method for robust IMPT that uses worst case optimization. The results show that worst case optimization on the form of (9) provides the treatment planner with more control over the resulting plan than objective-wise worst case optimization similar to (10). Moreover, the empirical results of Paper B that favors worst case optimization instead of expected value optimization are supported by biological arguments.

Improvement of treatment plans

Paper E generalizes deliverable Pareto surface navigation to the case where some apertures are shared between plans and some are individual. It provides empirical evidence that such partial sharing can be beneficial with respect to the tradeoff between plan quality and delivery time. This paper moreover provides the first method for deliverable Pareto surface navigation that can converge to the ideal (non-navigable) Pareto surface of plans for which all apertures are individual.

Paper F provides a solution to the problem of suboptimal plans that result when the treatment planner uses too weak requirements in the optimization. The constraints based on reference DVH and dose introduced in the paper can provide more stringent assurance than conventionally that the dose distribution does not deteriorate for lexicographic ordering methods for IMRT [16, 49].

Theoretical contributions

Paper D extends the theory for robust MCO. Specifically, the concept of *convex hull efficiency* is defined. The set of convex hull efficient solutions is shown to be a proper subset of the robust efficient solutions. It is shown that a necessary condition for convex hull efficiency is optimality with respect to a strictly increasing convex scalarizing function, and that a sufficient condition is optimality with respect to a strongly increasing convex scalarizing function.

5.3   Contributions by co-authors

For Papers A and C, Anders Forsgren and Björn Hårdemark acted as advisors, suggesting directions of the research and supervising the work. Papers D and E were written jointly with Rasmus Bokrantz. The design of the computational experiments and the theoretical work of these papers were performed in close collaboration. The respective first author—Rasmus Bokrantz for Paper D and I for Paper E—was principally responsible for conducting the computational experiments.

6 Bibliography

[1] A.-K. Ågren, A. Brahme, and I. Turesson. Optimization of uncomplicated control for head and neck tumors. *Int. J. Radiat. Oncol. Biol. Phys.*, 19(4):1077–1085, 1990.

[2] 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.

[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] P. Artzner, F. Delbaen, J. M. Eber, and D. Heath. Coherent measures of risk. *Math. Finance*, 9(3):203–228, 1999.

[5] E. Astreinidou, A. Bel, C. P. Raaijmakers, C. H. Terhaard, and J. J. Lagendijk. Adequate margins for random setup uncertainties in head-and-neck IMRT. *Int. J. Radiat. Oncol. Biol. Phys.*, 61(3):938–944, 2005.

[6] J. F. Barrett and N. Keat. Artifacts in CT: Recognition and avoidance. *Radiographics*, 24(6):1679–1691, 2004.

[7] V. Batumalai, M. G. Jameson, D. F. Forstner, P. Vial, and L. C. Holloway. How important is dosimetrist experience for intensity modulated radiation therapy? A comparative analysis of a head and neck case. *Pract. Radiat. Oncol.*, In Press., 2012.

[8] A. Ben-Tal and A. Nemirovski. Robust convex optimization. *Math. Oper. Res.*, 23(4):769–805, 1998.

[9] D. Bertsimas and M. Sim. The price of robustness. *Oper. Res.*, 52(1):35–53, 2004.

[10] 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.

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

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

[13] A. Brahme. Dosimetric precision requirements in radiation therapy. *Acta Oncol.*, 23(5):379–391, 1984.

[14] A. Brahme. Optimization of stationary and moving beam radiation therapy techniques. *Radiother. Oncol.*, 12(2):129–140, 1988.

[15] 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.

[16] S. Breedveld, P. R. M. Storchi, M. Keijzer, A. W. Heemink, and B. J. M. Heijmen. A novel approach to multi-criteria inverse planning for IMRT. *Phys. Med. Biol.*, 52(20):6339–6353, 2007.

[17] K. R. Britton, Y. Takai, M. Mitsuya, K. Nemoto, Y. Ogawa, and S. Yamada. Evaluation of inter- and intrafraction organ motion during intensity modulated radiation therapy (IMRT) for localized prostate cancer measured by a newly developed on-board image-guided system. *Radiat. Med.*, 23(1):14–24, 2005.

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

[19] 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.

[20] T. C. Y. Chan. *Optimization under Uncertainty in Radiation Therapy*. PhD thesis, Massachusetts Institute of Technology, 2007.

[21] T. C. Y. Chan, T. Bortfeld, and J. N. Tsitsiklis. A robust approach to IMRT optimization. *Phys. Med. Biol.*, 51(10):2567–2583, 2006.

[22] 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.

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

[24] M. Chu, Y. Zinchenko, S. G. Henderson, and M. B. Sharpe. Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty. *Phys. Med. Biol.*, 50(23):5463–5477, 2005.

[25] H. T. Chung, B. Lee, E. Park, J. 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. Biol. Phys.*, 71(4):1167–1174, 2008.

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

[27] D. Craft, T. Halabi, H. A. Shih, and T. Bortfeld. An approach for practical multiobjective IMRT treatment planning. *Int. J. Radiat. Oncol. Biol. Phys.*, 69(5):1600–1607, 2007.

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

[29] C. Cromvik and M. Patriksson. On the robustness of global optima and stationary solutions to stochastic mathematical programs with equilibrium constraints, part 2: Applications. *J. Optim. Theory Appl.*, 144(3):479–500, 2010.

[30] K. Deb and H. Gupta. Introducing robustness in multi-objective optimization. *Evol. Comput.*, 14(4):463–494, 2006.

[31] D. F. Dubois, B. R. Prestidge, L. A. Hotchkiss, J. J. Prete, and W. Bice. Intraobserver and interobserver variability of MR imaging- and CT-derived prostate volumes after transperineal interstitial permanent prostate brachytherapy. *Radiology*, 207(3):785–789, 1998.

[32] J. Dupacǒvá. Minimax stochastic programs with nonseparable penalties. In *Optimization Techniques, Lecture Notes in Control and Information Sciences*, pages 157–163. Springer-Verlag, Berlin, Germany, 1980.

[33] M. Ehrgott. *Multicriteria Optimization*. Lectures notes in economics and mathematical systems. Springer, Berlin-Heidelberg, Germany, 2005.

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

[35] L. El Ghaoui and H. Lebret. Robust solutions to least-squares problems with uncertain data. *SIAM. J. Matrix Anal. A.*, 18(4):1035–1064, 1997.

[36] C. Fiorino, M. Reni, A. Bolognesi, G. M. Cattaneo, and R. Calandrino. Intra- and inter-observer variability in contouring prostate and seminal vesicles: implications for conformal treatment planning. *Radiother. Oncol.*, 47(3):285–292, 1998.

[37] 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.

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

[39] B. Glimelius, A. Ask, G. Bjelkengren, T. Björk-Eriksson, E. Blomquist, B. Johansson, M. Karlsson, and B. Zackrisson. Number of patients potentially eligible for proton therapy. *Acta Oncol.*, 44(8):836–849, 2005.

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

[41] E. Hall and A. Giaccia. *Radiobiology For The Radiologist*. Lippincott Williams & Wilkins, Philadelphia, PA, 2006.

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

[43] C. W. Hurkmans, P. Remeijer, J. V. Lebesque, and B. J. Mijnheer. Set-up verification using portal imaging; review of current clinical practice. *Radiother. Oncol.*, 58(2):105–120, 2001.

[44] International Commission on Radiation Units and Measurements. *ICRU Report 50: Prescribing, recording, and reporting photon beam therapy*. International Commission on Radiation Units and Measurements, Bethesda, MD, 1993.

[45] 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.

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

[47] 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.

[48] O. Jäkel and P. Reiss. The influence of metal artefacts on the range of ion beams. *Phys. Med. Biol.*, 52(3):635–644, 2007.

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

[50] P. Keall, J. Siebers, R. Jeraj, and R. Mohan. The effect of dose calculation uncertainty on the evaluation of radiotherapy plans. *Med. Phys.*, 27(3):478–484, 2000.

[51] D. Kuroiwa and G. M. Lee. On robust multiobjective optimization. *Vietnam J. Math.*, 40(2&3):305–317, 2012.

[52] G. J. Kutcher and C. Burman. Calculation of complication probability factors for non-uniform normal tissue irradiation: The effective volume method. *Int. J. Radiat. Oncol. Biol. Phys.*, 16(6):1623–1630, 1989.

[53] K. Langen and D. Jones. Organ motion and its management. *Int. J. Radiat. Oncol. Biol. Phys.*, 50(1):265–278, 2001.

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

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

[56] J. Löf, B. K. Lind, and A. Brahme. Optimal radiation beam profiles considering the stochastic process of patient positioning in fractionated radiation therapy. *Inverse Probl.*, 11(6):1189–1209, 1995.

[57] J. Löf, B. K. Lind, and A. Brahme. An adaptive control algorithm for optimization of intensity modulated radiotherapy considering uncertainties in beam profiles, patient set-up and internal organ motion. *Phys. Med. Biol.*, 43(6):1605–1628, 1998.

[58] J. Löf and H. Rehbinder. Inverse planning optimization with RayOptimizer® in Pinnacle³®. White paper, RaySearch Laboratories, Stockholm, Sweden, 2003.

[59] A. J. 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.

[60] A. J. 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.

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

[62] J. Ludbrook, P. B. Greer, P. Blood, Y. D'yachkova, A. Coldman, W. A. Beckham, J. Runkel, and I. A. Olivotto. Correction of systematic setup errors in prostate radiation therapy: how many images to perform? *Med. Dosim.*, 30(2):76–84, 2005.

[63] J. T. Lyman and A. B. Wolbarst. Optimization of radiation therapy, iii: A method of assessing complication probabilities from dose-volume histograms. *Int. J. Radiat. Oncol. Biol. Phys.*, 13(1):103–109, 1987.

[64] J. Meyer, B. Cavanaugh, J. Purdy, and R. Timmermann. *IMRT, IGRT, SBRT: Advances in the Treatment Planning and Delivery of Radiotherapy*. Frontiers of Radiation Therapy and Oncology Series. Karger Verlag, Basel, Switzerland, 2007.

[65] K. Miettinen. *Nonlinear Multiobjective Optimization*, volume 12 of *International Series in Operations Research and Management Science*. Kluwer Academic Publishers, Dordrecht, The Netherlands, 1999.

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

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

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

[69] A. Niemierko. Reporting and analyzing dose distributions: a concept of equivalent uniform dose. *Med. Phys.*, 24(1):103–110, 1997.

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

[71] J. Nocedal and S. Wright. *Numerical optimization*. Springer Series in Operations Research Series. Springer, New York, NY, 1999.

[72] O. Nohadani, J. Seco, B. C. Martin, and T. Bortfeld. Dosimetry robustness with stochastic optimization. *Phys. Med. Biol.*, 54(11):3421–3432, 2009.

[73] A. Ólafsson and S. J. Wright. Efficient schemes for robust IMRT treatment planning. *Phys. Med. Biol.*, 51(21):5621–5642, 2006.

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

[75] E. Rietzel and G. T. Chen. 4D imaging and treatment planning. In *New Technologies in Radiation Oncology*, pages 81–97. Springer, Berlin-Heidelberg, Germany, 2006.

[76] R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *J. Risk*, 2(3):21–41, 2000.

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

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

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

[80] B. Schaffner and E. Pedroni. The precision of proton range calculations in proton radiotherapy treatment planning: experimental verification of the relation between CT-HU and proton stopping power. *Phys. Med. Biol.*, 43(6):1579–1592, 1998.

[81] U. Schneider, E. Pedroni, and A. Lomax. The calibration of CT Hounsfield units for radiotherapy treatment planning. *Phys. Med. Biol.*, 41(1):111–124, 1996.

[82] M. Schwarz. Treatment planning in proton therapy. *Eur. Phys. J. PLUS*, 126(7):1–10, 2011.

[83] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on stochastic programming: modeling and theory*, volume 9. Society for Industrial and Applied Mathematics, Philadelphia, PA, 2009.

[84] A. Shapiro and A. Kleywegt. Minimax analysis of stochastic problems. *Optim. Meth. Softw.*, 17(3):523–542, 2002.

[85] A. R. Smith. Proton therapy. *Phys. Med. Biol.*, 51(13):R491–R504, 2006.

[86] B. Sobotta, M. Söhn, and M. Alber. Robust optimization based upon statistical theory. *Med. Phys.*, 37(8):4019–4028, 2010.

[87] A. L. Soyster. Convex programming with set-inclusive constraints and applications to inexact linear programming. *Oper. Res.*, 21(5):1154–1157, 1973.

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

[89] S. E. Tenn. *Characterizing geometric accuracy and precision in image guided gated radiotherapy*. PhD thesis, University of California, Los Angeles, 2007.

[90] J. Unkelbach, T. Bortfeld, B. C. 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.

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

[92] J. Unkelbach and U. Oelfke. Inclusion of organ movements in IMRT treatment planning via inverse planning based on probability distributions. *Phys. Med. Biol.*, 49(17):4005–4029, 2004.

[93] J. Unkelbach and U. Oelfke. Incorporating organ movements in inverse planning: assessing dose uncertainties by bayesian inference. *Phys. Med. Biol.*, 50(1):121–139, 2005.

[94] M. van Herk. Errors and margins in radiotherapy. *Semin. Radiat. Oncol.*, 14(1):56–64, 2004.

[95] M. van Herk, P. Remeijer, C. Rasch, and J. V. Lebesque. The probability of correct target dosage: dose-population histograms for deriving treatment margins in radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 47(4):1121–1135, 2000.

[96] S. Webb. Motion effects in (intensity modulated) radiation therapy: a review. *Phys. Med. Biol.*, 51(13):R403–R425, 2006.

[97] S. Webb and A. Nahum. A model for calculating tumour control probability in radio-therapy including the effects of inhomogeneous distributions of dose and clonogenic cell density. *Phys. Med. Biol.*, 38(6):653–666, 1993.

[98] D. Yan and D. Lockman. Organ/patient geometric variation in external beam radio-therapy and its effects. *Med. Phys.*, 28(4):593–602, 2001.

[99] D. Yan, F. Vicini, J. Wong, and A. Martinez. Adaptive radiation therapy. *Phys. Med. Biol.*, 42(1):123–132, 1997.

Mästaren hade en viss benägenhet för nedstämdhet och det var ibland mycket svårt att muntra upp honom. I fråga om skinnet brukade jag säga åt honom, att det ju inte var enbart av ondo om skinnet var fullt av hål och öppningar liksom en trasig rock, eftersom man nu engång inte kunde ta den rocken av sig. En trasig rock är mindre het i solen och den torkar fortare efter regn. Och när två människor ligger tryckta mot varann, då är det inte så oävet med hål och öppningar och jag för min del skulle inte alls vilja ha skinnet annorlunda än vad det är.

- Willy Kyrklund, *Mästaren Ma*