Fully automated treatment planning solutions for robotic radiotherapy

Bas Schipaanboord

Fully automated treatment planning solutions for robotic radiotherapy

Bas Schipaanboord

Fully Automated Treatment Planning Solutions for Robotic Radiotherapy

Volledig geautomatiseerde optimalisatie van bestralingsplannen voor robotische radiotherapie

Thesis

to obtain the degree of Doctor from the
Erasmus University Rotterdam
by command of the
Rector Magnificus

Prof. dr. A.L. Bredenoord

and in accordance with the decision of the Doctorate Board.
The public defence shall be held on

Tuesday 7 June 2022 at 13:00 hrs
by

Bastiaan Wilhelmus Klaas Schipaanboord

born in Zeewolde, the Netherlands.

Erasmus University Rotterdam

Doctoral Committee

Promotor: Prof. dr. B.J.M. Heijmen

Other members: Prof. dr. U. Oelfke
                              Prof. dr. R.A. Nout
                              Assoc. prof. dr. J. Matos Dias

Co-promotor: Dr. ir. S. Breedveld

Contents

Introduction

In 2020, an estimated 19.3 million new cancer cases were diagnosed worldwide, of which more than 132.000 new cases in the Netherlands (World Health Organization: Sung et al. (2021)), making cancer a leading cause of death. Radiation therapy, or radiotherapy, is one of the main treatment modalities for treating cancer. In external beam radiotherapy, ionizing radiation is used to damage cancer cells by directing a beam of radiation through the patient towards the tumor. However, when the radiation passes through the patient's body, not only damage is done to the cancer cells, but also to the healthy tissue surrounding the tumor. The goal of radiotherapy is to inflict sufficient damage to the cancer cells to eradicate the tumor while limiting the damage to the healthy tissue as much as possible to prevent unnecessary radiation-induced side effects.

1.1 Stereotactic radiotherapy using the CyberKnife

Stereotactic Body Radiation Therapy (SBRT) is a radiotherapy approach in which high radiation doses are delivered to the tumors in a limited number of daily fractions (typically 3-5 fractions). High-precision dose delivery, generally based on image-guidance, is required to limit dose delivery to healthy tissue as much as possible.

The CyberKnife® robotic radiotherapy device (figure 1.1) is a system that can deliver SBRT. It features a linear accelerator (linac), which produces the ionizing radiation used for treatment, mounted on a robotic arm. The robotic arm allows for easy delivery of (non-coplanar) beam directions, without the need for manual couch shifts. The radiation beam can be shaped with a multi-leaf collimator (MLC), see figure 1.2. The CyberKnife's InCise™ 2 MLC consists of 26 leaf pairs of high-density material (tungsten). Each leaf can be moved individually in or out of the radiation beam, featuring the beam shaping. Patients are irradiated from

Figure 1.1: The CyberKnife® radiotherapy treatment device at the Erasmus MC.

multiple directions, and per direction multiple beam shapes (called segments), each with its own X-ray intensity, are used to allow for intensity-modulated radiation therapy (IMRT).

High-quality radiotherapy dose distributions can be delivered using the CyberKnife, with excellent clinical outcomes (Fuller et al., 2018, Meier et al., 2018, van der Voort van Zyp et al., 2009). However, the enhanced degrees of freedom of fully non-coplanar robotic radiotherapy increases the complexity of finding optimal patient-specific beam angles and beam segments.

1.2   Radiotherapy treatment planning

The process of patient-specific selection of the settings of the radiotherapy treatment unit that will result in the best possible, or anyway a high-quality, patient dose distribution is called treatment planning. For each patient, a personalized radiotherapy treatment plan is generated before treatment. This treatment plan consists of the delivery parameters that are used to irradiate the patient (beam directions, segment shapes, and segment intensities). Plan generation is generally performed in a commercial software package called Treatment Planning System (TPS). The planning procedure starts with the acquisition of a computed
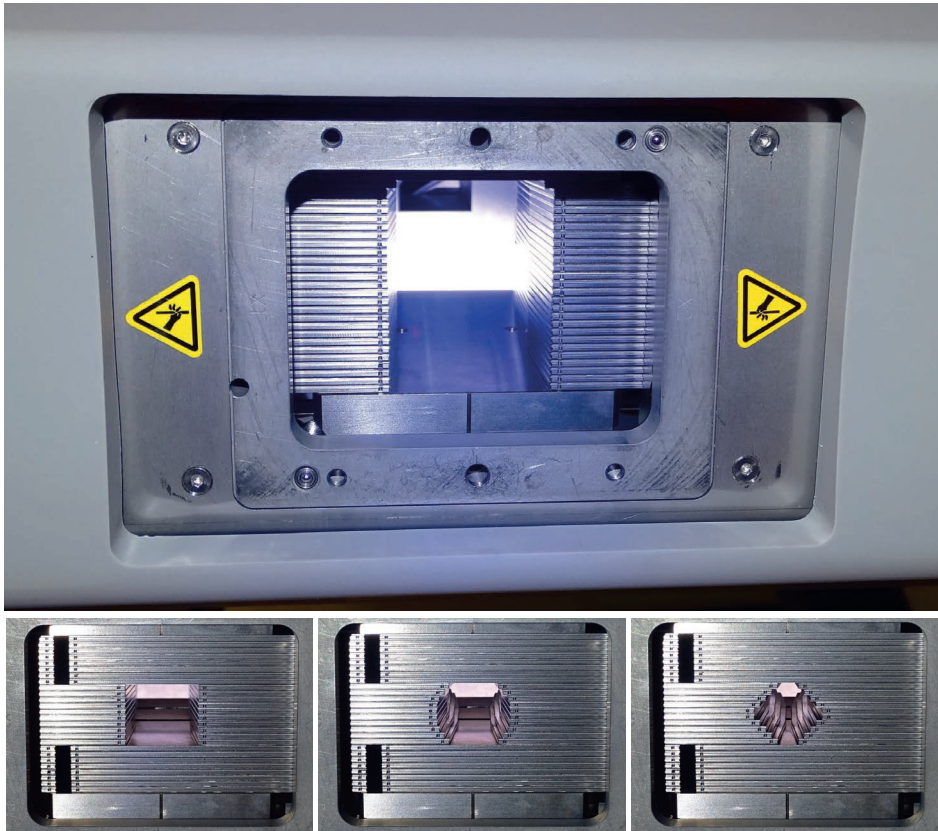
Figure 1.2: The InCise™ 2 multi-leaf collimator (MLC) can be used to shape the radiation beam.

tomography scan (planning CT) to provide a 3-dimensional representation of the patient's anatomy. Next, structures of interest, such as the tumor and nearby organs at risk (OAR), are segmented in the CT-scan. These segmented structures are used to optimize delivery parameters, aiming at high dose in the tumor and low doses in OARs. Based on the clinical planning protocol for the tumor site to be treated, a mathematical optimization problem is formulated (i.e. cost functions for tumor/OAR doses with weights). This optimization problem is subsequently solved to obtain the machine delivery parameters, and the resulting 3-dimensional patient dose distribution projected on the planning CT-scan. The planner evaluates whether the resulting dose distribution satisfies the requirements and objectives as stated in the clinical planning protocol. However, due to the anatomical variations per patient, the trade-offs between planning objectives may vary per patient and are unknown prior to planning. If the planner is not yet satisfied with the resulting

plan quality, the planner can adapt the mathematical optimization, e.g. by changing cost functions weights, and then rerun the optimization. In this way, the planner can balance the delicate trade-offs between tumor coverage, OAR, and plan delivery efficiency. This adaptation step is repeated until the planner is satisfied with the plan quality or if the planner thinks that no significant further improvements in plan quality can be achieved with further adaptation of the optimization problem. The plan quality of clinical treatment plans can vary drastically depending on the skills and ambition of the planner, the complexity of the case, and the time available for planning (Berry et al., 2016, Giglioli et al., 2016, Marino et al., 2015, Nelms et al., 2012).

1.3   Automated radiotherapy treatment planning

Automated treatment planning has been proposed to improve the quality and consistency of radiotherapy treatment plans, and can also substantially reduce the treatment planning workload (e.g. Breedveld et al. 2019b, Fogliata et al. 2014, Giglioli et al. 2020, Hansen et al. 2016, 2017, Heijmen et al. 2018, Hussein et al. 2016, 2018, Marrazzo et al. 2019, Oud et al. 2020, Purdie et al. 2014, Tol et al. 2015, Zarepisheh et al. 2019). In our center, Erasmus-iCycle has been developed for automated multi-criterial optimization (MCO) of beam fluence profiles and beam angles (FMO + BAO) (Breedveld et al., 2012). Erasmus-iCycle automatically generates a single Pareto-optimal radiotherapy treatment plan for each patient. Plan generation is based on a planning protocol specific 'wish-list' that is used for all patients treated according to the protocol. The wish-list contains hard planning constraints and prioritized planning objectives. Many validation studies have demonstrated that with an appropriate wish-list, quality of automatically generated plans supersedes that of plans generated with iterative planning by a planner (above) (Buergy et al., 2017, Della Gala et al., 2017, Heijkoop et al., 2014, Heijmen et al., 2018, Rossi et al., 2018, 2019, Sharfo et al., 2015, 2018, Voet et al., 2013a, 2014). As Erasmus-iCycle only optimizes pencil-beam intensities, the system was originally integrated with the commercial Monaco TPS (Elekta AB, Stockholm, Sweden) to convert generated plans into clinically deliverable plans for C-arm linacs (Voet et al., 2013a, 2014). This coupling to a commercial TPS also prevents regulatory issues with CE (Conformitè Europëenne) or FDA (U.S. Food & Drug Administration) approval.

1.4   Clinically deliverable treatment plans

A widely used optimization approach for radiotherapy treatment planning is to first optimize the beam fluence profiles (Fluence Map Optimization, FMO), and then convert the optimized fluences into deliverable multi-leaf collimator (MLC) segments and their corresponding segment weights. However, the conversion of an FMO plan into a clinically deliverable plan can degrade the plan quality, since delivery restrictions and the dosimetric impact of the MLC cannot be fully accounted for during the FMO phase.

At the start of this thesis, a planning workflow was available that could automatically generate clinically deliverable plans for CyberKnife robotic radiotherapy using the Iris™ variable aperture collimator (Rossi et al., 2018). However, no system was available for clinically deliverable plans for the CyberKnife equipped with the InCise™ 2 MLC. Additionally, shorter treatment times can be achieved with the InCise™ 2 MLC than with the fixed cones or with the Iris™ variable collimator (Tomida et al., 2017).   The ability to generate clinically deliverable treatment plans outside of the commercial TPS allows for an objective comparison between automated plans and manually generated plans in the clinical TPS, both on plan quality as on plan delivery efficiency.

1.5   Aim and outline of this thesis

The aim of this thesis was to develop and validate novel, fully automated treatment planning solutions for robotic radiotherapy, with emphases on clinical deliverability of generated treatment plans and on optimization of beam angles.

In chapter 2, a novel algorithm was proposed to convert Erasmus-iCycle FMO plans into plans that can be delivered at a CyberKnife with the InCise™ 2 MLC, using the same beam angles. The focus was on minimizing dosimetric plan quality loss between final segmented plans and initial FMO plans.   MLC segments were iteratively added to the segmented plan while considering all included beams simultaneously. The performance of the algorithm was investigated for three tumor sites: prostate, head and neck, and liver.

The exploratory analyses in chapter 2 were performed with a simplified dose calculation engine. Therefore, in chapter 3, the segmentation algorithm developed in chapter 2 was further improved by integration of a standalone version of the clinical dose engine. The challenge of this study was on how to cope with the fundamental differences between approximate dose delivered by pencil-beams and the accurate dose delivered by full segments (including MLC scatter and transmission effects). The improved segmentation algorithm was extensively tested

for prostate and liver cancer.

In chapter 4, a fully automated treatment planning workflow for robotic radiotherapy was proposed, and validated for prostate cancer. To this purpose, Erasmus-iCycle was coupled to the segmentation algorithm developed in chapter 3. Validation was performed by comparison of automatically generated plans with plans that were manually generated with the clinical TPS by an experienced medical physicist. Comparisons included dosimetric and delivery efficiency parameters.

Beam angle optimization (BAO) (as also used in the previous chapter) is a non-convex optimization problem and can therefore result in sub-optimal solutions. In chapter 5, a new BAO approach was proposed, allowing detection of sub-optimality of generated plans. For each patient, the basis of plan generation is a so-called total-beam-space (TBS) reference plan generated with Erasmus-iCycle providing a high-quality Pareto-optimal plan. TBS-FMO plans include all beam directions (i.e. no longer BAO as in Erasmus-iCycle), without enforcing any delivery restrictions. To convert TBS-FMO plans into deliverable plans, the segmentation algorithm developed in chapters 2 & 3 with slight modifications was used. For prostate cancer patients, the novel TBS-BAO approach was compared with the approach in chapter 4 and with manual planning. For this study, 594 clinically deliverable treatment plans were automatically generated for 33 patients.

In chapter 6, the proposed treatment planning solutions for automated planning for the robotic CyberKnife are discussed in a wider context, together with their challenges, opportunities, and potential future research.

Automated prioritised 3D dose-based MLC segment generation for step-and-shoot IMRT

B.W.K. Schipaanboord[1]**, S. Breedveld**[1]**, L. Rossi**[1]**, M. Keijzer**[2]**, B. Heijmen**[1]

[1] Department of Radiation Oncology, Erasmus MC Cancer Institute, Rotterdam, the Netherlands
[2] Delft Institute of Applied Mathematics, Delft University of Technology, Delft, the Netherlands

Abstract

Segmentation can degrade a high-quality dose distribution obtained by fluence map optimisation (FMO). A novel algorithm is proposed for generation of MLC segments to deliver an FMO plan with step-and-shoot IMRT while minimising quality loss. All beams are considered simultaneously while generating MLC segments for reproducing the 3-dimensional FMO dose distribution. Segment generation is only steered by the 3D FMO dose distribution, i.e. underlying FMO fluence profiles are not considered. The algorithm features prioritised generation of segments, focusing on accurate reproduction of clinical objectives with the highest priorities. The performance of the segmentation algorithm was evaluated for 20 prostate patients, 15 head-and-neck patients, and 12 liver patients. FMO dose distributions were generated by automated multi-criteria treatment planning (Pareto-optimal plans) and subsequently segmented using the proposed method. Various segmentation strategies were investigated regarding prioritisation of objectives and limitation of the number of segments. Segmented plans were dosimetrically similar to FMO plans and for all patients a clinically acceptable segmented plan could be generated. Substantial differences between FMO and segmented fluence profiles were observed. Avoidance of the usual reconstruction of 2D FMO fluence profiles for segment generation, and instead simultaneously generating segments for all beams to directly reproduce the 3D FMO dose distribution is a likely explanation for the obtained results. For the strategies of limiting the number of segments large reductions in number of segments were observed with minimal impact on plan quality.

2.1 Introduction

In intensity modulated radiation therapy (IMRT) treatment planning, the optimisation problem may be split into a fluence map optimisation (FMO) phase and a segmentation phase to convert the optimised fluences into multi-leaf collimator (MLC) segments. An advantage of this approach is that the FMO problem can be modelled as a convex multi-criterial optimisation problem (Breedveld et al., 2019b) with a guaranteed globally optimal solution in optimisation. As treatment machine limitations (e.g. limitations of the MLC) are not fully accounted for in FMO, the deliverable plan resulting from the segmentation phase may have a quality loss compared to the FMO plan. Additionally, decisions on multi-criterial trade-offs, as made during FMO, are generally not explicitly taken into account in the segmentation phase (Salari and Unkelbach, 2013).

One approach for segmentation in static step-and-shoot IMRT is to separately segment the beam fluence profiles for the involved beams into sets of deliverable segments by stratifying fluences into discrete intensity levels and subsequently generate feasible segments for each beam that match the optimised fluence profile (Süss et al., 2007, Xia and Verhey, 1998). In general, the more segments are included the better the fluence profile can be replicated. To restrict the treatment delivery time, pre-defined trade-offs between plan quality and treatment time can be used to restrict the number of intensity levels and number of segments (Craft et al., 2007). To the best of our knowledge, published MLC segmentation approaches for static step-and-shoot IMRT plans are all based on independent segmentation of the 2-dimensional fluence profiles of all beams (Gören and Taşkin, 2015, Long et al., 2016, Luan et al., 2006, Sun and Xia, 2004, Süss et al., 2007). Sequencing the fluences for each beam separately excludes mutual dosimetric compensation of imperfect segmentations of the 2-dimensional beam fluence profiles to optimally reproduce the initial 3-dimensional FMO dose distribution.

Extensive research has been done to improve MLC segmentation with non-discretised intensity levels and leaf positions (Long et al., 2016), to investigate segmentation efficiency under various MLC constraints (Gören and Taşkin, 2015), to explore regularization in the dose domain before segmentation (Nguyen et al., 2015) and to minimise beam-on-time (Crooks et al. 2002, Ahuja and Hamacher 2005, Boland et al. 2004). However, none of the published methods explicitly account for differences in objective priorities during segmentation. Consequently, discrepancies between FMO fluence and sequenced fluence may potentially lead to dose deviations in the PTV and OARs with uncontrolled balances, i.e. without explicitly considering the clinical priorities.

In contrast to FMO followed by segmentation, Direct Aperture Optimisation (DAO) has been proposed to directly generate MLC segments (Men et al., 2007, Romeijn et al., 2005, Shepard et al., 2002). DAO operates under the "What you see is what you get" principle, meaning that at every stage of the optimisation process the treatment plan is directly feasible for delivery and no segmentation phase (with possible loss in plan quality) is needed. However, including the non-convex modelling of the (physical) constraints of the collimator and treatment device leads to a non-convex optimisation problem. The column generation (CG) approach has been proposed as heuristic in the field of DAO to solve the optimisation problem in radiotherapy (Carlsson, 2008, Cassioli and Unkelbach, 2013, Men et al., 2007). Research on CG approaches for DAO includes investigations on convergence (Carlsson and Forsgren, 2014), generation of segments under various MLC constraints (Men et al., 2007) and inclusion of pre-defined multi-criterial trade-offs (Salari and Unkelbach, 2013).

In our centre we have given preference to plan optimisation using FMO followed by segmentation, because of the guaranteed Pareto and global optimality of the FMO plans that we generate with *Erasmus-iCycle*, an algorithm for automated a priori Multi-Criterial treatment plan Optimisation (MCO) (Breedveld et al., 2012). For each patient, a single Pareto-optimal FMO plan is generated with clinically favourable trade-offs, considering all treatment objectives with explicitly assigned priorities. This differs from a posteriori MCO (e.g. Bokrantz and Miettinen 2015, Craft and Richter 2013) in which, for each patient, a set of Pareto-optimal plans is generated with automated planning, while selection of a clinically favourable plan is performed by a user.

In this study, a prioritised dose-based MLC segment generation method is proposed which minimises 3-dimensional plan quality loss compared to the FMO plan by placing extra consideration on high priority clinical objectives. To reconstruct the 3-dimensional FMO dose, a CG approach was implemented that simultaneously optimises the beam segments for all treatment beams, rather than replicating the 2-dimensional fluences separately as is done in other published MLC segmentation methods. Segment generation was only steered by the 3D FMO dose distribution, i.e. underlying FMO fluence profiles were not considered. In this sense, the term segmentation as applied in this paper has a slightly different meaning from that in the literature. The CG approach was chosen because of its intuitive mechanism of generating segments and proven effectiveness in IMRT treatment planning (Carlsson, 2008, Romeijn et al., 2005, Salari and Unkelbach, 2013).

The overall goal was to achieve high quality radiotherapy treatment plans by combining global optimality of the FMO plan with a fast prioritised "DAO-like" segmentation. The segmentation was tailored to individual patients by using the

prioritised approach in combination with personalised clinical objectives obtained from the FMO plan. The proposed technique was developed in the context of CyberKnife robotic radiotherapy, but can be applied for segmentation of any static step-and-shoot IMRT or stereotactic body radiation therapy (SBRT) plan. Possibilities for keeping calculation and delivery times low were included in the investigations. The segmentation performance was evaluated for prostate, head-and-neck and liver tumours.

2.2 Materials & Methods

This section starts with briefly describing the applied FMO (section 2.2.1). Next, CG is introduced in section 2.2.2, while the proposed segmentation with CG is described in section 2.2.3. Prioritised steering on personalised objectives is described in section 2.2.4, which includes approaches to minimise the number of segments. Finally, plan evaluation criteria and details on our computational study are presented in sections 2.2.5 and 2.2.6, respectively.

2.2.1 Patients, FMO treatment plans and dose calculation model

An overview of the clinical cases and FMO plans used to evaluate the performance of the MLC segmentation is presented in table 2.1. All FMO plans were generated with fully automated multi-criterial optimisation as implemented in Erasmus-iCycle (Breedveld et al., 2012, 2017). A pencil-beam approach was used to describe dose delivered to the patient, i.e. $d = Ax$ with d the vector containing the patient's voxel doses, A the dose deposition matrix, and x the pencil beam weights (see also Breedveld et al., 2006, 2017, Nguyen et al., 2015, Zhu et al., 2012). For CyberKnife plans the beamlet and segmentation resolutions were defined at 800 mm from the source, while this was 1000 mm for conventional linac plans. For all plans a beam energy of 6 MV was used. For the prostate treatments, FMO was performed using pencil beams with a 5 x 5 mm^2 beamlet resolution, while a 5 x 10 mm^2 beamlet resolution was used for liver and head-and-neck cancer. These FMO resolutions have shown to provide a good balance between plan quality and computational efficiency. Since MLC segmentation of a FMO dose distribution can result in degradation of plan quality, segmentation on a higher resolution than the FMO resolution can compensate for potential degradation in plan quality. We modelled the CyberKnife InCise2 MLC and performed all segmentations (including the conventional linac plans) for this MLC, which has 2 banks of 26 leaves with a leaf thickness of 3.85 mm defined at 800 mm SAD. Since our dose engine is limited to integer values of resolution only, the segmentation was performed on a resolution of a 1 x 4 mm^2.

Table 2.1: Overview of the clinical cases and FMO plans.

Treatment site	Cases	Treatment unit	Beams		Prescribed dose	Fractions
Prostate	20	CyberKnife	25	non-coplanar	38 Gy	4
Head-and-neck	15	Conventional linac	9	coplanar	46 Gy	23
Liver	12	CyberKnife	25	non-coplanar	60 Gy	3

2.2.2 Column generation

CG is generally used to solve large-scale problems. The large-scale optimisation problem is denoted as the *Master Problem* (MP). Instead of solving the MP directly, the MP is solved by iteratively solving a restricted version of the problem denoted as the *Restricted Master Problem* (RMP). The RMP only includes a subset of the original decision-variables (i.e. the beamlet intensities x). During each iteration of the CG method, the RMP is solved and the solution is projected onto the MP. The projection on the MP can be used to identify the next promising subset of decision variables, which will be added to the RMP in the subsequent CG iteration. This identification step is called the *Pricing Problem* (PP). If no new decision-variables can be identified, the MP is solved to optimality. For a detailed description of CG in RT, see Carlsson and Forsgren, 2014, Men et al., 2007, Romeijn et al., 2005.

2.2.3 Problem definition and segment generation

The CG workflow, as introduced by Romeijn et al. (2005) for DAO in radiotherapy, was in this study used to segment FMO dose distributions by iteratively identifying promising MLC segments, see figure 2.1. The CG workflow is denoted in grey, while the proposed prioritised workflow consists of the grey workflow, followed by red.

The MP for plan segmentation is formulated in section 2.2.3.1. In each iteration the most promising segment is identified by solving the PP (section 2.2.3.2) and then added to the RMP (section 2.2.3.3). The RMP only contains the segments identified so far and it is solved to optimality to determine the intensities of the segments. At the end of each iteration, segments for which the intensity falls below the minimum required Monitor Units (MU) are removed (section 2.2.3.4). Then, in the next iteration the PP is again solved to identify the next promising segment.

For the proposed prioritised MLC segmentation, the CG workflow is incorporated into an adaptive framework (section 2.2.4). If a segmented solution converges to a dose distribution that does not comply with one or more of the DV criteria, segmentation is re-started with extra emphasis on high priority objectives.
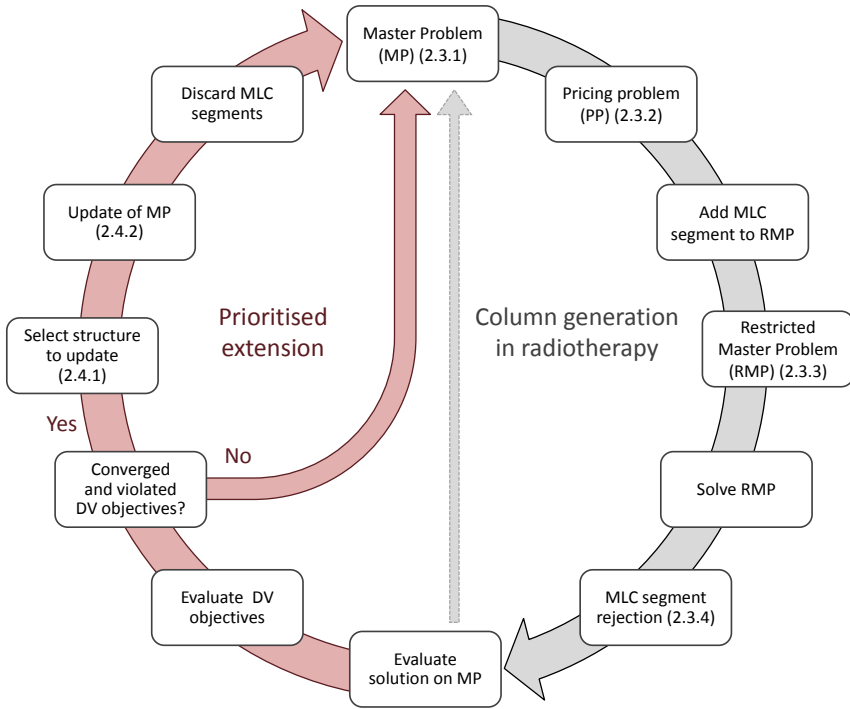
Figure 2.1: Schematic representation of the column generation (CG) approach for prioritised MLC segmentation as a combination of the grey and red workflows. The numbers between brackets refer to corresponding paragraphs.

2.2.3.1 Master Problem

With x denoting the fluence vector, our MP is formulated by:

$$\text{minimise} \quad f(x) + \omega p(x) \qquad \text{(MP)} \quad (2.1)$$
$$\text{subject to} \quad x \geq 0$$

$x \geq 0$ ensures non-negative fluences and $f(x)$ and $p(x)$ are given by:

$$f(x) = \sum_{v \in V} \|\eta_v^{1/2}(Ax - d_v^F)\|_2^2 \qquad (2.2)$$

$$p(x) = \frac{1}{N_{PTV}} \sum_{j=1}^{N_{PTV}} e^{-\alpha(A_j x - D^P)} \qquad (2.3)$$

The quadratic term $f(x)$ is the main driving force of the MP, similar to objective functions commonly used in treatment planning (Breedveld et al., 2006, Carlsson and Forsgren, 2014). The vector η_v contains voxel-dependent weight factors for each volume $v \in V$, which play a key role in the MP adaptation part of the algorithm

described in section 2.2.4. d_v^F defines the FMO dose distribution. $f(x)$ can be rewritten in the canonical form to increase computational efficiency (Breedveld et al., 2006).

The term $p(x)$ with weight ω in equation (2.1) puts extra emphasis on attaining an adequate PTV coverage by penalising under-dosage of the PTV, for which the Logarithmic Tumour Control Probability (LTCP) is used (equation (2.3), as proposed by Alber and Reemtsen (2007)). D^P denotes the prescribed dose for the PTV, N_{PTV} the number of sampled PTV voxels, and $A_j x$ the dose delivered to PTV voxel j. α is a constant related to cell survival (Alber and Reemtsen, 2007), which was tuned to achieve adequate coverage. An α equal to 0.90, 0.82 and 0.40 was used for prostate, head-and-neck and liver respectively.

2.2.3.2 Pricing problem and feasible segment generation

For the identification of the most promising feasible segment, an approach similar to Romeijn et al. 2005 has been implemented, in which the gradient from the MP is projected onto the beamlet grid for each of the beams. Beamlets with a negative gradient are favourable for inclusion into the next segment, as these indicate the most effective descent direction for the MP. These individual beamlets are grouped together into feasible MLC segments by constructing a layered graph (per beam direction) for possible combinations of adjacent negative beamlets. Mechanical restrictions of the MLC device are taken into account during construction of the graph. The following segment restrictions, similar to the restrictions of the CyberKnife InCise2 MLC, were enforced: the MLC segment contains only one contiguous opening, a minimum number of 2 leaf pairs open per segment (7.7 mm in total), a minimum opening size of 7.6 mm in the direction of the leaves, and interdigitation is allowed. Given that our dose engine operates on integer values of resolution, a field size restriction of two leaf pairs open per segment (8 mm in total) and a minimum opening size of 8 mm in the direction of the leaves were used instead. The graph is subsequently solved using a shortest path algorithm. The pricing problem was implemented in C++ using Boost Graph Libraries (v1.58) and solved using a Bellman-Ford shortest path algorithm. A post-processing step was implemented to guarantee feasible segments. Feasible segments are generated for all beam directions in parallel, but only the most promising segment (the one with the largest sum of negative contributions) is selected and added to the RMP.

2.2.3.3 Restricted Master Problem

The restricted version of the Master Problem is formulated in equation (2.4), where the quadratic (2.2) and LTCP (2.3) terms are now given by equations (2.5) and (2.6).

Variable x_r denotes segment intensities. In each iteration, a promising segment is identified and added to the RMP; a new column is added to the optimisation problem. The RMP is subsequently solved to optimality to determine the intensities x_r for included MLC segments. For this we use our in-house developed solver (Breedveld et al., 2017), specifically designed and tuned for solving radiotherapy optimisation problems, but any non-linear solver could in principle be used.

$$\begin{aligned}
\text{minimise} \quad & f_r(x_r) + \omega p_r(x_r) \qquad\qquad \text{(RMP)} \quad (2.4)\\
\text{subject to} \quad & x_r \geq 0
\end{aligned}$$

Where the quadratic and LTCP terms are now given by:

$$f_r(x_r) = \sum_{v \in V} \| \eta_v^{1/2}(A_r x_r - d_v^F) \|_2^2 \qquad\qquad (2.5)$$

$$p_r(x_r) = \frac{1}{N_{PTV}} \sum_{j=1}^{N_{PTV}} e^{-\alpha(A_{rj} x_r - D^P)} \qquad\qquad (2.6)$$

2.2.3.4   Segment rejection

For treatment delivery a minimum MU/segment is imposed because dose delivery for MU below this threshold may be inaccurate. Also, segments added to the RMP in an early stage of the segmentation process can decline in relevance due to addition of newer segments. We have chosen not to enforce the minimum MU/segment constraint while solving the RMP, in order to maintain the ability to identify and remove segments for which the contribution to the solution diminishes. When the intensity of a segment drops below the minimum the segment is removed from the RMP. Additionally, after segment removal, the intensities of the remaining segments are re-optimised and it is again verified whether they fulfil the minimum MU constraint. Discarding redundant columns (segments) from the RMP reduces the size of the problem which improves the computational efficiency of solving the RMP. For the hypo-fractionated SBRT plans (prostate and liver) in this study a minimum MU/segment per fraction of 5 was used and for the conventionally fractionated plans the minimum MU/segment per fraction was 3.

2.2.4   Prioritised MLC segmentation

The performance regarding the posed (personalised) objectives is tracked during segmentation (section 2.2.4.1). If a segmented solution converges towards a

solution that does not comply with one or more of the objectives, the MP is updated in an attempt to better reflect the requested trade-offs (section 2.2.4.2), taking into account the clinical priorities.

2.2.4.1   Prioritised objectives for segmentation

For the three investigated tumour sites, the tumour and OAR objectives with assigned priorities as used for the segmentation of FMO plans are presented in table 2.2a.   They are in line with the clinical treatment planning protocols at Erasmus MC. For each patient, the personalised goal values for the OAR objectives are obtained from the Pareto-optimal FMO plan generated with Erasmus-iCycle to obtain trade-offs in OAR sparing during the segmentation like they were made during the automated multi-criteria FMO. During segmentation there is no need to obtain a PTV coverage higher than requested in the clinical protocol, even if it is obtained during FMO. Therefore, to generate maximum space for OAR sparing during segmentation, the tumour objectives are enforced as provided in the treatment protocol.

When during a segmentation all objective functions have converged and one or more of the objective functions have not reached the goal value, an update of the MP is performed, i.e. the MP is modified to put more emphasis on the objective function with an unattained goal with the highest priority (table 2.2a) in a subsequent segmentation run. When multiple objectives with equal priority are not fulfilled, the structure with the largest deviation from the desired objective value is selected.

2.2.4.2   Updates of the MP during segmentation

Two mechanisms are used simultaneously for updates of the MP. The first approach is adjustment of the individual voxel weights η_v within the quadratic objective function (equation 2.2). Increasing the weight will magnify the difference between the attained dose in the segmented solution and the reference dose for that voxel, as obtained in the FMO plan. This will put more emphasis on that particular voxel for attaining its reference dose. The second option is to adjust the reference dose d_v^F within the quadratic objective function. By adjusting the reference dose for a voxel the difference in dose will be increased, thereby increasing the contribution to the MP objective function, but this option also favours deviations from the original FMO plan.

For each update, only the voxels in the selected structure that do not comply with the criteria contribute to adjustment of the MP cost function. A maximum of three MP updates was enforced to limit calculation time and to remain close to the FMO solution. The values of ω, η_v and d_v^F with the updates are presented in section 2.3.1.

Table 2.2: Personalised objectives for prioritised segmentation of a prostate, head-and-neck or liver FMO plan with assigned priorities (Pr.); for each patient, the goal values for the OAR objectives were the plan parameters in the corresponding FMO dose distributions, while the tumour objectives were always enforced as stated in the treatment planning protocol (a) and hard constraints to evaluate clinical acceptability of plans (b).

(a)

Prostate

Tumor objective		Pr.
PTV	$V_{D>38Gy} > 95\%$	1

Personalized objectives		Pr.
Rectum	D_{1cc}	2
Bladder	D_{1cc}	3
Urethra	$D_{5\%}$	4
	$D_{10\%}$	4
	$D_{50\%}$	4

Head-and-neck

Tumor objectives		Pr.
PTV	$V_{D>43.7Gy} > 98\%$	1
	$D_{1cc} < 49.2\,Gy$	2

Personalized objectives		Pr.
Spinal cord	D_{1cc}	3
Brainstem	D_{1cc}	3
Parotid L/R	D_{mean}	4
SMG L/R	D_{mean}	4
Larynx	D_{mean}	5
Cochlea L/R	D_{mean}	6
Esophagus	D_{mean}	6
Oral cavity	D_{mean}	6

Liver

Tumor objectives		Pr.
PTV	$V_{D>60Gy} > 95\%$	1
	$D_{1cc} < 75\,Gy$	2

Personalized objectives		Pr.
Liver - GTV	$V_{D>15Gy}$	3
Duodenum	D_{1cc}	4
Small bowel	D_{1cc}	4
Stomach	D_{5cc}	4
Spinal cord	D_{1cc}	5
Esophagus	D_{1cc}	5
Kidney L/R	$V_{D>15Gy}$	6

(b)

Prostate

Clinical constraints	
PTV	$V_{D>38Gy} > 95\%$
Rectum	$D_{1cc} < 32.3\,Gy$
Bladder	$D_{1cc} < 38\,Gy$

Head-and-neck

Clinical constraints	
PTV	$V_{D>43.7Gy} > 98\%$
	$D_{1cc} < 49.2\,Gy$
Spinal cord	$D_{1cc} < 50\,Gy$
Brainstem	$D_{1cc} < 60\,Gy$

Liver

Clinical constraints	
PTV	$V_{D>60Gy} > 95\%$
	$D_{1cc} < 75\,Gy$
Liver - GTV	$V_{D<15Gy} \geq 700\,cc$
Duodenum	$D_{1cc} < 30\,Gy$
Small bowel	$D_{1cc} < 30\,Gy$
Stomach	$D_{5cc} < 22.5\,Gy$
Spinal cord	$D_{1cc} < 18\,Gy$
Esophagus	$D_{1cc} < 27\,Gy$
Kidney L/R	$V_{D>15Gy} \leq 33\%$

If the update of the objective function yields a plan of inferior plan quality than before the update the segmentation falls back on the previous plan.

Segmentations are eventually terminated when the MP objective function has converged with all clinical objectives (table 2.2a) met, or if the maximum number of MP resets (three) has been reached. The convergence criterion was defined as the objective value being within 10% of its current value over the last 10 iterations for segment additions. This criterion was relaxed to 12.5% when all clinical objectives were met. As a result, fewer segments were included when an adequate plan has already been achieved.

2.2.4.3   Segment reduction

An important contributor to treatment delivery time is the number of segments. To investigate possibilities for active steering on the number of segments, we have implemented and evaluated six segmentation approaches: three prioritised segmentation (PS) methods and three non-prioritised (noPS) methods:

1. *PS_full*: Full prioritised segmentation as described in sections 2.2.4.1 and 2.2.4.2.
2. *PS_remove*: PS_full, followed by stepwise removal of segments. For every removal step, the segments are ranked based on their relative contribution to the PTV mean dose. Subsequently, the segment with the lowest contribution is removed and the intensities x_r of the remaining segments (equation (2.4)) are re-optimised. This process is continued until a tumour objective is violated.
3. *PS_terminate*: Start prioritised segmentation like in PS_full, but terminate as soon as all tumour objectives are met after the last update of the MP.
4. *noPS_full*: Start segmentation like in PS_full, but do not update the MP, i.e. the prioritised list of personalised objectives is not used during segmentation.
5. *noPS_remove*: noPS_full, followed by stepwise removal of segments (see 2 for details).
6. *noPS_terminate*: Start segmentation like in noPS_full, but terminate as soon as all tumour objectives are met.

2.2.5   Plan evaluation criteria

Plans segmented with the various approaches were mutually compared and compared with FMO. The analyses focused on clinical acceptability, dosimetric quality, number of segments, MU and segmentation time. Criteria for clinical acceptability are summarised in table 2.2b. Dose-volume parameters and the *Conformation Number (CN)* as proposed by van 't Riet et al. (1997) were used to quantify dosimetric quality. Additionally, visual inspections of the dose distributions were performed.

2.2.6   Computation times

Segmentations were performed on a dual CPU system, consisting of 2 octocore Intel Xeon E5-2690 CPUs, running at 2.90 GHz and with 128 GB of memory. For the various segmentation approaches, calculation times were recorded.

2.3 Results

Segmentation parameters found to be suitable for prioritised segmentation are presented in section 2.3.1. Prior to presenting the overall performance results for the segmentation approaches in section 2.3.3, one head-and-neck case is discussed in detail in section 2.3.2.

2.3.1 Prioritised segmentation parameters

A weight ω of 10^3 for the LTCP term of the objective function (2.1) was found to work adequately for all tumour sites and was kept fixed throughout the investigations. As mentioned in section 2.2.3.1, the contribution of the LTCP term diminished when an adequate PTV coverage was attained. The value of 10^3 provided an appropriate trade-off between steering on sufficient PTV coverage (when necessary) and reconstruction of the FMO dose distribution provided. The voxel weights η_v of the quadratic part of the objective function (2.2) were all set to 1 at the start of the segmentation. For updates of the MP cost function, the weights for selected voxels were increased from 1 to 5, to 10 and to 15 for subsequent updates. For OAR objectives, the voxel reference doses d_v^F, initially obtained from the FMO dose distribution were decreased by 0.33 Gy at the same time.

2.3.2 Segmentation performance - example patient

Figure 2.2 shows for an example patient axial and sagittal slices through the PTV for the FMO plan (a) and the segmented PS_full plan (b). FMO and segmented dose distributions were similar, though small deviations were noticeable. As required by the clinical protocol (table 2.2b), in both plans more than 98% of the PTV was covered by at least 95% of the prescribed dose (yellow isodose line in figure 2.2, 43.7 Gy=95%). Isodose lines of the segmented plan were somewhat smoother, which can be explained by the difference in fluence modulation. In FMO, the fluence is modulated per beamlet, although some form of regularisation (smoothing) is applied. In segmented plans, collections of connecting beamlets are irradiated with equal intensity, which generally results in smoother fluence profiles. In principle, the higher modulation in FMO could have resulted in a more conformal plan. However, in this case a minor increase in CN was observed in the segmented plan; CN=0.77 for FMO and 0.79 for the segmented plan. FMO and segmented fluence profiles for all beams are shown in figure 2.3. Even though FMO and segmented dose distributions were similar, substantial differences in fluence profiles per beam were observed.
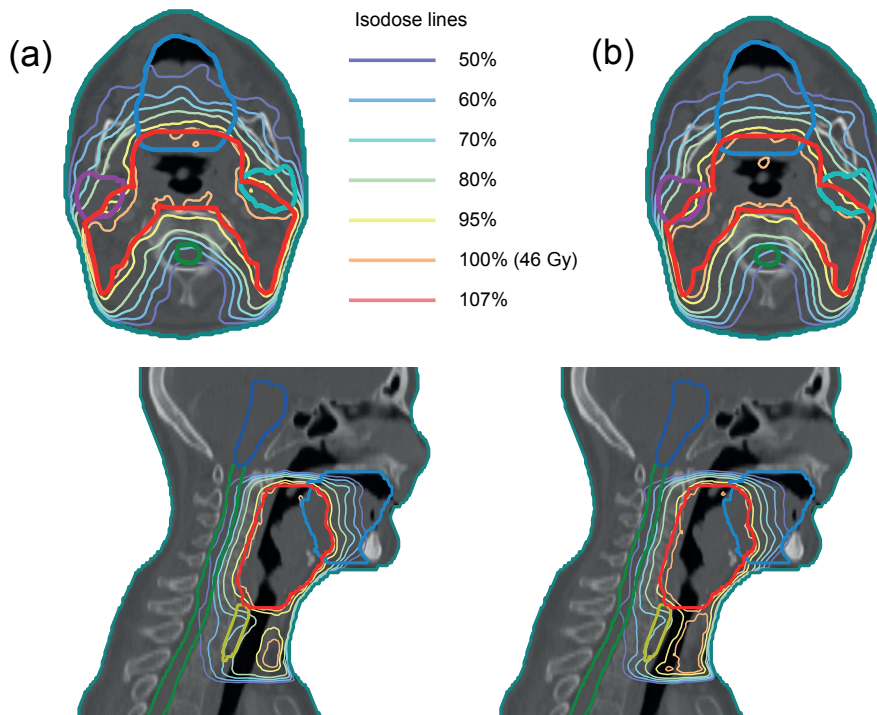
Figure 2.2: Similar FMO (a) and PS_full segmented (b) dose distributions for the example HN patient discussed in section 2.3.2. Depicted structures: PTV (red), spinal cord (green), brainstem (blue), parotid L/R (cyan/purple), oesophagus (yellow) and oral cavity (light blue).

DVHs for the FMO plan and all segmented plans are presented in figure 2.4. For all three PS methods the goal values for both PTV objectives were attained (table 2.2), with PTV coverages of 98.9%, 98.8% and 98.8% and PTV D_{1cc} of 49.0 Gy, 49.1 Gy and 49.2 Gy for PS_full, PS_remove and PS_terminate respectively. In comparison, in none of the noPS plans the PTV D_{1cc} goal was achieved (PTV D_{1cc} of 50.1 Gy). Since the PTV D_{1cc} was already violated for noPS_full, no segments could be removed without violating one of the tumour objectives in the reduction step and therefore all noPS plans were equal. Figure 2.5 shows for the example HN case the PTV objective functions ($V_{D>43.7Gy}$ and $V_{D>49.2Gy}$) as a function of iteration number during the segmentation process. In order to meet for both functions the goal values three MP updates for the PTV were necessary, to place more emphasis on crucial voxels that contributed to PTV overdose.
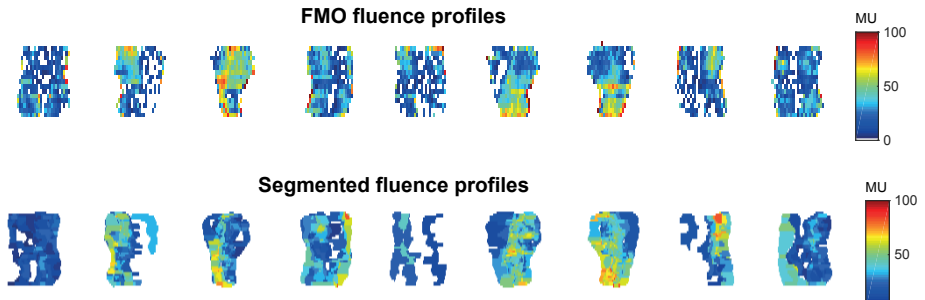
Figure 2.3: Fluence profiles to realise the FMO and segmented plans for the example patient discussed in section 2.3.2, plotted per beam. Even though the FMO and segmented dose distributions are similar (figure 2.2), substantial differences between FMO and segmented fluence profiles per beam were observed.
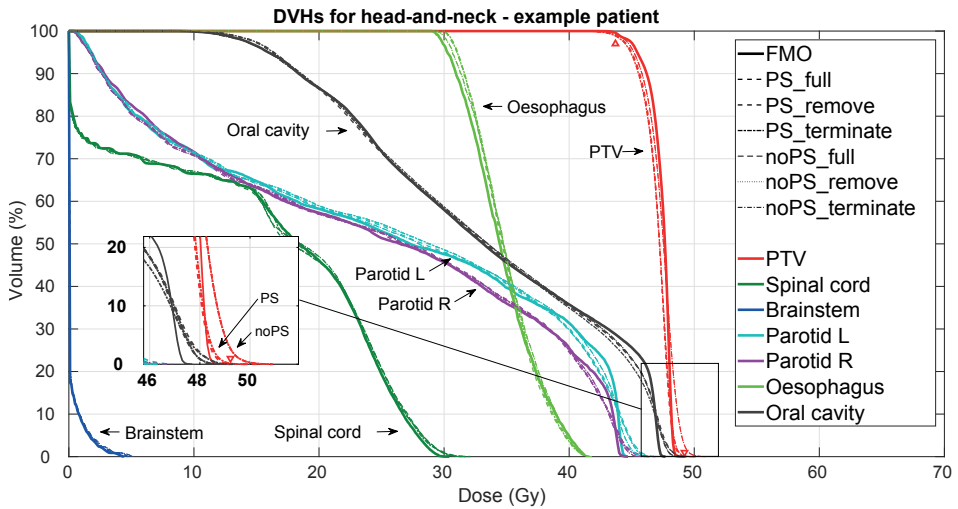
Figure 2.4: DVHs for FMO and segmented plans for the example head-and-neck case presented in section 2.3.2. Tumour objectives are denoted with red triangles. Inset: close-up of the DVHs around maximum dose. All three PS plans met both tumour objectives compared to none of the noPS plans (none of them fulfilled the PTV D_{1cc} objective).
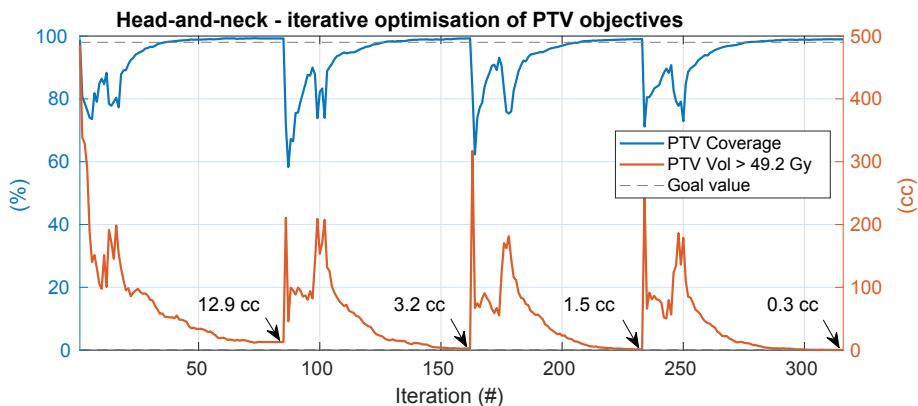
Figure 2.5: For the example head-and-neck patient discussed in section 2.3.2, tumour objective functions' values during the PS_full prioritised segmentation plotted against iteration number. In order to meet the goal values (table 2.2a), three updates of the objective function were necessary. Updates were performed at iteration 85 (PTV $V_{D>49.2Gy} = 12.9\ cc$), at iteration 162 (PTV $V_{D>49.2Gy} = 3.2\ cc$) and at iteration 233 (PTV $V_{D>49.2Gy} = 1.5\ cc$). At the end of PS_full, PTV $V_{D>49.2Gy} = 0.3\ cc$, which is within the goal value ($< 1\ cc$, table 2.2a).

2.3.3   Segmentation performance - all patients

Figures 2.6, 2.7 and 2.8 show population mean DVHs for the three patient groups. DVHs per individual patient can be found in the supplementary materials. For prostate and liver the mean DVHs for the six segmentation approaches were very close to those of the FMO plans. For the OARs this also held for head-and-neck cancer, but for the three noPS segmentation approaches the mean $V_{D>49.2Gy}$ for the PTV exceeded the clinical dose constraint (table 2.2b), see inset of figure 2.7. Figure 2.9 shows the number of clinically acceptable plans, which for head-and-neck indeed shows that for each of the three noPS approaches only 1 out of 15 plans was clinically acceptable.

Figures 2.10, 2.11 and 2.12 show details on the performance of the six segmentation approaches for the three investigated tumour sites. The subfigure "PTV D1cc" of figure 2.11 shows that acceptability issues with noPS plans for head-and-neck cancer were indeed indeed related to too large PTV volumes receiving high dose. Among the three PS approaches, differences between achieved dosimetric parameters for head-and-neck were clinically irrelevant (figure 2.11). Overall the best segmentation approach for this group seems PS_terminate as it has the lowest # segments (median: 75), the lowest # iterations (median: 334) and the lowest MU (median per fraction: 979 MU). Also for prostate and liver cancer, PS_terminate is often a good choice, mainly because of the relatively low # segments, low # iterations and low total MU.
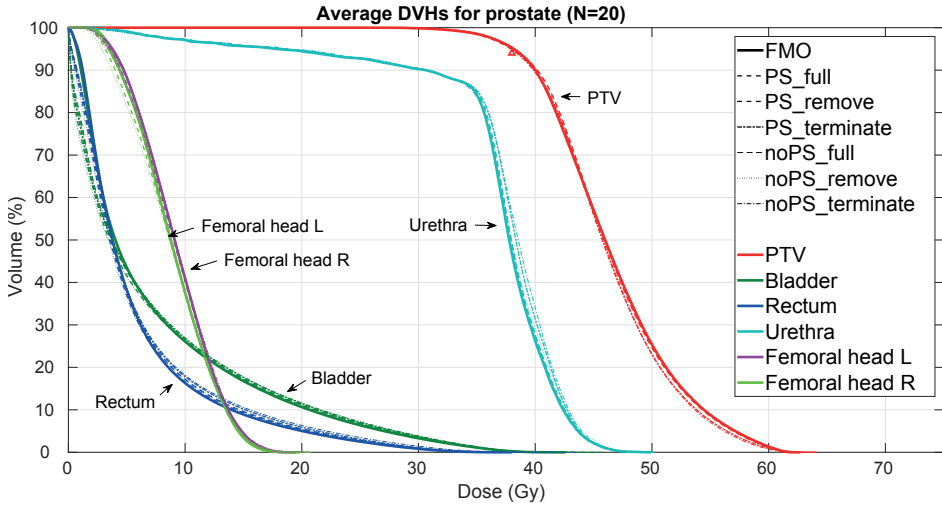
Figure 2.6: Average DVHs for the prostate cases (N=20) for the FMO plans (thick solid lines) and the segmented plans (dashed). The PTV coverage objective is denoted with a red triangle.

Another interesting observation are the differences in deviation observed per segmentation method over various objectives with different priorities (table 2.2a). For example, for prostate cancer the smallest deviations were observed for the most important OAR (i.e. rectum). With decreasing importance (increasing priority number) the deviations in dose from the FMO plan increased (compare objective subplots in figure 2.10). No similar trend was observed for the noPS approaches, which indicates that the proposed extension of the segmentation technique is able to reduce dose deviations in a prioritised manner.

2.4   Discussion

MLC segmentation approaches were investigated that aim at accurate reconstruction of optimised FMO distributions, while complying with the mechanical limitations of the treatment device. To reconstruct 3-dimensional FMO dose distributions, a column generation approach was implemented that simultaneously optimised the beam segments for all treatment beams, while ignoring the underlying FMO fluence profiles.

For OARs the proposed segmentation method uses a convex quadratic objective function to minimise the voxel-wise differences between the intended FMO dose and the segmented dose. During prioritised segmentation, this objective function may iteratively be adapted to maximally reproduce OAR dose parameters, while

Figure 2.7: Average DVHs for the head-and-neck cases (N=15) for the FMO plans (thick solid lines) and the segmented plans (dashed). Inset: close-up of the DVHs around maximum dose. PTV objectives are denoted with red triangles.

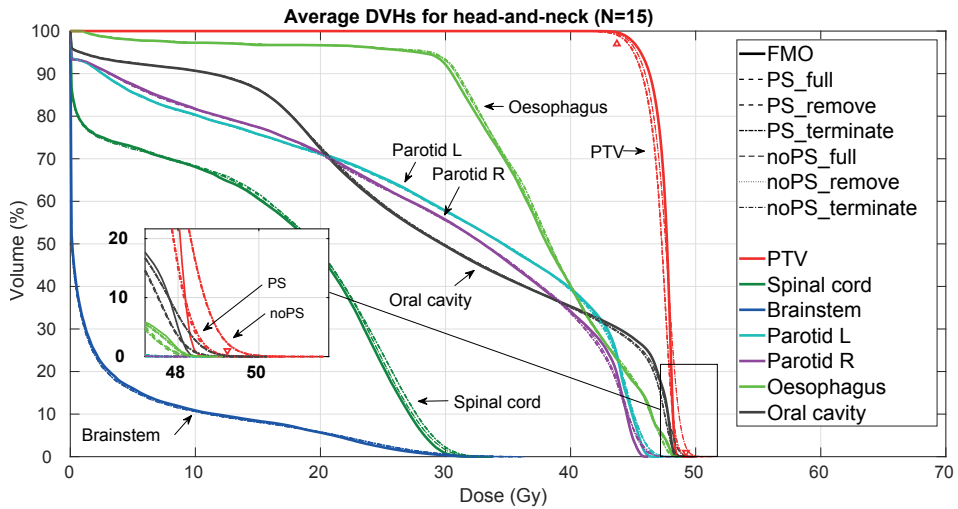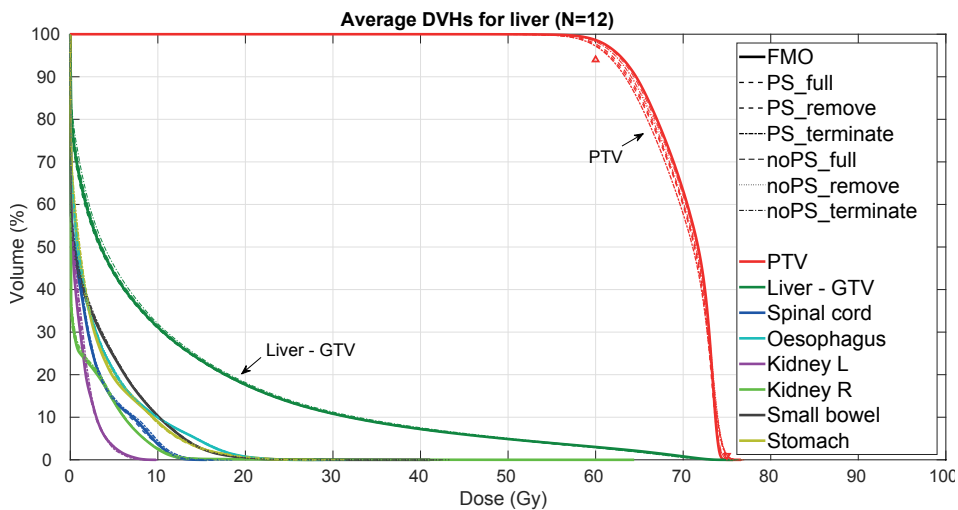**Figure 2.8:** Average DVHs for the liver cases (N=12) for the FMO plans (thick solid lines) and the segmented plans (dashed). PTV objectives are denoted with red triangles.
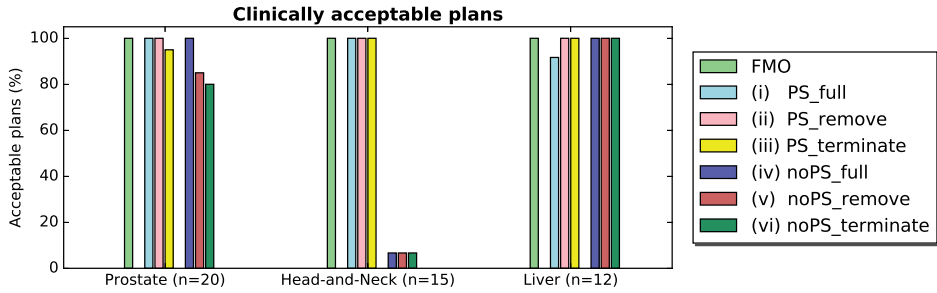
Figure 2.9: Percentage of clinically acceptable plans. A plan was considered acceptable if for all constraints the obtained values were within 0.25 Gy or 0.25% of imposed values (table 2.2b). Plans segmented using the prioritised methods (PS) outperformed the non-prioritised methods (noPS). Especially for head-and-neck where more emphasis was needed on crucial voxels, in 14 out of the 15 cases, in order to meet the PTV D_{1cc} constraint.

considering the clinical priorities. In initial attempts, we tried to also use for the PTV only a quadratic cost function, similar to the OARs. However, achieving clinically acceptable PTV coverages and maximum doses often failed. Therefore, the quadratic function was supplemented with an LTCP term. Also for FMO plan generation, the LTCP cost function is often used for obtaining adequate PTV dose (Alber and Reemtsen, 2007, Breedveld et al., 2012, 2017).

In our study, segmented plans were in good agreement with the FMO plans (section 2.3.3), and the number of segments and MU of the generated plans are in line with our clinical experience. Quantitative comparison with published segmentation methods is difficult due to large variations in clinical cases and plans, and in applied segmentation objectives, parameters and quality measures. Also, in published studies the number of evaluated cases is generally relatively low (typically 2 to 10 compared to 47 in this study). However, qualitative comparisons indicate a superior FMO plan reproduction with the proposed approach. Explanations for this could be i) the direct reconstruction of the 3D FMO dose distribution with total ignorance of obtained FMO fluences, and simultaneous segment generation for all beam directions, instead of the generally observed focus on reconstruction of separate fluence profiles, and ii) the prioritised approach with an explicit drive to avoid plan quality losses for the highest clinical priorities. With this approach, segmentation is fully focused on maintaining the quality of the FMO plan with minimal limitations in selecting optimal segments and preservation of the clinical trade-offs. We have indeed observed substantial differences between FMO- and segmented fluence profiles for similar dose distributions (see figures 2.2 and 2.3 for an example patient). Possibly, the featured large freedom in beam segment selection has contributed to the high quality of the reconstructed plans with
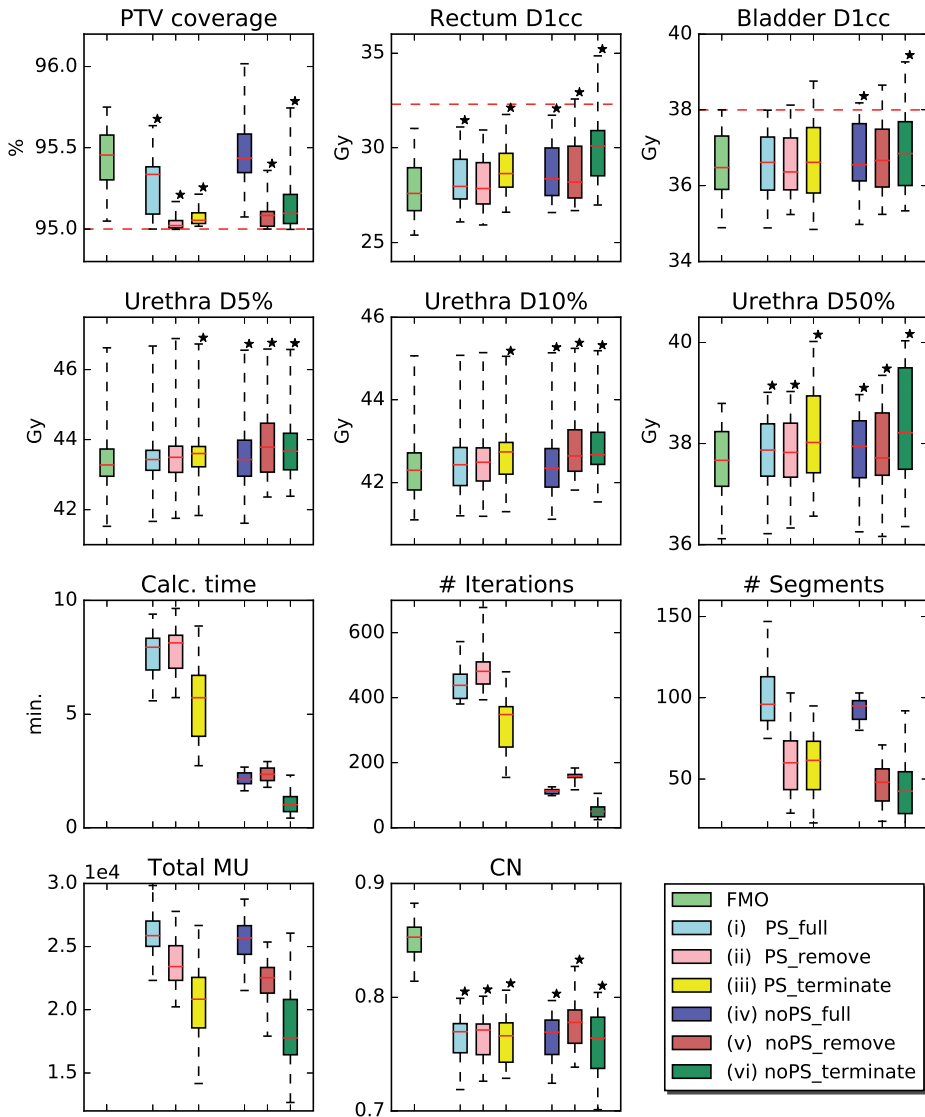
Figure 2.10: Comparison of treatment plan characteristics for FMO and the six segmentation approaches for prostate cancer. The whiskers in the boxplots denote the range of the data, i.e. minimum and maximum observed values and the coloured bars the 25-75% percentile range with the median value depicted with a horizontal line. Horizontal red dashed lines denote clinical constraints as summarised in table 2.2b. A plan was considered acceptable if for all constraints the obtained values were within 0.25 Gy or 0.25% of imposed values. MU = Monitor Unit, CN = Conformation Number, an asterisk (*) denotes a significant ($p < 0.05$) difference compared to FMO (Wilcoxon Signed Rank test).
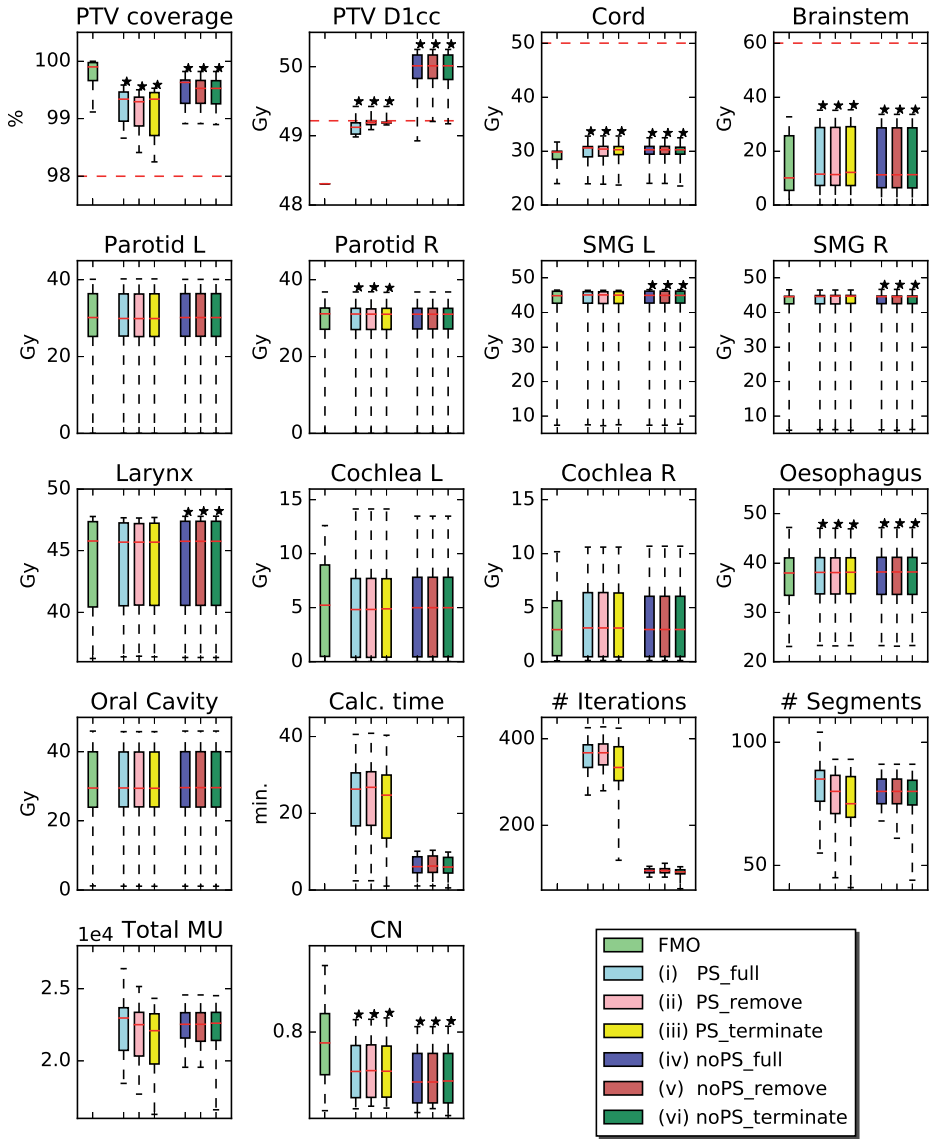
Figure 2.11: Comparison of treatment plan characteristics for FMO and the six segmentation approaches for head-and-neck cancer. The whiskers in the boxplots denote the range of the data, i.e. minimum and maximum observed values and the coloured bars the 25-75% percentile range with the median value depicted with a horizontal line. Horizontal red dashed lines denote clinical constraints as summarised in table 2.2b. A plan was considered acceptable if for all constraints the obtained values were within 0.25 Gy or 0.25% of imposed values. MU = Monitor Unit, CN = Conformation Number, an asterisk (*) denotes a significant ($p < 0.05$) difference compared to FMO (Wilcoxon Signed Rank test).
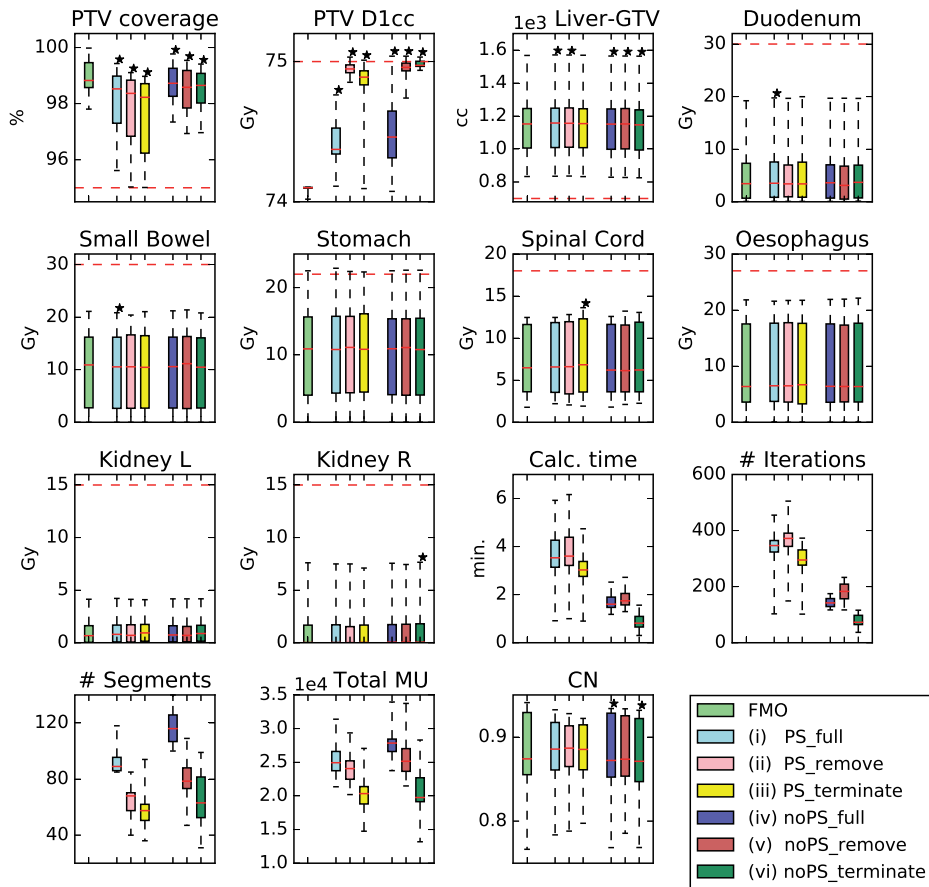
Figure 2.12: Comparison of treatment plan characteristics for FMO and the six segmentation approaches for liver cancer. The whiskers in the boxplots denote the range of the data, i.e. minimum and maximum observed values and the coloured bars the 25-75% percentile range with the median value depicted with a horizontal line. Horizontal red dashed lines denote clinical constraints as summarised in table 2.2b. A plan was considered acceptable if for all constraints the obtained values were within 0.25 Gy or 0.25% of imposed values. MU = Monitor Unit, CN = Conformation Number, an asterisk (*) denotes a significant ($p < 0.05$) difference compared to FMO (Wilcoxon Signed Rank test).

clinically acceptable numbers of segments (see figures 2.10, 2.11 and 2.12). Nguyen et al. 2015 investigated dose domain regularisation for MLC segmentation and observed that the segmented fluence profiles with and without regularisation could be substantially different. A direct comparison with published or commercially available segmentation approaches would be very interesting. For such studies it would be important to eliminate potential bias, e.g. originating from difference in dose calculation models, segmentation objectives or evaluation measures.

In this paper, FMO treatment plans were generated with Erasmus-iCycle (Breedveld et al., 2012), which has been successfully implemented in clinical practice for fully automated multi-criterial generation of clinically deliverable plans for head-and-neck, prostate, advanced lung cancer and advanced cervical cancer. For these tumour sites, Erasmus-iCycle is used for FMO plan generation, while the Monaco TPS (Elekta AB, Stockholm, Sweden) is effectively used for segmentation (Della Gala et al., 2017, Heijmen et al., 2018, Sharfo et al., 2015, Voet et al., 2013a, 2014). For automated offline treatment planning calculation time is not crucial, but in other scenarios it could be. Therefore, we have investigated calculation times for the segmentation approaches. When plan quality is the most important aspect, the PS_full is most suitable; fully converged segmentation with the best plan quality. When calculation time is more important, for example in the case of online-adaptive treatment, the PS_terminate could be a more suitable option. It provides a reduction in number of segments compared to the PS_full, with only minimal impact on plan quality. Additionally, the *terminate* and *remove* plans generally have a lower number of MU and so these plans are more efficient to deliver compared to the fully converged plans. A possible drawback of the terminate and remove approaches could be that the personalised objectives obtained from the FMO dose distributions are too challenging to reconstruct under the mechanical limitations. In that case, a plan would be returned with a high number of segments while fulfilling the highest feasible objectives. Regarding the terminate approach, since the segmentation minimises the difference between the FMO dose and segmented dose, terminating the segmentation when the tumour objectives are met does not necessarily mean that none of the remaining objectives are met. The segmentation primarily works on all dose points simultaneously with an extra emphasis on PTV coverage due to the LTCP term, prior to placing extra emphasis on other prioritised objectives. Investigations on further reduction of calculation times using GPU are on-going; preliminary results indicate a potential reduction in calculation time with a factor of 5.

The prioritised extension of the segmentation has been observed to be effective, but the impact is patient dependent and it increases calculation time. An option for clinical practice would be to always perform a full prioritised segmentation and let

the user decide which plan is most appropriate per patient, given that the intermediate (e.g. PS_terminate, noPS_full, noPS_terminate) plans are also available when PS_full is performed.

Interesting opportunities for further improving the proposed MLC segmentation approach would be to integrate published improvements in DAO techniques into this MLC segmentation method. One interesting approach in particular would be to integrate the Aperture Shape Optimization (ASO) algorithm as proposed by Cassioli and Unkelbach (2013), which optimises the shapes of the included segments in between CG iterations. However, this will also increase computation time. Another useful improvement might be to integrate a clinical dose engine in order to account for MLC scatter effects or to include a fuzzy controller to reduce numerical noise on the gradient maps as proposed by Yang et al. (2018).

2.5   Conclusions

Novel MLC segmentation approaches have been proposed for accurate reconstruction of high-quality FMO dose distributions, while complying with the mechanical limitations of the treatment device. 3-dimensional FMO dose distributions are reconstructed with total ignorance of underlying FMO fluences, and simultaneous segment generation for all beam directions. Due to the proposed prioritised approach, plan reconstruction has an accent on high priority planning objectives. Clinically acceptable segmented dose distributions could be generated for all cases with a plan quality that was in good agreement with the FMO plan and clinically acceptable numbers of segments.

Acknowledgments

Accurate 3D-dose-based generation of MLC segments for robotic radiotherapy

B.W.K. Schipaanboord[1], B. Heijmen[1], S. Breedveld[1]

[1] Department of Radiation Oncology, Erasmus MC Cancer Institute, Rotterdam, the Netherlands

Abstract

Radiotherapy treatment planning requires accurate modeling of the delivered patient dose, including radiation scatter effects, MLC leaf transmission, interleaf-leakage, etc. In fluence map optimization (FMO), a simple dose model is used to first generate an intermediate plan based on pencil-beams. In a second step (segmentation phase), this intermediate plan is then converted into a deliverable treatment plan with MLC segments. In this paper, we investigate novel approaches for the use of a clinical dose engine (CDE) for segmentation of FMO plans in robotic radiotherapy. Segments are sequentially added to the plan. Generation of each next segment is based on the total 3D dose distribution, resulting from already selected segments and the desired FMO dose, considering all treatment beams as candidates for delivery of the new segment. Three versions of the segmentation algorithm were investigated with differences in the integration of the CDE. The combined use of pencil-beams and segments in a segmentation method is non-trivial. Therefore, new methods were developed for the use of segment doses calculated with the CDE in combination with pencil-beams, used for the selection of new segments. For 20 patients with prostate cancer and 12 with liver cancer, segmented plans were compared with FMO plans. All three versions of the proposed segmentation algorithm could well mimic FMO dose distributions. Segmentation with a fully integrated CDE provided the best plan quality and lowest numbers of Monitor Units and segments at the cost of increased calculation time.

3.1   Introduction

In intensity-modulated radiation therapy (IMRT), dose is delivered with a multi-leaf collimator (MLC), which comes with radiation scatter effects, MLC leaf transmission, and interleaf-leakage. There are generally two optimization approaches to construct an IMRT plan.

The first approach splits the optimization problem into a fluence map optimization (FMO) phase and a segmentation phase. In FMO, each beam is discretized into a two-dimensional grid of beamlets, confined by the beam shape. Beamlet intensities are optimized using a pencil-beam formulation that relates intensities to delivered patient dose. Pencil-beam dose distributions are approximative, as final MLC configurations are not known in the FMO phase. In the subsequent segmentation phase, the optimized beamlet intensities are converted into MLC segments to create a deliverable plan. This conversion may degrade the FMO plan quality as the dosimetric impact of the MLC is not fully accounted for during FMO. Nevertheless, an advantage of this approach is that the FMO problem can be modeled as a convex multi-criterial optimization problem (Breedveld et al., 2019b) with a guaranteed globally optimal solution in the FMO phase.

The second approach for IMRT plan generation, called Direct Aperture Optimization (DAO), directly generates MLC segments (Men et al., 2007, Romeijn et al., 2005, Shepard et al., 2002). The non-convex modeling of the (physical) constraints of the treatment device in DAO leads to a non-convex optimization problem with a risk to get trapped in a local minimum. Moreover, the identification of each new segment is based on a simplified dose model, with similar shortcomings as FMO. In our center, we have given preference to treatment plan optimization using FMO followed by segmentation, because of the Pareto and global optimality of the FMO plans that are generated with Erasmus-iCycle (Breedveld et al., 2012).

In a previously published paper, we proposed a segmentation algorithm that reconstructs the 3D FMO dose distribution considering all beams simultaneously (Schipaanboord et al., 2019c), rather than separate reconstruction of the FMO beam fluence profiles as done in published segmentation algorithms (Gören and Taşkin, 2015, Long et al., 2016, Luan et al., 2006, Sun and Xia, 2004, Süss et al., 2007). This segmentation algorithm used a column generation (CG) approach, similar to DAO approaches, that iteratively identified segments and added them to the treatment plan. Segments were constructed using pencil-beams generated with a simplified MLC model, assuming that a following re-optimization of the segment weights with an accurate dose engine could (in part) compensate for the simplified dose model used during segmentation. Yet, this was not verified.

In this paper, we investigated approaches for integration of a clinical dose engine (CDE) into the segmentation algorithm proposed by Schipaanboord et al. (2019c). A CDE can accurately calculate dose delivered by MLC segments, taking into account scatter and transmission effects. When using an accurate (clinical) dose engine, there is a discrepancy between dose delivered by pencil-beams and by its corresponding segment, which results in challenges for the previously used column generation approach. Three versions of the segmentation algorithm with different degrees of integration of the CDE were investigated. For comparison of the segmentation approaches, calculation times, ability to dosimetrically mimic corresponding FMO plans, numbers of Monitor Units (MU), and numbers of segments were considered.

3.2   Materials & Methods

First, the column generation algorithm for MLC segmentation is summarized in section 3.2.1. Then, the integration of the CDE into the segmentation algorithm is described in section 3.2.2. Details on the FMO plans, plan evaluation criteria, and the validation of plan segmentation methods are presented in section 3.2.3. Section 3.2.4 describes the used computer hardware.

3.2.1   MLC segmentation using column generation

The use of column generation for the segmentation is previously described in detail in Schipaanboord et al. (2019c). In short, the segmentation problem is defined as the master problem (MP), which is iteratively solved by repeatedly solving a restricted version of the problem, denoted as the RMP. The RMP only holds a subset of the original decision variables, which in the perspective of MLC segmentation represents the so far selected segments. The RMP is solved to obtain the segment weights. The solution of the RMP is then converted into pencil-beam weights, which are used to calculate the gradient of the objective function of the MP. In the subsequent pricing problem (PP) step, the gradient is used to identify promising beamlets, which are then grouped to from a segment to be added to the RMP. After the PP step, the RMP is solved again to obtain the segment weights and the process is repeated. When no new segments can be identified that improve the MP, the MP is solved to optimality. The formulation of MP, PP and RMP used in this study are described in sections 3.2.1.1, 3.2.1.2 and 3.2.1.3, respectively.

3.2.1.1 Master Problem

For the formulation of the MP, a pencil-beam approach was used to describe dose delivered to the patient: $d = Ax$, where x denotes the pencil-beam weights (fluence intensities), A denotes the dose deposition matrix, and d denotes the vector containing the patient's voxel doses. The segmentation problem is formulated by:

$$\underset{x}{\text{minimize}} \quad f(x) + \omega_1 p(x) + \omega_2 g(x), \qquad \text{(MP)} \quad (3.1)$$

$$\text{subject to} \quad x \geq 0.$$

$x \geq 0$ ensures non-negative fluences and $f(x)$, $p(x)$, and $g(x)$ are given by:

$$f(x) = \sum_{v \in V} \|\eta_v^{1/2}(Ax - d_v^F)\|_2^2, \qquad (3.2)$$

$$p(x) = \frac{1}{N_{PTV}} \sum_{j=1}^{N_{PTV}} e^{-\alpha(A_j x - D^P)}, \qquad (3.3)$$

$$g(x) = \sum_{x \in X} x. \qquad (3.4)$$

The MP consists of three terms: 1) a quadratic term $f(x)$ that minimizes dose differences between input FMO dose and segmented dose, 2) a term $p(x)$ based on logarithmic tumor control probability (LTCP), with weight ω_1, to put extra emphasis on attaining adequate PTV coverage (Schipaanboord et al., 2019c), and 3) a configurable Monitor Units (MU) penalty term $g(x)$, with weight ω_2, to invoke reduction of the beam-on time.

In $f(x)$, the vector η_v contains voxel-dependent weight factors for each volume $v \in V$ and d_v^F defines the FMO distribution. In $p(x)$, D^P denotes the prescribed dose for the PTV, N_{PTV} the number of sampled PTV voxels, $A_j x$ the dose delivered to PTV voxel j, and α is a constant related to cell survival and this constant was tuned to achieve adequate PTV coverage (Schipaanboord et al., 2019c). In $g(x)$, X denotes the set of all pencil-beams.

3.2.1.2 Pricing Problem

New segments are identified in the PP step. For the identification, the gradient of the objective function of the MP is calculated and reshaped from the vectorized format to the 2-dimensional beamlet grid format for each of the beams. Beamlets with negative gradients are favorable for inclusion in the next segment, as they move the MP towards its minimum. A segment is formed by grouping a set of beamlets with negative gradients under the mechanical restrictions of the

treatment device. By grouping beamlets with negative gradients and adding them to the RMP, the CG approach utilizes a greedy gradient descent method to solve the MP. For identifying the most promising group of connected beamlets with negative gradients, a layered graph approach is used, similar to the approach in Romeijn et al. (2005) and also explained in detail in Schipaanboord et al. (2019c). The restrictions are taken into account while constructing a layered graph for each of the beam directions and the graphs are subsequently solved using a Bellman-Ford shortest path algorithm. Each identified segment i is then represented by a binary vector of open and closed beamlets, denoted by S_i.

The mechanical restrictions enforced during identification of new segments were similar to the restrictions of the CyberKnife InCise2 MLC: a segment can contain only one contiguous opening, a minimum field opening size of 7.5 x 7.7 mm² (57.75 mm²), a minimum number of 2 leaf pairs open per segment (7.7 mm in total), a minimum opening size of 5 mm in the directions of the leaves, and interdigitation is allowed. All segmentations were performed on a resolution of 5×3.85 mm². The leaf width matches that of the CyberKnife InCise 2 MLC (Accuray Inc.), which has 2 banks of 26 leaves with a leaf width of 3.85 mm defined at 800 mm SAD. To reduce inter-leaf leakage the CyberKnife InCise2 MLC is tilted by 0.5° which reduces the projected segment size and causes an asymmetry of approximately ±0.25 mm in the penumbral width at the top and bottom of the segment, analogous to the tongue-and-groove effect observed with interlocking leaves (Asmerom et al., 2016). In the finite sized pencil-beam algorithm of the CDE the projected aperture reduction is applied during segment dose calculation, and the penumbra asymmetry is ignored.

3.2.1.3 *Restricted Master Problem*

For each iteration of the CG, the identified segment is added to the RMP by adding a new column to the optimization problem. This column represents the dose deposition of the segment and is calculated by multiplying the pencil-beam matrix A with the binary open-closed vector representation of the segment S_i, the segment dose deposition matrix for all identified segments is denoted by A_r. The updated RMP is then solved to optimality to obtain the segment weights (segment intensities) x_r. For this, we use our in-house developed solver (Breedveld et al., 2017), specifically designed and tuned for solving radiotherapy optimization problems, but any non-linear solver can be used. The RMP is formulated by:

$$\underset{x_r}{\text{minimize}} \quad f_r(x_r) + \omega_1 p_r(x_r) + \omega_2 g_r(x_r), \qquad \text{(RMP)} \quad (3.5)$$

$$\text{subject to} \quad x_r \geq 0.$$

$x_r \geq 0$ ensures non-negative segment weights and $f_r(x_r)$, $p_r(x_r)$, and $g_r(x_r)$ are given by:

$$f_r(x_r) = \sum_{v \in V} \|\eta_v^{1/2}(A_r x_r - d_v^F)\|_2^2, \tag{3.6}$$

$$p_r(x_r) = \frac{1}{N_{PTV}} \sum_{j=1}^{N_{PTV}} e^{-\alpha(A_{rj}x_r - D^P)}, \tag{3.7}$$

$$g_r(x_r) = \sum_{x_r \in X_r} x_r, \tag{3.8}$$

i.e. identical to equations (3.2-3.4), but then in segment-space instead of beamlet-space. At last, to continue the CG loop, the segment weights are expressed in pencil-beam weights using (3.9). To convert the segment weights x_r into pencil-beam weights x, the open beamlets S_i of a segment i are multiplied by their corresponding segment weight x_{ri} and then accumulated over all identified segments I.

$$x = \sum_{i=1}^{I} S_i x_{ri} \tag{3.9}$$

The RMP solution expressed in pencil-beam weights can then be used to calculate the gradient of the objective function of the MP, which is used to identify the next segment.

3.2.2   Integration of clinical dose engine in MLC segmentation

The CDE can accurately calculate segment doses, taking into account scatter and transmission effects (not accounted for in the pencil-beam dose calculations), by modeling the MLC segment as one physical object. For the integration of the CDE in the column generation loop, every time a segment is identified in the PP the corresponding leaf configuration of the MLC is provided to the CDE. The CDE then returns a dose deposition vector for unit intensity of this segment. This dose deposition vector can then be used in the RMP to calculate the weights of the segments.

One of the challenges for integrating a CDE into the CG approach is that a segment dose calculated with the CDE differs from the segment dose calculated as the sum of its pencil-beams. Since the segments shapes, i.e. the leaf positions, are not known when the pencil-beams are calculated, an approximation of the field size, scatter and transmission effects are used. While, for the segment dose calculations, the leaf positions are known and thus the scatter and transmission effects are more

accurately predicted. This can cause problems when the RMP solution is converted back to the individual pencil-beam weights (using equation (3.9)). The MP expects a dose deposition per pencil-beam as given in A_{pb} ($d = A_{pb}x$), but x_r is calculated using the dose deposition per segment as given in A_{seg}. When converting x_r back to x, this causes a discrepancy in dose calculated in the MP as compared to the dose calculated in the RMP. Because of this dose discrepancy between the two dose deposition matrices, the dose resulting from x_r ($d = A_r x_r$) is not an accurate representation for the MP, since the MP is based on $A \neq A_r$.

In practice, due to the scatter and transmission effects accounted for in the RMP and not in the MP, this would result in an over-prediction of the delivered dose in the MP. As a consequence, fewer beamlets have a negative gradient and therefore fewer beamlets are predicted to be favorable for inclusion in the next segment. This would cause the segmentation to terminate prematurely. In the conventional CG approach, this problem does not occur because both the MP and the RMP use the same pencil-beam basis for dose calculation, i.e. $A_{pb}x = A_r x_r$.

To mitigate this difference in segment dose deposition between the MP and the RMP, we have introduced two extra versions of the RMP:

RMP$_{seg}$:

This version of the RMP accurately models the segment dose using the CDE. The RMP$_{seg}$ is formulated by:

$$\underset{x_{seg}}{\text{minimize}} \quad f_{seg}(x_{seg}) + \omega_1 p_{seg}(x_{seg}) + \omega_2 g_{seg}(x_{seg}), \qquad \text{(RMP}_{seg}\text{)} \quad (3.10)$$

$$\text{subject to} \quad x_{seg} \geq 0.$$

$x_{seg} \geq 0$ ensures non-negative segment weights and $f_{seg}(x_{seg}), p_{seg}(x_{seg})$, and $g_{seg}(x_{seg})$ are given by:

$$f_{seg}(x_{seg}) = \sum_{v \in V} \|\eta_v^{1/2}(A_{seg}x_{seg} - d_v^F)\|_2^2, \qquad (3.11)$$

$$p_{seg}(x_{seg}) = \frac{1}{N_{PTV}} \sum_{j=1}^{N_{PTV}} e^{-\alpha(A_{seg,j}x_{seg} - D^P)}, \qquad (3.12)$$

$$g_{seg}(x_{seg}) = \sum_{x_{seg} \in X_{seg}} x_{seg}. \qquad (3.13)$$

Here, A_{seg} denotes the accurate segment dose deposition matrix calculated with the CDE, each column of A_{seg} denotes the dose deposition per identified segment for unit intensity, x_{seg} denotes the segment weights for the CDE segments and X_{seg} the set of CDE segment weights.

RMP$_{pb}$:

This version of the RMP models the segment dose as the sum of its pencil-beams, similar to the RMP from section 3.2.1.3. The RMP$_{pb}$ is formulated by:

$$\underset{x_r}{\text{minimize}} \qquad f_{pb}(x_r) = \sum_{v \in V} \|\eta_v{}^{1/2}(A_{pb,r}x_r - A_{seg}x_{seg})\|_2^2, \qquad (\text{RMP}_{pb}) \quad (3.14)$$

$$\text{subject to} \qquad x_r \geq 0.$$

This RMP approximates the accurate dose from the RMP$_{seg}$ ($A_{seg}x_{seg}$) in pencil-beams. It finds the segment weights x_r that minimize the dose difference between the accurately modeled segments in the RMP$_{seg}$ and the segments modeled as the sum of its pencil-beams in the RMP$_{pb}$. This minimization step attempts to account for differences in segment intensities due to MLC scatter effects.

We have investigated three variants for the integration of the CDE in the column generation algorithm, see figure 3.1:

No integration (NO-int):

The first method does not use the CDE in the intermediate iterations and resembles a conventional CG model for MLC segmentation. Throughout the segmentation, the segment dose is modeled as the sum of its pencil-beams (RMP). The CDE is only applied for a final re-optimization of the segment weights using the RMP$_{seg}$ and for the final dose calculation, after all segments shapes have been identified. This method functions as a baseline performance.

Medium integration (MED-int):

The second method uses the CDE to calculate the accurate segment dose A_{seg} for each newly identified segment in the PP, and scales the segment as a whole. In the RMP$_{seg}$ (3.10) the weights x_{seg} for the accurate CDE segments are calculated. The dose corresponding to this solution, $A_{seg}x_{seg}$, is converted in pencil-beams by the RMP$_{pb}$ (3.14). For this, the RMP$_{pb}$ uses an open and closed beamlet representation S_i (3.9) to model the segment dose in $A_{pb,r}$. In this case, each non-zero element of S has the same value, namely 1. The conversion from accurate segment dose to pencil-beams is necessary to calculate the gradient on the pencil-beams for the next CG iteration.

Full integration (FULL-int):

The third method uses the CDE to calculate the accurate segment dose A_{seg} for each newly identified segment in the PP, similar to MED-int. Also, the RMP_{seg} (3.10) calculates the weights x_{seg} for the CDE segments. Additional to the conversion step using the RMP_{pb} (3.14), this method features an additional segment approximation step because the MLC scatter effects are not uniformly distributed over the open area of a segment but occur mostly at the edges of the opening. Instead of an open and closed beamlet representation, as used for MED-int, the pencil-beams are allowed to have partial contributions to the segment dose. I.e. S_i for the identified segment i is no longer binary (open-closed), but may now contain continuous values (3.9). This additional approximation step is formulated by:

$$\underset{S_i}{\text{minimize}} \qquad \sum_{v \in V} \|\eta_v^{1/2}(A_{pb}S_i - A_{seg,i})\|_2^2, \tag{3.15}$$

$$\text{subject to} \qquad S_i \geq 0. \tag{3.16}$$

This results in a closer representation of the accurate segment dose when converting the CDE segment weights to pencil-beam weights. Note that this approach physically still represents a deliverable segment (RMP_{seg}), but the re-weighting takes place in the pencil-beam space (RMP_{pb}).

3.2.3   Plan details and segmentation validation

The segmentation performance of the proposed methods was evaluated on a cohort of 32 patients treated with robotic radiotherapy, consisting of 20 SBRT prostate cases (38 Gy in 4 fractions, with highly heterogeneous PTV dose distributions to mimic HDR brachytherapy dose distributions) and 12 liver cases (60 Gy in 3 fractions). All FMO plans were generated with fully automated multi-criterial optimization as implemented in Erasmus-iCycle (Breedveld et al., 2012, 2017), generating Pareto-optimal plans, and with a 25 non-coplanar beam geometry (Rossi et al., 2015). These FMO plans were identical to the plans used in Schipaanboord et al. (2019c). The clinical objectives used to evaluate the segmentation performance are given in table 3.1.

To visualize the trade-off in plan quality versus MU, five segmentations per case were performed with increasing weights of the MU penalty term ($\omega_2 \in \{1, 2.5, 5, 7.5, 10\}$). A one-sided plan quality index (PQI) was defined as the mean difference in achieved criterion values (in Gy) between an FMO plan and its corresponding segmented plan (3.17), with n the number of criteria for that

No integration (NO-int)
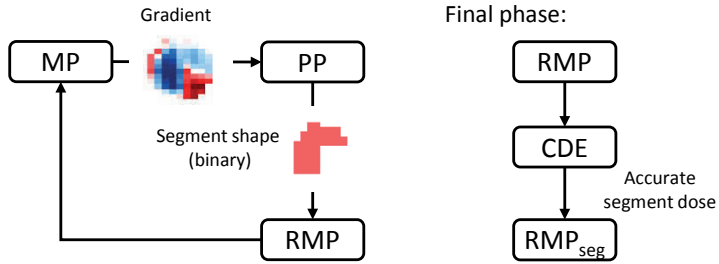
Medium integration (MED-int)
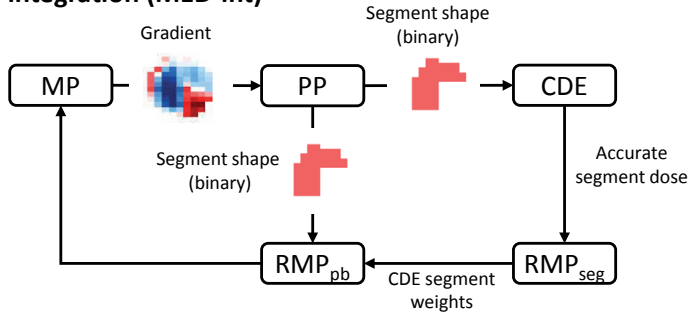
Full integration (FULL-int)
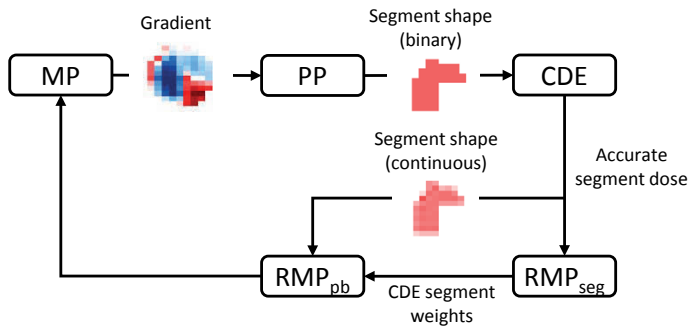
Figure 3.1: Integrations of the CDE into the column generation algorithm. The NO-int method only uses the CDE after all segments have been identified. The MED-int method uses the CDE during the segmentation for identification of new segments. For this, it uses a binary (open and closed) beamlet intensity representation in the RMP_{pb} to model the accurate segment dose from the RMP_{seg}. The FULL-int method uses the CDE during the segmentation for identification of new segments and uses a continuous beamlet intensity representation in the RMP_{pb} to model the accurate segment dose from the RMP_{seg} more closely than the MED-int. See section 3.2.2 for the formulations of the RMP_{pb} and the RMP_{seg}. Abbreviations: MP: Master Problem, PP: Pricing Problem, RMP: Restricted Master Problem, CDE: Clinical Dose Engine.

Table 3.1: Clinical objectives for prostate and liver cases.

Prostate		Liver	
PTV	$V_{38Gy} > 95\%$	PTV	$V_{60Gy} > 95\%$
Urethra	$D_{5\%} < 45\,Gy$		$D_{1cc} < 75\,Gy$
Urethra	$D_{10\%} < 42\,Gy$	Liver - GTV	$V_{15Gy} \geq 700\,cc$
Urethra	$D_{50\%} < 40\,Gy$	Duodenum	$D_{1cc} < 30\,Gy$
Rectum	$D_{1cc} < 32.3\,Gy$	Small bowel	$D_{1cc} < 30\,Gy$
Bladder	$D_{1cc} < 38\,Gy$	Stomach	$D_{5cc} < 22.5\,Gy$
		Spinal cord	$D_{1cc} < 18\,Gy$
		Esophagus	$D_{1cc} < 27\,Gy$

patient, FMO_i the value of criterion i for the FMO plan and SEG_i the achieved criterion value for the corresponding segmented plan.

$$PQI\,[\Delta Gy] = \frac{1}{n}\sum_{i=1}^{n}\begin{cases}|FMO_i - SEG_i| & \text{if } FMO_i \leq SEG_i \\ 0 & \text{if } FMO_i > SEG_i\end{cases} \tag{3.17}$$

Like the FMO plans, all segmented treatment plans were linearly scaled to the PTV coverage objective for a fair comparison in plan quality.

3.2.4 Computer hardware

Segmentations were performed on an Intel Core i7-3820 with 4 physical cores running at 3.60 GHz, with 32 GB of memory and equipped with an NVIDIA GeForce GTX Titan Black GPU card. The segmentation involves many (computationally intensive) matrix-vector multiplications, these are performed on the GPU to speed up the process. All calculation times are including data transfer between our research environment and the CDE.

3.3 Results

In section 3.3.1, we first present the segmentation performance on one prostate case. Overall performance on a patient population basis is then discussed in section 3.3.2.

3.3.1 Segmentation performance – example patient

Figure 3.2 shows the DVHs of the FMO plan and the three segmented plans for an example prostate case. Overall, there is a high degree of similarity between FMO and segmented doses. Also, differences between the three segmentations are minor. In figure 3.3, the input FMO dose distribution is compared with the dose distribution generated with the FULL-int segmentation method.
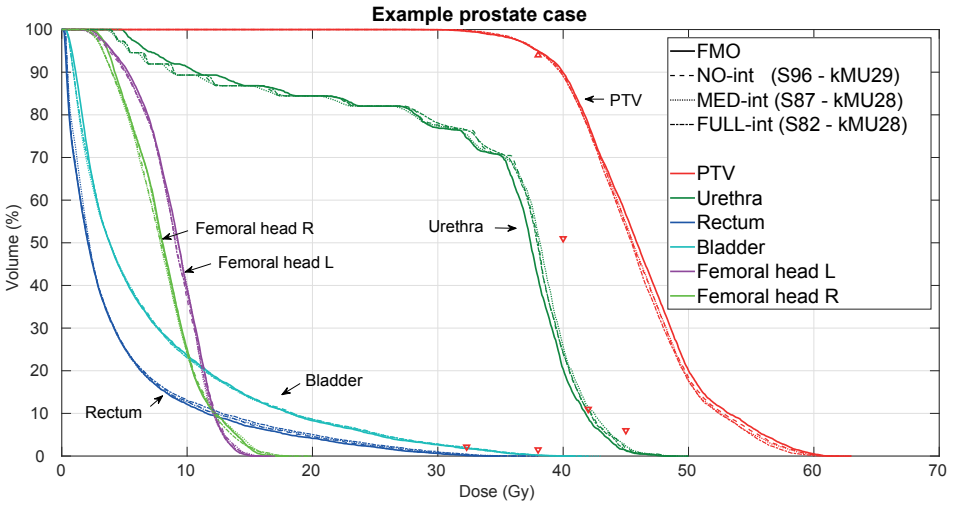
Figure 3.2: DVHs for FMO and segmented plans ($\omega_2 = 1$) for the example prostate case discussed in section 3.3.1. All plans were normalized to have a PTV coverage of 95% at 38 Gy. Sxx: number of segments, kMUxx: kMU (kilo MU). Clinical objectives are denoted with red triangles, see table 3.1.
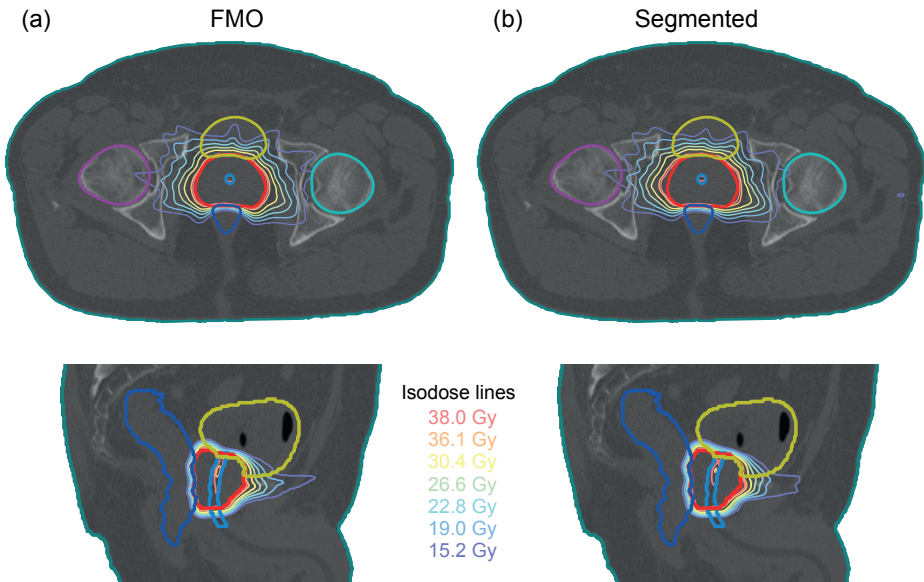
Figure 3.3: Axial (top) and sagittal (bottom) slices through the PTV for the FMO dose (a) and the segmented dose obtained with the fully integrated CDE (FULL-int) method (b) for the example prostate case discussed in section 3.3.1. Depicted structures: PTV (red), rectum (blue), bladder (yellow), urethra (light blue), and femoral heads L/R (cyan/purple).

3.3.2   Segmentation performance – patient population

Population averaged DVHs are shown in figure 3.4 for both the prostate plans and the liver plans. Evaluated on a population basis, for prostate, increasing the weight from 1 to 2.5 resulted in a violation of clinical constraints on either the rectum or the bladder. For clinically acceptable plans, a ω_2 equal to 1 was chosen for the prostate plans to be presented throughout this paper. For liver, ω_2 could be increased from 1 to 5 without compromising on any clinical constraints (see table 3.1). Therefore, for the liver plans a ω_2 equal to 5 was chosen to provide the best balance between clinically acceptable plans and deliverability of the plans. Unless otherwise stated, these ω_2 weights are used in the remainder of this paper. In general, the DVHs of the segmented plans are similar to the DVHs of the FMO plans for both treatment sites. For prostate, segmented urethra and rectum doses are slightly higher than in FMO. For liver, the segmented plans have a higher maximum dose than the FMO plans (less homogeneous PTV dose), yet all plans still meet the PTV $D_{1cc} \leq 75 \, Gy$ objective.

Dosimetric plan parameter values for segmented and FMO plans are compared in figure 3.5 (prostate) and in figure 3.6 (liver). For prostate, increases in urethra $D_{50\%}$ and rectum D_{1cc} were observed in the segmented plans, but occasional violations of clinical constraints or objectives were minor. For the other parameters, the FMO values were closely reproduced. Also for liver, differences between segmented and FMO plan parameter values were generally small and no constraint violations in segmented plans were observed. For both tumor sites, differences in the three segmentation approaches were small.

Figure 3.7 shows for the three segmentation approaches, PQI as a function of the number of MU and the number of segments. For data generation, segmentations were performed for various weights of the MU penalty term ($\omega_2 \in \{1, 2.5, 5, 7.5, 10\}$). A clear trade-off between PQI and the number of MU or the number of segments is observed. Plans generated with the highest MU penalty weight ($\omega_2 = 10$) have the lowest number of MU and the lowest number of segments, but also the highest PQI. The FULL-int method shows overall the best PQI values for equal MU or equal numbers of segments.

Table 3.2 shows segmentation times. The NO-int method is able to identify the segments within 2 minutes, but then needs an additional 12 minutes to recalculate the segment doses using the CDE. The average identification time for the MED-int and FULL-int methods was 49 minutes (over all plans, both prostate and liver). The additional recalculation step is not needed for MED-int and FULL-int, since it is integrated into the segmentation routine. For all methods, a final forward dose calculation can be performed to obtain the 3D dose distribution.
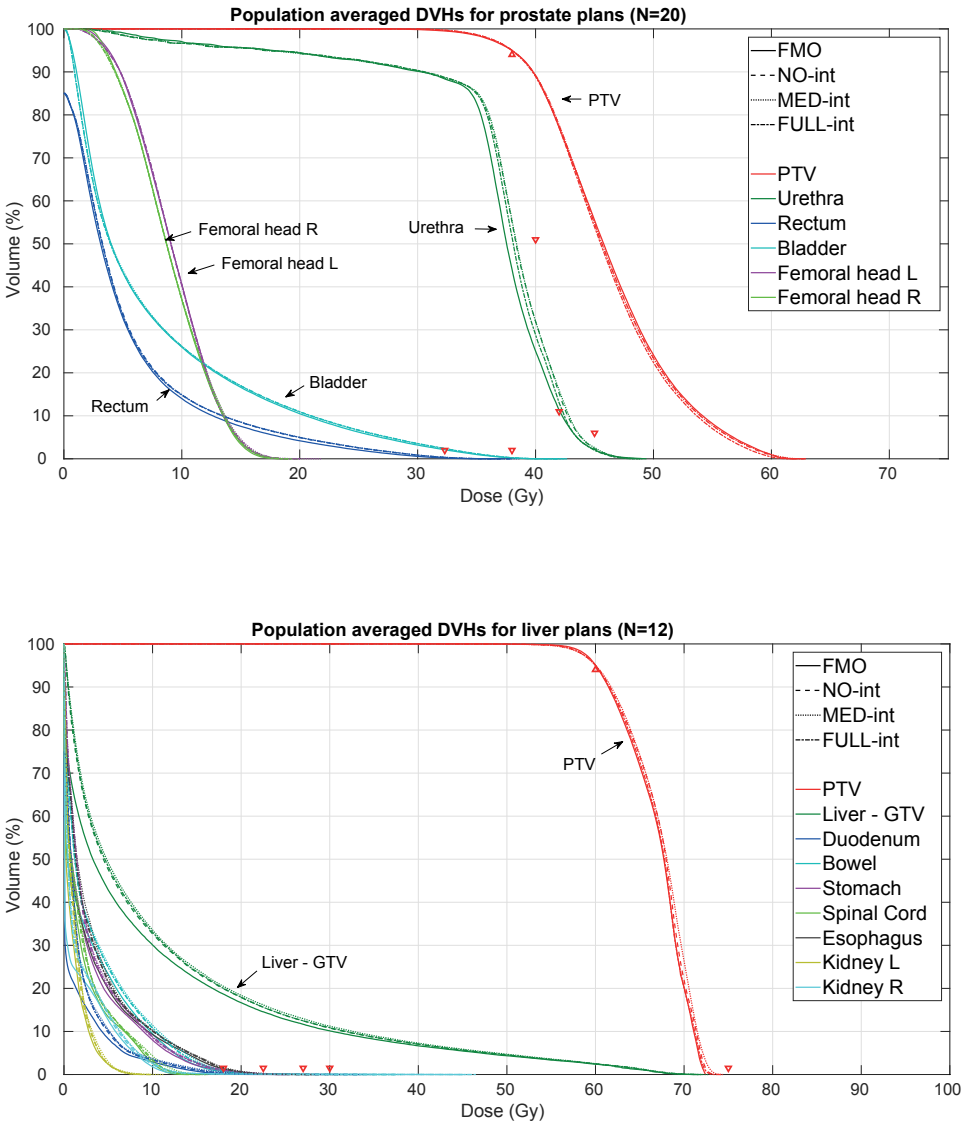
Figure 3.4: Population averaged DVHs. Upper panel: prostate, lower panel: liver. All prostate plans were normalized to have a PTV coverage of 95% at 38 Gy and all liver plans to a PTV coverage of 95% at 60 Gy. Clinical objectives are denoted with red triangles, see table 3.1.
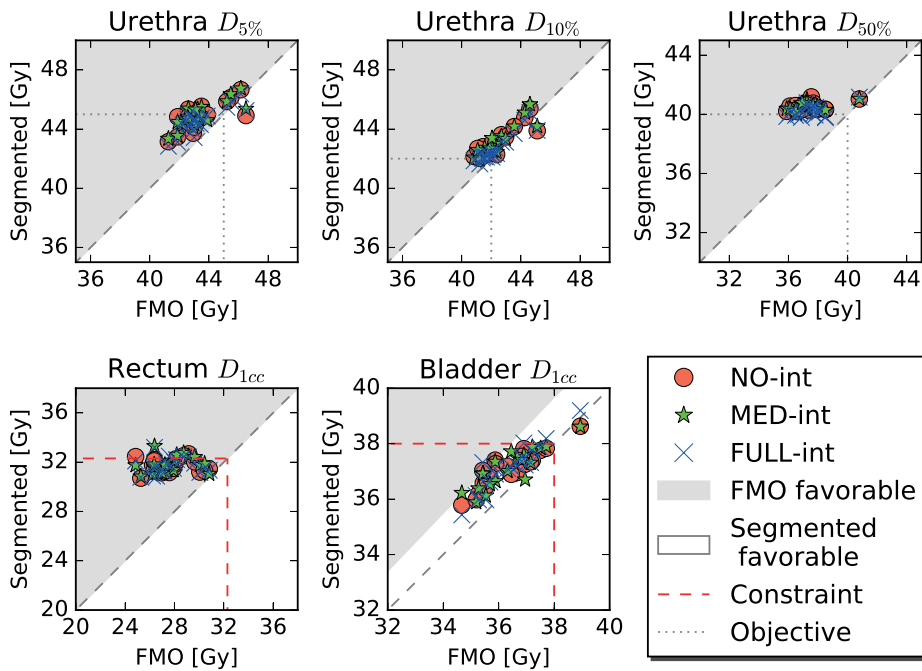
Figure 3.5: Segmentation results for the prostate plans. Segmented values for clinical objectives are plotted against the values of the input FMO plan. All plans were normalized to have a PTV coverage of 95% at 38 Gy.

Table 3.2: Average segmentation times needed to generate deliverable treatment plans. CDE = Clinical Dose Engine.

	Method	Segment identification [min]	CDE calculation & re-optimization [min]	**Subtotal** (identification) [min]	Total (incl. 3D dose) [min]
Prostate	NO-int	1.6 $\pm$ 0.2	12.6 $\pm$ 2.2	**14.2 $\pm$ 2.3**	20.0 $\pm$ 2.9
	MED-int	43.8 $\pm$ 7.9	-	**43.8 $\pm$ 7.9**	49.2 $\pm$ 8.3
	FULL-int	52.3 $\pm$ 8.4	-	**52.3 $\pm$ 8.4**	57.9 $\pm$ 9.0
Liver	NO-int	0.7 $\pm$ 0.1	11.5 $\pm$ 3.2	**12.3 $\pm$ 3.3**	15.5 $\pm$ 4.0
	MED-int	49.7 $\pm$ 14.6	-	**49.7 $\pm$ 14.6**	52.4 $\pm$ 15.4
	FULL-int	54.0 $\pm$ 13.7	-	**54.0 $\pm$ 13.7**	57.6 $\pm$ 14.7
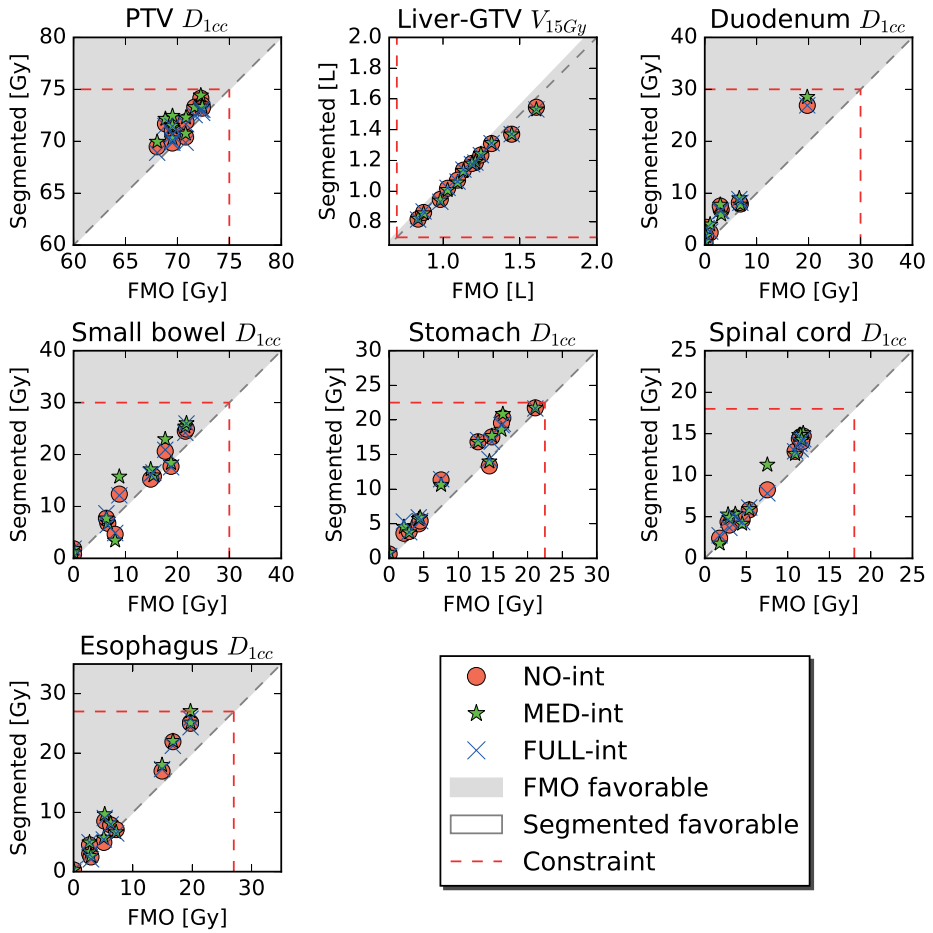
Figure 3.6: Segmentation results for the liver plans. Segmented values for clinical objectives are plotted against the values of the input FMO plan. All segmented plans met the constraints and also the FMO objective values were closely replicated. All plans were normalized to have a PTV coverage of 95% at 60 Gy.
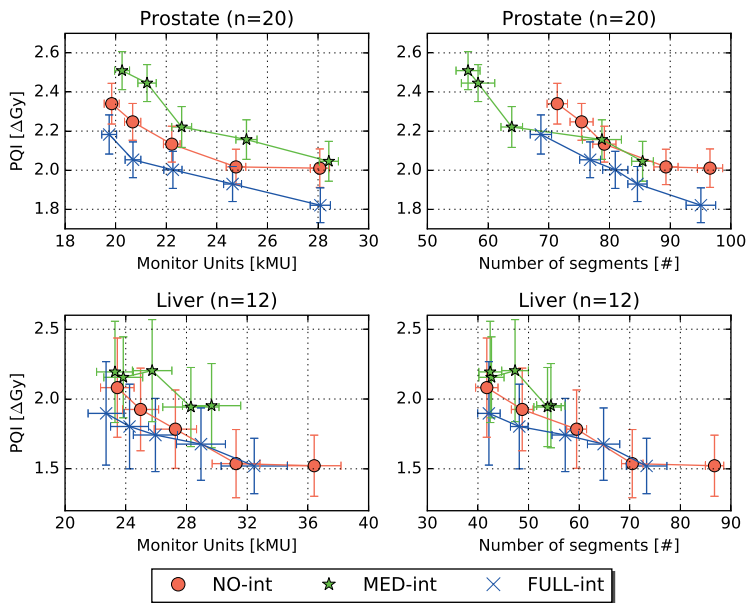
Figure 3.7: For various weights of the MU penalty term ($\omega_2 \in \{1, 2.5, 5, 7.5, 10\}$), PQI is plotted against the number of kMU (left panels) and the number of segments (right panels) for each of the three investigated segmentation approaches. Each marker represents a ω_2– "segmentation approach" combination plotted against the mean kMU or the mean segments over all patients; error bars in y-direction denote the standard error of the mean in PQI and error bars in x-direction denote the standard error of the mean in kMU or in segments. Mean kMU and mean segments decrease with increasing ω_2. Plans generated with the highest MU penalty weight (ω_2=10) have lowest kMU or segments.

3.4   Discussion

Accurate modeling of the segment dose, including leakage and scatter effects of the MLC, is essential when converting an FMO dose distribution into a deliverable radiotherapy treatment plan. In this study, we found that calculation of accurate segment doses at the end of the segmentation, together with a re-optimization of the segment weights, provides similar plan quality to that of treatment plans generated with a fully integrated CDE (figures 3.5 & 3.6). Yet, the fully integrated (FULL-int method) achieved this plan quality with fewer number of MUs or segments (figure 3.7). Therefore, if calculation times are not an issue, the FULL-int method would be the best segmentation method for the patients studied in this paper.

Published MLC segmentation methods are mainly based on independent segmentation of the 2D FMO beam fluence profiles (Gören and Taşkin, 2015, Long et al., 2016, Luan et al., 2006, Sun and Xia, 2004, Süss et al., 2007). In such

approaches, it is infeasible to fully account for MLC scatter effects. In literature, there are a few published papers that also use a CDE to model the dose. For example, Kontaxis et al. (2015a,b) use the Monte Carlo dose engine (GPUMCD) from Elekta AB for solving the segmentation problem by iteratively subtracting a segment dose from an ideal input dose. However, in their approach, the identification of new segments is done by segmentation of the fluence profiles per beam. For CyberKnife, Bedford et al. (2019b) used Accuray's CDE to calculate pencil-beam matrices for beam selection. This CDE was also used for dose calculation for dynamic arc delivery (Bedford et al., 2019a). To the best of our knowledge, this is the first paper to integrate the CDE into a column generation formalism to fully account for MLC scatter effects during segmentation, not only for MLC segmentation with column generation, but also for DAO approaches with column generation.

One of the obstacles for integrating the CDE in segmentation is that for each iteration the solution in the dose domain needs to be expressed in terms of pencil-beams. This conversion is needed to calculate the gradient of the objective function, which is then used to identify the next promising segment. Since the beam characteristics are different for each segment, the solution cannot be exactly expressed in terms of the generic pencil-beams. As a consequence, there is a discrepancy between the exact segment dose and the segment dose calculated with pencil-beams. This impacts the accuracy of the gradient used for identification of the next segment. In this study, we have minimized the impact of this discrepancy by introducing an additional optimization step within each iteration of the column generation loop. This optimization step minimizes the difference between the CDE dose and the dose from the pencil-beam solution. This additional optimization step reduces the discrepancy, yet does not fully resolve the issue. The remnant discrepancy is smaller for the FULL-int method than for the MED-int method, since the FULL-int method has more freedom for minimizing the dose differences per segment. This resulted in overall somewhat better plan quality for the FULL-int method (figure 3.7). Another interesting observation is that, when comparing NO-int with MED-int, the faster NO-int method performs generally better than the MED-int. A reason for this could be that, generally, at the start of the segmentation larger segments are added to the plan than at the later stages of the segmentation, where smaller segments are added that only subtly improve the MP per segment. We hypothesize that the prediction of smaller segments is more affected by the dose discrepancy between the RMP and the MP than the prediction of larger segments. In the NO-int method, the prediction by the PP of the segment dose to be added to the RMP is exactly what is added to the RMP. Due to this exact prediction, the NO-int method could continue to identify smaller segments that could be

beneficial to improve the MP for longer before stalling. The additional smaller segments in the NO-int plan, could then help to minimize the degradation when converted to accurate CDE doses. This could also explain why the NO-int method requires more segments than the FULL-int method for similar plan quality. Since the MED-int method can only linearly scale the segment weights as a whole to minimize the dose discrepancy between the RMP and the MP, the accuracy of the prediction is less than with the FULL-int method. Which will lead for the NO-int method to stall earlier than the FULL-int method. This suggests that, to exploit the added benefit of the CDE, the FULL-int method should be used, otherwise it is better and faster to use the NO-int method.

Since segments can be iteratively replaced by better ones during segmentation, more segments are generated than used in the final treatment plan. As a consequence, the MED-int and FULL-int methods have a larger overhead due to the many CDE dose calculations and, therefore, it takes around 3.5 times longer to calculate a treatment plan than with the NO-int method. Identification of the segments with the NO-int method takes 1-2 minutes. The subsequent serial recalculation of the segment doses using the CDE for re-optimization of the segment weights takes another 12 minutes and an extra 5 minutes for final 3D dose calculations. Since all the segments are already identified, these steps could be performed in parallel. This could reduce the calculation times to less than two minutes and make it potentially interesting for online adaptive treatment planning. Total segmentation times presented in table 3.2 are based on serial 3D calculations of the segment doses. On average, a CDE dose calculations costs 8 seconds. If these CDE calculations would be performed in parallel, this would result in a potential time reduction for prostate of 18 minutes for NO-int and 5 minutes for MED-int and FULL-int. For liver, potential time reductions are 14 minutes for NO-int and 3 minutes for MED-int and FULL-int.

Recently, many improvements for column generation in radiotherapy treatment planning have been published (e.g. Cassioli and Unkelbach 2013, Nguyen et al. 2016, 2017). In particular, two improvements would be very interesting in combination with the proposed algorithms. The first improvement would be to include a fuzzy controller to reduce numerical noise on the gradient maps as proposed by Yang et al. (2018). The second improvement would be to include the gradient of the previous iteration to improve convergence as proposed by Zhang et al. (2019).

3.5 Conclusions

With our segmentation algorithm, segments are sequentially added to the plan. Generation of each next segment is based on the total 3D dose distribution,

resulting from already selected segments, and the desired 3D FMO dose, considering all treatment beams as candidates for delivery of each new segment. When combined with a high-accuracy clinical dose engine, this algorithm could accurately reconstruct FMO plans optimized with pencil-beams into clinically deliverable treatment plans. Plan quality, MU, number of segments and calculation time depended on the level of integration of the clinical dose engine in the segmentation algorithm. The combined use of pencil-beams and segments in a segmentation method is non-trivial. Therefore, new methods were developed for the use of intermediate segment doses calculated with the CDE in combination with pencil-beams, used for selection of new segments.

Acknowledgments

3

Fully automated treatment planning for MLC-based robotic radiotherapy

B.W.K. Schipaanboord[1], **M.K. Giżyńska**[1], **L. Rossi**[1], **K.C. de Vries**[1], **B.J.M. Heijmen**[1], **S. Breedveld**[1]

[1] Department of Radiotherapy, Erasmus MC Cancer Institute, Rotterdam, the Netherlands

4

Abstract

Purpose: To propose and validate a fully automated multi-criterial treatment planning solution for a CyberKnife® equipped with an InCise™ 2 multi-leaf collimator.

Methods: The AUTO BAO plans are generated with fully-automated prioritized multi-criterial optimization (AUTO MCO) of pencil-beam fluence maps with integrated non-coplanar beam angle optimization (BAO), followed by MLC segment generation. Both the AUTO MCO and segmentation algorithms have been in-house developed. AUTO MCO generates for each patient a single, high-quality Pareto-optimal IMRT plan. The segmentation algorithm then accurately mimics the AUTO MCO 3D dose distribution, while considering all candidate beams simultaneously, rather than replicating the fluence maps. Pencil-beams, segment dose depositions, and final dose calculations are performed with a stand-alone version of the clinical dose calculation engine. For validation, AUTO BAO plans were generated for 33 prostate SBRT patients and compared to reference plans (REF) that were manually generated with the commercial treatment planning system (TPS), in absence of time pressure. REF plans were also compared to AUTO RB plans, for which fluence map optimization was performed for the beam angle configuration used in the REF plan, and the segmentation could use all these beams or only a subset, depending on the dosimetry.

Results: AUTO BAO plans were clinically acceptable and dosimetrically similar to REF plans, but had on average reduced numbers of beams ((beams in AUTO BAO)/(beams in REF) (relative improvement): 24.7/48.3 (-49%)), segments (59.5/98.9 (-40%)), and delivery times (17.1/22.3 min. (-23%)). Dosimetry of AUTO RB and REF were also similar, but AUTO RB used on average fewer beams (38.0/48.3 (-21%)) and had on average shorter delivery times (18.6/22.3 min. (-17%)). Delivered Monitor Units (MU) were similar for all three planning approaches.

Conclusions: A new, vendor-independent optimization workflow for fully automated generation of deliverable high-quality CyberKnife® plans was proposed, including BAO. Compared to manual planning with the commercial TPS, fraction delivery times were reduced by 5.3 min. (-23%) due to large reductions in beam and segment numbers.

4.1   Introduction

Quality of clinical treatment plans can vary drastically (Berry et al., 2016, Giglioli et al., 2016, Marino et al., 2015, Nelms et al., 2012), for example, depending on the skills and ambition of the planner, the complexity of the case, and the time available for planning. Automated treatment planning can be used to improve the quality and consistency of treatment plans (e.g. Breedveld et al. 2019b, Fogliata et al. 2014, Giglioli et al. 2020, Hansen et al. 2016, 2017, Heijmen et al. 2018, Hussein et al. 2016, 2018, Marrazzo et al. 2019, Oud et al. 2020, Purdie et al. 2014, Tol et al. 2015, Zarepisheh et al. 2019), and can also substantially reduce the treatment planning workload.

In our center, Erasmus-iCycle has been developed for automated multi-criterial optimization (MCO) of IMRT fluence profiles and beam angles (FMO + BAO) (Breedveld et al., 2012).   Erasmus-iCycle automatically generates a single Pareto-optimal radiotherapy treatment plan. As it only optimizes pencil-beam intensities, the system was originally integrated with the commercial Monaco TPS (Elekta AB, Stockholm, Sweden) to convert generated plans into clinically deliverable plans for C-arm linacs (Voet et al., 2013a, 2014).

In this paper, we propose and validate a novel, fully automated treatment planning solution for a CyberKnife® equipped with the InCise™ 2 MLC (Accuray Inc., Sunnyvale, USA). Plan generation is performed fully outside the clinical treatment planning system, using Erasmus-iCycle for pencil-beam based FMO and BAO (Breedveld et al., 2012, Voet et al., 2013a, 2014), followed by MLC segment generation aimed at close reproduction of the pencil-beam optimized 3D dose distributions. The applied segmentation algorithm is fully compatible with all characteristics of the InCise™ 2 MLC (Schipaanboord et al., 2019c, 2020). All pencil-beam, MLC segment, and final dose distributions were calculated with a stand-alone version of the commercial dose calculation engine. As automated plan generation includes BAO, the plans are denoted 'AUTO BAO' in the remainder of the paper. Generated BAO plans could in principle be delivered on a CyberKnife®, as the commercial dose calculation is used and the InCise™ 2 MLC is modeled accordingly.   However, the applied software has no FDA clearance, and also according to the recently issued MDR (Medical Device Regulations), the system can currently not be applied for clinical treatment.

The main aim of this paper was to develop a new automated treatment planning pipeline, independent of the CyberKnife® supplier, for generation of deliverable plans and to evaluate whether it could in principle replace the current manual planning, with the well-known plan quality issues of the latter and involved workload. Apart from dosimetric plan quality also delivery efficiency (Monitor

Table 4.1: Clinical dose-volume constraints

Structure	Constraint	
PTV	D_{max}	≤ 62.5 Gy
Urethra	$D_{5\%}$	≤ 45.5 Gy
	$D_{10\%}$	≤ 42 Gy
	$D_{50\%}$	≤ 40 Gy
Rectum	$V_{32.3Gy}$	≤ 1.0 cc (or 1.2 cc)
	D_{max}	≤ 38 Gy
Rectum mucosa	D_{max}	≤ 28.5 Gy
Bladder	V_{38Gy}	≤ 1.0 cc (or 1.5 cc)
	D_{max}	≤ 41.8 Gy
Femoral heads	D_{max}	≤ 24 Gy

Units (MU), number of beams, number of segments) and delivery times were evaluated.

For validation, the novel autoplanning workflow was first configured for prostate SBRT. For a group of 33 prostate SBRT patients, their AUTO BAO plan was compared to a reference plan ('REF') that was manually generated with the commercial TPS. Each REF plan was also compared to a corresponding plan with automated FMO for the (fixed, patient-specific) beam angles in the REF plan ('AUTO RB': AUTO Reference Beams).

4.2   Materials & Methods

4.2.1   Patient data and planning protocol

In this study, contoured planning CT-scans of 33 patients treated with robotic radiotherapy for low- to intermediate-stage prostate cancer were used. Patients were scanned in head-first, supine position, with an average slice thickness of 1.55 mm [min: 1.50 mm, max: 3.00 mm], and an average pixel spacing of 0.97 mm [min: 0.80 mm, max: 0.98 mm] in X and Y directions. The patients were irradiated with a hypofractionated SBRT protocol, delivering 38 Gy in 4 fractions, and featuring highly heterogeneous PTV dose distributions (mimicking HDR brachytherapy) (Aluwini et al., 2010, 2013). For the PTV, a uniform volume expansion of 3 mm of the CTV was used. The average PTV volume was 70.2 cm^3 (41.7 - 128.5 cm^3). The PTV coverage objective was defined as 95% of the PTV volume should receive the prescribed dose. The clinical dose-volume constraints for this protocol are listed in table 4.1. The intention for rectum and bladder was to keep the near-maximum doses (D_{1cc}) below 32.3 Gy and 38 Gy, respectively. However, when considered infeasible, the 1 cc constraint could be relaxed to 1.2 cc and 1.5 cc for rectum and bladder, respectively.

4.2.2   Reference plans (REF)

Attention was paid to using high-quality, manually-generated REF plans for validation of the new autoplanning pipeline. All REF plans were generated by a single experienced medical physicist, using manual iterative trial-and-error planning with the VOLO™ optimizer as introduced in Precision v. 2.0.0.0. (Accuray Inc., Sunnyvale, USA), in the absence of time pressure. Prior to generation of final plans, the physicist spent ample time (in the order of weeks) to develop a strategy for efficient generation of acceptable, high-quality REF plans.   The option to pre-select a randomized and spatially distributed subset of nodes prior to optimization was not used for the generation of the REF plans.   Instead, optimization started with all possible nodes available in the prostate robot motion path to use all available degrees of freedom for obtaining highest quality REF plans.

The REF plans were originally generated for planning study to validate the new VOLO™ optimizer by comparisons with clinical plans generated with the sequential optimization approach, both implemented in the commercial TPS. This study, by Giżyńska et al. 2021, showed that the REF plans were highly superior to the clinical plans, both in terms of dosimetric plan quality as in plan deliverability. Also, several other studies reported enhanced plan quality when using the VOLO™ optimizer instead of the sequential optimization approach (Calusi et al., 2020, Schüler et al., 2020, Zeverino et al., 2019).

4.2.3   AUTO plans

In contrast to manual Pareto navigation based MCO, Erasmus-iCycle based MCO entails automated generation of a single Pareto-optimal plan for each patient. Plan generation is based on a planning protocol specific 'wish-list' that is used for all patients treated according to the protocol. The wish-list contains hard planning constraints and prioritized planning objectives (Breedveld et al., 2012). For this study, a dedicated wish-list was constructed for the clinical prostate SBRT planning protocol, considering the planning constraints in table 4.1. The wish-list can be found in Appendix A. Erasmus-iCycle can handle DVH constraints directly. Originally the DVH constraints could only be used by approximation (Breedveld et al., 2017), but more recently also with high accuracy (Breedveld et al., 2019a). Due to the induced complexity of using DVH criteria, only hard clinical DVH constraints are included in the FMO wish-list. For the wish-list configuration, five patients (out of the 33 included patients) were used for training, and five extra patients for testing (fine-tuning of the wish-list).   In this study, BAO meant generation of patient-specific 25-beam configurations.

Both FMO and segmentation were performed using a pencil-beam resolution of

3 mm in the direction of the leaves and 3.85 mm perpendicular to the leaves at 800 mm SAD (leaf width of the InCise™ 2 MLC). Pencil-beam and segment dose depositions were calculated with a stand-alone version of CyberKnife's dose engine, provided by Accuray Inc. For the segmentation phase of this study, a MU penalty weight equal to 3 (Eq. 2 in Schipaanboord et al. 2019c) was used. A degradation tolerance of 0.25% per objective (table 4.1) was used for the segment reduction method, which iteratively removes low contribution segments after the segmentation as long as the objective degradation is within the specified tolerance. The segmentation algorithm does not necessarily utilize all candidate beams provided for FMO, therefore the final deliverable AUTO RB plans may have fewer beams than the provided candidate beamset of the corresponding REF plan.

4.2.4   Plan evaluation and comparison

A PTV coverage of 95% was aimed for, however, for some patients this was not feasible due to limiting OAR constraints (table 4.1). Prior to comparing treatment planning strategies (REF vs. AUTO BAO and AUTO RB), all three plans of a patient were normalized to exactly the same PTV coverage to minimize bias in dose delivery comparisons for healthy tissues, generally 95%. If 95% coverage was not feasible due to limiting OAR constraints for one or more treatment plans of a patient, all plans for that patient were normalized to the plan with the lowest PTV coverage to avoid inducing OAR constraint violations by normalizing to a higher PTV coverage.

AUTO plans were compared with REF plans using dosimetric plan parameters applied in clinical practice, Dose Volume Histograms (DVHs), visual inspection of the dose distributions, Conformation Number (van 't Riet et al., 1997), numbers of beam directions, numbers of MLC segments, numbers of MU per fraction, and estimated treatment delivery times, calculated with a stand-alone treatment time estimator provided by Accuray Inc. The estimated treatment time (ETT) includes beam-on time, robot movements, changing of apertures, and imaging, while excluding patient setup time.

Wilcoxon signed-rank tests for paired data were performed to assess statistical significance (p-value < 0.05) of differences between AUTO plans and manually generated REF plans.

4.2.5   Computation times

Computation times for AUTO plans were measured for 10 patients on an Intel Xeon Gold 6248 @ 2.5 GHz, containing 40 cores and with 386 GB of memory.
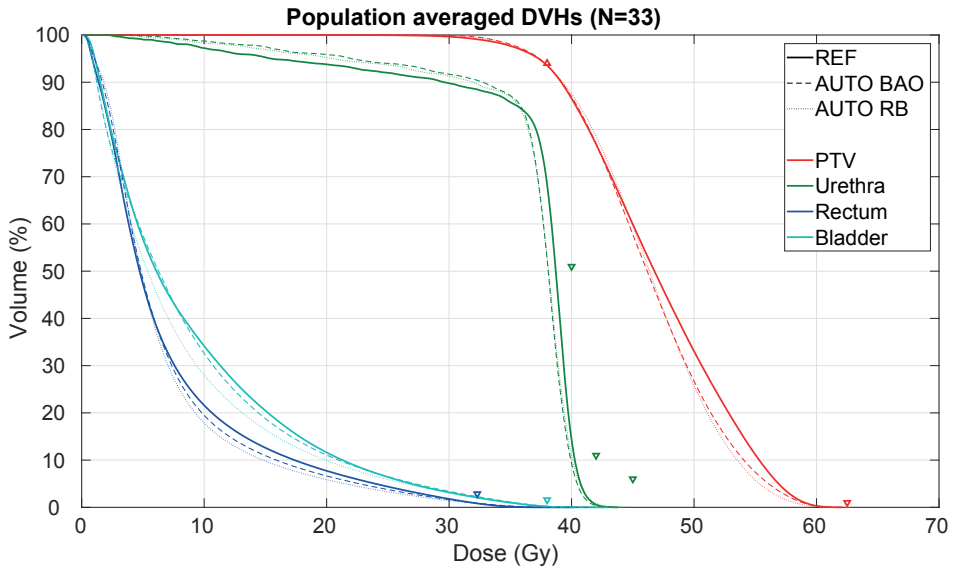
Figure 4.1: Population averaged PTV, rectum, bladder, and urethra Dose Volume Histograms (DVHs) for the reference plans (REF), the AUTO plans with beam angle optimization (AUTO BAO), and the AUTO plans with the reference beam geometry (AUTO RB). Clinical constraints are denoted with triangles, see table 4.1.

4.3   Results

4.3.1   Plan comparisons

The population averaged DVHs in figure 4.1 show high similarity for the three planning approaches, with small advantages for autoplanning compared to REF for the higher urethra doses and in the intermediate dose range for rectum and bladder (especially for AUTO RB). An example dose distribution for REF and AUTO BAO is shown in figure 4.3.1.

The dosimetric plan parameter comparisons presented in figure 4.3 confirm the overall similarity between AUTO plans and REF plans. Depending on the parameter, small overall advantages for AUTO plans or REF plans were seen. Some patients demonstrate differences that could possibly be clinically relevant, sometimes in favor of AUTO, sometimes in favor of REF. The upper parts of table 4.2 and 4.3 present overviews of differences between dosimetric plan parameters in REF plans and AUTO BAO plans (table 4.2) or AUTO RB plans (table 4.3). Although many of the differences in dosimetric plan parameters are statistically significant, they are small from the clinical point of view, sometimes in favor of AUTO and for other parameters in favor of REF.

Results for non-dosimetric parameters are presented in the bottom sections of
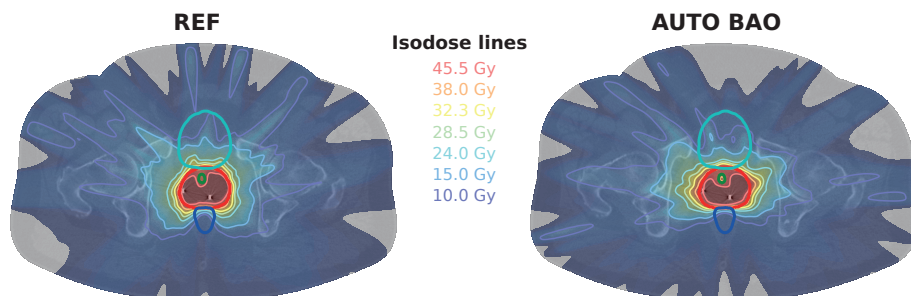
Figure 4.2: An example dose distribution for REF (left) and AUTO BAO (right). Depicted structures: PTV (red), Urethra (green), Rectum (dark blue), and Bladder (light blue).

tables 4.2 and 4.3 and in figure 4.4. A clinically relevant reduction in estimated treatment time was observed for the AUTO plans compared to the REF (AUTO BAO: -5.3 min. [-15.6, 1.0], AUTO RB: -3.7 min. [-11.7, 4.0]). This was related to reductions in the numbers of beams (AUTO BAO: -23.6 [-34, -12], AUTO RB: -10.4 [-23, -2]) and the numbers of segments (AUTO BAO: -39.4 [-115, -12], AUTO RB: -37.5 [-118, -37]), whereas no significant differences for the number of MU per fraction were observed.

4.3.2   Computation times

Table 4.4 shows autoplanning mean computations times with ranges, divided in 1) generation of the pencil-beam matrices together with the optimization of the FMO dose distributions (PB + FMO) and, 2) the subsequent segmentation of the FMO dose distributions (Segmentation).

4.4   Discussion

In this study, we have proposed a fully automated treatment planning workflow for a robotic CyberKnife® unit equipped with the InCise™ 2 MLC, and validated it for prostate SBRT by comparison of generated AUTO BAO plans with high-quality, manually generated reference plans (REF). The AUTO BAO plans are optimized fully independently of the commercial TPS, using in-house developed applications for integrated multi-criterial pencil-beam-based beam angle and fluence map optimization (Breedveld et al., 2012, Voet et al., 2013a, 2014), and subsequent generation of MLC segments (Schipaanboord et al., 2019c, 2020). Although practically not yet possible because of plan import restrictions, the plans are in principle deliverable at clinical CyberKnife® units. The AUTO BAO prostate SBRT
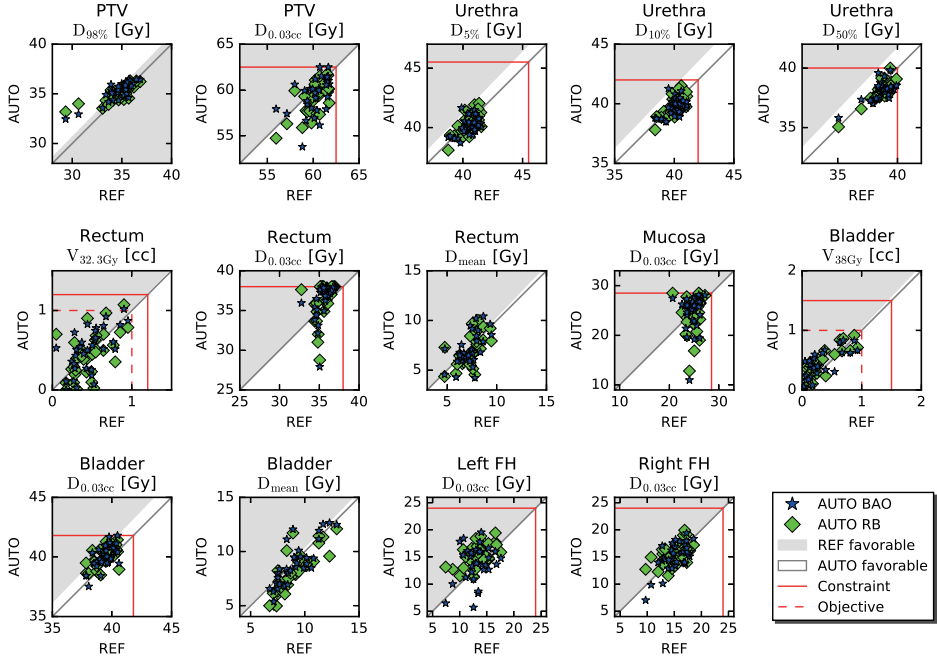
Figure 4.3: Comparisons of REF, AUTO BAO, and AUTO RB regarding dosimetric plan parameters. Every marker represents a plan parameter comparison for one of the 33 study patients. Red lines show treatment planning aims. FH: Femoral head. See section 4.2.1 for differences between constraint and objective levels.
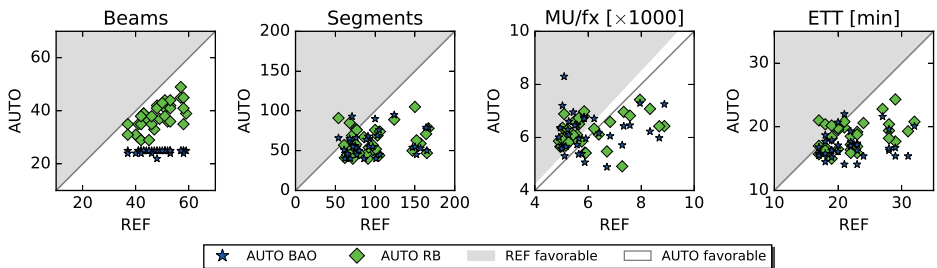
Figure 4.4: Non-dosimetric plan parameters of the AUTO plans compared to corresponding REF plan parameters. Every marker represents a plan parameter comparison for one of the 33 study patients. MU/fx: Monitor Units per fraction, ETT: Estimated treatment time.

Table 4.2: Plan parameter comparisons for AUTO BAO and REF plans. Doses (Dxx) are presented in Gy and volume (Vxx) in cc. Entrance dose was evaluated using D_{1cc} of a ring structure of 3 cm thickness inside the patient's external contour. FH: Femoral head, MU/fx: Monitor Units per fraction, ETT: Estimated treatment time.

		REF		AUTO BAO		AUTO BAO - REF		
		Mean	(min, max)	Mean	(min, max)	Δ	(min, max)	p
PTV	$D_{98\%}$	**34.7**	(29.3, 36.8)	**35.3**	(32.5, 36.5)	**0.6**	(-1.1, 3.1)	**<.001**
	$D_{0.03cc}$	**60.2**	(56.0, 61.8)	**59.5**	(53.8, 62.5)	**-0.7**	(-5.0, 2.6)	**.03**
Urethra	$D_{5\%}$	**40.5**	(38.8, 41.5)	**40.1**	(38.7, 41.7)	**-0.4**	(-1.5, 1.8)	**.005**
	$D_{10\%}$	**40.0**	(38.4, 41.1)	**39.7**	(38.5, 41.3)	**-0.4**	(-1.5, 1.7)	**.005**
	$D_{50\%}$	**38.6**	(35.1, 39.9)	**38.1**	(35.8, 39.8)	**-0.6**	(-1.8, 1.2)	**<.001**
Rectum	$V_{32.3Gy}$	**0.5**	(0.0, 1.0)	**0.5**	(0.0, 1.0)	**0.0**	(-0.5, 0.5)	.5
	$D_{0.03cc}$	**35.5**	(32.8, 37.0)	**36.4**	(27.9, 38.0)	**0.8**	(-7.2, 3.2)	**<.001**
	D_{mean}	**7.2**	(4.7, 9.5)	**7.1**	(4.2, 10.4)	**-0.2**	(-3.6, 2.4)	.4
Mucosa	$D_{0.03cc}$	**24.8**	(20.7, 27.2)	**24.7**	(11.0, 28.1)	**-0.0**	(-13.1, 6.8)	.3
Bladder	V_{38Gy}	**0.3**	(0.0, 0.9)	**0.4**	(0.0, 0.8)	**0.1**	(-0.3, 0.4)	**.001**
	$D_{0.03cc}$	**39.4**	(37.8, 40.6)	**40.0**	(37.5, 41.8)	**0.6**	(-1.1, 2.0)	**<.001**
	D_{mean}	**9.1**	(6.8, 12.9)	**8.9**	(5.4, 12.6)	**-0.2**	(-2.1, 3.1)	.1
Entrance	D_{1cc}	**17.3**	(14.2, 19.9)	**18.0**	(15.3, 20.1)	**0.7**	(-2.8, 4.2)	**.02**
Left FH	$D_{0.03cc}$	**13.2**	(7.4, 17.6)	**13.9**	(5.7, 19.6)	**0.7**	(-6.9, 7.8)	.2
Right FH	$D_{0.03cc}$	**15.1**	(9.7, 18.6)	**15.0**	(7.0, 19.4)	**-0.1**	(-3.3, 4.3)	.8
Conformality		**0.82**	(0.75, 0.88)	**0.79**	(0.75, 0.87)	**-0.03**	(-0.07, 0.02)	**<.001**
Beams		**48.3**	(37.0, 59.0)	**24.7**	(22.0, 25.0)	**-23.6**	(-34.0, -12.0)	**<.001**
Segments		**98.9**	(54.0, 167.0)	**59.5**	(40.0, 95.0)	**-39.4**	(-115.0, 22.0)	**<.001**
MU/fx [×1000]		**6.1**	(4.9, 8.9)	**6.2**	(4.9, 8.3)	**0.1**	(-2.7, 3.2)	.5
ETT [min]		**22.3**	(17.0, 32.0)	**17.1**	(14.1, 22.0)	**-5.3**	(-15.6, 1.0)	**<.001**

plans were clinically acceptable with a quality equivalent to the REF plans. However, as no manual fine-tuning of AUTO plans was needed, the quality of these plans was independent of manual planners and the workload was virtually zero. AUTO BAO also had shorter delivery times (23% reduction, 5.3 min.) and vastly reduced numbers of beams (24.7 vs. 48.3) and segments (59.5 vs. 98.9).

To the best of our knowledge, this is the first fully-independent system for automated generation of deliverable plans, that combines automated non-coplanar BAO, FMO, and segmentation all together. Automated generation of clinically deliverable plans was developed for other systems, but final plan generation was always performed using the commercial TPS. Or, deliverable plans were created but the planning workflow required manual tweaking during plan optimization or did not include non-coplanar BAO. For example, Erasmus-iCycle was integrated with Monaco (Elekta AB, Stockholm, Sweden) (Voet et al., 2013a, 2014) for treatment with C-arm linacs, and to MultiPlan (Accuray Inc., Sunnyvale, USA) (Rossi et al., 2018) for treatment with a CyberKnife® in combination with the IRIScollimator. The 4π planning approach relies on importing optimized beam angles into Eclipse (Varian Medical System, Palo Alto, USA), followed by conventional treatment planning (Murzin et al., 2018, Tran et al., 2017, Woods et al., 2016, Yu et al., 2018).

Table 4.3: Plan parameter comparisons for AUTO RB and REF plans. Doses (Dxx) presented in Gy and volumes (Vxx) in cc. Entrance dose was evaluated using D_{1cc} of a ring structure of 3 cm thickness inside the patient's external contour. FH: Femoral head, MU/fx: Monitor Units per fraction, ETT: Estimated treatment time.

		REF		AUTO RB			AUTO RB - REF	
		Mean	(min, max)	Mean	(min, max)	Δ	(min, max)	p
PTV	$D_{98\%}$	34.7	(29.3, 36.8)	35.1	(33.2, 36.3)	0.4	(-1.1, 3.8)	.02
	$D_{0.03cc}$	60.2	(56.0, 61.8)	59.1	(54.7, 61.8)	-1.1	(-3.4, 1.9)	<.001
Urethra	$D_{5\%}$	40.5	(38.8, 41.5)	40.3	(38.1, 42.0)	-0.2	(-1.5, 1.3)	.1
	$D_{10\%}$	40.0	(38.4, 41.1)	39.8	(37.8, 41.5)	-0.2	(-1.2, 1.0)	.03
	$D_{50\%}$	38.6	(35.1, 39.9)	38.1	(35.1, 40.0)	-0.6	(-1.4, 0.6)	<.001
Rectum	$V_{32.3Gy}$	0.5	(0.0, 1.0)	0.4	(0.0, 1.1)	-0.0	(-0.5, 0.7)	.3
	$D_{0.03cc}$	35.5	(32.8, 37.0)	36.4	(28.8, 38.0)	0.8	(-6.3, 4.9)	.004
	D_{mean}	7.2	(4.7, 9.5)	6.8	(4.4, 10.0)	-0.4	(-3.3, 2.5)	.03
Mucosa	$D_{0.03cc}$	24.8	(20.7, 27.2)	24.0	(12.8, 28.5)	-0.8	(-11.2, 7.8)	.3
Bladder	V_{38Gy}	0.3	(0.0, 0.9)	0.4	(0.1, 0.9)	0.1	(-0.2, 0.3)	<.001
	$D_{0.03cc}$	39.4	(37.8, 40.6)	40.1	(38.5, 41.4)	0.7	(-1.6, 2.2)	<.001
	D_{mean}	9.1	(6.8, 12.9)	8.2	(5.0, 12.1)	-0.9	(-2.7, 2.8)	<.001
Entrance	D_{1cc}	17.3	(14.2, 19.9)	17.2	(12.9, 21.4)	-0.0	(-4.4, 3.4)	.9
Left FH	$D_{0.03cc}$	13.2	(7.4, 17.6)	15.4	(11.3, 19.4)	2.2	(-1.7, 5.7)	<.001
Right FH	$D_{0.03cc}$	15.1	(9.7, 18.6)	15.6	(11.5, 19.9)	0.5	(-3.3, 4.2)	.2
Conformality		0.82	(0.75, 0.88)	0.78	(0.72, 0.84)	-0.04	(-0.07, -0.01)	<.001
Beams		48.3	(37.0, 59.0)	38.0	(29.0, 49.0)	-10.4	(-23.0, -2.0)	<.001
Segments		98.9	(54.0, 167.0)	61.5	(40.0, 105.0)	-37.5	(-118.0, 37.0)	<.001
MU/fx [×1000]		6.1	(4.9, 8.9)	6.2	(4.7, 7.4)	0.1	(-2.5, 1.8)	.2
ETT [min]		22.3	(17.0, 32.0)	18.6	(15.0, 24.3)	-3.7	(-11.7, 4.0)	<.001

Table 4.4: Planning and computation times in minutes calculated over the first 10 patients. The reported REF planning time is manual hands-on time, while the reported AUTO times are fully automated calculation times without manual interaction. PB + FMO: Calculation of pencil-beams + Fluence Map Optimization.

	PB + FMO		Segmentation		Total	
Approach	Mean	(min, max)	Mean	(min, max)	Mean	(min, max)
REF (manual)					50	(10, 170)
AUTO BAO	621	(529, 742)	16	(13, 19)	637	(542, 761)
AUTO RB	80	(53, 169)	18	(14, 24)	98	(67, 193)

The Expedited Constrained Hierarchical Optimization (ECHO) system also uses Eclipse to generate a final plan, the optimized fluence maps are imported and then leaf sequencing is performed within Eclipse (Zarepisheh et al., 2019). The ASEQ method used for online replanning creates deliverable plans with the use of the Monte Carlo dose engine (GPUMCD) from Elekta AB, but requires manual tweaking during plan optimization and uses predefined clinical beam configurations (Kontaxis et al., 2015a,b, 2017a,b).

The use of column-generation in radiotherapy treatment planning was first proposed by Romeijn et al. 2005 to solve the Direct Aperture Optimization (DAO) problem. The 4π planning approach (Dong et al., 2013a,b) also uses a

column-generation approach, based on the formulation proposed by Romeijn et al., but uses column-generation to solve the BAO/FMO problem. In the proposed autoplanning workflow for CyberKnife® with MLC, the BAO/FMO problem is solved using a multi-criterial optimization as implemented in Erasmus-iCycle (Breedveld et al., 2012). Then, a column-generation method, inspired by the approach by Romeijn et al., was used to solve the segmentation problem by mimicking the 3D FMO dose distribution (Schipaanboord et al., 2019c, 2020). This two-step approach of BAO/FMO followed by segmentation turns out to be more flexible for the multi-criteria optimization.

Computation times for AUTO BAO plans were on average 637 min. While this may seem long, important to realize is that no manual hands-on planning or manual corrections of the computed plans were applied. In a clinical context, this would mean that a deliverable plan is ready within a day of contour approval by the treating physician. In our clinical practice, this has always been a requirement for application of Erasmus-iCycle based automated plan generation for C-arm linacs (although current plan generation is much faster). Contributing most to the computation time was the applied integrated, iterative BAO which is computationally expensive for large numbers of beams. This BAO approach was chosen because it has been shown to provide good quality treatment plans in previous studies (Breedveld et al., 2012, Rossi et al., 2018, 2021, Sharfo et al., 2021, Voet et al., 2012). Alternatively, the proposed workflow could also be combined with a BAO approach that selects beam angles prior to FMO optimization. This would avoid the need for multiple FMO iterations, which would reduce the computation time. Or, instead of the patient-specific beam angle optimization, a pre-defined set of beam angles (class-solution) could be used for all patients. For example, the non-coplanar beam angle class-solution proposed by Rossi et al. 2015 Furthermore, our current research implementation utilizes serial calculation of the pencil-beams and the final dose calculations (included in both FMO and segmentation). Parallelization of these calculation steps could reduce the computation time substantially.

The presented comparison of AUTO RB plans with REF plans allowed us to compare autoplanning with manual planning without a bias of different beam angles; the AUTO RB plans indicate what performance could be achieved with the CyberKnife® system for each patient when using a different optimization approach for the same input beam angles (as used in the REF plan). As shown in table 4.3, not all input angles were always used in the final AUTO RB plans (on average 38.0 of the 48.3 beams in REF). Nevertheless, quality of the AUTO RB plans was similar to the REF plan quality. Apart from the reduced number of beams, this was obtained with also a substantially lower number of segments and a significantly reduced delivery

time (table 4.3). In the intermediate dose range, AUTO RB plans are slightly favorable over AUTO BAO plans (figure 4.1). This is attributed to the lower numbers of beams used in AUTO BAO (24.7 vs. 38.0). For AUTO BAO our aim was to obtain a quality that was comparable to REF. Going to e.g. 38 beams in AUTO BAO would have further enhanced the quality of the AUTO BAO plans, but at the cost of large increases in calculation time.

In this study, we have introduced a novel automated treatment planning pipeline for CyberKnife® SBRT and we validated it for prostate cancer. In an ongoing study, the new workflow is being investigated for lung SBRT, another type of treatment that is frequently performed with the CyberKnife®.

Recently, the RATING framework with guidelines for performing high-quality treatment planning studies has been published (Hansen et al., 2020). There is also a score sheet attached to the framework to get a quantitative impression on the quality of treatment planning papers. According this sheet, our study scored 92/100%. The filled-out sheet is provided in Appendix B.

4.5   Conclusions

A new vendor-independent workflow for fully automated generation of deliverable, high-quality CyberKnife plans was proposed, including patient-specific beam angle optimization (BAO). Compared to manual planning with the commercial TPS in absence of time pressure, dosimetric plan quality for prostate SBRT was similar, while fraction delivery times reduced by 5.3 min (from 22.3 to 17.1 min) due to large reductions in beam and segment numbers.

Acknowledgments

Conflict of interest

The authors have no relevant conflicts of interest to disclose.

Appendix A

Table A1: The wish-list used to automatically generate the FMO plans using Erasmus-iCycle, containing a list of constraints and a list of prioritized planning objectives. The optimization starts with optimizing on the PTV to attain sufficient PTV coverage, while respecting all constraints from the constraints list. Once an objective is optimized, it is added as an additional constraint to the constraint list and then the next objective will be optimized. The relative ranking of the objectives is based on clinical importance of the structures. A detailed description of the Erasmus-iCycle workflow is provided in Breedveld et al., 2012 (Breedveld et al., 2012). UrethraPlan is defined as the volume of the urethra within the PTV.

Constraints

Priority	Volume	Dose metric	Limit
	PTV	Dmax	61.5 Gy
	UrethraPlan	D5%	45 Gy
		D10%	42 Gy
		Dmean	40 Gy
	Rectum	Dmax	38 Gy
		D1cc	32.3 Gy
	Rectum mucosa	Dmax	27 Gy
	Bladder	Dmax	41.8 Gy
		D1cc	38 Gy
	Penile bulb	Dmax	1.5 Gy
	Ring PTV 2cm - 3cm	Dmax	25 Gy
	Skin dose	Dmax	20 Gy
	Shell 3mm	Dmax	38 Gy
	Shell 3cm	Dmax	20 Gy
	Shell 5cm	Dmax	20 Gy

Objectives

Priority	Volume	Dose metric	Goal	Sufficient	Parameters
1	PTV - OARs	↓ LTCP	0.01	0.01	D^p = 37 Gy, α = 0.90
2	PTV	↓ LTCP	0.25	0.25	D^p = 37 Gy, α = 0.60
3	UrethraPlan	↓ Dmean	39 Gy	38 Gy	
4		↓ LTCP	0.50		D^p = 40 Gy, α = -0.50
5	Rectum	↓ LTCP	0.00		D^p = 27 Gy, α = -0.20
6	Bladder	↓ LTCP	0.00		D^p = 31 Gy, α = -0.20
7	Dose bath	↓ Dmax	15 Gy		
8	Femoral heads	↓ Dmax	24 Gy		

Appendix B

Filled-out RATING score sheet can be found online.

TBS-BAO: Fully automated beam angle optimization for IMRT guided by a Total-Beam-Space reference plan

B.W.K. Schipaanboord[1], B.J.M. Heijmen[1], S. Breedveld[1]

[1] Department of Radiotherapy, Erasmus MC Cancer Institute, Rotterdam, the Netherlands

Abstract

Properly selected beam angles contribute to the quality of radiotherapy treatment plans. However, the Beam Angle Optimization (BAO) problem is difficult to solve to optimality due to its non-convex discrete nature with many local minima. In this study, we propose TBS-BAO, a novel approach for solving the BAO problem, and test it for non-coplanar robotic CyberKnife radiotherapy for prostate cancer. First, an ideal Pareto-optimal reference dose distribution is automatically generated using a priori multi-criterial fluence map optimization (FMO) to generate a plan that includes all candidate beams (total-beam-space, TBS). Then, this ideal dose distribution is reproduced as closely as possible in a subsequent segmentation/beam angle optimization step (SEG/BAO), while limiting the number of allowed beams to a user-selectable preset value. SEG/BAO aims at a close reproduction of the ideal dose distribution. For each of 33 prostate SBRT patients, 18 treatment plans with different pre-set numbers of allowed beams were automatically generated with the proposed TBS-BAO. For each patient, the TBS-BAO plans were then compared to a plan that was automatically generated with an alternative BAO method (Erasmus-iCycle) and to a high-quality manually generated plan. TBS-BAO was able to automatically generate plans with clinically feasible numbers of beams ($\sim$25), with a quality highly similar to corresponding 91-beam ideal reference plans. Compared to the alternative Erasmus-iCycle BAO approach, similar plan quality was obtained for 25-beam segmented plans, while computation times were reduced from 10.7 hours to 4.8/1.5 hours, depending on the applied pencil-beam resolution in TBS-BAO. 25-beam TBS-BAO plans had similar quality as manually generated plans with on average 48 beams, while delivery times reduced from 22.3 to 18.4/18.1 min. TBS reference plans could effectively steer the discrete non-convex BAO.

5

5.1   Introduction

An important aspect contributing to radiotherapy treatment plan quality is the use of appropriate beam angles.   Especially, for non-coplanar Stereotactic Body RadioTherapy (SBRT), in which high doses are delivered in a limited number of fractions, selection of the appropriate beam angles is crucial for high-quality treatment (Bedford et al., 2019b, Dong et al., 2013a,b, Rossi et al., 2012, 2018, 2021, Rwigema et al., 2015, Sharfo et al., 2017, 2021, Woods et al., 2016, Yu et al., 2018). The Beam Angle Optimization (BAO) problem is formally a discrete, combinatorial optimization problem that is difficult to solve to optimality (Bangert et al., 2012), since the optimization space is highly non-convex with many local minima (Craft, 2007, Södertröm and Brahme, 1993).

Various BAO approaches have been investigated.   For example, sequential addition of new beam angles by ranking candidate beams (Bangert and Unkelbach, 2016, Breedveld et al., 2012). The 4π approach, proposed by Dong et al. 2013b, uses a computer-assisted design (CAD) model to eliminate candidate beams that may cause collisions between the gantry and the couch or patient. Then, the beams are iteratively selected from the set of remaining candidate beams by predicting the benefit of each beam using the first-order information and selecting the beam with the lowest predicted objective function value. Furthermore, hybrid approaches that combine heuristics to explore the optimization space for finding the global minimum with local gradient-based search algorithms, have been explored. For example, Aleman et al. 2008 combined simulated annealing with a deterministic local search algorithm. Bertsimas et al. 2013 combined simulated annealing with gradient information to find the local minimum. Rocha et al. 2016 proposed a derivative-free multi-start framework with a pattern search algorithm, and Bedford et al. 2019b proposed an evolutionary algorithm to select non-coplanar beam angle sets for the CyberKnife. Amit et al. 2015 proposed a model-based approach based on the relationship between beam angles and anatomical features. Lately, also deep-learning-based BAO has been proposed (Sadeghnejad-Barkousaraie et al., 2020).

In this work, we introduce the novel TBS-BAO approach for solving the BAO problem. This approach integrates BAO with segmentation, rather than solving the BAO problem prior to or during the FMO phase. TBS-BAO first generates an 'ideal' reference dose distribution using all available candidate beams (total-beam-space, TBS). This plan is generated with a priori multi-criterial FMO, providing for each patient a high-quality Pareto-optimal plan with clinically favorable trade-offs, and represents the ideal plan quality for a patient that could be achieved when including all beam directions and when no delivery restrictions are taken into

account. This ideal FMO dose distribution is then used to guide the following IMRT segmentation with integrated BAO (SEG/BAO) for constructing a deliverable plan with a preset maximum number of beam directions. For SEG/BAO, the segmentation approach presented in Schipaanboord et al. 2019c, 2020 with some modifications is used as a starting point. MLC segments are sequentially added to the plan; each segment is added to the beam in the TBS that minimizes the dosimetric differences between the segmented plan and the ideal reference plan. The purpose of this study is to investigate to what extent the proposed integrated non-convex SEG/BAO can closely reproduce ideal reference plans. To investigate this, TBS-BAO prostate SBRT plans with variable pre-selected beam numbers were compared to plans generated with a previously published alternative BAO approach (Breedveld et al., 2012) and with high-quality manually generated plans.

5.2 Materials & Methods

5.2.1 Patient data

In this study, contoured CT scans of 33 prostate cancer patients, previously treated in our institution with the CyberKnife robotic treatment unit, were used. The patients were treated with a hypofractionated SBRT protocol, delivering 38 Gy in 4 fractions, and featuring highly heterogeneous PTV dose distributions (Aluwini et al., 2010, 2013). The average PTV volume was 70.2 cm^3 (41.7 - 128.5 cm^3). The PTV coverage objective was defined as 95% of the PTV volume should receive the prescribed dose. However, for some patients, the requested PTV coverage was not feasible due to limiting OAR constraints. If this was the case, the PTV coverage was compromised to respect the OAR constraints, in line with clinical practice. Additionally, if a plan featured a PTV coverage larger than 95%, the plan was re-normalized to exactly 95% PTV coverage to maximize OAR sparing, as is also done in clinical practice. The clinically used PTV objective and the applied dose-volume constraints are listed in table 5.1. This patient cohort was also used in previous studies Giżyńska et al. 2021, Schipaanboord et al. 2021.

5.2.2 Generation of ideal total-beam-space (TBS) reference plans

Ideal total-beam-space reference plans were fully automatically generated with our in-house developed Erasmus-iCycle optimizer for a priori Multi-Criterial plan Optimization (MCO) (Breedveld et al., 2009, 2012, 2017, 2019b). With Erasmus-iCycle, a single Pareto-optimal FMO plan is automatically generated for each patient. With proper configuration of the so-called wish-list with hard constraints and prioritized objectives, the generated Pareto-optimal plan also has

Table 5.1: Tumor objective and clinical dose-volume constraints

Structure	Objective / Constraint	
PTV	V_{38Gy}	≥ 95% (objective)
	D_{max}	≤ 62.5 Gy
Urethra	$D_{5\%}$	≤ 45.5 Gy
	$D_{10\%}$	≤ 42 Gy
	$D_{50\%}$	≤ 40 Gy
Rectum	$V_{32.3Gy}$	≤ 1.2 cc
	D_{max}	≤ 38 Gy
Rectum mucosa	D_{max}	≤ 28.5 Gy
Bladder	V_{38Gy}	≤ 1.5 cc
	D_{max}	≤ 41.8 Gy
Femoral heads	D_{max}	≤ 24 Gy

clinically favorable trade-offs. Many studies have demonstrated superiority of automatically generated plans over conventional manual planning (Bijman et al., 2020, 2021, Breedveld et al., 2012, Buschmann et al., 2018, Heijmen et al., 2018, Redapi et al., 2021, Rossi et al., 2018, 2021, Sharfo et al., 2016, 2015, 2017, 2018, 2021, Voet et al., 2012, 2013a). In contrast, in a posteriori MCO, (e.g. Bokrantz and Miettinen 2015, Craft and Richter 2013) for each patient, a set of Pareto-optimal plans is generated with automated planning, while selection of a clinically favorable plan is performed by a user.

In this study, we used a previously developed wish-list for the generation of the ideal reference plans for this patient cohort (Schipaanboord et al., 2021). FMO was performed using all 91 available CyberKnife candidate beams from the prostate MLC path set (Accuray Inc., 2015). To investigate the balance between calculation speed and plan quality, ideal plans were generated with a high pencil-beam resolution of 3 mm in the direction of the leaves and 3.85 mm perpendicular to the leaves (leaf width of the InCise™ 2 MLC), denoted as 'Ideal HR', and with low resolution pencil-beams (6 x 7.7 mm²), denoted as 'Ideal LR'. Pencil-beams were calculated with a stand-alone version of CyberKnife's dose engine, provided by Accuray Inc.

5.2.3   Integrated MLC segmentation and BAO (SEG/BAO) using column generation

A column generation (CG) approach was used for segmentation of ideal reference plans with integrated BAO. The approach was originally developed for FMO IMRT plans, using pre-selected favorable beam directions (Schipaanboord et al., 2019c), and in this paper extended for TBS-BAO, demanding that the algorithm does not only generate favorable segments for clinically optimal beam directions, but inherently also selects the optimal beam directions from the large total-beam-space, up to the pre-set maximally allowed number of beams.

Our published CG-based segmentation method (Schipaanboord et al., 2019c,

2020) generates MLC segments for FMO plans, while considering the available pre-selected beam angles simultaneously. It uses quadratic dose differences between the reference FMO dose distribution and the segmented dose distribution, supplemented with terms to focus on target coverage and total MU minimization, to define the master problem (MP). The MP is solved by iteratively solving the restricted master problem (RMP), which is restricted to feasible segment shapes only and is solved to obtain the corresponding segment weights. These segment weights (solution of the RMP) are then used to calculate the gradient on the objective function of the MP. In the subsequent pricing problem (PP) phase, a promising segment is identified and added to the RMP. In this phase, the gradient on MP is used to identify promising beamlets that can be grouped into a feasible segment given the mechanical restrictions of the MLC. Feasible candidate segments are generated for all pre-selected candidate beams simultaneously, however, only the most promising segment (the segment with the largest sum of negative gradients) is added to the RMP per CG iteration. Subsequently, the RMP is solved again to obtain the segment weights and the process is repeated until no feasible segment can be found that improves the MP.

The basis for SEG/BAO with CG is that segments are sequentially added to the plan, while considering all available candidate beams simultaneously. In the published papers, the input FMO distribution was generated with limited sets of pre-established favorable beam directions. Therefore, optimality of beam angles was upfront guaranteed. In this study, the ideal FMO reference plan is generated with all 91 candidate beams and the selection of the beams is performed during MLC segmentation instead. A heuristic was integrated into the column generation workflow, to limit the maximum allowed number of beams per plan to avoid clinically unfeasible numbers, while initially being able to select segments from all 91 candidate beams. Once the preset number of beam directions is reached, MLC segments could only be added to the already selected beams. This was achieved by only generating feasible segments during the PP phase for the already selected beam directions, once the limit was reached. The most promising segment to be added to the RMP was then selected from this subset of segments. Additionally, segments that are no longer favorable can be removed during the segmentation process. As a result, beam directions are also dynamically added and removed throughout the segmentation process. If the number of beams in the current solution is reduced to below the preset maximum allowed number of beams, all 91 candidate beams are again available during the PP phase for the identification of the next segment (until the preset number of beam directions is reached again).

Clinical hard constraints were not explicitly enforced during SEG/BAO segment generation. To accommodate such constraints in SEG/BAO, the published

Table 5.2: Wish-list for final weight re-optimization of selected MLC segments. ROV = Residual Objective Value at the end of segmentation (see section 5.2.3).

Constraints

	Volume	Dose metric	Limit
	PTV	D_{max}	≤ 62.5 Gy
	Urethra	$D_{5\%}$	≤ 45.5 Gy
		$D_{10\%}$	≤ 42 Gy
		D_{mean}	≤ 40 Gy
	Rectum	$V_{32.3Gy}$	≤ 1.2 cc
		D_{max}	≤ 38 Gy
	Rectum mucosa	D_{max}	≤ 28.5 Gy
	Bladder	V_{38Gy}	≤ 1.5 cc
		D_{max}	≤ 41.8 Gy
	Femoral heads	D_{max}	≤ 24 Gy
	Skin	D_{max}	≤ 20 Gy
	Dose	Quadratic differences	200% ROV
	MU/segment	linear	≥ 5

Objectives

Priority	Volume	Metric	Goal	Sufficient
1	PTV	↑ V_{38Gy}	95%	95%
2	Dose	↓ Quadratic differences	0	
3	Total MU	↓ Linear	0	

segmentation approach was extended with a segment weight re-optimization step including hard dose-volume constraints as applied in clinical planning (table 5.2). For this weight re-optimization, which was performed after selection of all segments, a fixed dedicated wish-list was used for all plans of all patients (table 5.2). Apart from the dose-volume constraints from table 5.1, additional constraints were added to the objective function: a dose-volume constraint to limit skin dose, a constraint on the quadratic dose difference with the input dose (based on the residual objective value found after segmentation), and a constraint to enforce the minimum required amount of MU per segment. The constraint on the quadratic dose differences aimed at limiting possible degradation of the dose distribution when maximizing the first objective of the wish-list (PTV coverage). For this term, a constraint value of 200% of the objective value at the end of the segment identification was used (prior to the segment weight re-optimization step with hard constraints).

In the remainder of this paper, final segmented patient plans, generated with Ideal HR or LR reference plans for a chosen maximum number of beams are denoted with 'TBS-BAO HR (maximum number of beams)' and 'TBS-BAO LR (maximum number of beams)', respectively. Both ideal reference plans and deliverable plans were fully automatically generated, i.e. no manual fine-tuning of

automatically generated plans was performed. During generation of the TBS-BAO plans, all delivery restrictions of the CyberKnife were taken into account (including the leaf width of 3.85 mm, both for Ideal HR and Ideal LR), and the applied dose calculation engine was the same as used in the clinical TPS. Therefore, generated plans could in principle be clinically delivered. For the integration of the clinical dose engine into the column generation formalism, the fully integrated version of the RMP as proposed in Schipaanboord et al. 2020 was used. In some cases, the final deliverable plan did not utilize the maximum allowed number of beams: either new segments did not further improve the score of the MP (and is thus converged), or no new feasible segment shapes could be identified.

5.2.4   Validation of the proposed novel TBS-BAO approach

For each patient, 9 TBS-BAO HR plans and 9 TBS-BAO LR plans were generated with maximum numbers of final beams from 10 to 50 in steps of 5. The first validation of these plans was based on comparisons with plans that were manually created ('Manual') with the clinical CyberKnife treatment planning system (Precision v2.0) by an experienced medical physicist (Giżyńska et al., 2021, Schipaanboord et al., 2021). Secondly, the TBS-BAO plans were compared to deliverable 25-beam 'iCycle-BAO(25)' plans generated with the Erasmus-iCycle option for integrated beam angle and beam fluence optimization (Breedveld et al., 2012), followed by segmentation with the approach as proposed in (Schipaanboord et al., 2019c, 2020). The Erasmus-iCycle BAO approach has been extensively validated and applied for various treatment sites (Breedveld et al., 2012, Rossi et al., 2018, 2021, Sharfo et al., 2017, 2021, Voet et al., 2012, 2013b). In Erasmus-iCycle, beam angles are sequentially selected. For the selection of each next beam direction, each not yet selected beam direction is evaluated by solving an IMRT optimization problem for a beam arrangement consisting of the already selected beam directions plus the candidate beam direction. Once all candidate beams are evaluated, the most favorable beam direction is added to the treatment plan. This routine is continued until the desired number of beam directions is reached or if further improvement in the patient dose distribution is considered clinically irrelevant.

5.2.5   Plan evaluations and comparisons

Plan evaluations and comparisons were based on obtained PTV coverages, plan parameters in table 5.1, patient body doses, Dose Volume Histograms (DVH), number of included beams, number of MLC segments, Weighted Segment Sizes (WSS (Calusi et al., 2020)), MUs per fraction, Conformation Numbers (van 't Riet et al., 1997), estimated treatment delivery times, and visual inspection of dose

distributions. Delivery times were calculated with a standalone treatment time estimator provided by Accuray Inc, considering beam-on times, robot movements, changing of apertures, and imaging. Patient setup time was not included. To assess statistical significance ($P < 0.05$), Wilcoxon signed-rank tests for paired data were performed on important clinical parameters.

5.2.6 Computation times

For 10 arbitrary selected patients, computation times for TBS-BAO(25) plans were compared to computation times for generation of corresponding iCycle-BAO(25) plans, and to manual planning times for the corresponding Manual plans. Plan generations were performed on an Intel Xeon Gold 6248 @ 2.5 GHz, containing 40 cores and with 386 GB of memory.

5.3   Results

5.3.1 Plan comparisons

Figure 5.1 shows generated dose distributions for an example patient. Dosimetric plan parameters presented in figure 5.2 show that there were no violations of hard constraints in any of the plans generated with TBS-BAO, iCycle-BAO, or manual planning. Subplot 5.2.a shows that for all six plan types (TBS-BAO HR, TBS-BAO LR, Ideal HR, Ideal LR, iCycle-BAO(25) and Manual), the mean PTV Coverage was very close to the clinically desired 95%, with similar outliers with lower coverage.

In general, mean dosimetric plan parameters for urethra, rectum, and bladder in segmented TBS-BAO HR and TBS-BAO LR plans improved with increasing preset maximum beam numbers, moving towards the value of the corresponding Ideal HR or Ideal LR reference plan with 91 beams (figure 5.2). However, improvements started to level off around a preset maximum of 25 beams. For maximum allowed beam numbers $\geq$ 25, the average delivery times and WSS gradually increased slightly, while beam segment numbers and MU per fraction remained relatively constant. Urethra plan parameters were favorable for TBS-BAO LR, compared to TBS-BAO HR, which was reversed for rectum and bladder. This was fully in line with the plan parameters for the corresponding reference Ideal plans; lower urethra mean values for Ideal LR, lower rectum and bladder values for Ideal HR. Figure 5.3 shows a comparison of the population averaged DVHs of ideal reference plans and segmented TBS-BAO plans.

Achieved iCycle-BAO(25) plan parameters for urethra, rectum, and bladder were highly similar to corresponding values in the TBS-BAO plans, sometimes closer to TBS-BAO HR, sometimes closer to TBS-BAO LR (figure 5.2). No significant
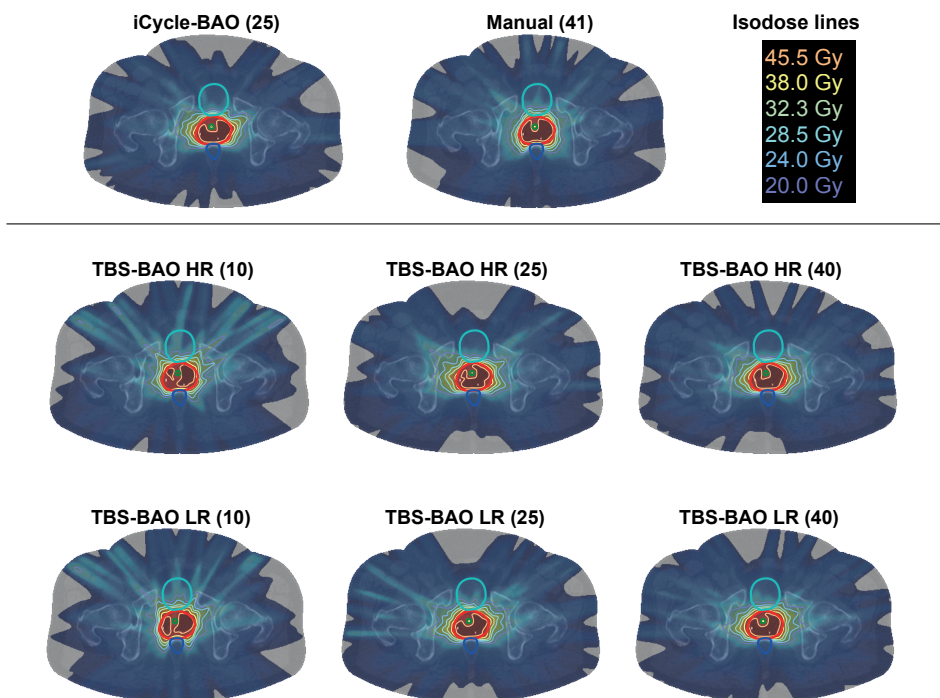
Figure 5.1: Dose distributions for an example patient. iCycle-BAO(25) and Manual (41) are validation plans. TBS-BAO-HR(XX) and TBS-BAO-LR(XX) are segmented plans generated with the proposed TBS-BAO approach for preset maximum numbers of beams of 10, 25 and 40. Depicted structures: PTV (red), Urethra (green), Rectum (dark blue), and Bladder (light blue).

5

differences in delivery times were observed between the iCycle-BAO(25) plans and TBS-BAO plans (figure 5.2.t): on average 18.0 minutes (min: 14.1, max: 22.0) for iCycle-BAO(25) and 18.4 minutes (min: 13.9, max: 28.9) for TBS-BAO HR(25). The Manual plans with on average 48 beams performed dosimetrically similar to TBS-BAO plans with a maximum of 25 selected beams, while delivery times of the former plans were substantially higher (figure 5.2.t): on average 22.3 minutes (min: 17.0, max: 32.0) for Manual and 18.4 minutes (min: 13.9, max: 28.9) for TBS-BAO HR(25). Figure 5.4 shows population averaged DVH comparisons between TBS-BAO plans and validation iCycle-BAO(25) and Manual plans. See Appendix A for more details on plan parameter comparisons.
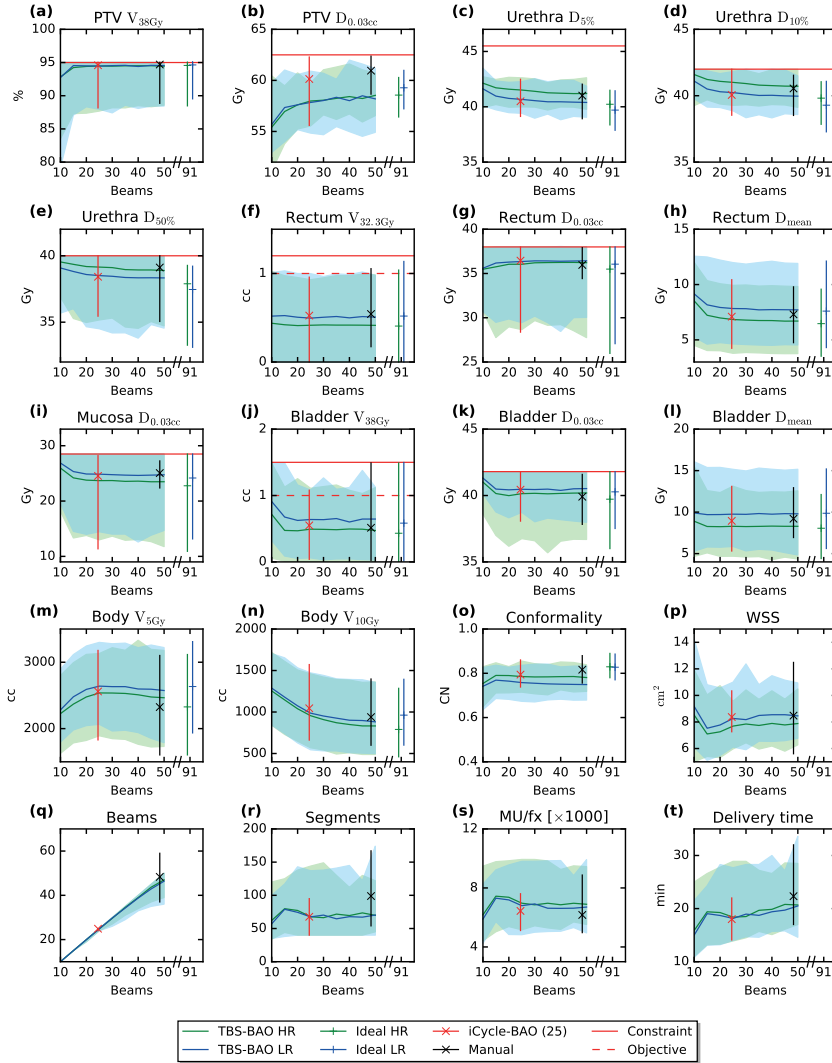
Figure 5.2: Overview of plan parameters for all segmented plans (TBS-BAO HR, TBS-BAO LR) together with the ideal FMO reference plans (Ideal HR, Ideal LR) and the validation plans (Manual, iCycle-BAO(25)). Each subplot shows results for one plan parameter. The lines representing TBS-BAO HR and TBS-BAO LR show for each maximum allowed number of beams (along the x-axis) the population mean parameter value, with colored areas depicting ranges. Plan parameter values for the Ideal HR and Ideal LR FMO plans and both validation plans are depicted with vertical line plots, positioned along the x-axis at the average number of realized beams per approach. The marker within each vertical line depicts the average value and the vertical line the range. Manual plans had on average 48 included beams. WSS = Weighted Segment Size.
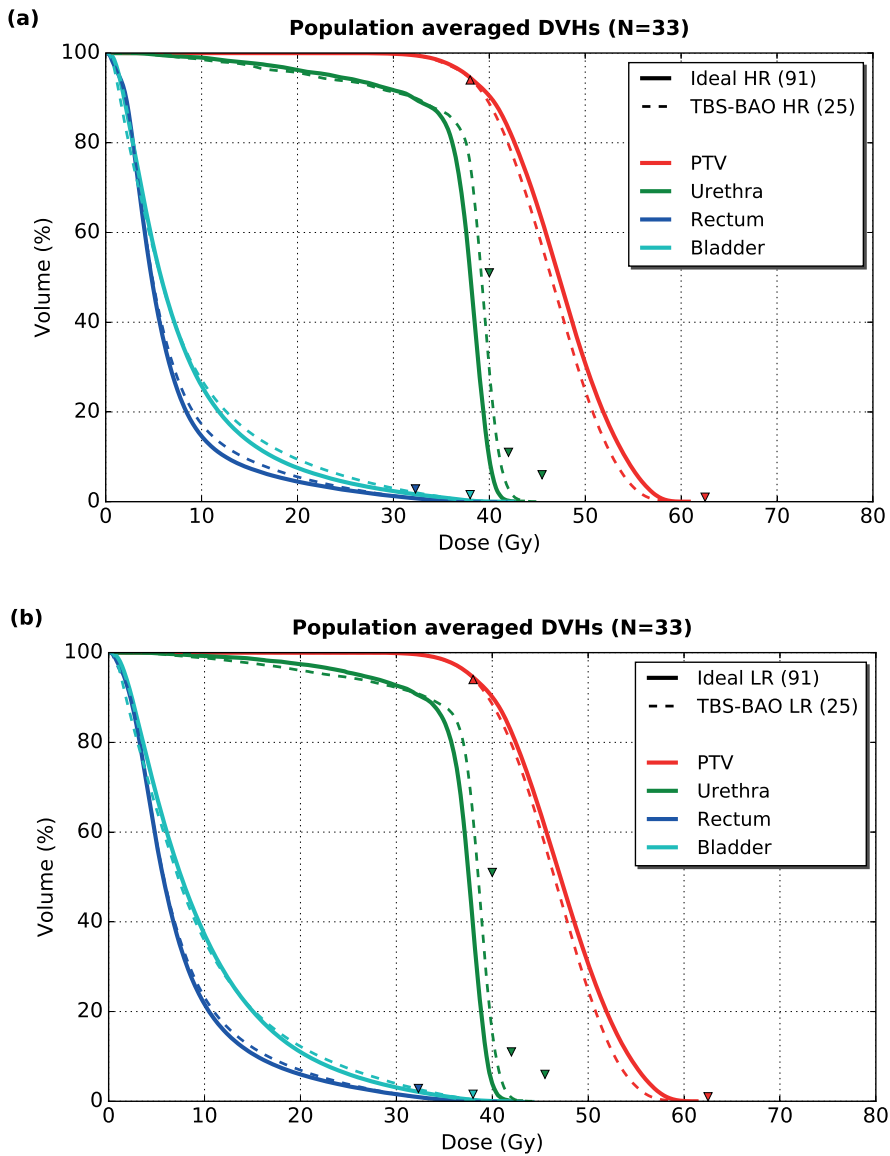
Figure 5.3: Population averaged DVHs for ideal reference FMO plans with 91 beams and segmented TBS-BAO plans with 25 beams. a) comparisons for Ideal HR reference plans, b) comparisons for Ideal LR reference plans. Clinical constraints are denoted with triangles, see table 5.1.
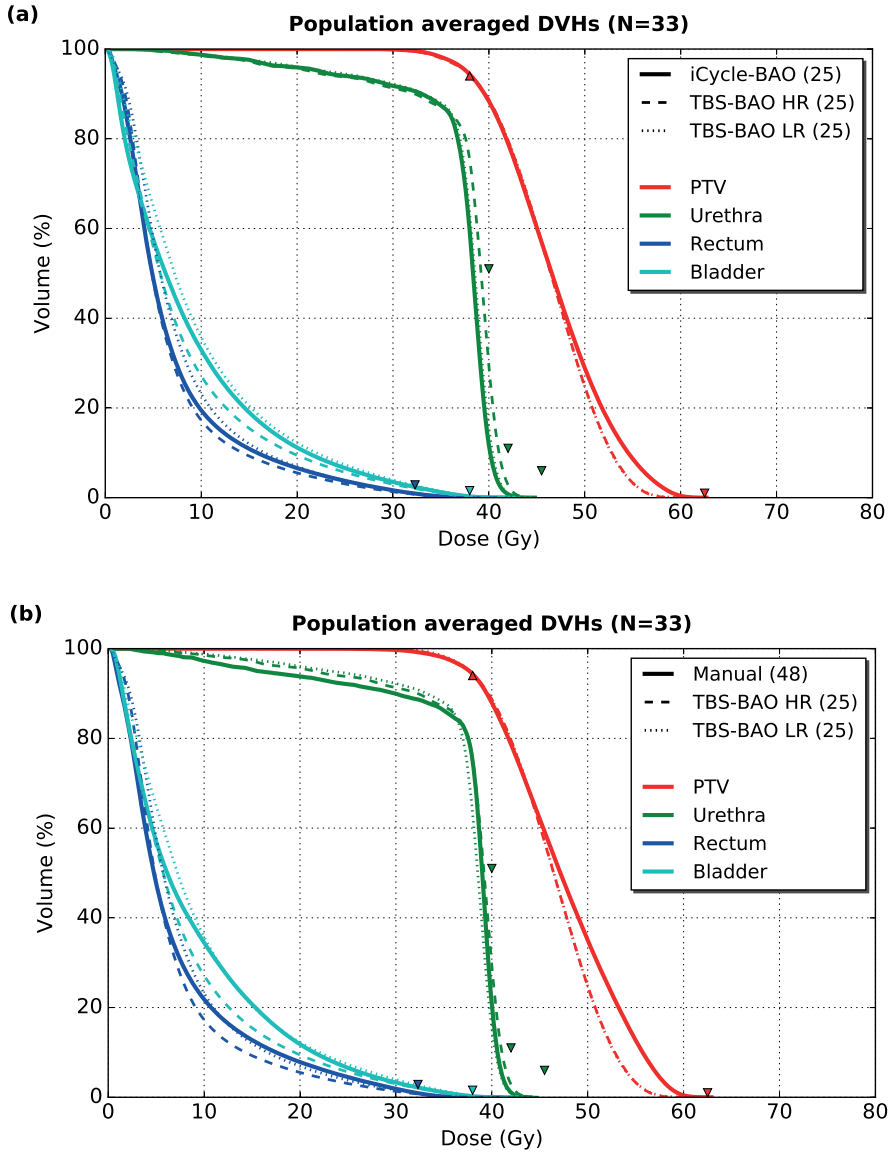
Figure 5.4: Population averaged DVHs for TBS-BAO HR(25), TBS-BAO LR(25) and validation plans. a) validation plans are iCycle-BAO(25) plans, b) validation plans are Manual plans with on average 48 non-coplanar beams. Clinical constraints are denoted with triangles, see table 5.1.

Table 5.3: Computation times in hours for 10 arbitrarily selected patients. The reported Manual planning time is manual hands-on time, while the other reported times are fully automated calculation times without manual interaction. PB + FMO: Calculation of pencil-beams + Fluence Map Optimization.

	PB + FMO		Segmentation		Total	
Approach	Mean	(min, max)	Mean	(min, max)	Mean	(min, max)
TBS-BAO HR (10)	**4.3**	(2.9, 7.2)	**0.4**	(0.3, 0.5)	**4.7**	(3.2, 7.5)
(25)			**0.5**	(0.3, 0.6)	**4.8**	(3.3, 7.5)
(50)			**0.5**	(0.3, 0.6)	**4.8**	(3.3, 7.5)
TBS-BAO LR (10)	**1.1**	(0.8, 2.2)	**0.4**	(0.3, 0.6)	**1.5**	(1.1, 2.7)
(25)			**0.4**	(0.3, 0.6)	**1.5**	(1.1, 2.7)
(50)			**0.5**	(0.3, 0.6)	**1.6**	(1.1, 2.7)
iCycle-BAO (25)	**10.4**	(8.8, 12.4)	**0.3**	(0.3, 0.4)	**10.7**	(9.2, 12.7)
Manual					**0.8**	(0.2, 2.8)

5.3.2 Computation times

Table 5.3 shows for 10 arbitrarily selected patients, measured computation times for automatic generation of TBS-BAO plans, times for automatic generation of corresponding iCycle-BAO(25) plans, and times for manual generation of the corresponding Manual plans. TBS-BAO has a clear advantage in computation time over iCycle-BAO, especially when Ideal LR plans are used as the reference plan.

5.4 Discussion

In this study, the novel TBS-BAO approach for solving the BAO problem was proposed, and validated for non-coplanar CyberKnife SBRT for prostate cancer patients. The essence for solving the BAO problem is that plan generation starts with creation of an 'ideal' reference dose distribution, using all available candidate beams (total-beam-space, TBS). This ideal reference plan is generated with multi-criterial FMO, providing for each patient a Pareto-optimal plan. The ideal FMO dose distribution then acts as a strong guidance in the following non-convex IMRT segmentation with integrated BAO (SEG/BAO). This SEG/BAO for a preset maximum number of included beams aimed at realizing the ideal reference plan as closely as possible. MLC segments are sequentially added to the beams, aiming at minimization of dosimetric differences between the segmented plan and the ideal reference plan. With the globally optimal reference plan at hand, there is always a quality measure for the final segmented plan with pre-defined maximum number of beams.

Dosimetrically, the clinically deliverable plans with 25 beams (TBS-BAO HR and TBS-BAO LR) were highly similar to the 91-beam ideal reference plans (Ideal HR and LR). Especially, for rectum and bladder, ideal dose distributions could be closely reproduced. For urethra sparing, a slight increase of approximately 1 Gy was

observed between the ideal reference plans and the segmented plans. For most plan parameters, improvements in TBS-BAO plans leveled off for allowed maximum numbers of non-coplanar beams ≥ 25.

While all plan parameters in Ideal HR and Ideal LR plans were within clinical constraints, there were slight trade-off differences between Ideal HR and Ideal LR in OAR sparing. Ideal HR performed slightly better for rectum and bladder, while Ideal LR was slightly better for urethra sparing (figure 5.2, panels c-l). The differences between Ideal HR and Ideal LR plans were reproduced in the segmented TBS-BAO HR and TBS-BAO LR plans. This indicates that depending on the pencil-beam resolution used during FMO, a slightly different trade-off in OAR sparing was achieved. Possibly, with the coarser pencil-beam resolution of Ideal LR, a somewhat larger area around the urethra was needed to be spared to fulfill the urethra constraints, at the price of a little less sparing for bladder and rectum (but still within constraints). As indicted in table 5.3, for a maximum of 25 beams, the total plan generation time for TBS-BAO HR of 4.8 hours reduced to 1.6 hours for TBS-BAO LR. This may also be considered in choosing between the HR or LR approach. Appropriate beamlet resolution may have to be tuned per tumor site. At the moment, our research system uses serial calculations for the pencil-beams and for the final dose calculations. Parallelization of these dose calculation steps could substantially reduce the computation times for both TBS-BAO HR and TBS-BAO LR.

Plan parameters of iCycle-BAO(25) plans were overall similar to those of TBS-BAO HR(25) and TBS-BAO LR(25) (figures 5.2 & 5.4) with small advantages of one or the other, depending on parameter. This also implies that plan parameters of the iCycle-BAO(25) plans were close to Ideal plan parameters. This observed high quality of iCycle-BAO plans is in line with observations in other studies (Rossi et al., 2018, 2021, Sharfo et al., 2017, 2021). With the availability of the Ideal plans, this could now be objectively verified. A clear advantage of the proposed TBS-BAO over iCycle-BAO is calculation time (table 5.3). With iCycle-BAO, calculation time substantially increases with beam number (Breedveld et al., 2012). A test for a single patient resulted in a computation time for TBS-BAO LR(40) of 1.1 hours while generating the corresponding iCycle-BAO(40) plan took 29.3 hours. Increasing the number of allowed beams per plan only had a marginal impact on the total calculation time for the TBS-BAO approach, since the ideal reference plan is always performed on all available candidate beams, and in the first part of the segmentation phase, all candidate beams are taking into account for the identification of the next MLC segment, irrespective of the maximum number of beams per plan. On average, the calculation time for the segmentation phase of the TBS-BAO plans increased from 23 min. (0.4 hours) for a plan with a maximum of 10 beams to 28 min. (0.5 hours) for a plan with a maximum of 50 beams.

Comparisons between Manual and TBS-BAO plans demonstrated that the Manual plans that were made by an experienced physicist in the absence of time pressure (Giżyńska et al., 2021) were indeed of high quality (figure 5.2). On the other hand, this Manual plan quality can possibly not be reached in clinical routine with planners with various skill sets and with time pressure. The manual plans had on average 48 included beams, while the quality was similar to TBS-BAO plans with a maximum of 25 beams. Delivery times of Manual plans were on average 3.9 min longer. Generation of TBS-BAO plans required virtually zero manual workload, while generation of the Manual plans took on average 0.8 hours.

The proposed TBS-BAO approach was developed for generating step-and-shoot IMRT. Extending the method for dynamic delivery would be an interesting option. This would require extra restrictions on segment shapes for consecutive segments, yet the requirement on the minimum MU per segment is no longer needed since the dose delivery is continuous. Additionally, Men et al. (2010) showed that column generation can be used to generate radiotherapy treatment plans for dynamic delivery, Kearney et al. (2017, 2018) demonstrated the feasibility of continuous arc delivery for the CyberKnife, and Bedford et al. (2020) demonstrated that sufficient delivery accuracy for clinical application can be accomplished with dynamic arc delivery on the CyberKnife. Currently, the CyberKnife does not support dynamic delivery, however, the proposed TBS-BAO approach could in principle be applied for generating treatment plans for conventional linacs.

5.5 Conclusions

The novel Total-Beam-Space-BAO approach (TBS-BAO) uses for each plan generation an 'ideal' dose distribution as reference. This ideal dose distribution is generated with fluence map optimization using all available candidate beams. The following integrated BAO and segmentation for a preset maximum number of included beams then aims at realizing the ideal reference plan as closely as possible. The full planning pipeline is automated; there is virtually no manual workload. For CyberKnife prostate SBRT, generation of TBS-BAO plans was much faster than plan generation with an alternative BAO approach. Compared to high-quality Manual planning, there was no involved workload and TBS-BAO plans with 25 beams had similar plan quality as the Manual plans with on average 48 beams. Delivery times for the TBS-BAO plans were reduced by on average 3.9 min. The ideal TBS reference FMO plans could effectively steer the discrete non-convex MLC segmentation with integrated BAO (SEG/BAO).

Acknowledgments

5

Appendix A

Table A1: Plan parameter comparisons for Manual and TBS-BAO HR plans with a maximum of 25 beams. Doses (Dxx) presented in Gy and volumes (Vxx) in cc. Entrance dose was evaluated using D_{1cc} of a ring structure of 3 cm thickness inside the patient's external contour. FH: Femoral head, MU/fx: Monitor Units per fraction, DT: Delivery time.

		Manual		TBS-BAO HR(25)		TBS-BAO HR(25) – Manual		
		Mean	(min, max)	Mean	(min, max)	Δ	(min, max)	p
PTV	$D_{98\%}$	**34.9**	(31.2, 36.8)	**35.0**	(33.5, 36.4)	**0.1**	(-1.2, 2.7)	.8
	$D_{0.03cc}$	**60.7**	(57.9, 62.4)	**57.9**	(55.5, 59.7)	**-2.8**	(-5.5, -0.5)	**<.001**
Urethra	$D_{5\%}$	**40.8**	(38.9, 41.9)	**41.4**	(40.0, 42.6)	**0.6**	(-1.0, 2.0)	**<.001**
	$D_{10\%}$	**40.4**	(38.5, 41.4)	**40.9**	(39.3, 42.0)	**0.6**	(-0.9, 1.9)	**<.001**
	$D_{50\%}$	**38.9**	(35.1, 40.0)	**39.1**	(35.2, 40.0)	**0.1**	(-1.0, 0.9)	.1
Rectum	$V_{32.3Gy}$	**0.5**	(0.2, 1.1)	**0.4**	(0.0, 0.9)	**-0.1**	(-0.5, 0.5)	**.004**
	$D_{0.03cc}$	**35.8**	(34.4, 37.6)	**36.0**	(28.5, 38.0)	**0.2**	(-6.6, 3.0)	.1
	D_{mean}	**7.3**	(4.7, 9.7)	**6.8**	(3.8, 9.9)	**-0.4**	(-2.7, 2.3)	**.01**
Mucosa	$D_{0.03cc}$	**25.0**	(22.0, 27.0)	**23.6**	(12.9, 28.2)	**-1.3**	(-11.2, 6.0)	.1
Bladder	V_{38Gy}	**0.4**	(0.0, 1.2)	**0.5**	(0.0, 1.0)	**0.1**	(-0.4, 0.4)	.1
	$D_{0.03cc}$	**39.7**	(37.8, 41.6)	**40.1**	(36.8, 41.5)	**0.3**	(-1.7, 1.5)	**.006**
	D_{mean}	**9.2**	(6.9, 12.9)	**8.3**	(4.7, 13.1)	**-0.9**	(-2.7, 1.8)	**<.001**
Entrance	D_{1cc}	**17.4**	(14.1, 19.9)	**17.3**	(12.9, 20.2)	**-0.1**	(-4.0, 3.5)	.8
Left FH	$D_{0.03cc}$	**13.3**	(7.6, 17.7)	**15.2**	(11.8, 19.0)	**1.9**	(-1.5, 7.2)	**<.001**
Right FH	$D_{0.03cc}$	**15.2**	(9.9, 18.6)	**15.4**	(9.8, 18.7)	**0.1**	(-2.6, 2.5)	.5
Conformality		**0.82**	(0.75, 0.88)	**0.79**	(0.71, 0.86)	**-0.03**	(-0.07, 0.00)	**<.001**
Beams		**48.3**	(37.0, 59.0)	**24.8**	(24.0, 25.0)	**-23.5**	(-34.0, -12.0)	**<.001**
Segments		**98.9**	(54.0, 167.0)	**67.8**	(40.0, 144.0)	**-31.2**	(-96.0, 31.0)	**<.001**
MU/fx [×1000]		**5.9**	(4.2, 8.9)	**6.7**	(4.6, 9.8)	**0.8**	(-2.0, 4.7)	**.002**
Delivery time [min]		**22.3**	(17.0, 32.0)	**18.4**	(13.9, 28.9)	**-3.9**	(-14.0, 6.6)	**<.001**

5

Table A2: Plan parameter comparisons for Manual and TBS-BAO LR plans with a maximum of 25 beams. Doses (Dxx) presented in Gy and volumes (Vxx) in cc. Entrance dose was evaluated using D_{1cc} of a ring structure of 3 cm thickness inside the patient's external contour. FH: Femoral head, MU/fx: Monitor Units per fraction, DT: Delivery time.

		Manual		TBS-BAO LR(25)		TBS-BAO LR(25) – Manual		
		Mean	(min, max)	Mean	(min, max)	Δ	(min, max)	p
PTV	$D_{98\%}$	**34.9**	(31.2, 36.8)	**35.5**	(33.2, 36.5)	**0.5**	(-0.5, 3.2)	**<.001**
	$D_{0.03cc}$	**60.7**	(57.9, 62.4)	**57.8**	(54.8, 60.9)	**-2.9**	(-5.6, -0.4)	**<.001**
Urethra	$D_{5\%}$	**40.8**	(38.9, 41.9)	**40.6**	(39.4, 42.2)	**-0.2**	(-2.0, 1.5)	.1
	$D_{10\%}$	**40.4**	(38.5, 41.4)	**40.2**	(38.9, 41.7)	**-0.2**	(-1.8, 1.5)	.2
	$D_{50\%}$	**38.9**	(35.1, 40.0)	**38.5**	(34.5, 39.8)	**-0.5**	(-1.7, 0.7)	**<.001**
Rectum	$V_{32.3Gy}$	**0.5**	(0.2, 1.1)	**0.5**	(0.0, 1.0)	**-0.0**	(-0.5, 0.5)	.3
	$D_{0.03cc}$	**35.8**	(34.4, 37.6)	**36.3**	(29.5, 38.0)	**0.5**	(-5.6, 2.2)	**.008**
	D_{mean}	**7.3**	(4.7, 9.7)	**7.8**	(4.6, 11.9)	**0.5**	(-2.7, 4.3)	.1
Mucosa	$D_{0.03cc}$	**25.0**	(22.0, 27.0)	**24.8**	(14.5, 28.5)	**-0.1**	(-9.6, 6.5)	.7
Bladder	V_{38Gy}	**0.4**	(0.0, 1.2)	**0.6**	(0.1, 1.1)	**0.2**	(-0.3, 0.8)	**<.001**
	$D_{0.03cc}$	**39.7**	(37.8, 41.6)	**40.4**	(38.7, 41.8)	**0.7**	(-1.0, 2.8)	**<.001**
	D_{mean}	**9.2**	(6.9, 12.9)	**9.7**	(5.9, 15.1)	**0.6**	(-1.5, 4.3)	**.04**
Entrance	D_{1cc}	**17.4**	(14.1, 19.9)	**17.5**	(13.7, 20.1)	**0.0**	(-3.7, 3.4)	.7
Left FH	$D_{0.03cc}$	**13.3**	(7.6, 17.7)	**15.9**	(12.5, 21.4)	**2.7**	(-1.0, 7.2)	**<.001**
Right FH	$D_{0.03cc}$	**15.2**	(9.9, 18.6)	**15.8**	(11.2, 19.7)	**0.6**	(-3.3, 5.0)	.1
Conformality		**0.82**	(0.75, 0.88)	**0.76**	(0.67, 0.84)	**-0.06**	(-0.11, -0.01)	**<.001**
Beams		**48.3**	(37.0, 59.0)	**24.9**	(24.0, 25.0)	**-23.5**	(-34.0, -12.0)	**<.001**
Segments		**98.9**	(54.0, 167.0)	**67.5**	(40.0, 137.0)	**-31.4**	(-114.0, 37.0)	**<.001**
MU/fx [×1000]		**5.9**	(4.2, 8.9)	**6.6**	(4.3, 9.8)	**0.6**	(-2.9, 4.7)	**.009**
Delivery time [min]		**22.3**	(17.0, 32.0)	**18.1**	(14.0, 27.7)	**-4.2**	(-15.8, 6.4)	**<.001**

5

Table A3: Plan parameter comparisons for iCycle-BAO and TBS-BAO HR plans with a maximum of 25 beams. Doses (Dxx) presented in Gy and volumes (Vxx) in cc. Entrance dose was evaluated using D_{1cc} of a ring structure of 3 cm thickness inside the patient's external contour. FH: Femoral head, MU/fx: Monitor Units per fraction, DT: Delivery time.

		iCycle-BAO(25)		TBS-BAO HR(25)		TBS-BAO HR – iCycle-BAO		
		Mean	(min, max)	Mean	(min, max)	Δ	(min, max)	p
PTV	$D_{98\%}$	**35.4**	(33.4, 36.5)	**35.0**	(33.5, 36.4)	**-0.4**	(-1.9, 0.5)	**<.001**
	$D_{0.03cc}$	**60.0**	(54.9, 62.3)	**57.9**	(55.5, 59.7)	**-2.1**	(-3.9, 1.3)	**<.001**
Urethra	$D_{5\%}$	**40.4**	(39.1, 42.5)	**41.4**	(40.0, 42.6)	**1.1**	(-0.8, 3.5)	**<.001**
	$D_{10\%}$	**39.9**	(38.5, 42.0)	**40.9**	(39.3, 42.0)	**1.0**	(-0.9, 3.3)	**<.001**
	$D_{50\%}$	**38.3**	(35.5, 39.8)	**39.1**	(35.2, 40.0)	**0.8**	(-0.9, 2.1)	**<.001**
Rectum	$V_{32.3Gy}$	**0.5**	(0.0, 0.9)	**0.4**	(0.0, 0.9)	**-0.1**	(-0.3, 0.1)	**<.001**
	$D_{0.03cc}$	**36.4**	(28.1, 38.0)	**36.0**	(28.5, 38.0)	**-0.4**	(-2.2, 0.7)	**.02**
	D_{mean}	**7.1**	(4.2, 10.3)	**6.8**	(3.8, 9.9)	**-0.2**	(-1.6, 0.8)	**.04**
Mucosa	$D_{0.03cc}$	**24.5**	(11.3, 27.9)	**23.6**	(12.9, 28.2)	**-0.8**	(-3.7, 1.6)	**.002**
Bladder	V_{38Gy}	**0.5**	(0.0, 0.9)	**0.5**	(0.0, 1.0)	**-0.0**	(-0.2, 0.4)	.1
	$D_{0.03cc}$	**40.3**	(38.1, 41.8)	**40.1**	(36.8, 41.5)	**-0.3**	(-2.8, 0.8)	.1
	D_{mean}	**8.9**	(5.3, 13.0)	**8.3**	(4.7, 13.1)	**-0.6**	(-3.9, 1.3)	**.008**
Entrance	D_{1cc}	**17.9**	(15.8, 19.5)	**17.3**	(12.9, 20.2)	**-0.6**	(-5.6, 2.9)	.1
Left FH	$D_{0.03cc}$	**13.8**	(5.7, 19.8)	**15.2**	(11.8, 19.0)	**1.4**	(-5.9, 7.6)	**.002**
Right FH	$D_{0.03cc}$	**14.9**	(7.6, 19.5)	**15.4**	(9.8, 18.7)	**0.5**	(-4.8, 4.9)	.2
Conformality		**0.80**	(0.74, 0.86)	**0.79**	(0.71, 0.86)	**-0.01**	(-0.05, 0.03)	**.03**
Beams		**24.9**	(24.0, 25.0)	**24.8**	(24.0, 25.0)	**-0.1**	(-1.0, 1.0)	.4
Segments		**67.7**	(40.0, 95.0)	**67.8**	(40.0, 144.0)	**0.1**	(-45.0, 51.0)	1.0
MU/fx [×1000]		**6.2**	(4.9, 7.4)	**6.7**	(4.6, 9.8)	**0.5**	(-0.8, 2.6)	**.003**
Delivery time [min]		**18.0**	(14.1, 22.0)	**18.4**	(13.9, 28.9)	**0.4**	(-4.8, 7.6)	.6

Table A4: Plan parameter comparisons for iCycle-BAO and TBS-BAO LR plans with a maximum of 25 beams. Doses (Dxx) presented in Gy and volumes (Vxx) in cc. Entrance dose was evaluated using D_{1cc} of a ring structure of 3 cm thickness inside the patient's external contour. FH: Femoral head, MU/fx: Monitor Units per fraction, DT: Delivery time.

		iCycle-BAO(25)		TBS-BAO LR(25)		TBS-BAO LR – iCycle-BAO		
		Mean	(min, max)	Mean	(min, max)	Δ	(min, max)	p
PTV	$D_{98\%}$	**35.4**	(33.4, 36.5)	**35.5**	(33.2, 36.5)	**0.0**	(-0.8, 0.8)	.4
	$D_{0.03cc}$	**60.0**	(54.9, 62.3)	**57.8**	(54.8, 60.9)	**-2.2**	(-4.7, 1.0)	**<.001**
Urethra	$D_{5\%}$	**40.4**	(39.1, 42.5)	**40.6**	(39.4, 42.2)	**0.2**	(-1.1, 2.4)	**.04**
	$D_{10\%}$	**39.9**	(38.5, 42.0)	**40.2**	(38.9, 41.7)	**0.2**	(-1.1, 2.2)	**.03**
	$D_{50\%}$	**38.3**	(35.5, 39.8)	**38.5**	(34.5, 39.8)	**0.1**	(-1.0, 1.1)	.1
Rectum	$V_{32.3Gy}$	**0.5**	(0.0, 0.9)	**0.5**	(0.0, 1.0)	**-0.0**	(-0.3, 0.1)	.1
	$D_{0.03cc}$	**36.4**	(28.1, 38.0)	**36.3**	(29.5, 38.0)	**-0.0**	(-0.8, 1.4)	.5
	D_{mean}	**7.1**	(4.2, 10.3)	**7.8**	(4.6, 11.9)	**0.7**	(-0.7, 2.3)	**<.001**
Mucosa	$D_{0.03cc}$	**24.5**	(11.3, 27.9)	**24.8**	(14.5, 28.5)	**0.4**	(-2.2, 3.2)	**.02**
Bladder	V_{38Gy}	**0.5**	(0.0, 0.9)	**0.6**	(0.1, 1.1)	**0.1**	(-0.1, 0.7)	**<.001**
	$D_{0.03cc}$	**40.3**	(38.1, 41.8)	**40.4**	(38.7, 41.8)	**0.1**	(-0.7, 1.7)	.6
	D_{mean}	**8.9**	(5.3, 13.0)	**9.7**	(5.9, 15.1)	**0.8**	(-1.1, 3.1)	**<.001**
Entrance	D_{1cc}	**17.9**	(15.8, 19.5)	**17.5**	(13.7, 20.1)	**-0.5**	(-3.1, 2.8)	.1
Left FH	$D_{0.03cc}$	**13.8**	(5.7, 19.8)	**15.9**	(12.5, 21.4)	**2.1**	(-2.5, 8.7)	**<.001**
Right FH	$D_{0.03cc}$	**14.9**	(7.6, 19.5)	**15.8**	(11.2, 19.7)	**0.9**	(-3.2, 5.2)	**.03**
Conformality		**0.80**	(0.74, 0.86)	**0.76**	(0.67, 0.84)	**-0.04**	(-0.12, 0.02)	**<.001**
Beams		**24.9**	(24.0, 25.0)	**24.9**	(24.0, 25.0)	**-0.0**	(-1.0, 1.0)	.7
Segments		**67.7**	(40.0, 95.0)	**67.5**	(40.0, 137.0)	**-0.2**	(-54.0, 57.0)	.7
MU/fx [×1000]		**6.2**	(4.9, 7.4)	**6.6**	(4.3, 9.8)	**0.4**	(-1.1, 2.7)	.3
Delivery time [min]		**18.0**	(14.1, 22.0)	**18.1**	(14.0, 27.7)	**0.1**	(-6.0, 7.4)	.8

5

Discussion

This thesis focused on developing and validating novel, fully automated treatment planning solutions for robotic radiotherapy, with emphases on clinical deliverability of generated treatment plans and on optimization of beam angles. Developed solutions are fully independent of the commercial TPS for manual plan generation. Automated plans were compared with manual plans both regarding plan quality and plan delivery efficiency. In this chapter, I will elaborate on the choice of algorithms, discuss the advantages and limitations of the methods proposed, and discuss potential future research.

6.1 FMO + segmentation vs. Direct Aperture Optimization (DAO)

In radiotherapy treatment planning, there are generally two optimization approaches for generating IMRT plans for pre-selected beam directions.

The first approach splits the optimization in two phases: a fluence map optimization (FMO) phase that optimizes the fluence profiles for all beam directions and a segmentation phase to convert the optimized fluences into a deliverable plan, consisting of multi-leaf collimator (MLC) segments and corresponding segment intensities (Breedveld et al., 2012, Luan et al., 2006, Sun and Xia, 2004, Süss et al., 2007, Xia and Verhey, 1998). The segmentation can degrade the quality of a dose distribution obtained with FMO.

The second approach is Direct Aperture Optimization (DAO) that optimizes directly on the MLC parameters (Men et al., 2007, Romeijn et al., 2005, Shepard et al., 2002). In this approach, the treatment plan is directly deliverable at every stage of the optimization process, therefore no segmentation phase (with a possible loss in plan quality) is needed. However, including the non-convex modeling of the (physical) constraints of the collimator and treatment device leads to a non-convex optimization problem with a risk to get trapped in local minima.

6

Both approaches can have their advantages and disadvantages. An advantage of FMO with our in-house Erasmus-iCycle, an algorithm for automated a priori Multi-Criterial treatment plan Optimization (MCO), is that it can handle hard dose constraints, whereas in most DAO approaches proposed in literature a weighted sum objective function is used without any hard constraints. In practice, the latter approach requires manual tweaking of the weights of the objective function per patient to arrive at a clinically desirable dose distribution, with favorable patient-specific trade-offs between objectives while respecting all hard constraints. An advantage of DAO is that the delivery constraints are respected during optimization, therefore no segmentation phase is required to generate a deliverable treatment plan.

In our center we have given preference to plan optimization using automated FMO with Erasmus-iCycle, followed by segmentation. The generated FMO plans have guaranteed Pareto and global optimality. Moreover, with proper configuration, the applied wish-list driven MCO results in plans of high clinical quality. In this thesis novel segmentation algorithms are proposed to reconstruct FMO dose distributions as accurately as possible, with a focus on most important dosimetric features of the FMO dose distribution.

6.2   Challenges in step-and-shoot segmentation

For the developed segmentation algorithm (chapters 2 & 3) a column generation approach was chosen with a quadratic dose minimization term as the main driving force to reconstruct input FMO dose distributions.

6.2.1   Column-generation

The column generation method is often used to solve large-scale optimization problems by iteratively solving a series of smaller, but increasingly growing, problems. This approach was chosen because of its intuitive mechanism of generating segments and proven effectiveness in IMRT treatment planning using DAO (Carlsson, 2008, Romeijn et al., 2005, Salari and Unkelbach, 2013).

During the work in this thesis, we observed that homogeneous tumor dose distributions were difficult to reconstruct within clinically acceptable numbers of segments using the column generation approach. Many segments with small segments weight were required to adequately reconstruct the homogeneous dose distribution. A limiting factor during optimization was the minimum MU per segment restriction for the robotic treatment unit to ensure accurate dose delivery per segment.

Potentially, this could be linked to how segments are handled in the current column generation formulation: all identified segments remain unchanged during the segmentation of the plan and each identified segment is based on the local gradient information at that specific iteration of the optimization. When the solution progresses and more segments are added to the plan, the existing segments would likely benefit from slight adaptions to adapt to the current dose distribution.

In the proposed implementation, the minimum MU segment restriction is not explicitly enforced during optimization of the segment weights; only the non-negative constraint of the segment weights is enforced (chapters 2 and 3, equations 2.1, 2.4, 3.4, and 3.8). Instead, during segmentation, when a segment weight falls below the minimum MU constraint, this segment is removed from the plan. The ability to dynamically remove identified segments and replace them with more suitable segments improves the delivery efficiency of the plans. An interesting option to adapt segment shapes is the Aperture Shape Optimization algorithm proposed by Cassioli and Unkelbach 2013. This algorithm uses a column generation approach and applies gradient information to further optimize the leaf positions for already identified segments, either during the optimization or only applied to the final deliverable plan. Possibly, the segmentation algorithm proposed in this thesis could be further refined with a similar type of segment shape fine-tuning.

6.2.2 Objective function – quadratic minimization term

Initially, the applied objective function for generation of MLC segments only consisted of a quadratic minimization term to minimize dose differences with the given input FMO dose distribution, and a non-negativity constraint (chapter 2, equations 2.1 & 2.2). The quadratic term was chosen because of its convex formulation and fast optimization. However, the quadratic term alone yielded plans with insufficient PTV coverage or too low PTV min doses. For this reason, the objective function was extended with an LTCP term to provide adequate PTV coverage (equation 2.3). In chapter 3, an MU penalty term was added to reduce beam-on-time (equation 3.4). Yet, the quadratic minimization term remained the main driving force of the segmentation.

Important to note is that in the developed approach, the quality of the input FMO dose distribution is vital. After all, the quadratic minimization term works towards the input dose distribution, without a drive to get better than the input dose distribution. At the same time, the input dose distribution must be physically achievable. Segmentation of a hypothetical dose distribution with the prescribed dose to the tumor and no dose to the healthy tissue does not work, since the

hypothetical dose distribution does not hold any information on the achievable trade-offs between PTV and OARs or between OARs. Throughout this thesis, Pareto and globally optimal FMO plans generated with Erasmus-iCycle were used as input dose distributions for the segmentation algorithm. These plans are based on pencil-beam dose depositions and therefore reflect (to a large extent) physically achievable patient dose.

In a so-far unpublished project, deliverable dose distributions for prostate cancer patients, generated with manual planning in the commercial Precision TPS (Accuray Inc., Sunnyvale, USA) were used as input for the segmentation, instead of the Erasmus-iCycle FMO dose distributions. Given the degeneracy of the treatment planning problem (Alber et al., 2002), it was investigated whether a different set of segments and weights could be found that would render the same dosimetric plan quality, but with improved plan delivery efficiency. Generally, plans generated with our segmentation approach had similar dosimetric plan quality as the initial manual plans with some slight deteriorations, while plan delivery efficiency was similar. As pointed out above, our segmentation approach focuses on reconstruction of the input plan, so no plan improvements were expected. The manual plans were made by an experienced medical physicist, who focused on both dosimetric plan quality and delivery efficiency. In this experiment, our segmentation approach could not beat the delivery efficiency obtained by the manual planner.

A novelty of our segmentation approach is that it reconstructs the 3-dimensional FMO dose distribution while considering all allowed beam directions simultaneously, instead of reconstructing 2-dimensional fluence profiles. Sequencing the fluences for each beam separately excludes mutual dosimetric compensation of imperfect segmentations of the 2D beam fluence profiles. The developed methodology also enabled to focus on reconstructing important dosimetric features of dose distributions, which would not have been possible with independent sequencing of 2-dimensional fluence profiles.

A challenge in the developed segmentation approach is full compliance with maximum or minimum dose constraints, since these constraints are dependent on a single or a few voxels, and are therefore difficult to enforce with a weighted sum of quadratic dose differences formulation. In chapter 2, this problem was mitigated to a certain extent with the use of the 'multi-criterial' segmentation, which dynamically adapted the weights of crucial voxels within the objective function to better respect the hard constraints (section 2.2.4). In chapter 5, section 5.2.3, a final segment weight re-optimization step was added using a wish-list optimization. This step was performed at the end of the segmentation and ensured compliance with hard constraints at the cost of a few minutes extra calculation time.

6.2.3 Integration of a clinical dose engine

To accurately model the delivered patient dose, including MLC scatter and leaf transmission effects, a standalone version of the clinical dose engine (CDE) was integrated into the column generation formalism. One of the challenges of the integration was that the accurate segment dose calculated with the CDE differed from the segment dose calculated as the sum of its (approximated) pencil-beams, the latter also calculated with the CDE. In the conventional CG formulation, pencil-beams are used to describe the beam intensity to dose deposition relation for both the Master Problem and the Restricted Master Problem, i.e. the dose prediction in MP is the same as the dose prediction in RMP. Therefore, the solution on the RMP can be directly converted into a solution on the MP, which is needed to calculate the gradients for identification of the next segment. However, when the segment dose prediction in the RMP is calculated with a CDE (to better reproduce the real dose delivery) dose differences are induced. Initially, this resulted in over-predictions of the delivered dose in MP. As a consequence, fewer beamlets had a negative gradient and therefore fewer beamlets were predicted to be favorable for inclusion in the segments. As a result, the segmentation would terminate prematurely. In chapter 3, two versions of the RMP were introduced to the CG workflow to mitigate this problem. One version of the RMP modeled the accurate segment dose using the CDE (RMP_{seg}), and another version modeled the segment dose using the approximated pencil-beams used in MP (RMP_{pb}). RMP_{seg} is solved to obtain the segment weights for the accurate dose delivery plan. For each identified segment added to RMP_{seg}, also a segment is added to RMP_{pb} that mimics the accurate segment dose in RMP_{seg}, but is constructed out of the pencil-beams used in MP. In the RMP_{pb}, each beamlet could partially (non-binary) contribute to the segment dose. Then, in an additional minimization phase added to each CG iteration, the individual beamlets and segment weights for RMP_{pb} were optimized to mimic the accurate segment dose from RMP_{seg} as accurately as possible. These individual beamlets and segment weights could then be used to calculate the gradient on MP for the identification of the next segment.

The proposed approach mitigated the problem to a certain extent so that the segmentation did not terminate prematurely, however, it did not solve the fundamental dose differences between MP and RMP when including a CDE. Other options for improving the dose discrepancy between MP and RMP could be explored. For example, better approximations of the pencil-beam characteristics. Pencil-beams that are located at the edge of the tumor projection are more likely to end up on the edge of an MLC segment. A possible improvement could be to include a heuristic that estimates the likelihood of a pencil-beam to end up on the edge of

6

an identified segment, based on the projection of the tumor and OARs. Depending on this estimation MLC scatter effects could be taken into account for the calculation of the pencil-beam dose deposition, e.g. with a Monte Carlo based pencil-beam algorithm.

6.3 Challenges in treatment delivery efficiency and plan complexity

An important aspect of radiotherapy treatment plans is delivery efficiency; shorter treatment times improve patient comfort and delivery accuracy (reduced risk on patient motion), and leakage dose is reduced with lower MU.

In chapter 3, a variable MU penalty term was added to the objective function to reduce beam-on-time (equation 3.4). The weight of this term was tuned per tumor site to balance the trade-off between plan quality and delivery efficiency. In chapter 4, we found that there was no significant difference in the total MU between manual planning and automated treatment planning. This might indicate that the total MU found for both approaches is needed to provide sufficient dose to the tumor. If that is the case, substantial improvements in minimizing beam-on time can be considered unlikely.

Alternatively, the total beam-off time could be reduced to improve delivery efficiency. When evaluating the individual segments per treatment plan, we observed that some treatment plans had multiple segments per beam direction that also had similar MU per segment. Potentially, these segments can be combined into one segment to reduce both the number of segments and the total MU per plan.

In the investigated robotic treatment unit, a Traveling Salesman Problem is solved for each plan to find the best route through the utilized node positions, given the constraints for cable management of the treatment device. Therefore, substantial improvements in travel path through the node positions are also considered unlikely.

Another possibility to enhance efficiency is to balance the number of segments and beam directions, since the next MLC segment can already be prepared when the treatment device moves from one node position to the next. If a plan could be generated that makes optimal use of this property, the delivery time could possibly be further reduced. As an indication, for the 25-beam plans from chapter 5 with on average 68 segments, the estimated time spent for robot movement was 7.0 min. (38%), for beam on time 6.9 min. (38%), and for changing apertures 2.5 min. (14%). For the 50-beam plans with on average 70 segments, these estimated times were: robot movement: 10.4 min. (50%), beam on time: 6.9 min. (33%), and changing apertures: 1.4 min. (7%). Utilizing one segment per beam direction could potentially speed up the delivery efficiency with 7%–14%.

Besides the focus on delivery efficiency, another important aspect of radiotherapy treatment plans is plan complexity. First of all, the proposed plans in this thesis respect all the MLC restrictions of the InCise™ 2 MLC. The Modulation Complexity Score (MCS) (McNiven et al., 2010) for the plans discussed in chapter 4 are slightly lower for the automated plans (0.41 [min: 0.31, max: 0.49]) than for manual plans (0.49 [min: 0.33, max: 0.60]). The MCS score accounts for the shape of the segments and the amount of modulation per beam direction (number of segments per beam direction). The MCS ranges from 0 to 1, where a score of 1 means the least amount of modulation (only rectangular segments with a single segment per beam), and a score of 0 means the most amount of modulation. The automated plans had fewer beam directions which in part can explain the increased complexity, since the plans had more segments per beam direction, resulting in more modulation per individual beam direction.

6.4   Opportunities for improving calculation time

Currently, the proposed automated planning solution is not fully optimized for calculation efficiency. The most time-consuming phases (based on the calculations in chapter 5 with a pencil-beam resolution of 3 x 3.85 mm^2) were: FMO optimization (219 min., 76%), pencil-beam calculation (41 min., 14%), segmentation (24 min., 8%), and final dose calculation (4 min., 1%). For a pencil-beam resolution of 6 x 7.7 mm^2, the calculations times were 53 min.(55%) for FMO optimization, 14 min. (15%) for pencil-beam calculation, 24 min. (25%) for segmentation, and 4 min. (4%) for the final dose calculation.

Both pencil-beam calculations and final dose calculations are currently performed in series, which, in principle, could be done in parallel. As a reference, pencil-beam calculations with a resolution of 5 x 3.85 mm^2 in the TPS take on average 2.5 seconds (patient-independent machine data needed for the pencil-beam calculations is pre-loaded when the patient data is loaded, this takes on average: ~14 seconds). During a plan optimization, the final segment dose calculations are performed during segmentation. To provide an estimate of the calculation time, recalculations of the final segment doses in the TPS take on average 59 seconds. Implementation of parallel dose calculations could potentially reduce our calculations times by 43 minutes (15%) for the high-resolution pencil-beam option and by 17 minutes (18%) for the low-resolution pencil-beam option.

Another opportunity to improve calculation time is to include deep-learning approaches. At the moment, full treatment planning using deep learning still has its challenges, for example in the generalizability of the models, overfitting (limited

6

data), or the deliverability of the predicted dose distribution (unpublished work at our center). However, deep-learning has already been shown to be suitable for speeding up or improving parts of the optimization process, for example, fast dose calculation (Kontaxis et al., 2020), boosting performance of conventional dose calculation algorithms (Bai et al., 2021, Xing et al., 2020), or beam angle optimization (Bohara et al., 2020, Sadeghnejad-Barkousaraie et al., 2020, 2021).

Our research system for automated treatment planning is the result of years of research and collaboration between various researchers. As a result, the current research implementation is a combination of multiple research applications written in various programming languages. For each plan generation, optimization data has to be transferred back and forth between applications, which contributes to total computation times. One coherent system that is fully optimized for calculation efficiency, similar to a commercial TPS, could substantially improve the calculation time. With improved calculation times the automated workflow could also be interesting for daily online re-planning.

6.5   Beam Angle Optimization

Properly selected beam angles contribute to the quality of radiotherapy treatment plans. The Beam Angle Optimization (BAO) problem, however, is difficult to solve to optimality (Bangert et al., 2012) due to its non-convex discrete nature with many local minima (Craft, 2007, Södertröm and Brahme, 1993).

In chapter 5, we proposed a method (TBS-BAO) that integrates BAO in plan segmentation (SEG/BAO), rather than solving the BAO problem prior to or during the FMO phase. The method first generates an ideal Total-Beam-Space (TBS) reference plan, which is then used to guide the following SEG/BAO phase that aims at reconstruction of the TBS reference plan. The TBS-BAO planning approach is fully automated; there is virtually no manual workload. The generation of TBS reference plans including all available beam directions is now technologically feasible due to the recent advancements in computational power.

When comparing the TBS-BAO approach to the previously proposed BAO method in Erasmus-iCycle and to manual planning, only minor differences in dosimetric plan quality were observed for prostate SBRT. However, the delivery times of the 25-beam TBS-BAO plans were 3.9 minutes shorter than for the manual plans, due to the inclusion of fewer beams and segments. So far, TBS-BAO has only been tested for prostate cancer. Studies on other tumor sites are on-going.

The observation that for prostate SBRT the dosimetric improvement in plan quality leveled off at around 25 to 30 beams per plan led to the hypothesis that the added value of BAO is reduced for high numbers of beams per plan. To verify this,
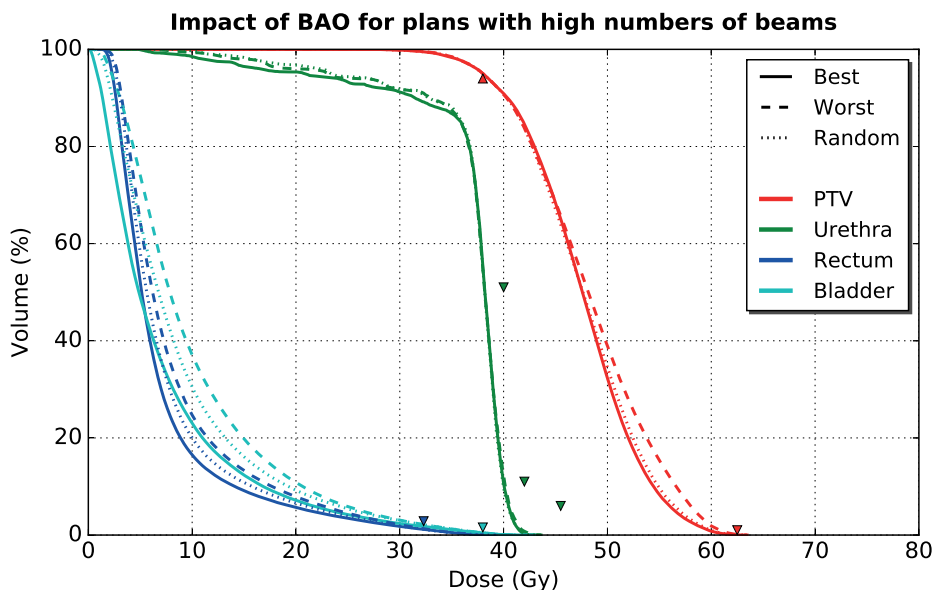
Figure 6.1: Population averaged DVHs for plans with high numbers of beams, comparing plans with the first 45 selected beams (best) out of the 91 candidate beams, plans with the complementary set of beams (worst), and plans with a set of 45 randomly selected beams (random).

an preliminary experiment was performed for 5 prostate SBRT patients for which we generated three treatment plans with high numbers of beams. For the first plan, the first 45 selected beams by TBS-BAO, out of the 91-candidate beam set, were used (best angles). For the second plan, the complementing 46 beams from the candidate beam set were used (worst). For the third plan, a random set of 45 beams from the candidate beam set were used (random). A clear difference in dosimetric plan quality was observed between the 'best' and 'worst' plans, whereas the plan quality of the 'random' plans was in between the plan quality of the 'best' and 'worst' plans, see figure 6.1. This suggests that even for plans with high numbers of beam directions per plan, there remains a dosimetric gain of using BAO.

6.6 Other future work and perspectives

6.6.1 Application of the proposed MLC segmentation for different modalities

An interesting option would be to adapt the developed MLC segmentation method for application in dynamic dose delivery, such as Volumetric Modulated Arc Therapy (VMAT). As mentioned earlier, the minimum MU per segment restriction was often a limiting factor for optimization of step-and-shoot plans. This can be

even more problematic with conventional fractionation schemes, where the total MU is distributed over many fractions. In dynamic delivery, when consecutive segments can be linked together, this MU per segment restriction is not an issue anymore. However, dynamic delivery requires more restrictions on segment shapes for consecutive segments.

In particular, dynamic delivery for the CyberKnife would be interesting; combining the non-coplanar beam space of the CyberKnife for high-quality treatment plans with the treatment efficiency of dynamic delivery. More specifically, to investigate options for integrating non-coplanar arc optimization in the TBS-BAO algorithm. This would allow for generating efficient treatment plans for the CyberKnife, while working towards the ideal plan quality of the total-beam-space reference plan. Men et al. (2010) showed that column generation can be used to generate treatment plans for dynamic delivery, Kearney et al. (2017, 2018) demonstrated the feasibility of continuous arc delivery for the CyberKnife, and Bedford et al. (2020) demonstrated that sufficient delivery accuracy for clinical application can be accomplished with dynamic arc delivery on the CyberKnife.

Existing MR-Linac systems only support step-and-shoot treatments, i.e. dynamic delivery is not available. The proposed TBS-BAO method could be used to generate deliverable IMRT treatment plans with integrated BAO. For this, the planning approach would need a Monte Carlo based dose calculation algorithm to account for the electron return effect to perform accurate pencil-beam and segment dose calculations.

6.6.2 Towards clinical application of the proposed methods

To use the developed workflow in clinical practice, the system should be MDR (Medical Device Regulations) or FDA (Food and Drug Administration) approved. For application as an independent self-functioning treatment planning system, it should comply to the highest risk class (MDR class III), since malfunctioning could cause death or an irreversible deterioration of a person's state of health. Also when applied as a QA-system, or when used as reference for conventional treatment planning, MDR-compliance is mandatory, although in a lower class (class IIb).

The advantages of having an independent self-functioning treatment planning system for the radiotherapy department would be to reduce the manual planning workload, reduce the treatment delivery time (chapters 4 & 5), or to use the tool as a QA-system to validate the clinical plans that are delivered to the patients (i.e. increase consistency of clinical plans). Currently, at our center, an automatic dose verification step is already performed for each treatment plan by recalculating the dose distribution using a third-party Monte Carlo dose engine (SciMoCa). This dose

verification step could be extended with a dose delivery verification step. For this, the final deliverable dose could be re-segmented with the proposed method from chapters 2 & 3. If the estimated delivery time of the reconstructed dose is substantially lower than the actual treatment plan, this plan could be flagged for re-evaluation by the medical physicist. Such a system could be run fully automatically in the background for all plans. Additionally, it would provide valuable feedback on achieved plan quality with the clinical TPS and insights on potential improvements to the vendor of the treatment device. Once such an automated planning system for the CyberKnife would be clinically available, the next steps would be to perform clinical validation studies to demonstrate the applicability in clinical practice and to explore other tumor sites.

6

Summary

Stereotactic Body Radiation Therapy (SBRT) is a radiotherapy approach in which high radiation doses are delivered to the tumors in a limited number of daily fractions (typically 3-5 fractions). High-precision dose delivery, generally based on image-guidance, is required to limit dose delivery to healthy tissue as much as possible. The CyberKnife® robotic radiotherapy device is a system that can deliver SBRT and allows for easy delivery of (non-coplanar) beam directions, without the need of manual couch shifts. High-quality radiotherapy dose distributions can be delivered using the CyberKnife, with excellent clinical outcomes (Fuller et al., 2018, Meier et al., 2018, van der Voort van Zyp et al., 2009). However, the enhanced degrees of freedom of fully non-coplanar robotic radiotherapy increases the complexity of finding optimal patient-specific beam angles and beam segments. This makes it more difficult to consistently generate high-quality treatment plans, especially under clinical circumstances with a variety of planner skill-sets and planning time constraints. The aim of this thesis was to develop and validate novel, fully automated treatment planning solutions for robotic radiotherapy, with emphases on clinical deliverability of generated treatment plans and on optimization of beam angles.

In radiotherapy treatment planning, there are generally two optimization approaches. The first approach splits the treatment plan optimization into two phases: a fluence map optimization (FMO) phase that optimizes the fluence profiles per beam direction, and a segmentation phase to convert the optimized fluences into a deliverable plan (consisting of multi-leaf collimator (MLC) segments and corresponding segment intensities). However, the segmentation phase can degrade the high-quality dose distribution obtained with FMO. The second approach is Direct Aperture Optimization (DAO) that optimizes directly on the MLC configurations. In this approach the treatment plan is directly deliverable at every

stage of the optimization process, so no segmentation phase (with a possible loss in plan quality) is needed. However, including the non-convex modeling of the (physical) constraints of the collimator and treatment device leads to a non-convex optimization problem. In our center, we have given preference to plan optimization using the first approach (FMO followed by segmentation), because of the guaranteed Pareto and global optimality of the FMO plans. The FMO plans are generated with Erasmus-iCycle, an algorithm for automated a priori Multi-Criterial treatment plan Optimization (MCO).

In chapter 2, a novel algorithm was proposed for converting a high-quality FMO dose distribution into a deliverable plan (consisting of MLC-segments and corresponding segment intensities). The algorithm focuses on minimizing plan quality loss by reproducing the 3-dimensional FMO dose distribution, rather than replicating the 2-dimensional fluences separately as is done in published MLC segmentation methods. All beams are considered simultaneously while generating MLC segments. The algorithm features prioritized generation of segments, focusing on accurate reproduction of clinical objectives with the highest priorities. The performance of the segmentation algorithm was evaluated for 20 prostate patients, 15 head-and-neck patients, and 12 liver patients. FMO dose distributions were generated by automated multi-criteria treatment planning (Pareto-optimal plans) and subsequently segmented using the proposed method. Segmented plans were dosimetrically similar to FMO plans and for all patients a clinically acceptable segmented plan could be generated. Substantial differences between FMO and segmented fluence profiles were observed. Avoidance of the usual reconstruction of 2D FMO fluence profiles for segment generation, and instead simultaneously generating segments for all beams to directly reproduce the 3D FMO dose distribution is a likely explanation for the obtained results.

In chapter 3, the proposed algorithm was further improved to accurately model the dose delivered to the patient with the integration of a clinical dose engine (CDE). Three versions of the algorithm were investigated with differences in the integration of the CDE. The combined use of pencil-beams and accurate segment doses in a segmentation method is non-trivial. Therefore, new methods were developed for the use of segment doses calculated with the CDE in combination with pencil-beams, used for the selection of new segments. For 20 patients with prostate cancer and 12 with liver cancer, segmented plans were compared with FMO plans. All three versions of the proposed segmentation algorithm could well mimic FMO dose distributions. Segmentation with a fully integrated CDE provided the best plan quality and lowest numbers of Monitor Units and segments at the cost

of increased calculation time.

In chapter 4, we proposed and validated a fully automated multi-criterial treatment planning solution for SBRT with a CyberKnife® equipped with an InCise™ 2 multi-leaf collimator. Automated treatment plans including non-coplanar Beam Angle Optimization (BAO) were generated fully outside the clinical TPS, using Erasmus-iCycle for pencil-beam based FMO and BAO, followed by the proposed algorithm from chapters 3 & 4 for MLC segment generation aimed at close reproduction of the FMO dose distribution. These plans are referred to as 'AUTO BAO'. For validation, AUTO BAO plans were generated for 33 prostate SBRT patients and compared to reference plans (REF) that were manually generated with the commercial TPS, in absence of time pressure. REF plans were also compared to AUTO RB plans, for which fluence map optimization was performed for the beam angle configuration used in the REF plan, and the segmentation could use all these beams or only a subset, depending on the dosimetry.
AUTO BAO plans were clinically acceptable and dosimetrically similar to REF plans, but had on average reduced numbers of beams ((beams in AUTO BAO)/(beams in REF) (relative improvement): 24.7/48.3 (-49%)), segments (59.5/98.9 (-40%)), and delivery times (17.1/22.3 min. (-23%)). Dosimetry of AUTO RB and REF were also similar, while AUTO RB used on average fewer beams (38.0/48.3 (-21%)) and had on average shorter delivery times (18.6/22.3 min. (-17%)). Delivered Monitor Units (MU) were similar for all three planning approaches. A new, vendor-independent optimization workflow for fully automated generation of deliverable high-quality CyberKnife® plans was proposed, including BAO. Compared to manual planning with the commercial TPS, fraction delivery times were reduced by 5.3 min. (-23%) due to large reductions in beam and segment numbers.

In chapter 5, we proposed TBS-BAO v, a novel approach for solving the BAO problem, and tested it for non-coplanar robotic CyberKnife radiotherapy for prostate cancer. Properly selected beam angles contribute to the quality of radiotherapy treatment plans, however, the BAO problem is difficult to solve to optimality due to its non-convex discrete nature with many local minima. The proposed BAO method, first, generates an 'ideal' Pareto-optimal reference dose distribution using multi-criterial fluence map optimization (FMO) for a plan that includes all candidate beams ('total-beam-space', TBS). The following integrated BAO and segmentation for a preset maximum number of included beams then aims at realizing the ideal reference plan as accurately as possible. For each of the 33 prostate SBRT patients, 18 treatment plans with different pre-set numbers of

allowed beams, were automatically generated with the proposed TBS-BAO. For each patient, the TBS-BAO plans were then compared to a plan that was automatically generated with an alternative BAO method (Erasmus-iCycle) and to a high-quality manually generated plan. TBS-BAO was able to automatically generate plans with clinically feasible numbers of beams (∼25), with a quality highly similar to corresponding 91-beam ideal reference plans. Compared to the alternative Erasmus-iCycle BAO approach, similar plan quality was obtained for 25-beam segmented plans, while computation times were reduced from 10.7 hours to 4.8/1.5 hours, depending on the applied pencil-beam resolution in TBS-BAO. 25-beam TBS-BAO plans had similar quality as manually generated plans with on average 48 beams, while delivery times reduced from 22.3 to 18.4/18.1 min. TBS reference plans could effectively steer the discrete non-convex BAO.

In chapter 6, the proposed treatment planning solutions are discussed in a wider context, together with their challenges, opportunities, and potential future research.

Samenvatting

Stereotactische radiotherapie is een vorm van bestralingstherapie waarbij de patiënt wordt bestraald met een hoge dosis in een beperkt aantal dagelijkse fracties (meestal tussen de 3 tot 5 fracties). Voor deze manier van bestralen is een hoge precisie vereist om de (onvermijdelijke) dosis aan het omliggende gezonde weefsel zo laag mogelijk te houden. De CyberKnife® is een robotisch bestralingssysteem dat zulke stereotactische bestralingsplannen kan afstralen. Daarnaast is het met de CyberKnife mogelijk om zogeheten 'niet-coplanaire' bundelhoeken af te stralen zonder dat de tafel (waarop de patiënt ligt) daarvoor handmatig versteld hoeft te worden. Mede hierdoor kan de CyberKnife bestralingsplannen van hoge kwaliteit afstralen en zijn er veelbelovende klinische resultaten mee behaald (Fuller et al., 2018, Meier et al., 2018, van der Voort van Zyp et al., 2009). Echter, door de extra mogelijkheden van de CyberKnife is het lastiger om bestralingsplannen te maken van een consistente hoge kwaliteit dan voor conventionele radiotherapie, met name in een klinische setting waarbij planners onder tijdsdruk werken. Het doel van dit proefschrift was om nieuwe methodes te ontwikkelen en valideren voor het volledig automatisch optimaliseren van bestralingsplannen voor robotbestraling. Hierbij lag een focus op het behalen van klinisch afstraalbare bestralingsplannen en op het optimaliseren van de bundelhoeken.

Een bestralingsplan beschrijft in essentie de configuratie van het bestralingsapparaat, en de daaruit volgende dosisverdeling in de patiënt. Een dergelijk bestralingsplan beschrijft de bundelhoeken, de vormen van de bundels, en de intensiteit van de bundels die gebruikt gaan worden voor het bestralen van de patiënt. In het algemeen zijn er twee optimalisatie technieken voor het genereren van een bestralingsplan. De eerste techniek splitst de optimalisatie in twee fases: een zogeheten fluence map optimalisatie (FMO) fase waarbij de intensiteitsprofielen per bundel worden geoptimaliseerd, en een segmentatiefase

waarin de geoptimaliseerde intensiteitsprofielen worden omgezet in een afstraalbaar plan (bestaande uit multi-leaf collimator (MLC) segmenten en bijbehorende intensiteiten per MLC segment). Echter, er kan bij deze segmentatiefase kwaliteitsverlies optreden bij het converteren van de geoptimaliseerde intensiteitsprofielen naar een afstraalbaar plan. De tweede techniek is Direct Apeture Optimization (DAO) en optimaliseert direct op de MLC configuraties. Met deze techniek is het bestralingsplan op elk moment van het optimalisatie proces afstraalbaar en is er daarom geen segmentatiefase nodig (met mogelijk kwaliteitsverlies). Voor deze techniek worden de (fysieke) voorwaarden en beperkingen van de MLC en het bestralingssysteem meegenomen in het optimalisatie proces. Een bijkomend nadeel is dat dit leidt tot een niet-convex optimalisatieprobleem, wat wiskundig moeilijk op te lossen is. Op onze afdeling gaat de voorkeur uit naar het optimaliseren van een bestralingsplan door middel van de eerste techniek (optimaliseren van intensiteitsprofielen per bundel gevolgd door het converteren naar een afstraalbaar plan). De reden hiervoor is dat wij op onze afdeling Pareto-optimale intensiteitsprofielen kunnen genereren met Erasmus-iCycle, een algoritme voor het automatisch genereren van een bestralingsplan door middel van het optimaliseren van bundelhoeken en intensiteitsprofielen op basis van meerdere doelfuncties.

In hoofdstuk 2 is er een nieuw segmentatie-algoritme geïntroduceerd voor het converteren van een hoge kwaliteit FMO dosisverdeling naar een afstraalbaar plan (bestaande uit MLC segmenten en bijbehorende intensiteiten per MLC segment). Het algoritme is gefocust op het minimaliseren van het kwaliteitsverlies door middel van het reconstrueren van de 3-dimensionale FMO dosisverdeling in plaats van het reconstrueren van de 2-dimensionale intensiteitsprofielen per bundelhoek, zoals wordt gedaan in bestaande gepubliceerde segmentatiemethodes. Voor het genereren van nieuwe MLC segmenten worden alle bundelhoeken tegelijkertijd meegenomen. Tevens wordt er bij het reconstrueren van de FMO dosisverdeling gefocust op het nauwkeurig reproduceren van de klinische doelstellingen met de hoogste prioriteiten. Het segmentatie-algoritme is getest op een dataset van 20 prostaatkanker patiënten, 15 hoofd-halskanker patiënten en 12 leverkanker patiënten. Hiervoor werden eerst automatisch Pareto-optimale FMO dosisverdelingen gegenereerd en vervolgens gesegmenteerd met het nieuwe segmentatie-algoritme. De gesegmenteerde plannen waren dosimetrisch gelijkwaardig aan de FMO dosisverdelingen, daarnaast kon er voor alle patiënten een klinisch acceptabel plan gegenereerd worden. Er werden substantiële verschillen gevonden tussen de FMO intensiteitsprofielen en de gesegmenteerde intensiteitsprofielen. Waarschijnlijk is dit een gevolg van het direct reconstrueren

van de 3-dimensionale dosisverdeling in plaats van het reconstrueren van de 2-dimensionale intensiteitsprofielen per bundelhoek.

In hoofdstuk 3 is het segmentatie-algoritme doorontwikkeld. Klinische software ('clinical dose engine', CDE) voor het nauwkeurig kunnen berekenen van de dosis afgifte per segment is hiervoor geïntegreerd in het segmentatie-algoritme. Voor het gezamenlijk gebruik van pencil-beam doses en segment doses in het segmentatie-algoritme waren nieuwe methodes nodig, hiervoor zijn drie implementaties van het algoritme met verschillende mates van integratie van de CDE ontwikkeld. De prestaties van de drie implementaties van het algoritme zijn met elkaar vergeleken voor een dataset van 20 prostaatkanker patiënten en 12 leverkanker patiënten. Hierbij werden de gesegmenteerde plannen met de FMO dosisverdelingen vergeleken. Alle drie de implementaties konden de FMO dosisverdelingen nauwkeurig reconstrueren. De implementatie met de volledig geïntegreerde CDE genereerde de hoogste kwaliteit dosisverdelingen met het laagste aantal monitor-eenheden en segmenten, ten koste van een langere rekentijd.

In hoofdstuk 4 is er een nieuw systeem geïntroduceerd en gevalideerd voor het volledig automatisch genereren van bestralingsplannen voor de CyberKnife® uitgerust met de InCise™ 2 MLC, inclusief niet-coplanaire bundelhoekoptimalisatie. De bestralingsplannen worden volledig onafhankelijk van het klinische TPS gegenereerd met behulp van Erasmus-iCycle voor de optimalisatie van de bundelhoeken en intensiteitsprofielen gevolgd door segmentatie met het algoritme ontwikkeld in de hoofdstukken 3 & 4. Deze automatisch gegenereerde plannen worden 'AUTO BAO' genoemd. Voor de validatie van het systeem zijn er AUTO BAO plannen gegenereerd voor 33 prostaatkanker SBRT patiënten en vergeleken met referentie plannen (REF), die handmatig en zonder tijdsdruk zijn gegenereerd in de commerciële TPS. Daarnaast zijn de REF plannen vergeleken met AUTO RB plannen, waarbij FMO optimalisatie was uitgevoerd voor de bundelhoeken uit de REF plannen. De daaropvolgende segmentatie kon alleen deze bundelhoeken gebruiken, of een gedeelte van de bundelhoeken afhankelijk van de dosimetrie van het plan. De AUTO BAO plannen waren klinisch acceptabel en dosimetrisch gelijkwaardig aan de REF plannen, maar hadden gemiddeld minder bundelhoeken ((bundelhoeken in AUTO BAO)/(bundelhoeken in REF) (relatief verschil): 24.7/48.3 (-49%)), segmenten (59.5/98.9 (-40%)) en behandeltijden (17.1/22.3 min. (-23%)). De AUTO RB en REF plannen waren ook dosimetric gelijkwaardig, echter AUTO RB gebruikte daarvoor gemiddeld minder bundelhoeken (38.0/48.3 (-21%)) en had gemiddeld kortere behandeltijden (18.6/22.3 min. (-17%)). Er was

geen significant verschil gevonden voor de benodigde monitor eenheden tussen de drie methodes. Een nieuw en onafhankelijk optimalisatiesysteem voor het automatisch genereren van afstraalbare plannen voor de CyberKnife met een hoge kwaliteit is ontwikkeld, inclusief niet-coplanaire bundelhoekoptimalisatie. Vergeleken met de handmatig gegenereerde plannen is de behandeltijd verkort met 5.3 min. (-23%) door een vermindering in het aantal bundelhoeken en segmenten.

Hoofdstuk 5 introduceert TBS-BAO (Total-Beam-Space – Beam Angle Optimization), een nieuwe methode voor het oplossen van het bundelhoek-optimalisatieprobleem. Het systeem is getest voor niet-coplanaire robotbestraling met de CyberKnife voor prostaatkanker. De kwaliteit van bestralingsplannen is mede afhankelijk van goed gekozen bundelhoeken, echter, het optimaal oplossen van het bundelhoek-optimalisatieprobleem is lastig doordat het probleem niet-convex en discreet is met veel lokale minima. De voorgestelde methode genereert eerst een 'ideaal' Pareto-optimale referentie dosisverdeling met behulp van FMO op basis van meerdere doelfuncties en gebruikt daarvoor alle mogelijke bundelhoeken ('total-beam-space', TBS). Daarna wordt deze ideale dosisverdeling zo nauwkeurig mogelijk nagebootst terwijl het aantal toegestane bundelhoeken wordt beperkt tot een door de gebruiker vooraf ingestelde waarde. Voor een dataset van 33 prostaatkanker SBRT patiënten wordt voor elke patiënt 18 bestralingsplannen gegenereerd, elk bestralingsplan met een ander vooraf ingesteld aantal toegestane bundelhoeken. Voor elke patiënt worden de TBS-BAO bestralingsplannen vergeleken met een automatisch gegenereerd bestralingsplan dat gebruikt maakt van een alternatieve bundelhoek-optimalisatiemethode (Erasmus-iCycle), en met een handmatig gegenereerd bestralingsplan. Automatisch gegenereerde TBS-BAO bestralingsplannen met een klinisch haalbaar aantal bundelhoeken (~25) waren dosimetrisch vergelijkbaar met het ideale referentie plan met 91 bundelhoeken. De TBS-BAO plannen en Erasmus-iCycle plannen hadden een vergelijkbare plan kwaliteit, terwijl de rekentijden werden verkort van 10.7 uur tot 4.8/1.5 uur, afhankelijk van de toegepaste pencil-beam resolutie in TBS-BAO. TBS-BAO plannen met 25 bundelhoeken waren vergelijkbaar in kwaliteit ten op zichten van handmatige gegenereerde plannen met gemiddeld 48 bundelhoeken, terwijl de behandeltijden werden verkort van 22.3 minuten to 18.4/18.1 minuten.

Hoofdstuk 6 bevat een algemene discussie over de voorgestelde methodes voor het optimaliseren van bestralingsplannen voor de CyberKnife, inclusief uitdagingen, kansen en suggesties voor toekomstig onderzoek.

References

Accuray Inc. *CyberKnife® Robotic Radiosurgery System - Physics Essential Guide*, 10.6.x/5.3.x/3.x edition, 09 2015. page: 436 (5-60).

R. K. Ahuja and H. W. Hamacher. A network flow algorithm to minimize beam-on time for unconstrained multileaf collimator problems in cancer radiation therapy. *Networks*, 45(1):36–41, 2005. ISSN 1097-0037. doi: 10.1002/net.20047.

M. Alber and R. Reemtsen. Intensity modulated radiotherapy treatment planning by use of a barrier-penalty multiplier method. *Optim. Methods Softw.*, 22:391–411, 2007. doi: 10.1080/10556780600604940.

M. Alber, G. Meedt, and F. Nüsslin. On the degeneracy of the IMRT optimization problem. *Med. Phys.*, 29:2584–2589, 2002. doi: 10.1118/1.1500402.

D. M. Aleman, A. Kumar, R. K. Ahunja, H. E. Romeijn, and J. Dempsey. Neighborhood search approaches to beam orientation optimization in intensity modulated radiation therapy treatment planning. *J. Glob. Optim.*, 42:587–607, 2008. doi: 10.1007/s10898-008-9286-x.

S. Aluwini, P. van Rooij, M. Hoogeman, C. Bangma, W. J. Kirkels, L. Incrocci, , and I.-K. Kolkman-Deurloo. Cyberknife stereotactic radiotherapy as monotherapy for low-to intermediate-stage prostate cancer: early experience, feasibility and tolerance. *J. Endourol.*, 24(5):865–869, 2010. doi: 10.1089/end.2009.0438.

S. Aluwini, P. van Rooij, M. Hoogeman, W. Kirkels, I. K. Kolkman-Deurloo, and C. Bangma. Stereotactic body radiotherapy with a focal boost to the MRI-visible tumor as monotherapy for low- and intermediate-risk prostate cancer: early results. *Radiat. Oncol.*, 8:84, Apr 2013. doi: 10.1186/1748-717X-8-84.

G. Amit, T. G. Purdie, A. Levinshtein, A. J. Hope, P. Lindsay, A. Marshall, D. A. Jaffray, and V. Pekar. Automatic learning-based beam angle selection for thoracic IMRT. *Med. Phys.*, 42(4):1992–2005, Apr 2015. doi: 10.1118/1.4908000.

G. Asmerom, D. Bourne, J. Chappelow, L. M. Goggin, R. Heitz, P. Jordan, W. Kilby, T. Laing, C. R. M. Jr, J. M. Noll, S. Sayeh, and A. Weber. The design and physical characterization of a multileaf collimator for robotic radiosurgery. *Biomed. Phys. Eng. Express*, 2(1):017003, 2016. doi: 10.1088/2057-1976/2/1/017003.

T. Bai, B. Wang, D. Nguyen, and S. Jiang. Deep dose plugin: Towards real-time monte carlo dose calculation through a deep learning-based denoising algorithm. *Machine Learning: Science and Technology*, 2(2), June 2021. ISSN 2632-2153. doi: 10.1088/2632-2153/abdbfe.

M. Bangert and J. Unkelbach. Accelerated iterative beam angle selection in IMRT. *Med. Phys.*, 43(3):1073–1082, Mar 2016. doi: 10.1118/1.4940350.

M. Bangert, P. Ziegenhein, and U. Oelfke. Characterizing the combinatorial beam angle selection problem. *Phys. Med. Biol.*, 57(20):6707–6723, Oct 2012. doi: 10.1088/0031-9155/57/20/6707.

J. L. Bedford, H. S. Tsang, S. Nill, and U. Oelfke. Treatment planning optimization with beam motion modeling for dynamic arc delivery of SBRT using Cyberknife with multileaf collimation. *Med. Phys.*, 46(12):5421–5433, Dec 2019a. doi: 10.1002/mp.13848.

J. L. Bedford, P. Ziegenhein, S. Nill, and U. Oelfke. Beam selection for stereotactic ablative radiotherapy using Cyberknife with multileaf collimation. *Med. Eng. Phys.*, 64:28–36, 02 2019b. doi: 10.1016/j.medengphy.2018.12.011.

J. L. Bedford, S. Nill, and U. Oelfke. Dosimetric accuracy of delivering SBRT using dynamic arcs on Cyberknife. *Med. Phys.*, 47(4):1533–1544, Apr 2020. doi: 10.1002/mp.14090.

S. L. Berry, A. Boczkowski, R. Ma, J. Mechalakos, and M. Hunt. Interobserver variability in radiation therapy plan output: Results of a single-institution study. *Pract. Radiat. Oncol.*, 6(6):442–449, 2016. doi: 10.1016/j.prro.2016.04.005.

D. Bertsimas, V. Cacchiani, D. Craft, and O. Nohadani. A hybrid approach to beam angle optimization in intensity-modulated radiation therapy. *Comput. Oper. Res.*, 40:2187–2197, 2013. doi: 10.1016/j.cor.2012.06.009.

R. Bijman, L. Rossi, A. W. Sharfo, W. Heemsbergen, L. Incrocci, S. Breedveld, and B. Heijmen. Automated Radiotherapy Planning for Patient-Specific Exploration

of the Trade-Off Between Tumor Dose Coverage and Predicted Radiation-Induced Toxicity-A Proof of Principle Study for Prostate Cancer. *Front Oncol*, 10:943, 2020. doi: 10.3389/fonc.2020.00943.

R. Bijman, A. W. Sharfo, L. Rossi, S. Breedveld, and B. Heijmen. Pre-clinical validation of a novel system for fully-automated treatment planning. *Radiother Oncol*, 158: 253–261, 05 2021. doi: 10.1016/j.radonc.2021.03.003.

G. Bohara, A. Sadeghnejad Barkousaraie, S. Jiang, and D. Nguyen. Using deep learning to predict beam-tunable pareto optimal dose distribution for intensity-modulated radiation therapy. *Med. Phys.*, 47(9):3898–3912, 2020. doi: 10.1002/mp.14374.

R. Bokrantz and K. Miettinen. Projections onto the Pareto surface in multicriteria radiation therapy optimization. *Med. Phys.*, 42:5862, 2015. doi: 10.1118/1.4930252.

N. Boland, H. W. Hamacher, and F. Lenzen. Minimizing beam-on time in cancer radiation treatment using multileaf collimators. *Networks*, 43(4):226–240, 2004. doi: 10.1002/net.20007.

S. Breedveld, P. Storchi, M. Keijzer, and B. Heijmen. Fast, multiple optimizations of quadratic dose objective functions in IMRT. *Phys. Med. Biol.*, 51:3569–3579, 2006. doi: 10.1088/0031-9155/51/14/019.

S. Breedveld, P. Storchi, and B. Heijmen. The equivalence of multi-criteria methods for radiotherapy plan optimization. *Phys. Med. Biol.*, 54:7199–7209, 2009. doi: 10.1088/0031-9155/54/23/011.

S. Breedveld, P. Storchi, P. Voet, and B. Heijmen. iCycle: Integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans. *Med. Phys.*, 39:951–963, 2012. doi: 10.1118/1.3676689.

S. Breedveld, B. van den Berg, and B. Heijmen. An interior-point implementation developed and tuned for radiation therapy treatment planning. *Comput. Optim. Appl.*, 68(2):209–242, Nov 2017. ISSN 1573-2894. doi: 10.1007/s10589-017-9919-4.

S. Breedveld, A. B. A. Bennan, S. Aluwini, D. R. Schaart, I. K. Kolkman-Deurloo, and B. J. M. Heijmen. Fast automated multi-criteria planning for HDR brachytherapy explored for prostate cancer. *Phys. Med. Biol.*, 64(20):205002, 10 2019a. doi: 10.1088/1361-6560/ab44ff.

S. Breedveld, D. Craft, R. van Haveren, and B. Heijmen. Multi-criteria optimization and decision-making in radiotherapy. *Eur. J. Oper. Res.*, 277(1):1 – 19, 2019b. ISSN 0377-2217. doi: 10.1016/j.ejor.2018.08.019.

D. Buergy, A. W. Sharfo, B. J. Heijmen, P. W. Voet, S. Breedveld, F. Wenz, F. Lohr, and F. Stieler. Fully automated treatment planning of spinal metastases - A comparison to manual planning of Volumetric Modulated Arc Therapy for conventionally fractionated irradiation. *Radiat. Oncol.*, 12(1):33, Jan 2017. doi: 10.1186/S13014-017-0767-2.

M. Buschmann, A. W. M. Sharfo, J. Penninkhof, Y. Seppenwoolde, G. Goldner, D. Georg, S. Breedveld, and B. J. M. Heijmen. Automated volumetric modulated arc therapy planning for whole pelvic prostate radiotherapy. *Strahlenther Onkol*, 194(4):333–342, 04 2018. doi: 10.1007/s00066-017-1246-2.

S. Calusi, R. Doro, V. Di Cataldo, S. Cipressi, G. Francolini, I. Bonucci, L. Livi, and L. Masi. Performance assessment of a new optimization system for robotic SBRT MLC-based plans. *Phys. Med.*, 71:31–38, Mar 2020. doi: 10.1016/j.ejmp.2020.02.009.

F. Carlsson. Combining segment generation with direct step-and-shoot optimization in intensity-modulated radiation therapy. *Med. Phys.*, 35(9):3828–3838, Sep 2008. doi: 10.1118/1.2964096.

F. Carlsson and A. Forsgren. On column generation approaches for approximate solutions of quadratic programs in intensity-modulated radiation therapy. *Ann. Oper. Res.*, 223(1):471–481, Dec 2014. doi: 10.1007/s10479-013-1360-1.

A. Cassioli and J. Unkelbach. Aperture shape optimization for IMRT treatment planning. *Phys. Med. Biol.*, 58(2):301, 2013. doi: 10.1088/0031-9155/58/2/301.

D. Craft. Local beam angle optimization with linear programming and gradient search. *Phys. Med. Biol.*, 52(7):N127–135, Apr 2007. doi: 10.1088/0031-9155/52/7/n02.

D. Craft and C. Richter. Deliverable navigation for multicriteria step and shoot IMRT treatment planning. *Phys. Med. Biol.*, 58(1):87–103, Jan 2013. doi: 10.1088/0031-9155/58/1/87.

D. Craft, P. Süss, and T. Bortfeld. The tradeoff between treatment plan quality and required number of monitor units in intensity-modulated radiotherapy. *Int. J. Radiat. Oncol. Biol. Phys.*, 67:1596–1605, 2007. doi: 10.1016/j.ijrobp.2006.11.034.

S. M. Crooks, L. F. McAven, D. F. Robinson, and L. Xing. Minimizing delivery time and monitor units in static IMRT by leaf-sequencing. *Phys. Med. Biol.*, 47(17):3105–3116, Sep 2002. doi: 10.1088/0031-9155/47/17/305.

G. Della Gala, M. L. P. Dirkx, N. Hoekstra, D. Fransen, N. Lanconelli, M. van de Pol, B. J. M. Heijmen, and S. F. Petit. Fully automated VMAT treatment planning for

advanced-stage NSCLC patients. *Strahlenther. Onkol.*, 193(5):402–409, May 2017. doi: 10.1007/s00066-017-1121-1.

P. Dong, P. Lee, D. Ruan, T. Long, E. Romeijn, D. A. Low, P. Kupelian, J. Abraham, Y. Yang, and K. Sheng. 4π noncoplanar stereotactic body radiation therapy for centrally located or larger lung tumors. *Int. J. Radiat. Oncol. Biol. Phys.*, 86(3):407–413, Jul 2013a. doi: 10.1016/j.ijrobp.2013.02.002.

P. Dong, P. Lee, D. Ruan, T. Long, E. Romeijn, Y. Yang, D. Low, P. Kupelian, and K. Sheng. 4Pi non-coplanar liver SBRT: A novel delivery technique. *Int. J. Radiat. Oncol. Biol. Phys.*, 85(5):1360–1366, 2013b. doi: 10.1016/j.ijrobp.2012.09.028.

M. Duijm, B. Schipaanboord, P. V. Granton, and J. Nuyttens. Local Reirradiation of Recurrent Non-small Cell Lung Carcinoma Resulting in Long Disease-free Survival, Although in the Presence of Osteonecrosis. *Cureus*, 10(10):e3471, Oct 2018. doi: 10.7759/cureus.3471.

A. Fogliata, F. Belosi, A. Clivio, P. Navarria, G. Nicolini, M. Scorsetti, E. Vanetti, and L. Cozzi. On the pre-clinical validation of a commercial model-based optimisation engine: application to volumetric modulated arc therapy for patients with lung or prostate cancer. *Radiother. Oncol.*, 113(3):385–391, Dec 2014. doi: 10.1016/j.radonc.2014.11.009.

D. B. Fuller, A. D. Falchook, T. Crabtree, B. L. Kane, C. A. Medbery, K. Underhill, J. R. Gray, A. Peddada, and R. C. Chen. Phase 2 Multicenter Trial of Heterogeneous-dosing Stereotactic Body Radiotherapy for Low- and Intermediate-risk Prostate Cancer: 5-year Outcomes. *Eur. Urol. Oncol.*, 1(6):540–547, 12 2018. doi: 10.1016/j.euo.2018.06.013.

F. R. Giglioli, L. Strigari, R. Ragona, G. R. Borzì, E. Cagni, C. Carbonini, S. Clemente, R. Consorti, R. El Gawhary, M. Esposito, M. D. Falco, D. Fedele, C. Fiandra, M. C. Frassanito, V. Landoni, G. Loi, E. Lorenzini, M. R. Malisan, C. Marino, E. Menghi, B. Nardiello, R. Nigro, C. Oliviero, G. Pastore, M. Quattrocchi, R. Ruggieri, I. Redaelli, G. Reggiori, S. Russo, E. Villaggi, M. Casati, and P. Mancosu. Lung stereotactic ablative body radiotherapy: A large scale multi-institutional planning comparison for interpreting results of multi-institutional studies. *Phys. Med.*, 32(4):600–606, Apr 2016. doi: 10.1016/j.ejmp.2016.03.015.

F. R. Giglioli, C. Garibaldi, O. Blanck, E. Villaggi, S. Russo, M. Esposito, C. Marino, M. Stasi, and P. Mancosu. Dosimetric Multicenter Planning Comparison Studies for Stereotactic Body Radiation Therapy: Methodology and Future Perspectives. *Int. J. Radiat. Oncol. Biol. Phys.*, 106(2):403–412, 02 2020. doi: 10.1016/j.ijrobp.2019.10.041.

M. K. Giżyńska, L. Rossi, W. den Toom, M. T. W. Milder, K. C. de Vries, J. Nuyttens, and B. J. M. Heijmen. Largely reduced OAR doses, and planning and delivery times for challenging robotic SBRT cases, obtained with a novel optimizer. *J. Appl. Clin. Med. Phys.*, 22(3):35–47, Mar 2021. doi: 10.1002/acm2.13172.

M. Gören and Z. C. Taşkin. A column generation approach for evaluating delivery efficiencies of collimator technologies in IMRT treatment planning. *Phys. Med. Biol.*, 60(5):1989–2004, Mar 2015. doi: 10.1088/0031-9155/60/5/1989.

C. R. Hansen, A. Bertelsen, I. Hazell, R. Zukauskaite, N. Gyldenkerne, J. Johansen, J. G. Eriksen, and C. Brink. Automatic treatment planning improves the clinical quality of head and neck cancer treatment plans. *Clin. Transl. Radiat. Oncol.*, 1:2–8, Dec 2016. doi: 10.1016/j.ctro.2016.08.001.

C. R. Hansen, M. Nielsen, A. S. Bertelsen, I. Hazell, E. Holtved, R. Zukauskaite, J. K. Bjerregaard, C. Brink, and U. Bernchou. Automatic treatment planning facilitates fast generation of high-quality treatment plans for esophageal cancer. *Acta Oncol.*, 56(11):1495–1500, Nov 2017. doi: 10.1080/0284186x.2017.1349928.

C. R. Hansen, W. Crijns, M. Hussein, L. Rossi, P. Gallego, W. Verbakel, J. Unkelbach, D. Thwaites, and B. Heijmen. Radiotherapy Treatment plannINg study Guidelines (RATING): A framework for setting up and reporting on scientific treatment planning studies. *Radiother. Oncol.*, 153:67–78, 12 2020. doi: 10.1016/j.radonc.2020.09.033.

S. T. Heijkoop, T. R. Langerak, S. Quint, L. Bondar, J. W. Mens, B. J. Heijmen, and M. S. Hoogeman. Clinical implementation of an online adaptive plan-of-the-day protocol for nonrigid motion management in locally advanced cervical cancer IMRT. *Int. J. Radiat. Oncol. Biol. Phys.*, 90(3):673–679, Nov 2014. doi: 10.1016/j.ijrobp.2014.06.046.

B. Heijmen, P. Voet, D. Fransen, J. Penninkhof, M. Milder, H. Akhiat, P. Bonomo, M. Casati, D. Georg, G. Goldner, A. Henry, J. Lilley, F. Lohr, L. Marrazzo, S. Pallotta, R. Pellegrini, Y. Seppenwoolde, G. Simontacchi, V. Steil, F. Stieler, S. Wilson, and S. Breedveld. Fully automated, multi-criterial planning for Volumetric Modulated Arc Therapy - An international multi-center validation for prostate cancer. *Radiother. Oncol.*, 128(2):343–348, 08 2018. doi: 10.1016/j.radonc.2018.06.023.

M. Hussein, C. P. South, M. A. Barry, E. J. Adams, T. J. Jordan, A. J. Stewart, and A. Nisbet. Clinical validation and benchmarking of knowledge-based IMRT and VMAT treatment planning in pelvic anatomy. *Radiother. Oncol.*, 120(3):473–479, 09 2016. doi: 10.1016/j.radonc.2016.06.022.

M. Hussein, B. J. M. Heijmen, D. Verellen, and A. Nisbet. Automation in intensity modulated radiotherapy treatment planning - a review of recent innovations. *Br. J. Radiol.*, 91(1092):20180270, Dec 2018. doi: 10.1259/bjr.20180270.

V. Kearney, J. P. Cheung, C. McGuinness, and T. D. Solberg. CyberArc: a non-coplanar-arc optimization algorithm for CyberKnife. *Phys. Med. Biol.*, 62(14):5777–5789, Jun 2017. doi: 10.1088/1361-6560/aa6f92.

V. Kearney, M. Descovich, A. Sudhyadhom, J. P. Cheung, C. McGuinness, and T. D. Solberg. A continuous arc delivery optimization algorithm for CyberKnife m6. *Med. Phys.*, 45(8):3861–3870, Jun 2018. doi: 10.1002/mp.13022.

C. Kontaxis, G. H. Bol, J. J. Lagendijk, and B. W. Raaymakers. Towards adaptive IMRT sequencing for the MR-linac. *Phys. Med. Biol.*, 60(6):2493–2509, 03 2015a. doi: 10.1088/0031-9155/60/6/2493.

C. Kontaxis, G. H. Bol, J. J. Lagendijk, and B. W. Raaymakers. A new methodology for inter- and intrafraction plan adaptation for the MR-linac. *Phys. Med. Biol.*, 60(19): 7485–7497, 10 2015b. doi: 10.1088/0031-9155/60/19/7485.

C. Kontaxis, G. H. Bol, L. G. W. Kerkmeijer, J. J. W. Lagendijk, and B. W. Raaymakers. Fast online replanning for interfraction rotation correction in prostate radiotherapy. *Med. Phys.*, 44(10):5034–5042, Oct 2017a. doi: 10.1002/mp.12467.

C. Kontaxis, G. H. Bol, B. Stemkens, M. Glitzner, F. M. Prins, L. G. W. Kerkmeijer, J. J. W. Lagendijk, and B. W. Raaymakers. Towards fast online intrafraction replanning for free-breathing stereotactic body radiation therapy with the MR-linac. *Phys. Med. Biol.*, 62(18):7233–7248, Aug 2017b. doi: 10.1088/1361-6560/aa82ae.

C. Kontaxis, G. H. Bol, J. J. W. Lagendijk, and B. W. Raaymakers. DeepDose: Towards a fast dose calculation engine for radiation therapy using deep learning. *Phys. Med. Biol.*, 65(7):075013, 04 2020. doi: 10.1088/1361-6560/ab7630.

T. Long, M. Chen, S. Jiang, and W. Lu. Continuous leaf optimization for IMRT leaf sequencing. *Med. Phys.*, 43(10):5403, Oct 2016. doi: 10.1118/1.4962030.

S. Luan, C. Wang, D. Z. Chen, X. S. Hu, S. A. Naqvi, X. Wu, and C. X. Yu. An improved MLC segmentation algorithm and software for step-and-shoot IMRT delivery without tongue-and-groove error. *Med. Phys.*, 33(5):1199–1212, May 2006. doi: 10.1118/1.2188823.

C. Marino, E. Villaggi, G. Maggi, M. Esposito, L. Strigari, E. Bonanno, G. R. Borzì, C. Carbonini, R. Consorti, D. Fedele, C. Fiandra, I. Ielo, T. Malatesta, M. R. Malisan, A. Martinotti, R. Moretti, B. Nardiello, C. Oliviero, S. Clemente, and P. Mancosu. A

feasibility dosimetric study on prostate cancer. *Strahlenther. Onkol.*, 191(7):573–581, Jul 2015. ISSN 1439-099X. doi: 10.1007/s00066-015-0822-6.

L. Marrazzo, I. Meattini, C. Arilli, S. Calusi, M. Casati, C. Talamonti, L. Livi, and S. Pallotta.   Auto-planning for VMAT accelerated partial breast irradiation. *Radiother. Oncol.*, 132:85–92, 03 2019. doi: 10.1016/j.radonc.2018.11.006.

A. L. McNiven, M. B. Sharpe, and T. G. Purdie.   A new metric for assessing IMRT modulation complexity and plan deliverability. *Med Phys*, 37(2):505–515, Feb 2010. doi: 10.1118/1.3276775.

R. M. Meier, D. A. Bloch, C. Cotrutz, A. C. Beckman, G. T. Henning, S. A. Woodhouse, S. K. Williamson, N. Mohideen, J. J. Dombrowski, R. L. Hong, D. G. Brachman, P. W. Linson, and I. D. Kaplan.   Multicenter Trial of Stereotactic Body Radiation Therapy for Low- and Intermediate-Risk Prostate Cancer: Survival and Toxicity Endpoints.   *Int. J. Radiat. Oncol. Biol. Phys.*, 102(2):296–303, 10 2018.   doi: 10.1016/j.ijrobp.2018.05.040.

C. Men, E. Romeijn, C. Taşkin, and J. Dempsey. An exact approach to direct aperture optimization in IMRT treatment planning. *Phys. Med. Biol.*, 52:7333–7352, 2007. doi: 10.1088/0031-9155/52/24/009.

C. Men, X. Jia, and S. Jiang.   GPU-based ultra-fast direct aperture optimization for online adaptive radiation therapy.   *Phys. Med. Biol.*, 55:4309–4320, 2010.   doi: 10.1088/0031-9155/55/15/008.

V. L. Murzin, K. Woods, V. Moiseenko, R. Karunamuni, K. R. Tringale, T. M. Seibert, M. J. Connor, D. R. Simpson, K. Sheng, and J. A. Hattangadi-Gluth. 4π plan optimization for cortical-sparing brain radiotherapy.   *Radiother. Oncol.*, 127(1):128–135, 04 2018. doi: 10.1016/j.radonc.2018.02.011.

B. E. Nelms, G. Robinson, J. Markham, K. Velasco, S. Boyd, S. Narayan, J. Wheeler, and M. L. Sobczak. Variation in external beam treatment plan quality: An inter-institutional study of planners and planning systems. *Pract. Radiat. Oncol.*, 2(4):296–305, 2012. doi: 10.1016/j.prro.2011.11.012.

D. Nguyen, D. O'Connor, V. Y. Yu, D. Ruan, M. Cao, D. A. Low, and K. Sheng. Dose domain regularization of MLC leaf patterns for highly complex IMRT plans. *Med. Phys.*, 42 (4):1858–1870, Apr 2015. doi: 10.1118/1.4915286.

D. Nguyen, D. Ruan, D. O'Connor, K. Woods, D. A. Low, S. Boucher, and K. Sheng. A novel software and conceptual design of the hardware platform for intensity modulated radiation therapy. *Med. Phys.*, 43(2):917–929, Feb 2016. doi: 10.1118/1.4940353.

D. Nguyen, D. O'Connor, D. Ruan, and K. Sheng. Deterministic direct aperture optimization using multiphase piecewise constant segmentation. *Med. Phys.*, 44 (11):5596–5609, Nov 2017. doi: 10.1002/mp.12529.

M. Oud, I. K. Kolkman-Deurloo, J. W. Mens, D. Lathouwers, Z. Perkó, B. Heijmen, and S. Breedveld. Fast and fully-automated multi-criterial treatment planning for adaptive HDR brachytherapy for locally advanced cervical cancer. *Radiother. Oncol.*, 148:143–150, Jul 2020. doi: 10.1016/j.radonc.2020.04.017.

T. G. Purdie, R. E. Dinniwell, A. Fyles, and M. B. Sharpe. Automation and intensity modulated radiation therapy for individualized high-quality tangent breast treatment plans. *Int. J. Radiat. Oncol. Biol. Phys.*, 90(3):688–695, Nov 2014. doi: 10.1016/j.ijrobp.2014.06.056.

L. Redapi, L. Rossi, L. Marrazzo, J. J. Penninkhof, S. Pallotta, and B. Heijmen. Comparison of volumetric modulated arc therapy and intensity-modulated radiotherapy for left-sided whole-breast irradiation using automated planning. *Strahlenther Onkol*, Aug 2021. doi: 10.1007/s00066-021-01817-x.

H. Rocha, J. Dias, T. Ventura, B. Ferreira, and M. Lopes. A derivative-free multistart framework for an automated noncoplanar beam angle optimization in IMRT. *Med. Phys.*, 43:5514–5526, 2016. doi: 10.1118/1.4962477.

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.*, 15:838–862, 2005. doi: 10.1137/040606612.

L. Rossi, S. Breedveld, B. J. M. Heijmen, P. W. J. Voet, N. Lanconelli, and S. Aluwini. On the beam direction search space in computerized non-coplanar beam angle optimization for IMRT - prostate SBRT. *Phys. Med. Biol.*, 57:5441–5458, 2012. doi: 10.1088/0031-9155/57/17/5441.

L. Rossi, S. Breedveld, S. Aluwini, and B. Heijmen. Non-coplanar beam angle class solutions to replace time-consuming patient-specific beam angle optimization in robotic prostate SBRT. *Int. J. Radiat. Oncol. Biol. Phys.*, 92:762–770, 2015. doi: 10.1016/j.ijrobp.2015.03.013.

L. Rossi, A. W. Sharfo, S. Aluwini, M. Dirkx, S. Breedveld, and B. Heijmen. First fully automated planning solution for robotic radiosurgery - comparison with automatically planned volumetric arc therapy for prostate cancer. *Acta Oncol.*, 57 (11):1490–1498, Nov 2018. doi: 10.1080/0284186x.2018.1479068.

L. Rossi, A. Méndez Romero, M. Milder, E. de Klerck, S. Breedveld, and B. Heijmen. Individualized automated planning for dose bath reduction in robotic radiosurgery for benign tumors. *PLoS One*, 14(2):e0210279, 2019. doi: 10.1371/JOURNAL.PONE.0210279.

L. Rossi, P. Cambraia Lopes, J. Marques Leitão, C. Janus, M. van de Pol, S. Breedveld, J. Penninkhof, and B. J. Heijmen. On the importance of individualized, non-coplanar beam configurations in mediastinal lymphoma radiotherapy, optimized with automated planning. *Front. Oncol.*, 11:951, 2021. ISSN 2234-943X. doi: 10.3389/fonc.2021.619929.

J. C. Rwigema, D. Nguyen, D. E. Heron, A. M. Chen, P. Lee, P. C. Wang, J. A. Vargo, D. A. Low, M. S. Huq, S. Tenn, M. L. Steinberg, P. Kupelian, and K. Sheng. 4π noncoplanar stereotactic body radiation therapy for head-and-neck cancer: potential to improve tumor control and late toxicity. *Int. J. Radiat. Oncol. Biol. Phys.*, 91(2):401–409, Feb 2015. doi: 10.1016/j.ijrobp.2014.09.043.

A. Sadeghnejad-Barkousaraie, O. Ogunmolu, S. Jiang, and D. Nguyen. A fast deep learning approach for beam orientation optimization for prostate cancer treated with intensity-modulated radiation therapy. *Med. Phys.*, 47(3):880–897, Mar 2020. doi: 10.1002/mp.13986.

A. Sadeghnejad-Barkousaraie, G. Bohara, S. B. Jiang, and D. Nguyen. A reinforcement learning application of a guided monte carlo tree search algorithm for beam orientation selection in radiation therapy. *Mach. Learn.: Sci. Technol.*, 2:035013, May 2021. doi: 10.1088/2632-2153/ABE528.

E. Salari and J. Unkelbach. A column-generation-based method for multi-criteria direct aperture optimization. *Phys. Med. Biol.*, 58(3):621–639, Feb 2013. doi: 10.1088/0031-9155/58/3/621.

B. Schipaanboord, D. Boukerroui, D. Peressutti, J. van Soest, T. Lustberg, A. Dekker, W. van Elmpt, and M. J. Gooding. An evaluation of atlas selection methods for atlas-based automatic segmentation in radiotherapy treatment planning. *IEEE Trans. Med. Imaging*, Apr 2019a. doi: 10.1109/tmi.2019.2907072.

B. Schipaanboord, D. Boukerroui, D. Peressutti, J. van Soest, T. Lustberg, T. Kadir, A. Dekker, W. van Elmpt, and M. Gooding. Can Atlas-Based Auto-Segmentation Ever Be Perfect? Insights From Extreme Value Theory. *IEEE Trans. Med. Imaging*, 38(1):99–106, 01 2019b. doi: 10.1109/tmi.2018.2856464.

B. W. K. Schipaanboord, S. Breedveld, L. Rossi, M. Keijzer, and B. Heijmen. Automated prioritised 3D dose-based MLC segment generation for step-and-shoot IMRT. *Phys. Med. Biol.*, 64(16):165013, Aug 2019c. doi: 10.1088/1361-6560/ab1df9.

B. W. K. Schipaanboord, B. Heijmen, and S. Breedveld. Accurate 3D-dose-based generation of MLC segments for robotic radiotherapy. *Phys. Med. Biol.*, 65(17): 175011, 08 2020. doi: 10.1088/1361-6560/ab97e7.

B. W. K. Schipaanboord, M. K. Giżyńska, L. Rossi, K. C. de Vries, B. J. M. Heijmen, and S. Breedveld. Fully automated treatment planning for MLC-based robotic radiotherapy. *Med. Phys.*, 48(8):4139–4147, 2021. doi: 10.1002/mp.14993.

B. W. K. Schipaanboord, B. J. M. Heijmen, and S. Breedveld. TBS-BAO: fully automated beam angle optimization for IMRT guided by a total-beam-space reference plan. *Phys. Med. Biol.*, 67(3):035004, Jan 2022. doi: 10.1088/1361-6560/ac4b37.

E. Schüler, A. Lo, C. F. Chuang, S. G. Soltys, E. L. Pollom, and L. Wang. Clinical impact of the volo optimizer on treatment plan quality and clinical treatment efficiency for cyberknife. *J. Appl. Clin. Med. Phys.*, 21(5):38–47, 2020. doi: 10.1002/acm2.12851.

A. Sharfo, S. Breedveld, P. Voet, S. Heijkoop, J.-W. Mens, M. Hoogeman, and B. Heijmen. Validation of fully automated VMAT plan generation for library-based plan-of-the-day cervical cancer radiotherapy. *PLoS ONE*, 11:e0169202, 2016. doi: 10.1371/journal.pone.0169202.

A. W. Sharfo, P. Voet, S. Breedveld, J. W. Mens, M. Hoogeman, and B. Heijmen. Comparison of VMAT and IMRT strategies for cervical cancer patients using automated planning. *Radiother. Oncol.*, 114:395–401, 2015. doi: 10.1016/j.radonc.2015.02.006.

A. W. Sharfo, M. Dirkx, S. Breedveld, A. Méndez Romero, and B. Heijmen. VMAT plus a few computer-optimized non-coplanar IMRT beams (VMAT+) tested for liver SBRT. *Radiother. Oncol.*, 123:49–56, 2017. doi: 10.1016/j.radonc.2017.02.018.

A. W. M. Sharfo, F. Stieler, O. Kupfer, B. J. M. Heijmen, M. L. P. Dirkx, S. Breedveld, F. Wenz, F. Lohr, J. Boda-Heggemann, and D. Buergy. Automated VMAT planning for postoperative adjuvant treatment of advanced gastric cancer. *Radiat. Oncol.*, 13(1):74, Apr 2018. doi: 10.1186/S13014-018-1032-Z.

A. W. M. Sharfo, L. Rossi, M. L. P. Dirkx, S. Breedveld, S. Aluwini, and B. J. M. Heijmen. Complementing Prostate SBRT VMAT With a Two-Beam Non-Coplanar IMRT Class Solution to Enhance Rectum and Bladder Sparing With Minimum Increase in Treatment Time. *Front. Oncol.*, 11:620978, 2021. doi: 10.3389/fonc.2021.620978.

D. M. Shepard, M. A. Earl, X. A. Li, S. Naqvi, and C. Yu. Direct aperture optimization: A turnkey solution for step-and-shoot IMRT. *Med. Phys.*, 29:1007–1018, 2002. doi: 10.1118/1.1477415.

X. Sun and P. Xia. A new smoothing procedure to reduce delivery segments for static MLC-based IMRT planning. *Med. Phys.*, 31(5):1158–1165, May 2004. doi: 10.1118/1.1713279.

H. Sung, J. Ferlay, R. L. Siegel, M. Laversanne, I. Soerjomataram, A. Jemal, and F. Bray. Global cancer statistics 2020: Globocan estimates of incidence and mortality worldwide for 36 cancers in 185 countries. *CA Cancer. J. Clin.*, 71(3):209–249, 2021. doi: 10.3322/caac.21660.

P. Süss, K.-H. Küfer, and C. Thieke. Improved stratification algorithms for step-and-shoot MLC delivery in intensity-modulated radiation therapy. *Phys. Med. Biol.*, 52: 6039–6051, 2007. doi: 10.1088/0031-9155/52/19/022.

S. Södertröm and A. Brahme. Optimization of the dose delivery in a few field techniques using radiobiological objective functions. *Med. Phys.*, 20(4):1201–1210, 1993. doi: 10.1118/1.596971.

J. P. Tol, A. R. Delaney, M. Dahele, B. J. Slotman, and W. F. Verbakel. Evaluation of a knowledge-based planning solution for head and neck cancer. *Int. J. Radiat. Oncol. Biol. Phys.*, 91(3):612–620, Mar 2015. doi: 10.1016/j.ijrobp.2014.11.014.

M. Tomida, T. Kamomae, J. Suzuki, Y. Ohashi, Y. Itoh, H. Oguchi, and T. Okuda. Clinical usefulness of MLCs in robotic radiosurgery systems for prostate SBRT. *J. Appl. Clin. Med. Phys.*, Jul 2017. doi: 10.1002/acm2.12128.

A. Tran, J. Zhang, K. Woods, V. Yu, D. Nguyen, G. Gustafson, L. Rosen, and K. Sheng. Treatment planning comparison of IMPT, VMAT and 4π radiotherapy for prostate cases. *Radiat. Oncol.*, 12(1):10, Jan 2017. doi: 10.1186/s13014-016-0761-0.

N. C. van der Voort van Zyp, J. B. Prévost, M. S. Hoogeman, J. Praag, B. van der Holt, P. C. Levendag, R. J. van Klaveren, P. Pattynama, and J. J. Nuyttens. Stereotactic radiotherapy with real-time tumor tracking for non-small cell lung cancer: clinical outcome. *Radiother. Oncol.*, 91(3):296–300, Jun 2009. doi: 10.1016/j.radonc.2009.02.011.

A. van 't Riet, A. C. Mak, M. A. Moerland, L. H. Elders, and W. van der Zee. A conformation number to quantify the degree of conformality in brachytherapy and external beam irradiation: application to the prostate. *Int. J. Radiat. Oncol. Biol. Phys.*, 37(3):731–736, Feb 1997. doi: 10.1016/s0360-3016(96)00601-3.

P. Voet, S. Breedveld, M. Dirkx, P. Levendag, and B. Heijmen. Integrated multi-criterial optimization of beam angles and intensity profiles for coplanar and non-coplanar head and neck IMRT and implications for VMAT. *Med. Phys.*, 39:4858–4865, 2012. doi: 10.1118/1.4736803.

P. Voet, M. Dirkx, S. Breedveld, D. Fransen, P. Levendag, and B. Heijmen. Towards fully automated multi-criterial plan generation: a prospective clinical study. *Int. J. Radiat. Oncol. Biol. Phys.*, 85:866–872, 2013a. doi: 10.1016/j.ijrobp.2012.04.015.

P. Voet, M. Dirkx, S. Breedveld, and B. Heijmen. Automated generation of IMRT treatment plans for prostate cancer patients with metal hip prostheses: Comparison of different planning strategies. *Med. Phys.*, 40:071704, 2013b. doi: 10.1118/1.4808117.

P. Voet, M. Dirkx, S. Breedveld, A. Al-Mamgani, L. Incrocci, and B. Heijmen. Fully automated VMAT plan generation for prostate cancer patients. *Int. J. Radiat. Oncol. Biol. Phys.*, 88:1175–1179, 2014. doi: 10.1016/j.ijrobp.2013.12.046.

K. Woods, D. Nguyen, A. Tran, V. Y. Yu, M. Cao, T. Niu, P. Lee, and K. Sheng. Viability of Non-Coplanar VMAT for Liver SBRT as Compared to Coplanar VMAT and Beam Orientation Optimized 4π IMRT. *Adv. Radiat. Oncol.*, 1(1):67–75, 2016. doi: 10.1016/j.adro.2015.12.004.

P. Xia and L. J. Verhey. Multileaf collimator leaf sequencing algorithm for intensity modulated beams with multiple static segments. *Med. Phys.*, 25(8):1424–1434, Aug 1998. doi: 10.1118/1.598315.

Y. Xing, Y. Zhang, D. Nguyen, M. H. Lin, W. Lu, and S. Jiang. Boosting radiotherapy dose calculation accuracy with deep learning. *J Appl Clin Med Phys*, 21(8):149–159, Aug 2020. doi: 10.1002/ACM2.12937.

J. Yang, Z. Gui, L. Zhang, and P. Zhang. Aperture generation based on threshold segmentation for intensity modulated radiotherapy treatment planning. *Med. Phys.*, 45(4):1758–1770, Apr 2018. doi: 10.1002/mp.12819.

V. Y. Yu, A. Landers, K. Woods, D. Nguyen, M. Cao, D. Du, R. K. Chin, K. Sheng, and T. B. Kaprealian. A Prospective 4π Radiation Therapy Clinical Study in Recurrent High-Grade Glioma Patients. *Int. J. Radiat. Oncol. Biol. Phys.*, 101(1):144–151, 05 2018. doi: 10.1016/j.ijrobp.2018.01.048.

M. Zarepisheh, L. Hong, Y. Zhou, J. Hun Oh, J. G. Mechalakos, M. A. Hunt, G. S. Mageras, and J. O. Deasy. Automated intensity modulated treatment planning: The expedited constrained hierarchical optimization (ECHO) system. *Med. Phys.*, 46(7):2944–2954, Jul 2019. doi: 10.1002/mp.13572.

M. Zeverino, M. Marguet, C. Zulliger, A. Durham, R. Jumeau, F. Herrera, L. Schiappacasse, J. Bourhis, F. O. Bochud, and R. Moeckli. Novel inverse planning optimization algorithm for robotic radiosurgery: First clinical

implementation and dosimetric evaluation.   *Phys. Med.*, 64:230–237, Aug 2019. doi: 10.1016/j.ejmp.2019.07.020.

L. Zhang, P. Zhang, J. Yang, J. Li, and Z. Gui.   Aperture shape generation based on gradient descent with momentum.   *IEEE Access*, 7:157623–157632, 2019.   ISSN 2169-3536. doi: 10.1109/ACCESS.2019.2949871.

L. Zhu, T. Niu, K. Choi, and L. Xing.   Total-variation regularization based inverse planning for intensity modulated arc therapy.   *Technol. Cancer Res. Treat.*, 11(2): 149–162, Apr 2012. doi: 10.7785/tcrt.2012.500244.

List of publications

First author publications

- B. W. K. Schipaanboord, B. J. M. Heijmen, and S. Breedveld. TBS-BAO: fully automated beam angle optimization for IMRT guided by a total-beam-space reference plan. *Phys. Med. Biol.*, 67(3):035004, Jan 2022. doi: 10.1088/1361-6560/ac4b37

- B. W. K. Schipaanboord, M. K. Giżyńska, L. Rossi, K. C. de Vries, B. J. M. Heijmen, and S. Breedveld. Fully automated treatment planning for MLC-based robotic radiotherapy. *Med. Phys.*, 48(8):4139–4147, 2021. doi: 10.1002/mp.14993

- B. W. K. Schipaanboord, B. Heijmen, and S. Breedveld. Accurate 3D-dose-based generation of MLC segments for robotic radiotherapy. *Phys. Med. Biol.*, 65(17): 175011, 08 2020. doi: 10.1088/1361-6560/ab97e7

- B. W. K. Schipaanboord, S. Breedveld, L. Rossi, M. Keijzer, and B. Heijmen. Automated prioritised 3D dose-based MLC segment generation for step-and-shoot IMRT. *Phys. Med. Biol.*, 64(16):165013, Aug 2019c. doi: 10.1088/1361-6560/ab1df9

- B. Schipaanboord, D. Boukerroui, D. Peressutti, J. van Soest, T. Lustberg, A. Dekker, W. van Elmpt, and M. J. Gooding. An evaluation of atlas selection methods for atlas-based automatic segmentation in radiotherapy treatment planning. *IEEE Trans. Med. Imaging*, Apr 2019a. doi: 10.1109/tmi.2019.2907072

- B. Schipaanboord, D. Boukerroui, D. Peressutti, J. van Soest, T. Lustberg, T. Kadir, A. Dekker, W. van Elmpt, and M. Gooding. Can Atlas-Based Auto-Segmentation Ever Be Perfect? Insights From Extreme Value Theory.

IEEE Trans. Med. Imaging, 38(1):99–106, 01 2019b.   doi: 10.1109/tmi.2018.2856464

Co-author publications

- M. Duijm, B. Schipaanboord, P. V. Granton, and J. Nuyttens. Local Reirradiation of Recurrent Non-small Cell Lung Carcinoma Resulting in Long Disease-free Survival, Although in the Presence of Osteonecrosis. *Cureus*, 10(10):e3471, Oct 2018. doi: 10.7759/cureus.3471

PhD Portfolio

Name PhD Student	B.W.K. Schipaanboord
PhD period	2016 – 2021
Erasmus MC Department	Radiotherapy
Research School	Molecular Medicine
Promotor	Prof. dr. B.J.M. Heijmen
Copromotor	Dr. ir. S. Breedveld

General courses	
Biomedical English Writing and Communication	2019
Presenting Skills for junior researchers	2018
Scientific Integrity	2018

Specific courses		
ESTRO - Physics workshop	Barcelona, ES	2018
ESTRO - Research Masterclass in Radiotherapy Physics	Florence, IT	2017
Programming with python	Rotterdam, NL	2017

Presentations at international conferences			
ESTRO 2021	Oral presentation	Madrid, ES	2021
WIMP 2021	Poster presentation	Online	2021
ESTRO 2020	Electronic poster	Online	2020
WIMP 2020	Poster presentation	Breckenridge, US	2020
ICCR 2019	Poster presentation	Montreal, CA	2019
ESTRO 37	Oral presentation	Barcelona, ES	2018

Presentations at national meetings

NVKF 2021	Oral presentation	Online	2021
NVKF 2019	Oral presentation	Rotterdam, NL	2019

In-house presentations

Journal Club	2017, 2019, 2020
Research Day	2017, 2018
Physics R&D and WD	2016 – 2021
Referee evenings & lunches	2016 – 2021

Conferences, seminars, and workshops

ESTRO 2021	Madrid, ES	2021
NVKF 2021	Online	2021
WIMP 2021	Online	2021
ESTRO 2020	Online	2020
NVRO / NVKF 2020	Online	2020
NVKF 2020	Online	2020
WIMP 2020	Breckenridge, US	2020
NVKF 2019	Rotterdam, NL	2019
ICCR 2019	Montreal, CA	2019
ESTRO 37	Barcelona, ES	2018
Educational sessions		2019 – 2020
Radiotherapy Research Days		2017 – 2020
Research Rounds		2016 – 2020

Teaching activities

Supervising master student	Remy de Jong	2019 – 2020

Others

Integration of clinical dose engines	2017 – 2021
Two reviews for Physica Medica	2018, 2019

Distinctions and awards

Early Career Poster Competition – WIMP 2021	Online	2021
Rising Stars Award – ICCR 2019	Montreal, CA	2019
Best Research Pitch – Radiotherapy Research Day	Rotterdam, NL	2018

Curriculum Vitae

Bas Schipaanboord was born on 14 June 1990, in Zeewolde, the Netherlands. After graduating high school in 2008, he moved to Groningen to study 'Life Science and Technology' with a major in 'Biomedical Engineering'. He obtained his bachelor's degree in 2011 and continued with the master 'Biomedical Engineering' with the specialization 'Clinical Physics'. During his master he developed a prime interest in imaging techniques and medical devices and in his search to fully understand the fundamental physics of these devices, he started a second master: 'Applied Physics'. His master's theses involved pattern recognition techniques on PET data for the early diagnosis of neurodegenerative brain diseases.

In June 2015, he moved to Oxford, United Kingdom, for an internship at Mirada Medical, a software company specialized in medical software for diagnostic imaging and radiation oncology. After finishing his internship, he continued to work for Mirada Medical focusing on image registration and auto-contouring techniques.

In November 2016, Bas started his PhD at the radiotherapy department of the Erasmus Medical Center – Cancer Institute in Rotterdam under the supervision of Prof. dr. Ben Heijmen and Dr. ir. Sebastiaan Breedveld. His work focused on improving (automated) treatment planning for robotic radiotherapy, which resulted in this thesis.

Currently, Bas is working as a Clinical Scientist at the Netherlands Cancer Institute Antoni van Leeuwenhoek in Amsterdam, focusing on translating the latest technical developments into clinical practice.

Acknowledgments

My PhD has been a journey to never forget, it has taken me from the fundamental basics of medical physics to snowboarding in the Rocky Mountains. I would like to thank everyone that supported me throughout the years immensely, this thesis would not have been possible without your support. I would like to thank a few people in particular:

First of all, I like to thank my promotor Ben Heijmen. Thanks for your enthusiasm and the many fruitful discussions throughout the PhD. With the help of your critical view, all papers of this thesis got accepted in the journal of our first choice. You also taught me to keep an eye on the bigger picture and, more importantly, to not get lost in the details. Also thank you and Alessandra for your hospitality.

Many thanks to my co-promotor Sebastiaan Breedveld. I still remember when we met at the social event of ICCR 2016. You sparked my interest in the field of automated treatment planning and not long after I started a PhD at the Erasmus MC with you as my supervisor. Thanks for the support throughout the years, the many ideas, your mathematical insight, and also many thanks for Erasmus-iCycle, which has been a valuable asset during my PhD.

Special thanks to my paranymphs: Emma, thanks for riding the PhD rollercoaster with me. I could always come to you for help, to express my frustrations, or just for a good cup of coffee. Thanks for everything and for always being positive even when eye was not. Thank you, Daniel, you always made the effort to try to understand my (sometimes unnecessarily overly complicated) attempts to explain my PhD projects in detail. Thanks for the great times, not only during the PhD but also for the years before.

Thanks to all my colleagues, fellow PhD candidates, and room-mates throughout the years. You brought my PhD experience to the next level. Marloes, thank you for being my PhD-buddy. I especially enjoyed all our conversations about food,

restaurants, and wines, many thanks for all the recipes and good tips. Without you, my PhD wouldn't have been the same. Thank you Alba, Linda, Dominique, Thyrza, Rens, Rik, Jesús, Marie, Iva, Michelle, Nienke H., Nienke S., Kemal, Ioannis, and Xander for the many coffees and cakes together, the after-work socials, and making my PhD even more enjoyable.

To Andras Zolnay and Erik Jan Tromp from the IT-team for their continuous support with Matterhorn and with the implementation of the dose library. To Erik de Klerck for sharing his pictures of the CyberKnife and thanks to Wilhelm den Toom for your help with photographing the MLC of the CyberKnife.

Thanks to Warren Kilby and John Dooley from Accuray Inc. for their help and support when I had questions regarding the CyberKnife or the dose engine.

Thanks to Mark Gooding for teaching me the ins and outs of medical imaging and to all other colleagues from Mirada Medical for the wonderful time in Oxford.

Many thanks to the "Vrijmiboys" (Rick, Lars, Tom, Sven, Evert-Jan, and Jeroen). As the only beta of the group, I have probably bored you with my technical stories on many occasions. Thanks for the many valuable moments we've shared together and for all the fun and laughter throughout the years.

Special thanks to Emiel and Rein for the much-needed distractions on the weekends. Whether we were enjoying Rein's BBQ skills, pulling off heists, or helping each other with construction work, I could completely switch off from work.

To JC platypus (Peter, Daniel, Jorrit, Joran, Mark, Reinier, Johannes, Peter, Erik, Rogier, and Freek), you were there at the start of my Medical Physics career. Thanks for an amazing student life and for providing me with many procrastination opportunities.

To all my teammates at Footy and in particular to James and Maurice. Keep up the great work as waterboy and coach! I'm looking forward to when we can play together on the pitch again.

I especially want to thank my family for their unconditional support and love: my parents Wim & Nelie; Anitra & Tom, Josephine, Hugo, and Juliette; Celesta & Sander; Diantha & Matthijs; Esther & Ivor; and in-memory Opa & Oma Schip. I am truly grateful to have such a loving family, every time we get together it's an extremely fun chaos. I would also like to wish my sister Diantha all the best in obtaining her PhD in the future!

Finally, my biggest thanks goes out to Anna. You have been a tremendous support for me. Your positivity, your inexhaustible source of energy, your amazing cooking skills, your talent to find new restaurants and activities to do, and so much more has helped me through the PhD. I feel very lucky to have met you and I'm looking forward to all our future travels together.