

Automatic Configuration of Fast Automated Multi-Objective Treatment Planning in Radiotherapy

Rens van Haveren

Automatic Configuration of Fast Automated Multi-Objective Treatment Planning in Radiotherapy

Rens van Haveren

Automatic Configuration of Fast Automated Multi-Objective Treatment Planning in Radiotherapy

© Rens van Haveren 2021

Author: Rens van Haveren

Cover design: Li Bromfield

Layout: Rens van Haveren

Printing: Ridderprint | www.ridderprint.nl

ISBN: 978-94-6416-547-0

Automatic Configuration of Fast Automated Multi-Objective Treatment Planning in Radiotherapy

Automatische configuratie voor het snel en automatisch optimaliseren van bestralingsplannen in de radiotherapie op basis van meerdere doelfuncties

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Erasmus Universiteit Rotterdam
op gezag van de rector magnificus

Prof.dr. F.A. van der Duijn Schouten

en volgens besluit van het College voor Promoties.
De openbare verdediging zal plaatsvinden op

dinsdag 25 mei 2021 om 15:30 uur

door

Rens van Haveren
geboren te Korendijk

Promotiecommissie

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

Overige leden: Prof.dr. R.A. Nout
Prof.dr. U. Oelfke
Prof.dr. K. Miettinen

Copromotor: Dr.ir. S. Breedveld

Table of contents

1	Introduction	1
1.1	Radiotherapy	1
1.2	Radiotherapy treatment plan optimisation	1
1.3	Thesis outline	3
2	Lexicographic extension of the reference point method applied in radiation therapy treatment planning	5
	Abstract	6
2.1	Introduction	7
2.2	Multi-objective optimisation and trade-offs	9
2.3	Radiation therapy treatment planning	9
2.4	Lexicographic reference point method	12
2.5	Illustrative example	17
2.6	Prostate cancer study	20
2.7	Discussion	25
2.8	Conclusions	28
	Appendix	29
2.A	Equivalent optimisation problem for the LRPM	29
3	Fast and fuzzy multi-objective radiotherapy treatment plan generation for head and neck cancer patients with the lexicographic reference point method (LRPM)	31
	Abstract	32
3.1	Introduction	33
3.2	Methods and materials	35

3.3	Results	41
3.4	Discussion	44
3.5	Conclusions	46
4	Automatically configuring the reference point method for automated multi-objective treatment planning	49
	Abstract	50
4.1	Introduction	51
4.2	Methods and materials	52
4.3	Results	64
4.4	Discussion	68
4.5	Conclusions	71
	Appendices	72
4.A	Non-Pareto-optimal database generation	72
4.B	Plan projection onto the Pareto front	72
5	Automatic configuration of the reference point method for fully automated multi-objective treatment planning applied to oropharyngeal cancer	75
	Abstract	76
5.1	Introduction	77
5.2	Methods and materials	78
5.3	Results	82
5.4	Discussion	86
5.5	Conclusions	89
	Appendix	90
5.A	Supporting information	90
6	Plan-library supported automated replanning for online-adaptive intensity-modulated proton therapy of cervical cancer	93
	Abstract	94
6.1	Introduction	95
6.2	Methods and Materials	96
6.3	Results	100
6.4	Discussion	102
7	Fast and exact Hessian computation for a class of nonlinear functions used in radiation therapy treatment planning	105
	Abstract	106
7.1	Introduction	107

7.2	Main result	109
7.3	Examples in radiation therapy	111
7.4	Discussion	113
7.5	Conclusions	115
	Appendices	115
7.A	Proofs and derivation	115
7.B	Parameters for commonly used functions	118
8	Discussion	121
8.1	Introduction	121
8.2	RPM versus LRPM for automatic planning	121
8.3	RPM versus 2p _{ec} method regarding automatic configuration	124
8.4	RPM for adaptive planning	125
8.5	Future work	126
8.6	Towards clinical introduction of the RPM	127
	References	131
	List of publications	143
	Summary	145
	Samenvatting	149
	PhD portfolio	153
	Curriculum Vitae	155

Introduction

1.1 RADIOTHERAPY

Radiotherapy is commonly used for treatment of cancer: more than half of all cancer patients receive radiotherapy alone or as part of the treatment. Radiotherapy is typically applied when the tumour cells are localised in a part of the body. Ionising radiation can damage the malignant cells and, in this way, eradicate the tumour. Unavoidably, however, this radiation also damages healthy tissues surrounded by the tumour mass which may lead to radiation-induced side-effects. Therefore, the aim of a radiotherapy treatment is to deliver a sufficiently high dose to the tumour while keeping the doses to the healthy surrounding tissues as low as possible. Radiotherapy treatment plan optimisation aims at realising this.

1.2 RADIOTHERAPY TREATMENT PLAN OPTIMISATION

For a high-quality radiotherapy treatment plan, the doses delivered to the tumour and to the radiation-sensitive surrounding healthy tissues (organs-at-risk, OARs) need to be balanced delicately. The goal of treatment plan optimisation, or treatment planning, is to find settings of the applied treatment unit that result in an optimal balance for the patient to be treated. The treatment planning process starts by making a computer tomography (CT) scan to obtain a three-dimensional representation of the part of the patient's anatomy where the tumour and OARs are located. Common treatment modalities with photon beams (X-rays) include intensity-modulated radiotherapy (IMRT) and volumetric-modulated arc therapy (VMAT). Another treatment modality is intensity-modulated proton therapy (IMPT), which uses protons (instead of photons) to damage cancerous tissue. These are all techniques that allow for a high-precision and personalised treatment, given a well-designed treatment plan.

Traditionally, radiotherapy treatment plans for patients are generated in an interactive trial-and-error procedure ("manual planning") using a (commercial) software applic-

ation, called the treatment planning system (TPS). In this interactive procedure, the TPS is stepwise steered towards an acceptable, and intentionally high-quality plan. In such a plan, the patient-specific treatment machine parameters (e.g. number of radiation beams and their angles, beam-on times, and beam-shaping parameters) should result in a three-dimensional dose distribution that well satisfies all dosimetric goals for the tumour and OARs, as defined in the clinical planning protocol. In this way, an as high as possible probability for cure with an as low as possible probability of developing radiation-induced side-effects can be ensured. However, the manual planning procedure can be time-consuming and the quality of the final plan is dependent on the allotted time, and skills and experience of the planner.

In order to avoid inconsistencies in plan quality as much as possible and to ensure that each patient is provided with a high-quality treatment plan, automation of the planning process is needed. In 2010, this led to the clinical introduction of the in-house developed Erasmus-iCycle algorithm for automatic multi-objective optimisation of beam angles and IMRT intensity profiles (Breedveld et al., 2012). For each patient, Erasmus-iCycle automatically generates a single *Pareto-optimal* IMRT plan based on a “wish-list”, which contains predefined planning constraints and prioritised planning objectives for steering the multi-objective plan generation. In a Pareto-optimal plan, none of the planning objectives can be improved without deterioration of at least one of the others. Separate wish-lists are created for separate patient populations (e.g. prostate cancer or head and neck cancer). When ready, they are applied for automated planning for new patients in the respective patient groups, without making any patient-specific modifications. Wish-lists are created in a tuning process based on manually generated treatment plans of previously treated patients and input from radiation oncologists, planning radiotherapy technologists (RTTs), and medical physicists. The best wish-list results in plans with the most favourable balances between all clinical treatment aims. The process starts with defining a simple (“initial guess”) wish-list, which is used to automatically generate plans for a small group of previously treated patients (typically five to ten). The plans for these patients are then evaluated by the clinical expert team, and their input is used to update the wish-list. Then, another round of automated plan generations for the group of patients follows. This iterative process of updating the wish-list finishes when there are no more insights for new updates that could further improve the plan quality. In this tuning process, the manually generated plans of the previously treated patients are (initially) used as a guideline, and the final aim is to improve upon these plans. An optimally tuned wish-list ensures that the plans generated are both Pareto optimal and clinically favourable. Several studies (Voet et al., 2013a, Heijmen et al., 2018) have demonstrated that the quality of automatically generated plans with Erasmus-iCycle is superior to that of manually generated plans. At the start of this PhD project, November 2014, Erasmus-iCycle had already been used in the clinic to automatically generate plans

for approximately 600 patients per year.

Based on a tuned wish-list, the core part of Erasmus-iCycle, the 2-phase ϵ -constraint (2p ϵ c) method (Breedveld et al., 2007, 2009), automatically generates high-quality and Pareto-optimal intensity profiles for each patient, given a fixed beam set-up. The iterative 2p ϵ c method requires solving several consecutive optimisation problems before the final intensity profiles are obtained. The required multiple optimisations have a cost in terms of optimisation time. Another challenge of the 2p ϵ c method is the interactive tuning of wish-lists. This trial-and-error configuration procedure, which has to be repeated for each patient group, is both labour-intensive and time-consuming.

This thesis investigates extending the Erasmus-iCycle optimisation suite with the novel *lexicographic reference point method* (LRPM) and the *reference point method* (RPM), a special case of the LRPM, for automatic optimisation of intensity profiles. This extension aims for (1) faster optimisation times, and (2) reducing the configuration workload of the planning algorithm. In contrast to the 2p ϵ c method, the LRPM and RPM generate Pareto-optimal intensity profiles by solving a *single* optimisation problem, allowing much faster plan generation than the 2p ϵ c method, particularly when many planning objectives are included in the plan optimisation. This makes the LRPM and RPM more eligible for (online) adaptive planning, where a plan needs to be adapted to the daily anatomy of the patient in a short amount of time. To avoid time-consuming and labour-intensive configuration procedures, as needed for high-quality plan generation with the 2p ϵ c method, automatic configuration of the RPM was investigated.

1.3 THESIS OUTLINE

The focus of this thesis is on (1) development and validation of the LRPM and RPM for fast automatic generation of Pareto-optimal and clinically favourable IMRT intensity profiles, (2) development and validation of automatic RPM configuration based on a set of training plans, and (3) to explore the use of the RPM for daily adaptive re-planning in IMPT.

Throughout this thesis, automatically generated plans with the 2p ϵ c method are used as a gold standard. The quality of the faster LRPM and RPM is therefore always evaluated by comparing the resulting plans to those generated with the 2p ϵ c method.

Chapters 2 and 3 focus on implementation and validation of the LRPM for multi-objective treatment plan optimisation in IMRT.

In chapter 2, the LRPM is introduced as a lexicographic extension of the RPM. In addition, part of the LRPM parameters is utilised to explicitly steer the trade-offs between plan objectives. Feasibility of the LRPM for automated plan generation is investigated for prostate cancer.

Chapter 3 further explores the applicability of the LRPM for automatic generation of high-quality treatment plans, but now for head and neck cancer patients. From a multi-

objective optimisation point of view, automatic plan generation for head and neck cancer is far more challenging than for prostate cancer due to the increased number of objectives that need to be optimised in a balanced way for generation of clinically favourable plans.

Manual configuration of the LRPM is highly complex and time-consuming. In chapters 4 and 5, automatic configuration of the LRPM is explored based on previously obtained dose distributions generated with Erasmus-iCycle.

Chapter 4 introduces an automatic configuration procedure based on user-specified tolerances for differences between RPM generated plans and previously obtained dose distributions. The automatic procedure was extensively tested for prostate cancer.

Chapter 5 extends and further explores the applicability of the automatic configuration procedure using a database of plans for head and neck cancer patients.

In chapter 6, a new automated strategy is presented for online-adaptive IMPT for cervical cancer patients. The proposed strategy restores spot positions of the planning CT, adds spots to sufficiently cover the tumour, and then applies an RPM optimisation of the spot weights to generate an IMPT plan for the daily anatomy.

In radiotherapy plan optimisation with an interior-point method ([Breedveld et al., 2017](#)), accurate and repeated computation of gradient vectors and Hessian matrices is required for solving the optimisation problem. Chapter 7 presents a new computationally efficient canonical form to hard-code the computation of gradients and Hessians for functions used in radiotherapy treatment planning.

Chapter 8 discusses the LRPM and RPM in a wider context, including a view on future research.

Lexicographic extension of the reference point method applied in radiation therapy treatment planning

R. van Haveren¹, S. Breedveld¹, M. Keijzer², P.W.J. Voet³, B.J.M Heijmen¹, and W. Ogryczak⁴

¹ Erasmus MC, University Medical Center Rotterdam, Department of Radiation Oncology, Rotterdam, The Netherlands

² Delft University of Technology, Delft Institute of Applied Mathematics, Delft, The Netherlands

³ Elekta Instrument AB, Stockholm, Sweden

⁴ Warsaw University of Technology, Institute of Control and Computation Engineering, Warsaw, Poland

European Journal of Operational Research, Volume 263, Issue 1, 247–257, 16 November 2017

DOI: [10.1016/j.ejor.2017.04.062](https://doi.org/10.1016/j.ejor.2017.04.062)

ABSTRACT

In radiation therapy treatment planning, generating a treatment plan is a multi-objective optimisation problem. The decision-making strategy is uniform for each group of cancer patients, e.g. prostate cancer, and can thus be automated. Predefined priorities and aspiration levels are assigned to each objective, and the strategy is to attain these levels in order of priority. Therefore, a straightforward lexicographic approach is sequential ϵ -constraint programming where objectives are sequentially optimised and constrained according to predefined rules, mimicking human decision making. The clinically applied 2-phase ϵ -constraint (2p ϵ c) method captures this approach and generates clinically acceptable treatment plans.

However, the number of optimisation problems to be solved for the 2p ϵ c method, and hence the computation time, scales linearly with the number of objectives. To improve the daily planning workload and to further enhance radiation therapy, it is extremely important to reduce this time. Therefore, we developed the lexicographic reference point method (LRPM), a lexicographic extension of the reference point method, for generating a treatment plan by solving a single optimisation problem. The LRPM processes multiple a priori defined reference points into modified partial achievement functions. In addition, a priori bounds on a subset of the partial trade-offs can be imposed using a weighted sum component.

The LRPM was validated for 30 randomly selected prostate cancer patients. While the treatment plans generated using the LRPM were of similar clinical quality to those generated using the 2p ϵ c method, the LRPM decreased the average computation time from 12.4 to 1.2 minutes, a speed-up factor of 10.

2.1 INTRODUCTION

In multi-objective optimisation problems, multiple functions, or *objectives*, are simultaneously optimised. Consequently, there usually exist multiple (or even infinitely many) *Pareto-optimal* solutions to such a problem. In this paper, we focus on applications where each objective is assigned at least one priority and aspiration level. The foremost interest of the decision maker (DM) is then to attain the aspiration level for the objective that corresponds to the highest priority. Attaining the aspiration levels for objectives with lower priorities may then not deteriorate the achieved values of the objectives corresponding to the higher priorities. However, the objectives may slightly deteriorate if this allows a large improvement for at least one of the others, i.e. partial *trade-offs* between objectives have an important role as well.

One of these applications is *radiation therapy treatment planning*. For patients diagnosed with cancer and selected for treatment by radiation therapy, a treatment plan has to be made. This is the task of the treating physician who is the DM in this application. The DM's most important objective is to sufficiently irradiate the tumour, in order to eradicate the malignant cells. The healthy organs, which have certain levels of *radiosensitivity*, need to be spared as much as possible. Lowering the radiation below these levels has little effect. The other objectives thus serve to attain these levels for the healthy organs, and have different priorities. However, attaining these levels for the healthy organs may never interfere with sufficiently irradiating the tumour, otherwise the patient will not be cured.

There are two main approaches for multi-objective treatment planning. The first approach involves interactive methods. An example of an interactive method is Pareto front navigation where a representation of the Pareto-optimal solutions is generated (Craft et al., 2006, Miettinen et al., 2008). Other examples are aspiration/reservation-based methods where the DM specifies aspiration levels for the desired solution. Based on these levels, a solution is generated. The DM can then either accept this solution or adjust the aspiration levels after which a new solution is generated (Korhonen and Wallenius, 1988, Granat and Makowski, 2000, Ogryczak and Kozłowski, 2011). Another interactive method is presented in Long et al. (2012) and was especially developed for treatment planning. The treatment plan is obtained by repeatedly generating part of the Pareto front for bi-objective optimisation problems. First, part of the Pareto front is generated for the objectives with the two highest priorities, and the DM selects the desired trade-off. The objective with the highest priority is then constrained to its corresponding objective value, and part of the Pareto front is generated for the objectives with the second and third most highest priority. The DM then selects the desired trade-off again. This process continues until all objectives are dealt with, and generally leads to desirable treatment plans. The benefit of this approach is that the DM can explore different treatment plans. However, the total computation times are long, and the plan quality is dependent on the experience of the DM.

The other main approach for treatment planning involves automated methods which aim to formalise the decision-making strategy of the DM. This strategy is then replicated algorithmically and thus allows automated treatment planning. We have developed the Erasmus-iCycle algorithm (Breedveld et al., 2012) for automated treatment planning. Erasmus-iCycle has been successfully validated and is in clinical use (Rossi et al., 2012, Van de Water et al., 2013, Voet et al., 2013a, 2014, Sharfo et al., 2015). The core part of Erasmus-iCycle is the multi-objective 2-phase ϵ -constraint (2p ϵ c) method (Breedveld et al., 2007, 2009). The 2p ϵ c method sequentially applies the ϵ -constraint method (Miettinen, 1999) to generate a treatment plan. For a certain treatment site (e.g. prostate cancer), the 2p ϵ c method uses a uniform set of parameters to generate a treatment plan for each patient of that site. The benefits of this approach are the short computation times, and the consistent physician-independent plan quality. However, the DM is presented with only a single treatment plan.

Although the 2p ϵ c method has acceptable computation times in the clinical workflow of Erasmus MC – Cancer Institute, faster plan computation enables new applications such as *adaptive treatment planning*. A treatment plan is delivered in up to 40 fractions (one fraction per day), because of the biological property that healthy cells recover faster than malignant cells. However, the patient’s anatomy differs every day. Adaptive treatment planning aims at generating a new plan each day, based on up-to-date anatomical information. The patient awaits treatment during the plan generation, and therefore, acceptable planning times are in the order of minutes. Another reason why short planning times enhance radiation therapy is that it enables physicians to start validating a treatment plan minutes after they finished the delineation of the tumour and healthy organs. In this way, the validation is done with the background of the patient fresh in mind. The method proposed in this paper is an important step towards strongly reducing the computation times, and thus towards improving the efficiency of the workflow and further enhancing radiation therapy as a treatment of cancer.

The aim of this paper was to introduce a multi-objective method able to automatically generate treatment plans for prostate cancer patients within the order of minutes while maintaining the plan quality of those generated using the clinically applied 2p ϵ c method. The proposed method is the lexicographic reference point method (LRPM) which extends the reference point method (RPM) presented in Wierzbicki (1986) by allowing the DM to a priori define more reference points. In contrast to the 2p ϵ c method, applying the LRPM only requires solving a *single* optimisation problem. We demonstrate the feasibility of the LRPM in treatment planning for 30 randomly selected prostate cancer patients. Similar to the 2p ϵ c method, the LRPM should use the same set of input parameters to generate the treatment plans for the 30 prostate cancer patients. This uniform set of input parameters represents a fixed decision-making strategy, i.e. the input parameters are not fine-tuned per

patient.

The rest paper is structured as follows. Section 2.2 summarises the notation used throughout the paper. Section 2.3 explains the problem of radiation therapy treatment planning. In section 2.4, the LRPM is introduced. In section 2.5, we use an example problem to illustrate our intentions with the LRPM. In section 2.6, the feasibility of the LRPM is demonstrated for the 30 prostate cancer patients. Section 2.7 discusses our findings, and section 2.8 concludes the paper.

2.2 MULTI-OBJECTIVE OPTIMISATION AND TRADE-OFFS

For $N \in \mathbb{N}$, we define the index set $[N] := \{1, 2, \dots, N\}$. For any two vectors $x = (x_1, \dots, x_N)$, $y = (y_1, \dots, y_N) \in \mathbb{R}^N$, we write $x = y$ if and only if $x_i = y_i$ for all $i \in [N]$, $x \leq y$ if and only if $x_i \leq y_i$ for all $i \in [N]$, and $x < y$ if and only if $x_i < y_i$ for all $i \in [N]$.

We reserve $m \in \mathbb{N}$ for the number of decision variables, and $n \geq 2$ for the number of objectives. Without loss of generality, each objective is to be minimised. They are denoted as $f_i : \mathbb{R}^m \rightarrow \mathbb{R}$ for $i \in [n]$ and assumed to be of the same order of magnitude, and convex and twice differentiable. A multi-objective optimisation problem is then denoted as

$$\underset{x \in X}{\text{minimise}} f(x), \quad (2.1)$$

where $X \subseteq \mathbb{R}^m$ is the *feasible set* and $f(x) = [f_1(x), f_2(x), \dots, f_n(x)]$. A number of $c \in \mathbb{N}$ constraints $g(x) = [g_1(x), \dots, g_c(x)] \leq 0$ define the feasible set, i.e. $X = \{x \in \mathbb{R}^m \mid g(x) \leq 0\}$. All functions $g_i, i \in [c]$, are also assumed to be convex and twice differentiable, and X is a non-empty convex and compact set. We say that \mathbb{R}^m is the *decision space*, and \mathbb{R}^n is the *objective space*.

For a multi-objective optimisation problem, we use the following notions of optimality: (1) $\hat{x} \in X$ is *weakly Pareto optimal* if there is no $x \in X$ such that $f(x) < f(\hat{x})$, and $\hat{y} \in f(X)$ is *weakly Pareto optimal* if its corresponding decision variable is weakly Pareto optimal; (2) $\hat{x} \in X$ is *Pareto optimal* if there is no $x \in X$ such that $f(x) \leq f(\hat{x})$ and $f_j(x) < f_j(\hat{x})$ for some $j \in [n]$, and $\hat{y} \in f(X)$ is *Pareto optimal* if its corresponding decision variable is Pareto optimal.

Finally, *partial trade-off rates* are introduced according to Miettinen (1999). For $x \in X$, the partial trade-off rate at x involving objectives f_i and f_j is

$$\lambda_{ij}(x) := \frac{\partial f_i(x)}{\partial f_j}. \quad (2.2)$$

2.3 RADIATION THERAPY TREATMENT PLANNING

Radiation therapy treatment planning involves solving a multi-objective optimisation problem. While the tumour must receive a sufficient dose, the dose delivered to the several sur-

rounding healthy organs should be decreased as much as possible to minimise radiation-induced complications. It is important to realise that there is not a single clinically optimal treatment plan. If several treatment plans are presented, different physicians (DMs) are likely to pick different treatment plans. This is not only due to the differences in experience, but also because of the complexity of the problem: there are between 10 and 30 objectives (Rossi et al., 2012, Van de Water et al., 2013, Van Haveren et al., 2017b) which are mutually correlated. Instead of a single optimal treatment plan, there is thus a set of clinically acceptable ones.

2.3.1 Fluence map optimisation

During treatment, several beams are focussed on the patient's tumour from a number of different beam directions. Each beam is subdivided into smaller sub-beams, and on each individual sub-beam, the intensity x_j is to be set. The intensities form the *fluence map* $x = (x_1, \dots, x_m)$, which is the decision vector in the problem (2.1) of radiation therapy treatment planning.

Since each patient has a unique anatomy, a computed tomography (CT) scan is made and used to identify the location of the tumour and surrounding healthy tissues. The region of interest is then discretised into a number of k *voxels* (volumetric pixels), i.e. small volume elements. The absorbed dose, measured in gray (Gy), in voxel i is denoted as d_i . Suppose there are $V \in \mathbb{N}$ voxels, then the relation between the nonnegative *dose distribution vector* $d = (d_1, \dots, d_V)$ and fluence map x satisfies $d(x) = Ax$ where A is the $V \times m$ *dose deposition matrix* consisting of nonnegative elements. This matrix is calculated using the algorithm described in Storchi and Woudstra (1996).

The objectives and constraints in problem (2.1) depend on the dose distribution vector (and thus on the fluence map). Typical objectives are the mean dose, generalised mean dose, and/or maximum dose delivered to an organ,

$$\text{mean}(x) = \frac{1}{|O|} \sum_{i \in O} d_i(x), \quad (2.3)$$

$$\text{gmean}_r(x) = \left(\frac{1}{|O|} \sum_{i \in O} d_i(x)^r \right)^{1/r}, \quad (2.4)$$

$$\text{max}(x) = \max_{i \in O} d_i(x). \quad (2.5)$$

In Equations (2.3)-(2.5), the set O consists of all voxels within the corresponding organ and $|O|$ denotes the number of voxels within that organ.

To realise a sufficient irradiation of the tumour, we use a different quantity: the Logarithmic Tumour Control Probability (LTCP)

$$\text{LTCP}_{D^p, \alpha}(x) = \frac{1}{|T|} \sum_{i \in T} \exp(\alpha[D^p - d_i(x)]), \quad (2.6)$$

Table 2.1: Example of a wish-list for two objectives.

Priority	Objective	Aspiration level
1	$f_1(x)$	10
2	$f_2(x)$	11
3	$f_1(x)$	2
4	$f_2(x)$	5

where T is the set of all voxels within the tumour volume, D^p is the prescribed dose in Gy, and α is the cell sensitivity (Alber and Reemtsen, 2007, Breedveld et al., 2009, 2012). Essentially, large penalties are given to voxels whose dose values are lower than the prescribed dose. Lower LTCP values (2.6) thus imply a higher overall tumour dose.

These objectives are convex, but for the twice differentiability, it is required that $r \geq 1$ for the generalised mean dose (2.4). For the maximum dose (2.5), an unbounded decision variable y is added together with the linear constraints $d_i(x) \leq y$ for all $i \in O$.

The set X of all feasible fluences is the intersection of a finite multi-dimensional interval, due to treatment device constraints, and the imposed inequality constraints, each containing a function of the form (2.3)-(2.6). With these settings, the problem (2.1) is thus in convex and twice differentiable form.

2.3.2 Clinically applied 2-phase ϵ -constraint method

Prior to applying the 2p ϵ c method (Breedveld et al., 2007, 2009), at least one priority and aspiration level is assigned to each objective. For convenience, this information is gathered in a *wish-list*, e.g. see table 2.1.

We explain the 2p ϵ c method using the wish-list in table 2.1. This method has two phases in which a sequence of ϵ -constraint optimisation problems is solved. In the first phase, the aim is to attain the aspiration levels in order of priority. The phase is initiated by minimising objective $f_1(x)$, after which an automated decision is made based on the value $1.03 \cdot f_1^{\min}$ where $f_1^{\min} = \min_{x \in X} f_1(x)$, i.e. the 3% relaxation of the minimum. If this value is less or equal than the aspiration level 10, the constraint $f_1(x) \leq 10$ is added to the set X ; otherwise, the constraint $f_1(x) \leq 1.03 \cdot f_1^{\min}$ is added to the set X (e.g. if the minimum is 15, the constraint $f_1(x) \leq 15.45$ is added). This procedure of gradually adding constraints is used to process all priorities in the wish-list. Thereafter, the second phase of the 2p ϵ c method initiates in which the objectives are minimised in order of priority to strive for a Pareto-optimal solution.

Essentially, for the wish-list in table 2.1, the 2p ϵ c method should be interpreted as a

perturbation of the lexicographic minimisation problem

$$\text{lexmin}_{x \in X} [\max(f_1(x), 10), \max(f_2(x), 11), \max(f_1(x), 2), \max(f_2(x), 5), f_1(x), f_2(x))]. \quad (2.7)$$

Solving the lexicographic minimisation problem (2.7) is essentially the same as applying the 2pec method but then without the 3% relaxation.

There are two reasons for the automated decision of the 3% relaxation. The first reason is to avoid instability regarding approximation errors in numerical calculations (Klepikova, 1985). The other reason is that attaining the aspiration levels for the objectives with lower priorities becomes more likely. This is essential for the medical preference model in which both the achieved objective values and trade-offs between objectives are important.

In practice, we use the same wish-list per patient group to represent a fixed decision-making strategy. Developing a wish-list is a complex iterative trial-and-error process in which physicians collaborate. The wish-list gradually evolves by generating and evaluating the resulting treatment plans (generated with the 2pec method) for a small group of patients (Voet et al., 2013a, 2014). This process has to be done for each treatment site since the healthy tissues surrounding the tumour and their possible complications vary for each site.

2.4 LEXICOGRAPHIC REFERENCE POINT METHOD

We aim for a different approach to perturb the lexicographic minimisation model (2.7). The approach should consist of only a single optimisation problem to decrease the required computation time for generating a treatment plan. Furthermore, both aspiration levels and trade-offs should be taken into account. Similar to the 2pec method, the LRPM should use the same set of input parameters per treatment site (e.g. prostate cancer). This uniform set of input parameters represents a fixed decision-making strategy.

With these requirements, a hybrid approach based on the reference point method (RPM) (Ogryczak, 1997, Granat and Makowski, 2000, Ogryczak and Kozłowski, 2011), and the weighted sum method (Miettinen, 1999) seems most suitable. In the original RPM (Wierzbicki, 1982), a single reference point is used to represent the preferences of the DM. One extension of that work is presented in Wierzbicki (1986) where two reference points, one consisting of optimistic aspiration levels and the other consisting of pessimistic aspiration levels, are used to generate the solution.

For the proposed lexicographic reference point method (LRPM), we extend the RPM for more than two reference points (section 2.4.1). In addition, a weighted sum component is included to affect the trade-offs (section 2.4.2).




Figure 2.1: A piecewise linear reference path through the circled reference points r^1 to r^4 . For each of the three problems, the thick curve represents the set of Pareto-optimal solutions, and the one generated by the LRPM is depicted as a square.

Table 2.2: General reference list: a summary of the reference points and priority levels.

Priority	Reference point	$f_1(x)$	$f_2(x)$...	$f_n(x)$
1	r^1	r_1^1	r_2^1	...	r_n^1
2	r^2	r_1^2	r_2^2	...	r_n^2
\vdots	\vdots	\vdots	\vdots	\vdots	\vdots
p	r^p	r_1^p	r_2^p	...	r_n^p

2.4.1 Reference points

The main idea of the LRPM is to extend the RPM by allowing the DM to a priori define more reference points in the objective space. This should be done so that it is possible to define a strictly increasing curve, which we refer to as a *reference path*, that goes through the defined reference points. An example is shown in figure 2.1 where two objectives and four reference points are used for three *problems*, i.e. cases that have the same objectives but a different feasible set X resulting in a different set of Pareto-optimal solutions. figure 2.1 also shows several *indifference curves*: the DM has no preference for one point on an such a curve over another point on the same curve.

In general, $p \in \mathbb{N}$ reference points, r^1, r^2, \dots, r^p , are defined in the objective space \mathbb{R}^n . The objectives are only allowed to improve, yielding $r^p < \dots < r^2 < r^1$. Therefore, the first priority is to attain the reference point r^1 , the second priority is to attain reference point r^2 , and so on. A *reference list* summarises these points, see table 2.2.

Using the reference list (table 2.2), a reference path can be parametrised, e.g. by linear interpolation as shown in figure 2.1. In general, a reference path $\gamma : \mathbb{R} \rightarrow \mathbb{R}^n$ is given by $\gamma(z) = (q_1(z), q_2(z), \dots, q_n(z))$ where the $q_i : \mathbb{R} \rightarrow \mathbb{R}$ define its parametric representa-




Figure 2.2: Convex PAFs for two objectives satisfying $s_i(r_i^j) = v_j$ for $i \in [2]$ and $j \in [3]$.

tion, and strictly increase by construction.

For every value $z \in \mathbb{R}$, the feasible points of interest (if any) satisfy $f_i(x) \leq q_i(z)$ for all $i \in [n]$, where lower values of z are more preferred than higher values. In terms of the RPM, the partial achievement functions (PAFs) $s_i : \mathbb{R} \rightarrow \mathbb{R}$ are exactly the inverse functions of q_i which are well defined since the q_i strictly increase. In case of linear interpolation (as shown in figure 2.1), typical PAFs are depicted in figure 2.2 where the vertical axis represents a uniform scale to measure the DMs dissatisfaction for each individual objective. This dissatisfaction is expressed in the *value levels* v_j for which it holds that $s_i(r_i^j) = v_j$ for all $i \in [n]$ and $j \in [p]$. The sequence $(v_j)_{j=1}^p \subseteq \mathbb{R}$ thus strictly decreases.

For linear interpolation, the PAFs are explicitly given by

$$s_i(f_i(x)) = \begin{cases} v_p + \alpha_1 w_i^p(f_i(x) - r_i^p), & f_i(x) \leq r_i^p \\ v_j + w_i^j(f_i(x) - r_i^j), & r_i^j < f_i(x) \leq r_i^{j-1}, \quad j \in [p] \\ v_1 + \alpha_2 w_i^2(f_i(x) - r_i^1), & r_i^1 < f_i(x), \end{cases} \quad [1] \quad (2.8)$$

where

$$w_i^j = \frac{v_{j-1} - v_j}{r_i^{j-1} - r_i^j}, \quad i \in [n], \quad j \in [p] \quad [1]. \quad (2.9)$$

Parameters α_1 and α_2 satisfy $0 < \alpha_1 \leq 1 \leq \alpha_2$. Parameter α_1 models the increase of the DMs satisfaction in case better outcomes than the reference point r^p are generated. Parameter α_2 represents an increase of the DMs dissatisfaction in case outcomes worse than the reference point r^1 are generated.

For convex programming, the PAFs need to be convex. While this is trivial in case $p = 2$ (choose $v_2 < v_1$), the PAFs (2.8) are not necessarily convex for $p > 2$. However, by choosing appropriate value levels, the convexity of all PAFs can be guaranteed. It suffices

to initially set $v_p < v_{p-1}$ and to ensure that the following inequalities hold

$$v_{j-1} \geq v_j + (v_j - v_{j+1}) \max_{i \in [n]} \frac{r_i^{j-1} - r_i^j}{r_i^j - r_i^{j+1}}, \quad j \in [p-1] \quad [1]. \quad (2.10)$$

Selecting the v_j according to condition (2.10) guarantees that the slopes w_i^j (2.9) are increasing ($w_i^2 \geq w_i^3 \geq \dots \geq w_i^p$) so that the PAFs are convex (2.8), e.g. as shown in figure 2.2. Note that the value levels can be scaled to any desired range.

With the convex and piecewise linear PAFs (2.8), the convex *scalarising achievement function* (SAF) to be minimised is similar as in Wierzbicki (1986). This leads to the minimisation problem

$$\text{minimise}_{x \in X} \quad \max_{i \in [n]} s_i(f_i(x)). \quad (2.11)$$

However, solving minimisation problem (2.11) generally leads to weakly Pareto-optimal solutions. Also, the control over the preferences for the trade-offs between objectives can be improved. Both of these issues can be overcome by adding a weighted sum component to the SAF whose effect we discuss in the next section.

2.4.2 Weighted sum component

Similar as in the approach of Kaliszewski (1994), Kaliszewski and Michalowski (1997), we add a weighted sum component to the SAF. This overcomes the issue of possibly generating weakly Pareto-optimal solutions, and allows to include preferences for the trade-offs between objectives. Adding the weighted sum component extends the minimisation problem (2.11) to

$$\text{minimise}_{x \in X} \quad \left[\max_{i \in [n]} s_i(f_i(x)) + \sum_{i \in [n]} \rho_i s_i(f_i(x)) \right], \quad (2.12)$$

where the nonnegative parameters $\rho = (\rho_1, \dots, \rho_n) \geq 0$ are referred to as the *trade-off parameters*. In appendix 2.A, we give an equivalent minimisation problem for (2.12) which can be used directly as input for mathematical solvers.

Although the trade-off parameters are usually chosen uniformly (i.e. $\rho_i = \rho$ for all $i \in [n]$) to guarantee ρ -properly Pareto optimality (Wierzbicki, 1986), we follow the approach presented in Kaliszewski and Michalowski (1997) where one parameter ρ_i is ascribed per PAF. Solving the convex minimisation problem (2.12) leads to a Pareto-optimal solution if $\rho_i > 0$ for all $i \in [n]$ (the solution is even *properly Pareto optimal* (Kaliszewski and Michalowski, 1997), i.e. the trade-offs between objectives are bounded), while only weak Pareto optimality can be guaranteed if there is at least one $j \in [n]$ with $\rho_j = 0$.

In two dimensions, the effect of the weighted sum component on the indifference curves is shown in figure 2.3. There are thus multiple bends in the indifference curves

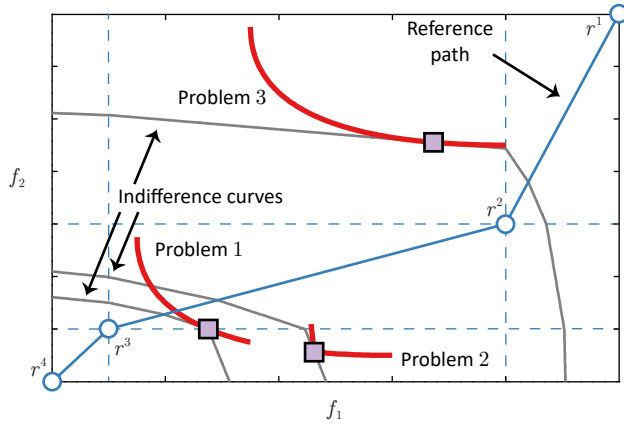


Figure 2.3: For each of the three problems, the thick curve represents the set of Pareto-optimal solutions, and the one generated by the LRPM including the weighted sum component is depicted as a square. Each indifference curve bends multiple times.




Figure 2.4: Indifference curve for large trade-off parameters. The curve is bent multiple times, in particular at its intersection with the lines $f_1 = r_1^2$ and $f_2 = r_2^2$.

caused by the weighted sum component. In figure 2.4, these bends are shown in more detail (the angles are exaggerated for demonstration purposes). In general, an indifference curve is additionally bent each time it intersects an affine hyperplane in \mathbb{R}^n of the form: $f_i = r_i^j$ where $i \in [n]$ and $j \in [p]$. For the angles θ_{ij}^{kl} (e.g. as shown in figure 2.4), it holds that for $i, j \in [n]$ with $i \neq j$, and $k, l \in [p]$ [1] with $k \geq l$, the tangent of the angle between

$f_i \in (r_i^{k-1}, r_i^k)$ and $f_j \in (r_j^{l-1}, r_j^l)$ is

$$\begin{aligned} \tan(\theta_{ij}^{kl}) &= \frac{w_i^k}{w_j^l} \cdot \frac{\rho_i}{1 + \rho_j} \\ &= \frac{v_{k-1} - v_k}{v_{l-1} - v_l} \cdot \frac{r_j^{l-1} - r_j^l}{r_i^{k-1} - r_i^k} \cdot \frac{\rho_i}{1 + \rho_j}. \end{aligned} \quad (2.13)$$

If $f_i < r_i^p$ for an objective $i \in [n]$, the quantity $\alpha_1 w_i^p$ should be used in equation (2.13) instead of w_i^p . Similarly, if $r_i^1 < f_i$, the quantity $\alpha_2 w_i^2$ should be used instead of w_i^2 . The angles (2.13) thus depend on the reference points, value levels (if $k \neq l$), and trade-off parameters.

For p fixed reference points, all angles (2.13) are thus determined by the value levels and trade-off parameters. If the DM wants to set a priori upper bounds b_{12} and b_{21} for the partial trade-offs λ_{12} and λ_{21} (2.2) in a specific domain $f_1 \in (r_1^2, r_1^1)$ and $f_2 \in (r_2^2, r_2^1)$ (e.g. see figure 2.4), then the trade-off parameters ρ_1 and ρ_2 should be determined so that the inequalities $\tan(\theta_{12}^{22}) \leq b_{12}$ and $\tan(\theta_{21}^{22}) \leq b_{21}$ hold. There is, however, no generalisation for more objectives and/or a larger number of a priori bounds for partial trade-offs per domain since this leads to an overdetermined or inconsistent system of equations.

2.5 ILLUSTRATIVE EXAMPLE

This section demonstrates the strategy we used to find suitable parameters for the LRPM. For three problems with the same objectives but different feasible sets, solutions are generated using the 2p ϵ c method with a uniform wish-list (i.e. uniform parameters). These solutions are to be matched as well as possible using the LRPM, also with uniform parameters.

For $i \in [3]$, we define problem i as the multi-objective optimisation problem

$$\underset{x \in X_i}{\text{minimise}} [f_1(x), f_2(x)],$$

where the objectives are

$$\begin{aligned} f_1(x_1, x_2) &= 1.5 + 0.125 \cdot (x_1 - 8)^2 + 0.25 \cdot x_2, \\ f_2(x_1, x_2) &= 1 + 0.125 \cdot (x_1 - 2)^2 + x_2. \end{aligned}$$

The problems' feasible sets are $X_1 = [4, 10] \times [0, 10]$, $X_2 = \{(x_1, x_2) \in [0, 10] \times [0, 1] \mid 4x_1 - x_2 \leq 12\}$, and $X_3 = [2, 10] \times [8, 10]$.

In figure 2.5, the diamonds are the solutions of the three problems that were generated using the 2p ϵ c method with the wish-list in table 2.1. We intend to find uniform parameters for the LRPM for generating solutions similar to these.




Figure 2.5: For each of the three problems, the thick curve represents the set of Pareto-optimal solutions, and the one generated using the 2pεC method is depicted as a diamond. The indifference curves shown for the LRPM correspond to $(\rho_1, \rho_2) = (0.5, 1)$.

Algorithm 2.5.1: Three-step strategy to obtain a reference list.

```

Input: wish-list used for the 2pεC method
Output: reference list used for the LRPM
/* Step 1: use aspiration levels of the wish-list. */
for Priority  $j = 1$  to  $k$  in the wish-list do
    Find corresponding objective  $f_i$  and aspiration level  $b$ ;
    Set  $r_i^j = b$ ;
end
Add a reference point with lower bounds of the objectives;
/* Step 2: Linear interpolation in between defined aspiration levels. */
for Objective  $i = 1$  to  $n$  do
    Find  $J = \{j_1, \dots, j_m\}$  for which  $r_i^J$  are defined;
    for  $l = 1$  to  $m - 1$  do
        if  $j_{l+1} - j_l > 1$  then
            for  $k = j_l + 1$  to  $j_{l+1} - 1$  do
                 $r_i^k = r_i^{j_{l+1}} + (r_i^{j_{l+1}} - r_i^{j_l}) \cdot (k - j_l) / (j_{l+1} - j_l)$ ;
            end
        end
    end
end
/* Step 3: Linear interpolation for remaining aspiration levels. */
for Objective  $i = 1$  to  $n$  do
    Find smallest number  $j_{\min}$  for which  $r_i^{j_{\min}}$  is defined;
    if  $j_{\min} > 1$  then
        for  $k = 1$  to  $j_{\min} - 1$  do
             $r_i^k = r_i^{j_{\min}} + (r_i^{j_{\min}} - r_i^{j_{\min}+1}) \cdot (j_{\min} - k)$ ;
        end
    end
end
end

```

Table 2.3: Reference list used for the LRPM resulting from applying Algorithm 2.5.1 to the wish-list in table 2.1, and its value levels.

Priority	Reference point	$f_1(x)$	$f_2(x)$	Value level
1	r^1	10	14	10
2	r^2	6	11	6
3	r^3	2	8	2
4	r^4	1	5	1
5	r^5	0	0	0

The approach used for finding the parameters for the LRPM consists of two steps. First, the reference points are determined using the different priorities defined in a wish-list, and second, the trade-off parameters are determined through a trial-and-error procedure.

We apply a three-step strategy to obtain the reference points using the wish-list in table 2.1 as input: (1) set $r_i^j = b$, for priority j , objective f_i and aspiration level b in the wish-list so that the aspiration levels in both the wish-list and reference list correspond to the same priority, and add a reference point with lower bounds for the objectives (i.e. $r_1^1 = 10$, $r_2^2 = 11$, $r_3^3 = 2$, $r_4^4 = 5$, and e.g. $r^5 = (0, 0)$); (2) for each objective, use linear interpolation in between aspiration levels that are already set in the reference list (i.e. $r_1^2 = 6$, $r_1^5 = 1$ for objective $f_1(x)$, and $r_2^3 = 8$ for objective $f_2(x)$); and (3) use linear interpolation to set the remaining aspiration levels (i.e. $r_2^1 = 14$). This strategy is generalised in Algorithm 2.5.1 to work for any wish-list. The number of reference points thus equals the number of priorities in the wish-list plus one. Applying Algorithm 2.5.1 to the wish-list in table 2.1 results in the reference list in table 2.3 which we used for the three problems.

The value levels in table 2.2 were chosen according to the convexity condition (2.10) where we used equality, and the initial pair $(v_5, v_4) = (0, 1)$. Equality was used to obtain only a small effect on the angles (2.13) of the indifference curves.

For the trade-off parameters, we apply a trial-and-error procedure which we initiate by setting these parameters to 0.01 so that the effect of the weighted sum component is small and Pareto optimality is guaranteed. As shown in table 2.4, the objective value of f_2 concerning problems 1 and 2 are worse for the LRPM than those for the 2p ϵ c method. Therefore, we set ρ_2 to a multiple of a fixed stepsize, e.g. 0.5. table 2.4 shows the further actions taken. This procedure led to the pair $(\rho_1, \rho_2) = (0.5, 1)$ which resulted into the indifference curves shown in figure 2.5. The trade-off parameters should not be too large since the weighted sum component should not be too dominant in the SAF. Therefore, an upper bound of 1 was used for these parameters.

All optimisation problems in this section were solved using the interior-point method implementation of the function `fmincon` in MATLAB R2013a.

Table 2.4: Objective values for the three problems using: (1) the 2p ϵ c method with the wish-list in table 2.1, and (2) the LRPM with the reference list in table 2.3 for different combinations of trade-off parameters. Also, the values of the PAFs s_1 and s_2 , defined by (2.8), are given.

2pϵc method					
		f_1	f_2		
	Problem 1	1.56	4.52		
	Problem 2	4.71	1.11		
	Problem 3	4.12	10.78		

LRPM					
(ρ_1, ρ_2)		f_1	f_2	s_1	s_2
(0.01, 0.01)	Problem 1	1.50	5.47	1.50	1.16
	Problem 2	4.57	2.20	4.57	0.44
	Problem 3	4.72	10.04	4.72	4.72
(0.01, 0.50)	Problem 1	1.54	4.73	1.50	1.16
	Problem 2	4.63	1.13	4.63	0.23
	Problem 3	4.72	10.04	4.72	4.72
(0.01, 1.00)	Problem 1	1.62	4.14	1.50	1.16
	Problem 2	4.63	1.13	4.63	0.23
	Problem 3	4.96	9.84	4.96	4.45
(0.50, 1.00)	Problem 1	1.56	4.50	1.56	0.90
	Problem 2	4.63	1.13	4.63	0.23
	Problem 3	4.72	10.04	4.72	4.72

2.6 PROSTATE CANCER STUDY

This study included 30 randomly selected prostate cancer patients. For each patient, we generated two Pareto-optimal treatment plans: one using the clinically applied 2p ϵ c method (section 2.3.2), and the other using the LRPM (section 2.4). For each method, we used uniform parameters to generate the plans for all patients. Our aim was to show that the LRPM can generate plans similar to those generated using the 2p ϵ c method. Since the latter are clinically acceptable (Voet et al., 2014), we can avoid discussing clinical preferences. In addition, alternative plans were generated using the LRPM to demonstrate the effect of the weighted sum component (section 2.4.2).

For all patients, there were 11 objectives involved which are specified in the wish-list in table 2.5 (Voet et al., 2014). For the 2p ϵ c method, this wish-list was used to generate the treatment plans for all patients. The healthy organs surrounding the tumour are the rectum, anus, bladder, and femoral heads. The prescribed dose for the tumour was 78 Gy. The wish-list has several fixed constraints, e.g. for all voxels inside the tumour (volume), the maximum dose (type) cannot exceed 81.12 Gy (upper bound) or 104% of the prescribed dose. The constraints on the tumour ensure a sufficient irradiation while avoiding necrosis. Furthermore, the wish-list contains objectives, e.g. the most important objective is the gmean₁₂ dose (type), evaluated using the voxels in the rectum (volume), which has 30 Gy as aspiration level. Additional objectives concern the external ring (edge of the treat-

Table 2.5: Uniform wish-list for prostate cancer patients.

Constraints			
Number	Volume	Type	Upper bound
1	Tumour	LTCP _{78,0.8}	0.5
2	Tumour	max	81.12 Gy
3	Rectum	max	78 Gy
Objectives			
Priority	Volume	Type	Aspiration level
1	Rectum	gmean ₁₂	30 Gy
2	Rectum	gmean ₈	20 Gy
3	Rectum	mean	10 Gy
4	External ring	max	31.2 Gy
5	Tumour shell 5 mm	max	72.5 Gy
6	Anus	mean	10 Gy
7	Tumour shell 15 mm	max	54.6 Gy
8	Tumour shell 25 mm	max	39 Gy
9	Bladder	mean	40 Gy
10	Right femoral head	max	20 Gy
	Left femoral head	max	20 Gy

ment area of interest) to decrease the entrance dose, and shells around the tumour at 5 mm, 15 mm, and 25 mm to increase the conformality of the dose distribution, i.e. to rapidly decrease the dose outside the tumour.

For the LRPM, we determined uniform parameters using the two-step approach described in section 2.5. The reference points were thus obtained by applying Algorithm 2.5.1 to the wish-list in table 2.5. Since the objectives concerning the femoral heads have an equal priority, the aspiration values in the wish-list were set into the same reference point. Furthermore, we set $\alpha_1 = 1$ since the last reference point, for which we used the zero vector, is always infeasible (physically impossible to attain). Also, we set $\alpha_2 = 1$ to not overrate the dissatisfaction of the DM for outcomes worse than the aspiration levels for the high priorities in table 2.5, which are too optimistic. Convexity of the PAFs was ensured by obeying the condition (2.10) in which we used equality, and initialised by the pair $(v_p, v_{p-1}) = (0, 1)$. The trade-off parameters ρ were established through a trial-and-error process similar as described in section 2.5, but in absence of a DM. For a training set, 5 out of 30 patients, we attempted to manually match the objective values of the plans generated using the LRPM to those of treatment plans generated using the 2p ϵ c method as well as possible. For the trade-off parameters, we used a stepsize of 1/12 and upper bound of 1 which resulted in the trade-off parameters in table 2.6.

For all patients, we used 23 coplanar beams which had an equal angular spacing. The numerical data used is available as part of The Radiotherapy Optimisation Test Set (TROTS) dataset (Breedveld and Heijmen, 2017). Generating a treatment plan requires solving a large-scale dense nonlinear convex optimisation problem. Each of these was

Table 2.6: Objective values of the two treatment plans generated for one patient.

Priority	Volume	Type	2pec method	LRPM	ρ
1	Rectum	gmean_{12}	62.2 Gy	61.7 Gy	2/3
2	Rectum	gmean_8	57.8 Gy	57.3 Gy	7/12
3	Rectum	mean	30.4 Gy	30.6 Gy	7/12
4	External ring	max	31.2 Gy	28.6 Gy	1/6
5	Tumour shell 5 mm	max	76.3 Gy	75.7 Gy	11/12
6	Anus	mean	19.9 Gy	20.4 Gy	5/12
7	Tumour shell 15 mm	max	57.0 Gy	56.3 Gy	5/12
8	Tumour shell 25 mm	max	47.5 Gy	47.8 Gy	1/6
9	Bladder	mean	33.0 Gy	33.9 Gy	1/12
10	Right femoral head	max	40.0 Gy	40.0 Gy	1/12
	Left femoral head	max	40.0 Gy	40.0 Gy	1/12

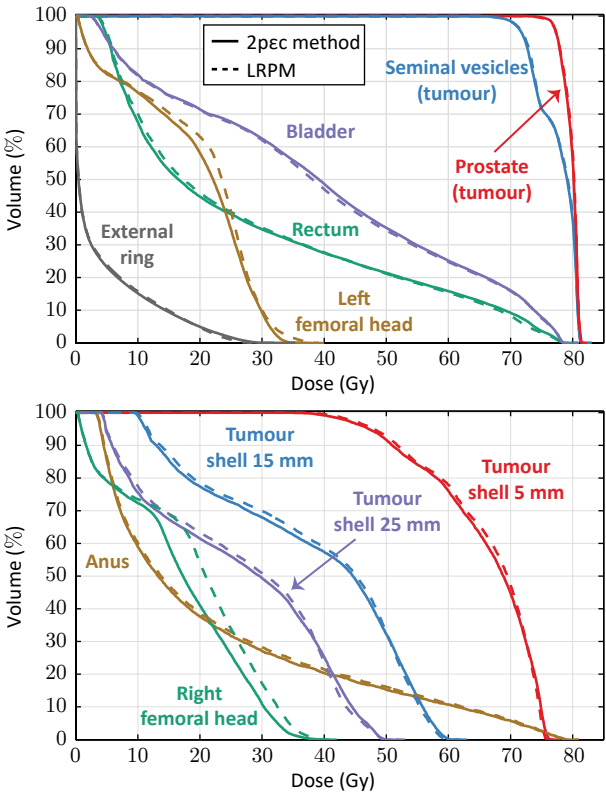


Figure 2.6: The DVHs of the two treatment plans generated for one patient. Each point on a curve corresponds to the volume that receives at least the related dose, e.g. less than 10% of the rectum receives more than 70 Gy in both treatment plans.

successfully solved using a custom implementation of a *primal-dual interior-point method* (Breedveld et al., 2017). For one particular patient who was not included in the training




Figure 2.7: The dose distribution of the plan generated using the LRPM for one patient. The thick solid curves are the contours of the tumour (prostate and seminal vesicles), and the surrounding organs (rectum, bladder, and femoral heads). The thinner solid lines represent *isodose curves*: the tissue on such a curve receives the same dose. The closely spaced isodose curves near the rectum represent a steep dose fall-off.

Table 2.7: The differences ($2p\epsilon c$ method — LRPM) for the evaluation criteria, for which we did two runs with the LRPM. SD stands for standard deviation.

2p ϵc method - LRPM	ρ according to table 2.6		$\rho = 0.01$	
	Mean \pm SD	Range	Mean \pm SD	Range
Rectum V _{75Gy} (%-point)	0.51 \pm 0.23	[0.04, 1.12]	0.84 \pm 0.46	[-0.14, 1.78]
Rectum V _{60Gy} (%-point)	0.53 \pm 0.32	[0.10, 1.35]	0.94 \pm 0.46	[0.12, 2.24]
Rectum mean dose (Gy)	-0.09 \pm 0.53	[-0.71, 1.19]	-1.65 \pm 1.73	[-4.40, 1.27]
Anus mean dose (Gy)	-0.18 \pm 0.49	[-1.50, 0.82]	-1.21 \pm 0.94	[-5.09, 0.12]
Bladder V _{65Gy} (%-point)	-1.40 \pm 1.30	[-4.50, 0.19]	-2.81 \pm 1.92	[-8.82, 0.04]
Bladder mean dose (Gy)	-1.93 \pm 1.79	[-6.48, 0.30]	-3.21 \pm 2.14	[-7.69, 0.74]
Tumour shell 5 mm maximum dose (Gy)	-0.02 \pm 0.95	[-2.26, 1.51]	-3.89 \pm 1.57	[-7.40, -1.68]
Tumour shell 15 mm maximum dose (Gy)	-0.19 \pm 2.07	[-4.82, 3.24]	-7.49 \pm 5.50	[-21.87, 1.02]
Tumour shell 25 mm maximum dose (Gy)	-1.38 \pm 2.44	[-8.96, 2.30]	-5.55 \pm 5.33	[-18.06, 1.50]
Right femoral head maximum dose (Gy)	-0.86 \pm 1.92	[-8.02, 2.08]	0.14 \pm 2.83	[-8.39, 7.38]
Left femoral head maximum dose (Gy)	-0.86 \pm 1.69	[-5.17, 1.20]	0.10 \pm 3.28	[-8.81, 6.23]

set, there were 3 376 beamlets (the decision variables), 93 945 linear constraints, and 27 nonlinear constraints. For this patient, the results are shown for the two treatment plans using the optimal objective values (table 2.6), and the *dose-volume histograms* (DVHs, see figure 2.6). DVHs essentially summarise a 3-dimensional dose distribution: the closer a curve of a healthy organ is to the origin, the better that organ is spared. In figure 2.7, the *dose distribution* is shown of the treatment plan generated using the LRPM.

For all treatment plans generated, the tumour received a sufficient dose. This means that the volume of the tumour receiving at least 95% of the prescribed dose (74.1 Gy) was more than 99%, which is of clinical importance, while the maximum dose was not higher than 81.12 Gy.




Figure 2.8: Comparison of evaluation criteria for the 30 prostate cancer patients, where one plan was generated using the 2pec method and the other using the LRPM with the trade-off parameters in table 2.6. Each marker represents the plan values of an evaluation criterion for one patient. Points close to the identity line represent a similar result while points below (above) this line are in favour of the plan generated using the LRPM (2pec method).

For the healthy tissues, the DM compares the evaluation criteria in figure 2.8 where, e.g. the notation rectum $V_{75\text{Gy}}$ stands for the volume of the rectum receiving a dose of at least 75 Gy. Some of the evaluation criteria differ from the objectives used the optimisations (table 2.5), e.g. the rectum $V_{75\text{Gy}}$ is not convex so we used several generalised means as decent convex substitutes for obtaining a clinically desired dose distribution.

In table 2.7, differences in evaluation criteria with respect to the 2p ϵ c method are shown for two LRPM runs. In the first run, we used the trade-off parameters obtained from the trial-and-error procedure (table 2.6) while we used uniform trade-off parameter $\rho = 0.01$ in the other run. The latter is in-line with commonly chosen values in the literature (Wierzbicki, 1986, Miettinen, 1999, Granat and Makowski, 2000, Ogryczak and Kozłowski, 2011).

All treatment plans generated using the LRPM with the trade-off parameters obtained from the trial-and-error procedure (table 2.6) were found clinically acceptable. Overall, these plans gave slightly better results for the higher priorities, while the 2p ϵ c method gave better results for the lower priorities. For some patients, relatively large differences in objective values were observed in low priorities which were in favour of the 2p ϵ c method. However, the ranges of the objective values are within the set of clinically acceptable plans. There are no formal prescriptions for these ranges, mainly because they are correlated and differ per patient. However, these ranges are generally small for high priorities but larger for low priorities. The differences observed were thus not clinically relevant, i.e. both the small differences in high priorities and the larger differences in low priorities were not clinically relevant.

In contrast, the treatment plans generated using the LRPM with all trade-off parameters set to 0.01 were not found clinically acceptable for all patients. Due to the large negative plan differences (table 2.7), the objective values for some plans were not within the set of clinically acceptable plans.

All computations were performed using 2x 2.90 GHz Intel Xeon processors (16 cores), and 128 GiB of memory. On average, the LRPM reduced the computation time relative to the 2p ϵ c method from 12.4 to 1.2 minutes, a speed-up factor of 10.

2.7 DISCUSSION

This paper introduced the LRPM as a multi-objective method, and demonstrated its feasibility to automatically generate clinically acceptable treatment plans for prostate cancer patients using a uniform set of parameters. We showed that the plans generated using the LRPM had similar plan quality to those generated using the clinically applied 2p ϵ c method, while the computation times were reduced by a factor of 10. This time gain was achieved because the LRPM only requires a single optimisation problem to be solved instead of multiple which is the case for the 2p ϵ c method.

For the prostate cancer study, the main assumption was that the plans generated using the 2p ϵ c method were clinically acceptable. This assumption is justified by the previous study of Voet et al. (2014), which showed that the plan quality of the automatically generated plans, using the 2p ϵ c method with the wish-list in table 2.5, was similar to that of manually generated plans. Since the same wish-list was used for all patients, using uniform parameters was also a requirement for the LRPM.

For the application in radiation therapy treatment planning, the most important benefit of the LRPM as opposed to the clinical method is that the average computation times were reduced from 12.4 to 1.2 minutes. This is of extreme importance since this time gain improves the daily planning workload, and is an important step towards applications such as adaptive radiotherapy where clinically acceptable and Pareto-optimal treatment plans can be generated based on up-to-date anatomical information.

Another benefit of the LRPM is that a priori preferences for the bounds on trade-offs can be set. In radiation therapy treatment planning, both the achieved objective values and trade-offs between objectives are important. In this application, the weighted sum component in the LRPM may thus be essential to obtain clinically acceptable treatment plans for all patients. The interactive approach presented in Long et al. (2012) also leads to desirable treatment plans in general. In this approach, Pareto-optimal sets are sequentially generated for two objectives so that the DM can visualise the possible trade-offs and corresponding objective values simultaneously. This method is effective in practice, but requires human interaction and multiple optimisations to generate the desired solution. Breedveld et al. (2007, 2009) present a similar approach, but with an automated DM. A priori aspiration levels for the objectives are set in a wish-list, and a fixed relaxation of 3% is used (section 2.3.2). Depending on the importance of attaining an aspiration level, and on the shape of the Pareto-optimal set, this may or may not result into a clinically desired trade-off.

In another study (Van Haveren et al., 2017b), it was demonstrated that using the LRPM with the same strategy as presented in this paper leads to high-quality treatment plans for 15 head and neck cancer patients. The plans generated using the LRPM were compared to those generated using the 2p ϵ c method. The LRPM generated plans that had equal or better quality than those generated using the 2p ϵ c method. In addition, the number of objectives used was 22, and the average computation time was reduced from 209.2 to 9.5 minutes, a speed-up factor of 22, in favour of the LRPM.

The LRPM was introduced as a multi-objective method whose application is not limited to radiation therapy treatment planning. In general, the LRPM extends the regular RPM Wierzbicki (1982, 1986) by allowing more reference points to be defined a priori by the DM. These points need to be chosen strictly monotonic coordinatewise. The weighted sum component of the LRPM can be used for partially controlling the bounds on trade-offs according to the relations (2.13), but can also be used as a regularisation term (small

Table 2.8: Reference list to approximate the lexicographic minimisation problem (2.7). Here, $\xi > 0$ is a small parameter, m_1 a lower bound for objective $f_1(x)$, and M_2 an upper bound for objective $f_2(x)$.

Priority	Reference point	$f_1(x)$	$f_2(x)$	Value level
				$\xi = 0.01, m_1 = 0, M_2 = 14$
1	r^1	$10 + \xi$	M_2	$1.44 \cdot 10^8$
2	r^2	10	$11 + \xi$	$4.79 \cdot 10^5$
3	r^3	$2 + \xi$	11	$6.00 \cdot 10^2$
4	r^4	2	$5 + \xi$	1
5	r^5	m_1	5	0

positive values for the trade-off parameters). For either use, the LRPM generates solutions that are Pareto optimal.

However, there are limitations for using the LRPM to perturb a lexicographic minimisation problem such as (2.7). Instead of the reference list in table 2.3, we could use the one in table 2.8 for a better approximation. As can be seen in table 2.8, however, is that the value levels become huge if the convexity condition (2.10) is satisfied. This can make it difficult to decide upon appropriate trade-off parameters ρ_i . Also, in particular if more objectives are involved, this likely leads to numerical problems for the solver. To reduce these problems, the SAF optimisation can be transformed into a multilevel lexicographic optimisation process. This, however, leads to increased solution complexity (solution time) similar to the 2pεc method.

In the conversion of a wish-list to a reference list in Algorithm 2.5.1, we do not distinguish between linear and nonlinear objectives. In particular, this implies that we apply the linear interpolation strategy for nonlinear objectives. This is likely to compromise reaching aspiration levels for the other objectives. An improvement of the algorithm could therefore be to apply a different approach to parametrise the corresponding PAF. The latter can also be achieved by approximating the PAF with a piecewise linear function, i.e. the linear interpolation strategy could still be applied, but to more reference points.

For radiation therapy treatment planning, the overall goal is to reduce the dose to the healthy organs in a prioritised order, with aspiration levels being put as population-based references. The amount of damage an organ can have before losing functionality is patient dependent, and impossible to know beforehand. Therefore, these selected aspiration levels do not need to be strictly met. We found that it suffices to perturb the lexicographic minimisation problem (2.7) by choosing the reference points according to Algorithm 2.5.1, selecting the value levels according to condition (2.10) using equality and with the initial pair $(v_p, v_{p-1}) = (0, 1)$, and using an iterative trial-and-error process to determine the weighted sum component. While this procedure results into clinically acceptable treatment plans, it may not be the optimal procedure. Therefore, we intend to develop analytical guidelines for determining the input parameters, i.e. the reference points, value levels, and

trade-off parameters. Determining these input parameters may be achieved via inverse modelling approaches. If such an approach would be successful, the LRPM would become more user-friendly and easier to apply to other treatment sites.

2.8 CONCLUSIONS

We proposed the LRPM as a multi-objective method extending the RPM by allowing the DM to define more reference points a priori for specifying the preferences. In addition, we used a weighted sum component to ensure Pareto optimality, and to affect the partial trade-offs. The feasibility of the LRPM was demonstrated in radiation therapy treatment planning, where we used uniform parameters to generate treatment plans for 30 randomly selected prostate cancer patients. The plan quality of these plans was equal to those generated using the clinically applied 2p ϵ c method. The LRPM reduced the average computation time from 12.4 to 1.2 minutes, a speed-up factor of 10. This time gain is an important step towards the application of adaptive radiotherapy.

ACKNOWLEDGEMENTS

The authors want to thank Elvera Boogaart for her preliminary work on implementing the LRPM in radiation therapy treatment planning.

APPENDIX 2.A EQUIVALENT OPTIMISATION PROBLEM FOR THE LRPM

The scalarising achievement function to be minimised in the convex optimisation problem (2.12) is not differentiable. However, using that the PAFs are convex and piecewise linear, we can conveniently introduce an additional $n + 1$ unbounded decision variables: a_i for $i \in [n]$ for the PAFs, and z for the maximum of the PAFs. The feasible set is then extended to

$$L := \{(x, a, z) \in \mathbb{R}^{m+n+1} \mid x \in X \subseteq \mathbb{R}^m, a = (a_1, \dots, a_n) \in \mathbb{R}^n, z \in \mathbb{R}\},$$

and an equivalent optimisation problem is

$$\begin{aligned} & \underset{(x, a, z) \in L}{\text{minimise}} && z + \sum_{i \in [n]} \rho_i a_i \\ & \text{subject to} && a_i \leq z, && i \in [n] \\ & && v_p + \alpha_1 w_i^p(f_i(x) - r_i^p) \leq a_i, && i \in [n] \\ & && v_j + w_i^j(f_i(x) - r_i^j) \leq a_i, && i \in [n], j \in [p] \quad [1] \\ & && v_1 + \alpha_2 w_i^2(f_i(x) - r_i^1) \leq a_i, && i \in [n]. \end{aligned}$$

This minimisation problem is convex, and its objective and constraints are twice differentiable. The optimisation problem can thus be solved by standard mathematical solvers, e.g. interior-point methods (Wright, 1997, Breedveld et al., 2017).

Fast and fuzzy multi-objective radiotherapy treatment plan generation for head and neck cancer patients with the lexicographic reference point method (LRPM)

**R. van Haveren¹, W. Ogryczak², G.M. Verduijn¹, M. Keijzer³, B.J.M. Heijmen¹, and
S. Breedveld¹**

¹ Erasmus MC, University Medical Center Rotterdam, Department of Radiation Oncology,
Rotterdam, The Netherlands

² Warsaw University of Technology, Institute of Control and Computation Engineering,
Warsaw, Poland

³ Delft University of Technology, Delft Institute of Applied Mathematics, Delft, The Netherlands

Physics in Medicine and Biology, Volume 62, Issue 11, 4318–4332, 5 May 2017

DOI: [10.1088/1361-6560/62/11/4318](https://doi.org/10.1088/1361-6560/62/11/4318)

ABSTRACT

Previously, we have proposed Erasmus-iCycle, an algorithm for fully automated IMRT plan generation based on prioritised (lexicographic) multi-objective optimisation with the 2-phase ϵ -constraint (2p ϵ c) method. For each patient, the output of Erasmus-iCycle is a clinically favourable, Pareto-optimal plan. The 2p ϵ c method uses a list of objective functions that are consecutively optimised, following a strict, user-defined prioritisation. The novel Lexicographic Reference Point Method (LRPM) is capable of solving multi-objective problems in a single optimisation, using a fuzzy prioritisation of the objectives. Trade-offs are made globally, aiming for large favourable gains for lower prioritised objectives at the cost of only slight degradations for higher prioritised objectives, or vice versa.

In this study, the LRPM is validated for 15 head and neck cancer patients receiving bilateral neck irradiation. The generated plans using the LRPM are compared with the plans resulting from the 2p ϵ c method.

Both methods were capable of automatically generating clinically relevant treatment plans for all patients. For some patients, the LRPM allowed large favourable gains in some treatment plan objectives at the cost of only small degradations for the others. Moreover, because of the applied single optimisation instead of multiple optimisations, the LRPM reduced the average computation time from 209.2 to 9.5 min, a speed-up factor of 21.9 relative to the 2p ϵ c method.

3.1 INTRODUCTION

The overall aim of radiotherapy treatment planning is to sufficiently irradiate the planning target volume (PTV) while reducing doses to the surrounding organs-at-risk (OARs) as much as possible, using a clinically desired prioritisation.

To encourage fast plan computation and to guarantee optimality, the selected treatment objectives should preferentially be convex and twice continuously differentiable, e.g. the *generalised equivalent uniform dose* (gEUD), *logarithmic tumour control probability* (LTCP), see Niemierko (1997) and Alber and Reemtsen (2007), respectively. Also, constrained optimisation is preferred since this allows to control the domains for the objectives, e.g. the dose delivered to the PTV should be within 95% and 105% of the prescribed dose. Fast and accurate algorithms (Mehrotra, 1992, Forsgren et al., 2002, Boyd and Vandenberghe, 2004, Nocedal and Wright, 2006, Breedveld et al., 2017) can then be applied to solve the corresponding constrained convex optimisation problems. Since the structure and properties of such optimisation problems tend to be quite specific in each application, we have developed an interior-point method tuned for radiotherapy treatment plan optimisation (Breedveld et al., 2017).

Several techniques may be used to approach the prioritised (lexicographic) multi-objective (multi-criteria) radiotherapy treatment plan optimisation problem, such as Pareto navigation tools (Craft et al., 2006, Miettinen et al., 2008) or interactive methods (Korhonen and Wallenius, 1988, Granat and Makowski, 2000, Ogryczak and Kozłowski, 2011, Long et al., 2012). These techniques feature (partial) exploration of the Pareto front to compare the possible trade-offs between the treatment objectives, but require interference of a physician to steer towards the final plan. This may be time consuming and the result is operator-dependent.

In our institution, we investigate an alternative approach: *automated multi-objective treatment planning* (Breedveld et al., 2007, Jee et al., 2007, Wilkens et al., 2007, Breedveld et al., 2009, 2012, Van Haveren et al., 2017a), in which the decision making is formalised and processed algorithmically, yielding a single clinically favourable, Pareto-optimal plan for each patient. This approach eliminates hands-on time and results are operator-independent.

We have developed Erasmus-iCycle (Breedveld et al., 2012), an algorithm for fully automated treatment plan generation, based on the *2-phase ϵ -constraint* (2p ϵ c) method for prioritised multi-objective optimisation (Breedveld et al., 2007, 2009). The system has been tested for several treatment sites and is now in full clinical use for treatment of prostate cancer (Voet et al., 2013b, 2014), cervical cancer (Sharfo et al., 2015) (adaptive approach, Heijkoop et al. (2014)), head and neck (HN) cancer (Voet et al., 2013a), liver cancer (Leinders et al., 2013), and advanced lung cancer (Della Gala et al., 2016). Apart from optimisation of beam intensity profiles, Erasmus-iCycle can also optimise beam orienta-

tions, e.g. for Cyberknife (Rossi et al., 2012, 2015), and intensity-modulated proton therapy (IMPT) plans (Van de Water et al., 2013).

The 2p ϵ c method optimises the beam intensity profiles for fixed beam directions, i.e. fluence map optimisation (FMO). Each plan generation is based on a so-called *wish-list*, containing planning constraints, that must be obeyed, and treatment objectives with assigned priorities that have to be attained as well as possible (in order of their priority). For each treatment site, a fixed wish-list is used to represent a uniform decision-making structure for all patients with the same tumour type. In the 2p ϵ c method, treatment objectives are sequentially optimised according to their priorities in the wish-list. After each objective optimisation, an appropriate constraint for the current treatment objective is added to the problem, and used in subsequent optimisations. For Pareto-optimal plan generation, the number of performed optimisations scales linearly with the number of treatment objectives to be optimised.

Recently, we introduced the *Lexicographic Reference Point Method* (LRPM) for fully automated FMO (Van Haveren et al., 2017a), as an alternative to the 2p ϵ c method. Similar to the 2p ϵ c method, input parameters are uniform for all patients with the same tumour type. In contrast to the 2p ϵ c method, the LRPM has a fuzzy objective prioritisation and only requires a single optimisation to generate a Pareto-optimal treatment plan. The fuzzy lexicographic scalarisation technique is an extension to the original reference point method (Wierzbicki, 1982, 1986), including a prioritised structure for the objectives. In contrast to other scalarisation techniques, e.g. a weighted-sum scalarisation (Miettinen, 1999), the LRPM considers both the objective values and the global trade-offs made between the objectives.

The challenge in using the same decision-making structure for different patients with the same type of cancer is that each patient has its own specific shape and location of the Pareto front due to the uniqueness of each patient's anatomy. Next to well-selected aims for objective values, sane trade-offs are required to arrive at clinically favourable plans. It is undesirable to fix trade-offs to a certain level, as some Pareto fronts are steep while others are gradual. Secondly, trade-offs should be made *global* rather than for only two (subsequent) objectives (contrary to the approaches in Breedveld et al. (2009), Long et al. (2012)). The LRPM is capable of solving a prioritised multi-objective problem featuring global trade-offs, i.e. large gains in lower prioritised objectives can be favoured if the degradation in higher prioritised objectives is small, or vice versa. As a result, the strict lexicographic ordering of the objectives (as applied in the 2p ϵ c method) becomes fuzzy.

HN cancer is one of the most complex tumour sites regarding multi-objective optimisation, requiring many constraints (10–20) and objectives (20–30) to optimally distribute unavoidable dose delivery between the various OARs. The aim of this study is to demonstrate the feasibility of the LRPM for generating high-quality plans for HN cancer patients.

The plans resulting from the LRPM and 2p ϵ c method will be compared both regarding quality and computation time.

In section 3.2, descriptions of both the 2p ϵ c method and LRPM are provided with their configurations for HN cancer patients. In section 3.3, we analyse automatically generated plans for 15 HN cancer patients. Sections 3.4 and 3.5 discuss our findings, and conclude the paper.

3.2 METHODS AND MATERIALS

In prioritised multi-objective optimisation problems, achieving a goal for higher prioritised objectives is more important than for lower prioritised objectives. For these problems, we assume that the prioritised objectives $f_i(x)$ for $i \in \{1, 2, \dots, n\}$ need to be minimised while obeying the imposed constraints.

The 2p ϵ c method and LRPM, both in-house developed for automated prioritised optimisation of radiotherapy treatment plans are discussed in sections 3.2.1 and 3.2.2, respectively. Details on the study of generated HN plans are discussed in section 3.2.3.

3.2.1 2-phase ϵ -constraint method

For this study, the automated treatment plan generation with the 2p ϵ c method is performed with the wish-list in table 3.1. This wish-list (Voet et al., 2013a) is the result of an iterative process in which physicians, dosimetrists and physicists collaborated. In each iteration, plans were generated and evaluated for a small fixed group of patients, and the wish-list was adjusted according to this evaluation.

The applied wish-list in table 3.1 shows that the first priority is to decrease the LTCP (Alber and Reemtsen, 2007),

$$\text{LTCP}(d; \alpha, D^p) = \sum_{j=1}^M \exp [\alpha (D^p - d_j)] . \quad (3.1)$$

to a value of 0.4 to ensure a sufficient coverage for the PTV. Here, M denotes the number of voxels in the PTV and parameters α and D^p are the cell sensitivity (set to 0.82) and the prescribed dose (46 Gy), respectively. Moreover, an LTCP-value of 0.4 is also sufficient: no effort is put into achieving lower LTCP-values than 0.4 (dose escalation), while there is no penalty involved if a high dose in the PTV is required for better OAR sparing (Petit et al., 2013). The LTCP is used instead of the *tumour control probability* (TCP) since the former is convex and the latter not.

After achieving a sufficient coverage for the PTV, the focus is on the OARs. First, the focus is to decrease the mean dose delivered to the salivary glands to 39 Gy (priorities 2 and 3), representing an NTCP-value of about 50% (Murdoch-Kinch et al., 2008, Dijkema et al., 2010). Before lowering these doses even further (priorities 5, 7, 16 and 17), the

Table 3.1: Wish-list used for all HN patients. The prescribed dose for the PTV is 46 Gy.

Constraints					
	Volume	Type	Limit		
	PTV	D_{\max}	48.3 Gy	(= 105% of D^p)	
	Parotid glands/SMGs	D_{\max}	48.3 Gy	(= 105% of D^p)	
	Oral cavity/Larynx	D_{\max}	48.3 Gy	(= 105% of D^p)	
	Unspecified tissue	D_{\max}	48.3 Gy	(= 105% of D^p)	
	PTV shell 0 cm	D_{\max}	46 Gy	(= D^p)	
	Spinal cord/Brainstem	D_{\max}	38 Gy		
	Cochleas	D_{\max}	30 Gy		
Objectives					
Priority	Volume	Type	Goal	Sufficient	Parameters
1	PTV	↓ LTCP	0.4	0.4	$D^p = 46$ Gy, $\alpha = 0.82$
2	Parotid glands	↓ D_{mean}	39 Gy		
3	SMGs	↓ D_{mean}	39 Gy		
4	PTV shell 0.5 cm	↓ D_{\max}	43.7 Gy	43.7 Gy	
5	Parotid glands	↓ D_{mean}	30 Gy		
6	PTV shell 1.5 cm	↓ D_{\max}	36.8 Gy		
7	Parotid glands	↓ D_{mean}	20 Gy		
8	PTV shell 3 cm	↓ gEUD	20.93 Gy		$a = 15$
	PTV shell 4 cm	↓ gEUD	16.1 Gy		$a = 15$
9	Oral cavity	↓ D_{mean}	39 Gy		
10	Spinal cord	↓ D_{\max}	30 Gy		
	Brainstem	↓ D_{\max}	30 Gy		
11	External ring	↓ D_{\max}	41.4 Gy		
12	Larynx	↓ D_{mean}	34.5 Gy		
13	Swallowing muscles	↓ D_{mean}	34.5 Gy		
14	Oesophagus	↓ D_{mean}	34.5 Gy		
15	Cochleas	↓ D_{\max}	23 Gy		
16	Parotid glands	↓ D_{mean}	10 Gy		
17	SMGs	↓ D_{mean}	10 Gy		
18	Unspecified tissue	↓ D_{mean}	–		

Abbreviations: PTV = planning target volume; SMG = submandibular gland; LTCP = logarithmic tumour control probability (3.1); gEUD = generalised equivalent uniform dose.

maximum dose/gEUD for the PTV shells is lowered (priorities 4, 6 and 8). The PTV shells are artificial structures at 0.5, 1.5, 3 and 4 cm distance from the PTV and serve to increase dose conformality, i.e. accomplish a steep dose fall-off outside the PTV. Hereafter, we aim to decrease the mean dose of the oral cavity (priority 9) and the maximum doses in spinal cord, brainstem and external ring (priorities 10 and 11). The external ring is an artificial structure of a 2 cm ring following the inside of the body contour and serves to control the entrance dose. Next, we decrease the mean doses in the larynx, swallowing muscles and oesophagus (priorities 12, 13 and 14) and the maximum dose in the cochleas (priority 15). The focus is then returned to the salivary glands (priorities 16 and 17), but now we aim for an even lower goal of 10 Gy. The second phase of the 2pec method then consecutively minimised all objectives again in order of priority, but now to their fullest ensuring a Pareto-optimal plan. Finally, the lowest priority 18 serves to lower the overall dose inside the patient and has no

goal value, meaning that this objective is only minimised at the very end of the algorithm (this final optimisation thus does not influence the attained values for the other objectives).

The wish-list has a *multi-level* structure (indicating complex decision making, see Breedveld et al. (2012)), meaning that some OARs (with the same type) appear multiple times to gradually lower the dose. For example, the mean dose of the SMGs can be found twice (priorities 3 and 17), but with another goal. The goal given to the high priority 3 has a relatively low demand compared to the goal for priority 17. A multi-level wish-list serves to prevent that for instance, the SMGs receive a low dose at the expense of an unacceptably high dose in the oral cavity.

We refer the reader to Breedveld et al. (2009) for a more in-depth description of the 2pec method for prioritised multi-objective optimisation.

3.2.2 Lexicographic reference point method

3.2.2.1 Algorithm description

As described above, for all patients with a specific tumour type, a uniform configuration (i.e. the same wish-list) is applied for automated multi-objective treatment plan optimisation with the 2pec method. Also for the LRPM, all plans for a specific tumour type are generated with a uniform configuration, consisting of a reference path and trade-off parameters (explained below). In contrast to the 2pec method, the LRPM is designed to consider all treatment objectives in a single optimisation using a fuzzy objective prioritisation. In this section, we illustrate the principles of plan optimisation with the LRPM for two objectives.

The basic idea of the LRPM is to represent the lexicographic ordering of the objectives by multiple reference points. A reference point assigns goal values to the objectives that are equally important to attain, e.g. the reference point $(f_1, f_2) = (30, 40)$ in figure 3.1(a) means that attaining the goal value of 30 for f_1 is as important as attaining 40 for f_2 . The lexicographic ordering is implemented by using multiple reference points, so that aimed improvements are allowed to vary. This principle is shown in figure 3.1(a), e.g. for subsequent reference points $r^1 = (50, 50)$ and $r^2 = (30, 40)$, the aimed improvements are 20 for f_1 and 10 for f_2 (i.e. more focus on improving f_1) while for subsequent reference points $r^2 = (30, 40)$ and $r^3 = (20, 10)$, the aimed improvements are 10 for f_1 and 30 for f_2 (i.e. more focus on improving f_2). The general rule is that the goal values for each objective may only improve for each pair of subsequent reference points.

After multiple reference points are selected in the configuration process (details in Van Haveren et al. (2017a)), a strictly monotonic *reference path* through these reference points is made, see figure 3.1(a). This path may be nonlinear in general, but we consider the piecewise linear case as in figure 3.1(a). The principle of the LRPM is as follows: the first priority is to meet the goal values in the first reference point r^1 . If r^1 is feasible (i.e. no constraints are violated), the LRPM will steer the solution to the second reference point




Figure 3.1: Principle LRPM for two objectives. (a) reference points/path for lexicographic ordering, (b)-(c) plan selection with non-fuzzy indifference curves, (d) effect fuzzy indifference curve on plan selection, (e) plan selection for a group of patients using a uniformly configured LRPM.

r^2 (second priority), following the reference path. This process continues until no further improvement is possible for any of the objectives (without violating at least one constraint), at which point the Pareto-optimal plan is found. In other words, the LRPM is designed to follow the reference path downwards while remaining feasible so that the generated plan is the intersection of the reference path and the Pareto front.

Technically, the LRPM minimises a single overall function depending on all objective

values, subject to the constraints imposed. This procedure is visualised with *indifference curves*. An indifference curve is a set of points where the overall function takes on a certain constant value. In figure 3.1(b), several indifference curves are depicted (partially, to improve visibility), where the corresponding constant values decrease when moving down the reference path. The optimal solution corresponds with the lowest value for the indifference curve while satisfying the constraints. For the final Pareto-optimal plan, the intersection of the area under the indifference curves and above the Pareto front is exactly a single point, e.g. the square in figure 3.1(b) corresponds to the generated solution.

However, following the reference path as explained above does generally not lead to clinically relevant plans, see figure 3.1(c), where it is intuitively clear that the square does not represent a well-balanced plan, as a large favourable gain for objective f_2 can be realised for only a small degradation of objective f_1 . The problem is that the indifference curves generate the plan solely based on the objective values on the reference path while completely ignoring the trade-offs made between objectives, i.e. these indifference curves are non-fuzzy since they strictly obey the imposed lexicographic ordering of the objectives. To address this issue so that global trade-offs between objectives are also considered, bends are introduced to the indifference curves to create fuzzy indifference curves as in figure 3.1(d). These bends are configured by specifying *trade-off parameters* (one for each objective) integrated in the LRPM. To demonstrate the effect, compare the plans generated in figure 3.1(d): the square is the result of using the non-fuzzy indifference curves (no trade-off parameters) and the diamond results from using fuzzy indifference curves. In some cases, e.g. for patient 1 in figure 3.1(e), the fuzzy and non-fuzzy indifference curves generate the same plan since the intersection of the reference path and the Pareto front happens to represent a well-balanced plan.

With fuzzy indifference curves, the LRPM can be uniformly configured to generate clinically relevant plans for a group of patients (figure 3.1(e) sketches the situation). The fuzziness is required to account for the variation in shape and location of the Pareto fronts, caused by differences in anatomy.

With constraints summarised in the vector $\mathbf{g}(x)$, for which we assume without loss of generality that each entry should be less or equal to zero, the mathematical model for the

LRPM is

$$\begin{aligned}
& \text{minimise} && z + \sum_{i=1}^n \rho_i a_i \\
& \text{subject to} && a_i \leq z, && i = 1, \dots, n \\
& && v_p + \alpha_1 w_i^p (f_i(x) - r_i^p) \leq a_i, && i = 1, \dots, n \\
& && v_j + w_i^j (f_i(x) - r_i^j) \leq a_i, && i = 1, \dots, n, j = 2, \dots, p \\
& && v_1 + \alpha_2 w_i^2 (f_i(x) - r_i^1) \leq a_i, && i = 1, \dots, n \\
& && \mathbf{g}(x) \leq \mathbf{0}.
\end{aligned}$$

Here, $r^j = (r_1^j, \dots, r_n^j)$ are the p reference points, and parameters v_j, w_i^j and α_1, α_2 concern the parametrisation of the piecewise linear reference path, see figure 3.1(a). The ρ_i are the trade-off parameters used to bend the indifference curves, see figure 3.1(d). The z and a_i are additional unbounded decision variables required for a convex and twice differentiable formulation of the optimisation problem. A more in-depth description of the LRPM for prioritised multi-objective plan generation can be found in [Van Haveren et al. \(2017a\)](#).

3.2.2.2 Technical issues

There are two practical issues with applying the LRPM: the sufficient parameter values in the wish-list (table 3.1) and numerical issues for the LTCP (3.1) as objective and/or constraint.

In the LRPM, objectives are always encouraged to improve. However, for an objective with a sufficient value, it is undesired to improve an objective beyond this value since this would deteriorate other objectives too severely. To address this issue, we have to replace each objective with a sufficient value. For example, if objective f_1 has a sufficient value of 0.4, we replace the objective function f_1 by the convex function

$$h_1(x) := \max[f_1(x), 0.4].$$

To implement this into the optimisation problem derived in [Van Haveren et al. \(2017a\)](#), the entries $f_1(x)$ are replaced by the newly introduced decision variable h_1 and the following constraints are added

$$\begin{aligned}
h_1 &\geq f_1(x), \\
h_1 &\geq 0.4.
\end{aligned}$$

In this way, the optimisation problem remains both smooth and convex. These constraints still allow f_1 to be below 0.4 (in case dose escalation is required ([Petit et al., 2013](#))), but this is not encouraged when optimising h_1 .

The other issue we encountered was that the LTCP can cause numerical problems. The problem is poorly scaled since the exponential terms in the LTCP can lead to large values compared to the mean, maximum/minimum and gEUD. To solve this issue, we introduced the *logarithmic LTCP* (LLTCP), i.e. $\text{LLTCP} = \ln(\text{LTCP})$: an equivalent convex objective which has a one-to-one correspondence with the (L)TCP.

3.2.3 Study set-up

In this study, we consider 15 HN cancer patients receiving bilateral neck irradiation, all with a prescribed dose of 46 Gy (no boost techniques were applied). The data is available as part of the TROTS (The Radiotherapy Optimisation Test Set) dataset (Breedveld and Heijmen, 2017). For each patient, we use both the LRPM and the 2pcc method to automatically generate a treatment plan. All plans were generated using a 23 equi-angular coplanar beam set-up to ensure achievable volumetric modulated arc therapy (VMAT) dose distributions (Voet et al., 2013a, Sharfo et al., 2015). To objectively compare the performance of both multi-objective methods, we did not apply VMAT segmentation to avoid a bias in the plan comparisons.

All optimisation problems were solved using the in-house developed primal-dual interior-point method (Breedveld et al., 2017), specifically tuned for the radiotherapy plan optimisation setting, with 2x 2.90 GHz Intel Xeon E5-2690 CPUs (total of 16 cores) and 128 GiB of memory running on Linux.

3.3 RESULTS

All generated plans showed the same value of 0.4 for the LTCP, resulting in a PTV coverage of at least 99% for 95% of the prescribed dose for each plan. On average, the LRPM plans even showed slightly increased PTV coverage of $0.02\% \pm 0.04\%$ -point (range $[-0.06 \ 0.08]$). The differences in plan parameters for the most relevant objectives and evaluation criteria for the individual patients are visualised in figure 3.2. As all plan values for the PTV shell 0.5 cm are equal for each patient (43.7 Gy), this plan objective is not shown in the figure.

In figure 3.2, the differences in the plan trade-offs can be seen for each patient. For example, the trade-offs in both plans for patient 10 are similar, whereas the trade-offs for patients 4 and 8 lead to noticeable plan differences. For patient 4, the LRPM plan significantly reduces the NTCP of the right SMG at the cost of a slight degradation of the NTCP for the left parotid gland and the mean doses in the larynx, swallowing muscles and oesophagus. A different trade-off is seen for patient 8: the LRPM plan significantly reduces the mean doses to the MCP and oesophagus, and also slightly reduced the NTCP for the right SMG at the cost of a slight degradation of the conformality (3 and 4 cm PTV shells). The LRPM aims for more balanced trade-offs (large favourable gains at the cost of small degradations), i.e. the strict lexicographic ordering of the objectives becomes fuzzy. The particular

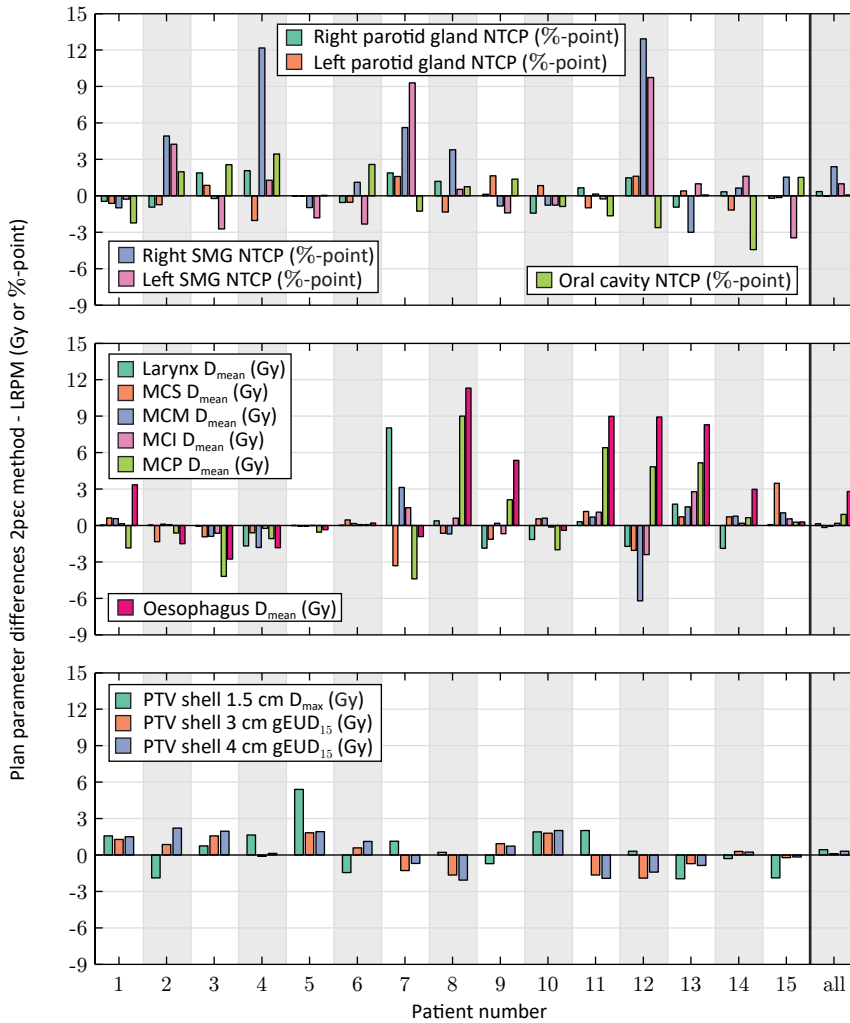


Figure 3.2: Plan differences (2pec method - LRP) per patient for the most relevant treatment objectives and evaluation criteria. Positive values are in favour of the LRP, whereas negative values indicate that the 2pec method performed better. All generated plans are Pareto optimal. SMG = submandibular gland, MCS = musculus constrictor superior, MCM = musculus constrictor medius, MCI = musculus constrictor inferior, MCP = musculus constrictor cricopharyngeus.

OARs that allow large favourable gains (without large degradations for other OARs) differ per patient (see figure 3.2) i.e. the LRP is not configured to improve certain OARs but to find a sane and balanced global trade-off. Differences in conformality (measured by the maximum dose/gEUD for the PTV shells) were minimal.

In figure 3.3, the distributions of differences in selected plan parameters are sketched using boxplots. Most medians are positive, and thus in favour of the LRP. For the dif-




Figure 3.3: Boxplots of the plan differences (2p ϵ c method - LRPM) for treatment objectives and evaluation criteria. The bottom and top edges of the boxes are the first and third quartiles respectively, the horizontal solid lines represent the medians, the whiskers are at 1.5 times the interquartile range and the pluses represent outliers. SMG = submandibular gland, MCS = musculus constrictor superior, MCM = musculus constrictor medius, MCI = musculus constrictor inferior, MCP = musculus constrictor cricopharyngeus.

ferences in maximum doses of the spinal cord and brainstem, the medians are even well above zero. The negative medians (in favour of the 2p ϵ c method) are only slightly below zero. From the 13 observed outliers in figure 3.3, there were 10 in favour of the LRPM. The boxplots show that the global trade-offs are generally better balanced for the plans generated with the LRPM in comparison with the plans generated with the 2p ϵ c method, i.e. relatively large favourable gains were possible for relatively small degradations.

Whereas the LRPM only needs a single optimisation for each patient, the 2p ϵ c method

requires multiple optimisations which scales linearly with the number of treatment objectives. Consequently, we observed a mean computation time for treatment plan generation of 9.5 ± 3.7 min (range [0.5 14.5]) for the LRPM and 209.2 ± 91.0 min (range [9.5 362.0]) for the 2p ϵ c method, an observed speed-up factor of the mean computation times of 21.9 for the LRPM relative to the 2p ϵ c method. This time gain allows a more effective and efficient clinical workflow (e.g. after the physician finished delineation of the target, the final plan approval can be done within minutes by the same physician) and is an important step towards the highly desired application in online adaptive radiotherapy.

3.4 DISCUSSION

The purpose of this study was to demonstrate that a uniformly configured LRPM is capable of fully automated generation of plans for HN cancer patients with at least similar clinical plan quality as the plans that were automatically generated with the clinically applied 2p ϵ c method, thereby significantly reducing the required plan computation time. Generally, it was observed that the LRPM is capable of better balancing the global trade-offs between the different OARs, resulting in more favourable plans. In a previous prospective clinical study (Voet et al., 2013a), we demonstrated for HN cancer fully automated plans with the 2p ϵ c method had significantly higher quality than the manually generated plans in clinical routine (in 97% of cases, the treating physician selected the automatically generated plan for treatment because of superior quality). Therefore, the plans generated with the 2p ϵ c method can be assumed to be clinically relevant, and a proper benchmark for quality comparisons with LRPM plan generation as performed in this study.

For the LRPM optimisation problems, LTCP functions (3.1) were replaced by equivalent LLTCP functions (section 3.2.2.2) to avoid numerical issues caused by the exponential terms in the LTCP. An LTCP function with goal value $b > 0$ can thus be replaced by an LLTCP function with goal value $\ln(b)$ without changing the result of the FMO and maintaining convexity. As the use of the LLTCP tends to reduce the number of iterations needed by the solver and thereby reducing the computation time, we also used the LLTCP (instead of the LTCP) in the 2p ϵ c method for a fair comparison.

The dense convex nonlinear optimisation problems (solved with the in-house developed algorithm described in Breedveld et al. (2017)) are of large-scale: the number of beamlets is in the order of $O(10^4)$ and the number of total voxels considered is in the order of $O(10^5)$. For the LRPM, a single optimisation problem needs to be solved which led to an observed average computation time of 8.6 min. The 2p ϵ c method needs to solve a sequence of optimisation problems to determine a clinically relevant Pareto-optimal plan. On average, the 2p ϵ c method required solving 28 (range [23 32]) optimisation problems to generate a plan. The 2p ϵ c method contains heuristics to reduce the overall runtime of the algorithm, resulting in the varying number of optimisation problems to be solved for

different patients. Firstly, the 2p ϵ c method checks whether or not an objective should be optimised. If the solution of the previous optimisation implies a lower objective value for the current objective than the specified goal in the wish-list (table 3.1), no optimisation is performed but the objective is simply constrained to the specified goal. Also, if the objective was previously optimised but unable to attain its specified goal, it is evident that a lower goal cannot be reached either, so the optimisation is skipped. This mostly depends on the patient's anatomy, e.g. if a parotid gland has a large overlap with the PTV, the higher goal cannot be reached, and the scheduled optimisations for the lower goals will be skipped. The second heuristic is that the solver does not always solve to optimality, but stops if the solution becomes feasible and lower than the specified goal for the current objective. Another property of the 2p ϵ c method is that due to the sequential addition of objectives, the first few optimisations are solved faster than the last optimisations. On the other hand, the LRPM solves a single optimisation problem, roughly of the same size as the last optimisation in the 2p ϵ c method, and is always solved to optimality (no early termination).

The input parameters for the LRPM consist of reference points and trade-off parameters. While the former are automatically converted from the wish-list (Van Haveren et al., 2017a), the latter were determined in an iterative manner, where differences with the 2p ϵ c method were analysed for a training set in each iteration. After the results were satisfying enough, the corresponding parameters of the LRPM were applied for all 15 HN cancer patients. This procedure of determining suitable trade-off parameters can be time consuming, especially with many (20–30) objectives. We plan to address this issue in the near future by developing knowledge-based algorithms to automatically configure the LRPM.

For all HN cancer patients considered in this study, no boost techniques were applied. From a technical point of view, the LRPM has no limitations on the number of boost volumes. Computational difficulties may arise due to the scaling of the LTCP function, however this limitation is rectified with the technique proposed in section 3.2.2.2. We have investigated this for generation of single boost prostate plans (Van Haveren et al., 2017a), and no difficulties were detected. Further investigation is required to determine if this also holds for more complex configurations with multiple boost volumes.

The comparison between the two automated planning approaches shows that the LRPM results in clinically more favourable trade-offs (compared to the 2p ϵ c method) for some patients. Both multi-objective methods have different unrelated mechanisms for the trade-offs. The 2p ϵ c method statically relaxes the minimum of an objective by 3% (in case the desired goal value was infeasible) to create some room for lower prioritised objectives. While this approach generally results into clinically satisfying trade-offs, it may sometimes cause a jump from one steep part of the Pareto front to another. Consequently, it is also not possible to define an acceptable dynamic relaxation in a sequential optimisation approach (Long et al., 2012), as it is impossible to predict the effect of the relaxation mechanism on

lower prioritised objectives. This issue can be overcome by defining global trade-offs where all objectives are weighted simultaneously, which is the case for the LRPM.

For the results (section 3.3), a clear distinction should be made between the plan comparison for individual cases (figure 3.2) and the distributions of plan differences for all patients (figure 3.3), since there are quite a few outliers while most medians are around zero. For example, the outlier of 12.9%-point improvement of NTCP for the right SMG is clinically significant, but the median of this plan parameter is close to zero. Also, although the medians for the differences in maximum doses of the spinal cord and brainstem are well above zero, there is little clinical significance (although a lower dose is always preferred), i.e. the differences do not have a significant impact on the quality of life of the patients. However, for re-irradiation of recurrent HN cancer, the additional sparing of the spinal cord and brainstem may lead to a better possible re-treatment.

Recently, we demonstrated the feasibility of the LRPM for generation of high-quality VMAT plans for prostate cancer (Van Haveren et al., 2017a). In this study, 30 randomly selected prostate cancer patients were considered. For each patient, treatment plans generated with the 2p ϵ c method were compared with the plans resulting from the LRPM. In a previous study (Voet et al., 2014), it was demonstrated that the plans generated with the 2p ϵ c method were of high clinical quality. For these prostate cancer patients, both the 2p ϵ c method and LRPM achieved almost identical results. This is because the trade-offs for the prostate site are much more straightforward compared to the HN site. Still, the average computation time for the LRPM was reduced from 12.4 to 1.2 min, a speed-up factor for the average computation time of 10 relative to the 2p ϵ c method.

3.5 CONCLUSIONS

In this paper, we investigated the use of the novel LRPM for automated multi-objective treatment plan generation with fuzzy objective prioritisation for HN cancer patients receiving bilateral neck irradiation. For the majority of treatment plans generated with the LRPM, quality was at least as good as the quality of the corresponding plan generated with the clinically applied non-fuzzy 2p ϵ c method (in Erasmus-iCycle) for automated plan generation. For individual cases, the fuzziness of the LRPM led to significant reductions of dose in certain OARs at the cost of small increments of dose for other OARs while maintaining a similar PTV coverage. For some cases, this resulted in a clearly favourable plan for the LRPM. Average computation times were reduced with the LRPM from 209.2 to 9.5 min, a speed-up factor of 21.9 relative to the 2p ϵ c method. This time gain improves the effectiveness and efficiency of the clinical workflow, and is an important step towards online adaptive radiotherapy while avoiding deteriorations in plan quality. The LRPM is suited for fast and high-quality automated plan generation for HN cancer patients.

ACKNOWLEDGEMENTS

The authors would like to thank Stevin Petit for his help in providing the patient data and Peter Voet for his work on developing the wish-list for the HN cancer site.

Automatically configuring the reference point method for automated multi-objective treatment planning

R. van Haveren, B.J.M. Heijmen, and S. Breedveld

Erasmus MC, University Medical Center Rotterdam, Department of Radiation Oncology, Rotterdam, The Netherlands

Physics in Medicine and Biology, Volume 64, Issue 3, 035002, 22 January 2019

DOI: [10.1088/1361-6560/aaf9fe](https://doi.org/10.1088/1361-6560/aaf9fe)

ABSTRACT

Automated treatment planning algorithms have demonstrated capability in generating consistent and high-quality treatment plans. Their configuration (i.e. determining the algorithm's parameters), however, can be a labour-intensive and time-consuming trial-and-error procedure.

Previously, we introduced the reference point method (RPM) for fast automated multi-objective treatment planning. The RPM generates a single Pareto-optimal plan for each patient. When the RPM is configured appropriately, this plan has clinically favourable trade-offs between all plan objectives. This paper proposes a new procedure to automatically generate a single configuration of the RPM per tumour site. The procedure was tested for prostate cancer.

Planning CT scans of 287 previously treated patients were included in a database, together with corresponding Pareto-optimal plans generated using our clinically applied 2-phase ϵ -constraint method (part of Erasmus-iCycle) for automated multi-objective treatment planning. The procedure developed acquires plan characteristics observed in a training set. Based on these, an RPM configuration is automatically generated according to user preferences which specify acceptable differences between training set plans and corresponding RPM generated plans. For example, compared to the training set plans, the RPM generated plans need to have similar PTV coverage, and preferably reduced high rectum dose while slight deteriorations in other objectives are allowed. Training sets of different sizes were tested, and the quality of the resulting RPM configurations was evaluated on the test set (subset of the database not used for training).

Using the new procedure, an RPM configuration was generated for each training set. The quality of RPM generated plans was similar or slightly better than that of the corresponding test set plans. The proposed automated configuration procedure greatly reduces the manual configuration workload, and thereby improves the efficiency and effectiveness of an automated clinical treatment planning workflow.

4.1 INTRODUCTION

In intensity-modulated radiation therapy (IMRT) treatment planning, the aim is to irradiate the planning target volume (PTV) with a sufficiently high dose while keeping the dose to the organs-at-risk (OARs) as low as physically possible. This is a multi-objective optimisation problem in which different priority levels for the various treatment plan objectives are set. Ideally, the plan generated for each patient is both on the Pareto front (i.e. none of the treatment plan objectives can be improved without deteriorating at least one of the other objectives) and of high clinical quality (i.e. plans that are hard to improve from a clinical point of view). Several approaches have been developed that explicitly aim for this (Hussein et al., 2018, Breedveld et al., 2019), e.g. automated Pareto front generation followed by manual plan selection (“a posteriori” multi-objective planning Monz et al. (2008), Craft and Monz (2010)), interactive lexicographic multi-objective methods (Long et al., 2012), matheuristic methods (Cabrera et al., 2018), or “a priori” automated multi-objective treatment planning (Breedveld et al., 2009, Van Haveren et al., 2017a).

This paper focusses on the latter, a priori automated multi-objective treatment planning, where the objectives and priorities are defined beforehand. This approach aims to automatically generate a single clinically favourable, Pareto-optimal treatment plan for each patient without any manual tuning.

One method that we developed for this purpose is the 2-phase ϵ -constraint (2pec) method (Breedveld et al., 2007, 2009). This method is used to automatically optimise the IMRT fluence map in Erasmus-iCycle (Breedveld et al., 2012). Various studies have demonstrated that the quality of automatically generated plans (using Erasmus-iCycle) is superior to that of manually generated plans (Voet et al., 2013a,b, 2014, Heijkoop et al., 2014, Sharfo et al., 2015, 2018, Della Gala et al., 2016, Buschmann et al., 2018, Heijmen et al., 2018). Erasmus-iCycle is clinically used for treatment of prostate cancer, head and neck cancer, locally advanced cervical cancer, and advanced non-small cell lung cancer.

Another method that we recently developed for automated multi-objective treatment planning is the lexicographic reference point method (LRPM). The LRPM is a collection of multi-objective methods (Van Haveren et al., 2017a) and its most basic form is equivalent to the reference point method (RPM) presented in Wierzbicki (1982, 1986). The RPM can consistently generate high-quality treatment plans using a single configuration for prostate cancer patients, both for photon therapy (Van Haveren et al., 2017a) and adaptive proton therapy (Jagt et al., 2018). The same was demonstrated for photon therapy treatment of head and neck cancer patients (Van Haveren et al., 2017b), but with the more general LRPM. In these studies, the plan quality of treatment plans generated using the LRPM was similar or slightly better than that of the plans generated using the 2pec method, while average plan computation times were up to 22 times shorter. In this paper, we focus on the use of the RPM for automated treatment plan generation.

However, it remains a major challenge to successfully configure either the 2p ϵ c method or the RPM for a tumour site. Current configuration approaches involve labour-intensive and time-consuming interactive tuning procedures that have to be repeated for each tumour site. To improve this process, we explore an automatic configuration approach.

Since each multi-objective method has its own set of parameters, automatic procedures to find these are specific to the method used. This paper focusses specifically on an automatic procedure for configuring the RPM. There are a few reasons to focus on the RPM instead of the 2p ϵ c method. First, all parameters in the RPM are combined in a single optimisation problem and thus have a direct effect on the plan generated. In contrast, the 2p ϵ c method involves solving consecutive optimisation problems, gradually using more of its parameters in each optimisation. A second reason is that the prioritisation of the objectives is less strict for the RPM than for the 2p ϵ c method. As a result, large improvements of lower prioritised objectives can be favoured if the deterioration in higher prioritised objectives remains small, or vice versa. In other words, an RPM generated plan tends to have better balanced trade-offs between the objectives than those for a plan generated using the 2p ϵ c method (Van Haveren et al., 2017a). Finally, the plan computation times are lower for the RPM than for the 2p ϵ c method while the plan quality is similar for both methods (Van Haveren et al., 2017a).

The aim of this paper is to explore the feasibility of automatically configuring the RPM. For this purpose, we developed a new configuration procedure for the RPM and tested this for prostate cancer. A database including planning CT scans of 287 previously treated prostate cancer patients together with the corresponding plans generated using the 2p ϵ c method was created. An automatic configuration of the RPM is based on a training set, a subset of the database. Evaluating the quality of the resulting RPM configuration is based on its performance on the test set, the subset of the database not used for training. Plan quality on the test set is analysed as a function of the training set size.

4.2 METHODS AND MATERIALS

Figure 4.1 gives an overview of the procedure for automatically configuring the RPM using a training set consisting of CT scans with corresponding treatment plans as input. The output is an RPM configuration which allows automatic generation of treatment plans (with the RPM) that have similar or better quality than that of the training set plans. Predefined user preferences specify how the user defines similar or better plan quality. The quality of the final RPM configuration is evaluated by comparing test set plans with RPM generated plans for an unbiased performance test.

First, we describe our plan database in section 4.2.1. Then, the principles of plan optimisation using the RPM are explained in section 4.2.2. This is followed by a description of the procedure for automatically configuring the RPM in section 4.2.3. Finally, our study




Figure 4.1: Schematic overview of the automated configuration procedure for the RPM.

set-up to evaluate the quality of the RPM configurations for different-sized training sets is described in section 4.2.4.

4.2.1 Patient database

The database used in this study included planning CT scans of 287 prostate cancer patients (Wang et al., 2016). The database consists of three groups of patients, ranked according to the risk of tumour involvement in the seminal vesicles (Voet et al., 2014): there were 60 patients who had low risk, 113 patients who had intermediate risk, and 114 patients who had high risk.

For each planning CT scan, a single plan was generated. For this generation, we used a fixed coplanar, equi-angular 23 beam set-up to mimic VMAT-like dose distributions. The plans were automatically generated using the multi-objective 2p ϵ c method (Breedveld et al., 2007, 2009) with the “wish-list” in table 4.1 (Sharfo et al., 2018). In this list, the PTV is defined as the prostate PTV for the low risk and intermediate risk patients. In addition,

Table 4.1: Wish-list used for generating the database plans with the the 2pcc method. The down-arrows (\downarrow) indicate that the objectives are to be minimised. The prescribed dose for the PTV is $D^p = 78$ Gy.

Constraints					
	Volume	Type	Limit		
	PTV	D_{\max}	81.4 Gy	(= 105% of D^p)	
	PTV	D_{mean}	78 Gy	(= D^p)	
	PTV shell 50 mm	D_{\max}	39 Gy	(= 50% of D^p)	
	Rectum/Anus	D_{\max}	79.56 Gy	(= 102% of D^p)	
	Femoral heads	D_{\max}	39 Gy	(= 50% of D^p)	
	Unspecified tissue	D_{\max}	81.4 Gy	(= 105% of D^p)	
Objectives					
Priority	Volume	Type	Goal	Sufficient	Parameters
1	PTV	\downarrow LTCP	0.5	0.5	$D^p = 78$ Gy, $\alpha = 0.8$
2	Rectum	\downarrow gEUD	20 Gy		$r = 12$
3	Rectum	\downarrow gEUD	10 Gy		$r = 8$
4	PTV shell 5 mm	$\downarrow D_{\max}$	62.4 Gy	62.4 Gy	
	External ring 20 mm	$\downarrow D_{\max}$	19.5 Gy		
5	Rectum	$\downarrow D_{\text{mean}}$	5 Gy		
6	Anus	$\downarrow D_{\text{mean}}$	5 Gy		
7	PTV shell 15 mm	$\downarrow D_{\max}$	39 Gy	39 Gy	
	PTV shell 25 mm	$\downarrow D_{\max}$	23.4 Gy	23.4 Gy	
8	Bladder	$\downarrow D_{\text{mean}}$	5 Gy		
9	Femoral heads	$\downarrow D_{\max}$	39 Gy		

Abbreviations: PTV = planning target volume; D^p = prescribed dose; LTCP = logarithmic tumour control probability; gEUD = generalised equivalent uniform dose.

intermediate risk patients receive a prescribed dose of 72.2 Gy to the seminal vesicles PTV. For high risk patients, the PTV in table 4.1 is defined as the union of the prostate PTV and the seminal vesicles PTV. For adequate dose conformity, the wish-list includes goal values for the maximum dose levels to the PTV shells at 5, 15, and 25 mm distance from the PTV. Also, for avoiding high entrance dose, a goal value for the maximum dose to the external ring (20 mm ring inside the body contour) is included. The prescribed dose of 78 Gy to the PTV is delivered in 39 fractions of 2 Gy.

To achieve a clinically acceptable tumour coverage ($V_{95\%} \geq 99\%$) in a convex optimisation framework, we used the logarithmic tumour control probability (LTCP) as an objective function (Alber and Reemtsen, 2007)

$$\text{LTCP}(d; \alpha, D^p) = \frac{1}{M} \sum_{j \in [M]} \exp [\alpha (D^p - d_j)],$$

where $d = (d_1, \dots, d_M)$ is the dose vector, M is the number of voxels in the PTV, D^p is the prescription dose, and α the cell sensitivity. Low LTCP values indicate high PTV coverage.

Similarly, to limit the rectum volume receiving high doses, we used the generalised

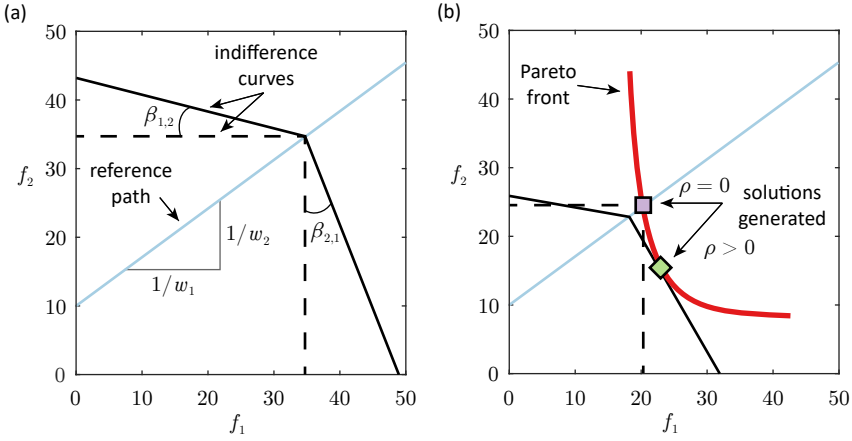


Figure 4.2: Effect of RPM parameters for two objective functions f_1 and f_2 : (a) The reference path parameters and indifference curves, see eq. (4.2); (b) RPM generated solutions on a Pareto front without trade-off parameters ($\rho = 0$, square), and with trade-off parameters ($\rho > 0$, diamond).

equivalent uniform dose (gEUD) as objective functions (Niemierko, 1997)

$$\text{gEUD}(d; r) = \left(\frac{1}{M} \sum_{j \in [M]} d_j^r \right)^{1/r}.$$

The gEUD reduces to the mean dose for $r = 1$, while it approximates the maximum dose for higher values of r .

4.2.2 Reference point method (RPM)

A detailed description of the RPM can be found in the literature (Ogryczak and Kozłowski, 2011, Van Haveren et al., 2017a). This section provides a short summary of its application in automated multi-objective optimisation for generating Pareto-optimal fluence maps. In general, there are n objectives involved, f_1, \dots, f_n , together with a feasible set X consisting of constraints, e.g. those in the wish-list in table 4.1.

The parameters of the RPM are characterised by two sets: the reference path parameters, and the trade-off parameters (figure 4.2). With the trade-off parameters set to zero, the RPM aims to find a Pareto-optimal plan on the reference path, which is a strictly increasing parametric curve (through the objective space). Technically, the reference path is defined by the so-called partial achievement functions $s_i : \mathbb{R} \rightarrow \mathbb{R}$,

$$s_i(f_i(x); w_i, c_i) = w_i f_i(x) + c_i, \quad i \in [n], \quad (4.1)$$

where $w_1, \dots, w_n > 0$ and $c_1, \dots, c_n \in \mathbb{R}$ are the reference path parameters, and the set notation $[n] := \{1, \dots, n\}$ is introduced for convenience throughout this paper. In

figure 4.2, the reference path is exactly the collection of points (f_1, f_2) for which the partial achievement functions (4.1) are equal, i.e. $s_1(f_1) = s_2(f_2)$.

The purpose of the other set of parameters, the trade-off parameters $\rho_i \geq 0$ for $i \in [n]$, is balancing the trade-offs between objectives, i.e. to what extent slight deteriorations of objective values are allowed if large improvements can be realised for other objective values. The trade-off parameters affect the indifference curves for a fixed reference path as is illustrated in figure 4.2. More specifically, the relation between the RPM parameters and the angle $\beta_{i,j}$ is

$$\beta_{i,j}(w_i, w_j, \rho_i, \rho_j) = \text{atan} \left(\frac{w_i}{w_j} \cdot \frac{\rho_i}{1 + \rho_j} \right). \quad (4.2)$$

In the new automated configuration procedure (section 4.2.3), we ensure using trade-off parameters $\rho_i > 0$ since this guarantees generation of Pareto-optimal fluence maps (Kaliszewski and Michalowski, 1997).

RPM generated fluence maps are obtained by solving the following minimisation problem,

$$\underset{x \in X}{\text{minimise}} \quad \left\{ \max_{i \in [n]} \left[s_i(f_i(x); w_i, c_i) \right] + \sum_{i \in [n]} \rho_i s_i(f_i(x); w_i, c_i) \right\}. \quad (4.3)$$

For convex objectives and a convex feasible set, problem (4.3) can be solved using a standard convex mathematical solver: we use our interior-point method tuned for solving fluence map optimisation problems in radiotherapy (Breedveld et al., 2017).

4.2.3 Automated configuration procedure

The aim is to find a single configuration of the RPM, i.e. a single set of RPM parameters to automatically generate plans for each patient without manual tuning, in such a way that the quality of the RPM generated plans is similar or better than that of the database plans. Finding RPM parameters that recreate a single plan in the database (generated using the 2p ϵ c method) is straightforward. First, the objective values achieved define a point in the objective space. Then, applying the RPM with reference path parameters through this point, and trade-off parameters set to zero results in the same fluence map since the plans generated using the 2p ϵ c method are Pareto optimal (note: if the database plan was not Pareto optimal, then the corresponding RPM generated plan, which is Pareto optimal by definition, would score better for at least one of the objective values, see figure 4.2(b) and appendix 4.B).

However, finding a single set of RPM parameters that leads to a desirable fluence map for each patient in a database is less straightforward as illustrated in figure 4.3. The trade-off parameters are essential to avoid generating clinically irrelevant fluence maps. In addition, user preferences are required to specify acceptable or desired population-based compromises, e.g. compared to the plans in the training set, the tumour coverage needs to be similar,

but the mean dose to the rectum is allowed to slightly increase to decrease the high rectum dose (dose levels of 60 Gy or higher delivered to the rectum).

Our strategy consists of two steps. First, we fit a reference path to all points in the objective space defined by the objective values observed in the training set plans, explained in section 4.2.3.1. Secondly, we develop an iterative process for determining the trade-off parameters according to the user preferences which is explained in section 4.2.3.2. No a priori knowledge of the test set was used in developing our strategy.

4.2.3.1 Step 1: reference path parameters

To develop a strategy for fitting a reference path using plans that were previously obtained, it is crucial to see that a reference path divides the objective space. For example, in figure 4.3(a) with two objectives f_1 and f_2 , the reference path itself is characterised by all points (f_1, f_2) for which the partial achievement functions (4.1) are equal, i.e. $s_1(f_1) = s_2(f_2)$. If, from any point on the reference path, only the value f_1 is increased then $s_1(f_1) > s_2(f_2)$. Similarly, if only f_2 is increased then $s_1(f_1) < s_2(f_2)$. In figure 4.3(b), a visualisation for three training plans with corresponding Pareto fronts is given. The Pareto-optimal plans are represented by the points $y^k = (y_1^k, y_2^k)$ for $k \in [3]$ (circles in figure 4.3(b)).

A first strategy could be to fit the reference path by applying linear regression. In figure 4.3(b), this means that the reference path should be fit to the points y^k , $k \in [3]$ as well as possible, i.e. we desire $s_1(y_1^k) = s_2(y_2^k)$ for all points $k \in [3]$. The resulting reference path is plotted in figure 4.3(c). However, in the following step to match the trade-off parameters for this reference path, we arrive at the situation in figure 4.3(d). This figure shows the differences between objective values for training plans (circles) and RPM generated plans (diamonds). In particular, these differences are relatively large for patient 3.

Our strategy attempts to avoid such large differences by already considering the trade-offs observed in the training plans when the reference path is determined. First, it should be noted that these trade-offs are characterised for each plan by its linear approximation of the Pareto front at y^k , see figure 4.3(e). Then, the direction of these linear approximations are matched to the indifference curves (figure 4.2). For example, the direction of the linear approximation of Pareto front 1 at point y^1 can be achieved by a relatively small value of $\beta_{1,2}$ (4.2). Similarly, relatively small values of $\beta_{2,1}$ can be used to match the directions of the other linear approximations. Therefore, the point y^1 should be in the part of the objective space where $s_1(f_1) < s_2(f_2)$, see figure 4.3(a). This means we desire $s_1(y_1^1) < s_2(y_2^1)$. Similarly, the points y^2 and y^3 should be in the part of the objective space where $s_1(f_1) > s_2(f_2)$. The resulting reference path is given in figure 4.3(f) for which the trade-off parameters can be matched better than in figure 4.3(d). This can be observed by comparing the RPM generated plans (diamonds) in figure 4.3(d) and figure 4.3(f).




Figure 4.3: Reference path fitting: (a) relation reference path and partial achievement functions, (b) Pareto fronts for three patients in the training set together with the objective values (circles), (c) reference path fitted by linear regression, (d) RPM generated plans (diamonds) after matching the trade-off parameters, (e) linearisations of the Pareto fronts at the optimal points in the training set (dashed lines), and (f) reference path fitted according to our strategy, and the RPM generated plans (diamonds) after matching the trade-off parameters.

To formalise our strategy, suppose that there are p plans in the training set. To obtain the linearisations, let the objective values observed in the training set plans be $y^k = (f_1(x^k), \dots, f_n(x^k))$ where x^k is the fluence map of training set plan $k \in [p]$. The linear approximations of the Pareto fronts at y^k are connected to the Lagrange multipliers $\lambda^k = (\lambda_1^k, \dots, \lambda_n^k)$, which are exactly those in the weighted sum optimisation (Miettinen, 1999, Breedveld et al., 2009)

$$\underset{x \in X}{\text{minimise}} \quad \sum_{i \in [n]} \lambda_i^k f_i(x) \quad (4.4)$$

for which x^k is the optimal solution. In optimisation problem (4.4), the set X^k is the feasible set corresponding to the k^{th} plan in the training set. The multipliers can usually be obtained explicitly from the Lagrangian (Miettinen, 1999). In Erasmus-iCycle, the Lagrange multipliers are a byproduct after solving the optimisation problem (Breedveld et al., 2009, 2017).

Per plan, there are n objectives and thus $n(n-1)/2$ pairwise comparisons between them. For each of these, we can check if the corresponding linear approximation is more similar to the situation for y^1 in figure 4.3(e), i.e. $\lambda_i^k < \lambda_j^k$, or more similar to the situation for y^2 or y^3 , i.e. $\lambda_i^k > \lambda_j^k$.

We use penalty functions for generating the reference path according to our strategy. For example, we want $s_i(y_i^k) < s_j(y_j^k)$ if $\lambda_i^k < \lambda_j^k$ since the corresponding angle $\beta_{i,j}$ is then less than 45 degrees. Therefore, we should penalise if $s_i(y_i^k) > s_j(y_j^k)$, which can be accomplished with the function $\phi_1^{i,j,k}$:

$$\begin{aligned} \phi_1^{i,j,k}(y_i^k, y_j^k; \theta) &= \max \{s_i(y_i^k; \theta) - s_j(y_j^k; \theta), 0\}, \\ \phi_2^{i,j,k}(y_i^k, y_j^k; \theta) &= \max \{s_j(y_j^k; \theta) - s_i(y_i^k; \theta), 0\}. \end{aligned}$$

Similarly, $\phi_2^{i,j,k}$ is used if $\lambda_i^k > \lambda_j^k$, i.e. if the angle $\beta_{j,i}$ is less than 45 degrees. For convenience, we introduced $\theta = (w, c)$ to denote the reference path parameters.

Define the sets

$$\begin{aligned} P_1 &:= \{(i, j, k) \in [n] \times [n] \times [p] \mid i < j, \lambda_i^k \leq \lambda_j^k\}, \\ P_2 &:= \{(i, j, k) \in [n] \times [n] \times [p] \mid i < j, \lambda_i^k > \lambda_j^k\}, \end{aligned}$$

where the condition $i < j$ is imposed to exclude redundant pairwise comparisons. Also, in order to obtain reference path parameters that are numerically stable in the RPM plan generation (4.3), we impose lower ($0 < w_{\text{low}}$ and c_{low}) and upper bounds (w_{up} and c_{up}) for the reference path parameters. Then, the reference path parameters are generated by

solving the following convex minimisation problem

$$\begin{aligned}
& \underset{\theta=(w,c)}{\text{minimise}} && \sum_{h \in [2]} \sum_{(i,j,k) \in P_h} \phi_h^{i,j,k}(y_i^k, y_j^k; \theta) + \delta \|\theta - \theta_{\text{low}}\|_2^2 \\
& \text{subject to} && 0 < w_{\text{low}} \leq w_i \leq w_{\text{up}} && i \in [n], \\
& && c_{\text{low}} \leq c_i \leq c_{\text{up}} && i \in [n].
\end{aligned} \tag{4.5}$$

Problem (4.5) thus minimises the sum of all penalty functions. Moreover, we added the regularisation term $\delta \|\theta - \theta_{\text{low}}\|_2^2$ with $\theta_{\text{low}} = (w_{\text{low}}, c_{\text{low}})$ in minimisation problem (4.5) to ensure a unique solution. Motivation for the values used for δ , and the lower (w_{low} and c_{low}) and upper bounds (w_{up} and c_{up}) of the reference path parameters are reported in section 4.2.3.3.

4.2.3.2 Step 2: trade-off parameters

Determining suitable trade-off parameters consists of two parts. The first part is an initialisation process in which the trade-offs parameters are matched to the linear approximations of the Pareto fronts for all training set plans (see figure 4.3(f)). The second part is an iterative procedure to guide the differences between objective values in the training set plans and those in the RPM generated plans towards clinically acceptable, predefined ranges. This second part is required since the first part is usually not sufficient: small deviations in the trade-offs can lead to large differences in objective values.

For the first part, trade-off information can be extracted by simplifying the RPM minimisation function (4.3). Using the reference path parameters acquired in section 4.2.3.1, the maximum of the partial achievement functions (4.1) for each plan $k \in [p]$ can be computed explicitly. If the maximum index is denoted by $m_k \in [n]$, the RPM minimisation function (4.3) reduces to

$$\begin{aligned}
& \underset{x \in X}{\text{minimise}} && (1 + \rho_{m_k}) w_{m_k} f_{m_k}(x) + \sum_{i \in [n] \setminus \{m_k\}} \rho_i w_i f_i(x) + \\
& && (1 + \rho_{m_k}) c_{m_k} + \sum_{i \in [n] \setminus \{m_k\}} \rho_i c_i,
\end{aligned} \tag{4.6}$$

which is an optimisation problem similar to (4.4). Since the term on the second line in (4.6) is constant (independent of x), it is irrelevant in solving the optimisation problem. From here, there are various options to initialise the trade-off parameters. One would be to build a linear but overdetermined system by matching the multiplicative quantities of the objective functions in (4.4) to those in (4.6) for each plan (Kaliszewski and Michalowski, 1997). However, our experience is that matching the arctangent for the maximum partial achievement β_{i, m_k} (4.2) to the actual arctangent observed in the training set plan leads to a better initialisation, see figure 4.3(f). The latter is given by the arctangent of the ratio

between the corresponding Lagrange multipliers. If we minimise the sum of the absolute differences between these arctangents for each plan in the training set, and also bound the trade-off parameters ($0 < \rho_{\text{low}} \leq \rho_i \leq \rho_{\text{up}}$), then initialising the trade-off parameters is accomplished by solving

$$\begin{aligned} & \underset{\rho}{\text{minimise}} \quad \sum_{k \in [p]} \sum_{i \in [n]} \left| \text{atan} \left(\frac{w_i}{w_{m_k}} \cdot \frac{\rho_i}{1 + \rho_{m_k}} \right) - \text{atan} \left(\frac{\lambda_i^k}{\lambda_{m_k}^k} \right) \right| \\ & \text{subject to} \quad 0 < \rho_{\text{low}} \leq \rho_i \leq \rho_{\text{up}} \quad i \in [n]. \end{aligned} \quad (4.7)$$

Since optimisation problem (4.7) is not convex, it is technically challenging to solve. The number of decision variables of the problem is, however, small (namely equal to the number of objectives n). We used the `fmincon` function in Matlab R2017a (Mathworks Inc.), which consistently found a local minimum for problem (4.7). Motivation for the values used for the lower (ρ_{low}) and upper bound (ρ_{up}) of the trade-off parameters are reported in section 4.2.3.3.

For the second part, the trade-off parameters are adjusted based on population-based differences for each objective. Let the Pareto-optimal fluence maps generated using the RPM be denoted by x_{RPM}^k for plan $k \in [p]$, then the differences in objective values for all training plans are calculated and combined in a vector,

$$\Delta f_i = [y_i^1 - f_i(x_{\text{RPM}}^1), \dots, y_i^p - f_i(x_{\text{RPM}}^p)], \quad i \in [n]. \quad (4.8)$$

Since lower objective values are preferred, a positive quantity in the vector (4.8) implies favourability for this objective in the RPM generated plan.

At this point, user preferences are required to continue the process. In this study, we use the preferences in table 4.2. Statistical population-based quantities together with desired lower bounds are specified for the differences observed in objective values between RPM generated plans and training set plans (4.8). In this way, the user can control unfavourable outliers while, for example, maintaining acceptable median values. The general aim of these quantities was to improve the database plans regarding high dose rectum sparing and mean dose bladder sparing, while maintaining a similar PTV coverage. Slight deteriorations were allowed for the PTV conformality and the mean doses to the rectum and anus. Larger deteriorations were allowed for the maximum dose to the femoral heads (but not above 39 Gy since this is a constraint in the wish-list in table 4.1).

Using the preferences in table 4.2, the trade-off parameters are adjusted heuristically according to algorithm 4.2.1. The general line of thought is that ρ_i is increased if objective f_i scored worse for the population than desired, but is decreased if objective f_i scored better for the population than desired. The heuristic of decreasing ρ_i serves to speed up the second part of determining suitable trade-off parameters.

Table 4.2: User preferences for the training set in terms of population-based quantities for the objective value differences between database plans and RPM generated plans (database — RPM). Median = 50th percentile, 1st quartile = 25th percentile.

Plan objective	Difference type	Lower bound
PTV LTCP	1 st percentile	0
Rectum gEUD ₁₂	1 st percentile	0
Rectum gEUD ₈	1 st percentile	0
PTV shell 5 mm D _{max}	median	-0.5
	1 st quartile	-1
	5 th percentile	-1.5
External ring 20 mm D _{max}	median	0
	1 st quartile	-0.75
	5 th percentile	-1
Rectum D _{mean}	median	-0.75
	1 st quartile	-1.5
	5 th percentile	-3
Anus D _{mean}	median	-0.75
	1 st quartile	-1.5
	5 th percentile	-3
PTV shell 15 mm D _{max}	median	-0.5
	1 st quartile	-1.25
	5 th percentile	-0.75
PTV shell 25 mm D _{max}	median	-0.5
	1 st quartile	-1.25
	5 th percentile	-1.75
Bladder D _{mean}	median	0
	1 st quartile	-0.25
	5 th percentile	-0.5
Femoral heads D _{max}	median	-5
	1 st quartile	-4
	5 th percentile	-3

Abbreviations: RPM = reference point method; PTV = planning target volume; LTCP = logarithmic tumour control probability; gEUD = generalised equivalent uniform dose.

4.2.3.3 Hyperparameters automated configuration procedure

The hyperparameters, i.e. the fixed parameters required to perform the actions described in the blocks in figure 4.1, are gathered in table 4.3. These hyperparameters were chosen heuristically. The bounds for the RPM parameters w_i , c_i , and ρ_i were chosen so that solving optimisation problem (4.3) is computationally stable, assuming that all objectives are of the same order of magnitude as reported in Van Haveren et al. (2017a) (if not, objectives should be scaled). We used a small value for the regularisation parameter δ to ensure a stable solution while having a small effect on the outcome. Hyperparameter values in algorithm 4.2.1 were chosen so that they give effective triggers for adjusting the trade-off parameters.

4.2.4 Treatment planning study set-up

The first and main part of the study aims to investigate the effect of using different-sized training sets on the quality of the RPM configuration generated. For this purpose, we ap-

Algorithm 4.2.1: Iterative trade-off parameter adjustment.

Input: Trade-off parameters ρ^k used in the last iteration k , and all corresponding differences Δf_i , $i \in [n]$ given by eq. (4.8).
Output: New set of trade-off parameters ρ^{k+1} .
 /* Hyperparameter initialisation (table 4.3) */
 $\rho_{\text{low}} > 0$: minimum value allowed for ρ ;
 $\rho_{\text{up}} > 0$: maximum value allowed for ρ ;
 $\delta^+ > 0$: increase factor;
 $\delta^{++} > 0$: large increase factor;
 $\delta^- > 0$: decrease factor;
 $\delta^{--} > 0$: large decrease factor;
 $\Delta_{\text{action}} > 0$: minimal value difference statistical population-based quantities for adjusting ρ^k ;
 /* Adjustment per objective */
for Objective $i \in [n]$ **do**
 Calculate predefined statistical population-based quantities for *Objective* i ;
 Pick the quantity that scored worst (largest negative difference with lower bound);
 Let u be a sample from the Uniform(0, 1) distribution;
 if The worst quantity scored lower than desired **then**
 /* Increase ρ_i but not beyond ρ_{up} */
 if The worst quantity scored lower than desired for all previous ρ **then**
 Set $\rho_i^{k+1} = \min(\rho_{\text{up}}, \rho_i^k + \delta^{++}u)$;
 else
 Set $\rho_i^{k+1} = \min(\rho_{\text{up}}, \rho_i^k + \delta^+u)$;
 end
 else
 if Difference for quantity $> \Delta_{\text{action}}$ **then**
 /* Decrease ρ_i but not below ρ_{low} */
 if The worst quantity scored better than desired for all previous ρ **then**
 Set $\rho_i^{k+1} = \max(\rho_{\text{low}}, \rho_i^k - \delta^{--}u)$;
 else
 Set $\rho_i^{k+1} = \max(\rho_{\text{low}}, \rho_i^k - \delta^-u)$;
 end
 end
 end
end

plied the automated configuration procedure (figure 4.1) to different-sized training sets of 9, 36, 72, 108, and 144 database plans. A variation of k -fold cross-validation was applied to each training set size ($k = 5$ for 9, $k = 4$ for 36, $k = 3$ for 72, $k = 2$ for both 108 and 144) to provide unbiased performance tests. For example, in case of 9 training set plans, we divided the database of 287 plans into 5 mutually exclusive training folds of 9 plans and a remainder fold of 242 plans. For each training fold of size p , we randomly selected approximately $60p/287$ plans from the low risk patient group, $113p/287$ plans from the intermediate risk patient group, and $114p/287$ plans from the high risk patient group. An RPM configuration was generated for each training fold and evaluated on the test set.

In the second part, a workflow for applying the automated configuration procedure to database plans that are not Pareto optimal is described and tested. We used a single

Table 4.3: Hyperparameters used in the automated configuration procedure.

Reference path parameters	
Description	Value/Range
Range for the reference path parameters	$[w_{\text{low}}, w_{\text{up}}] = [0.5, 10]$ $[c_{\text{low}}, c_{\text{up}}] = [0, 10^3]$
Regularisation parameter	$\delta = 10^{-16}$
Trade-off parameters	
Description	Value/Range
Range for the trade-off parameters	$[\rho_{\text{low}}, \rho_{\text{up}}] = [0.01, 50]$
Maximum number of iterations for trade-off tuning	$k_{\text{max}} = 40$
Increase factor	$\delta^+ = 0.35$
Large increase factor	$\delta^{++} = 1.2$
Decrease factor	$\delta^- = 0.03$
Large decrease factor	$\delta^{--} = 1.2$
Threshold for adjusting the trade-off parameters	$\Delta_{\text{action}} = 0.5$

training fold of 72 plans to demonstrate the workflow, and evaluated its performance on the corresponding test set.

For all automatically generated RPM configurations, the same user preferences in table 4.2 were used.

4.3 RESULTS

All database plans and all RPM generated plans in this study had sufficient PTV coverage, i.e. $V_{95\%} \geq 99\%$. Therefore, we focus on reporting the differences observed in plan objective values and other clinically relevant criteria values for the OARs.

All computations were done on a dual Intel Xeon E5-2690 Linux server. Most of the computation time required for obtaining the RPM configurations was spent in the “RPM plan generation” block in figure 4.1. The time required to process the actions in the other blocks in figure 4.1 was negligible. The average computation time for generating a single plan using the RPM was 1.1 min, so the average time required for one iteration of adjusting the trade-off parameters takes $1.1p$ min (where p is the number of training set plans). Total offline RPM configuration times were between 2 and 34 h. However, no user interaction was required. Also, in practice, the RPM configuration needs to be performed only once per treatment site. After the RPM configuration has been generated, clinically favourable and Pareto-optimal plan can be automatically generated in approximately 1.1 min for each patient of the corresponding treatment site.

Results for the Pareto-optimal plan database are given in section 4.3.1, and an example of a workflow for a non-Pareto-optimal plan database is given in section 4.3.2.

4.3.1 Pareto-optimal plan database

For each training fold, the user preferences in table 4.2 were satisfied. In figure 4.4, the differences observed in OAR objective values between Pareto-optimal database plans and Pareto-optimal RPM generated plans are shown for the corresponding test sets as a function of the training set size. For cross-validation, multiple mutually exclusive training folds were used per training set size to generate the RPM configurations. Paired two-sided Wilcoxon signed-rank tests were used to assess whether or not the plan objective differences have a median of zero. Differences in figure 4.4 marked with * are statistically significant ($p < 0.05$), indicating that the median is different from zero according to the test.

For all test sets, the median delivery of high dose to the rectum for the RPM generated plans was lower than that in the database plans (rectum gEUD_{12} , gEUD_8 , $V_{75\text{Gy}}$, and $V_{60\text{Gy}}$ in figure 4.4). The same observation can be made for the bladder mean dose. These improvements came at the cost of slight deteriorations in median values of other objectives, in particular the mean anus dose and the maximum doses to both the shell of 5 mm around the PTV and the femoral heads. This is in line with both our clinical priorities in the wish-list (table 4.1) and our user preferences (table 4.2).

For the training set size of 9, the several training folds show different median values for the objective differences related to the rectum and femoral heads. For larger training set sizes, the median values amongst the different training folds are more similar. Other differences occur in the severity of the outliers, e.g. for the training set size of 144, the test set for the first training fold shows no negative outliers for the rectum gEUD_{12} and gEUD_8 while the test set for the second training fold does.

Population median DVHs of the database plans and RPM generated plans for the test set patients are shown in figure 4.5, where the RPM configurations are based on 9 training plans figure 4.5(a), and 144 training plans figure 4.5(b). Both median DVHs show an almost identical median for the PTV $V_{95\%}$, while the medians of the mean bladder dose, and the rectum $V_{75\text{Gy}}$ and $V_{60\text{Gy}}$ are lower for the RPM generated plans than for the database plans.

4.3.2 Non-Pareto-optimal plan database

In this section, we give an example of a workflow where the starting point is a database of treatment plans that are not Pareto optimal. To automatically generate realistic non-Pareto-optimal plans similar to those in the Pareto-optimal database regarding the objective values, we applied the procedure described in appendix 4.A. Differences in objective values between the non-Pareto-optimal plans and corresponding Pareto-optimal plans are shown in figure 4.6, i.e. all deteriorations in objective values were less than 2 Gy.

Before the automated configuration procedure (figure 4.1) of the RPM can be applied, we require the objective values and Lagrange multipliers achieved in each plan in the database. For the objective values, we use those achieved in the non-Pareto-optimal plans. To




Figure 4.4: Boxplots of the differences in plan objective values (white background) and other clinically relevant evaluation criteria values (grey background) between database plans and RPM generated plans, evaluated on the test sets as a function of the training set size. Positive values are in favour of the RPM. The lines in the boxes represent medians, the boxes are between the first and third quartile, the whiskers are between the 2.5th and 97.5th percentile, and the circles are outliers. Note: there are variations in scales along the vertical axes.

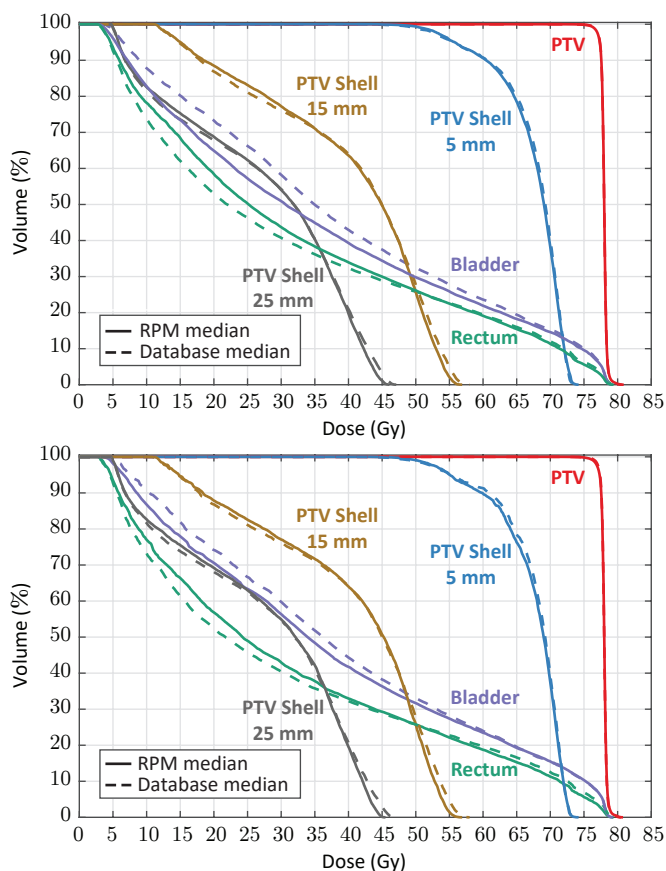


Figure 4.5: Population median DVHs for both the RPM generated plans of the test set patients (solid lines) and the corresponding database plans (dashed lines). The RPM configurations used were based on the set of (a) 9 training plans, first fold, and (b) 144 training plans, first fold.

obtain relevant Lagrange multipliers of a non-Pareto-optimal plan, we first project this plan onto the Pareto front according to the procedure described in appendix 4.B. From the resulting Pareto-optimal fluence map, the Lagrange multipliers can again be obtained as a byproduct of the optimisation.

For the automated configuration procedure of the RPM, we used the same 72 training set plans as in the first training fold in figure 4.4. The differences observed in objective values for OARs between database plans and corresponding Pareto-optimal RPM generated plans are shown in figure 4.7 for both the non-Pareto-optimal database and the Pareto-optimal database. Differences in figure 4.7 marked with * are statistically significant ($p < 0.05$) indicating that the median of the objective differences is not zero according to the Wilcoxon signed-rank test.

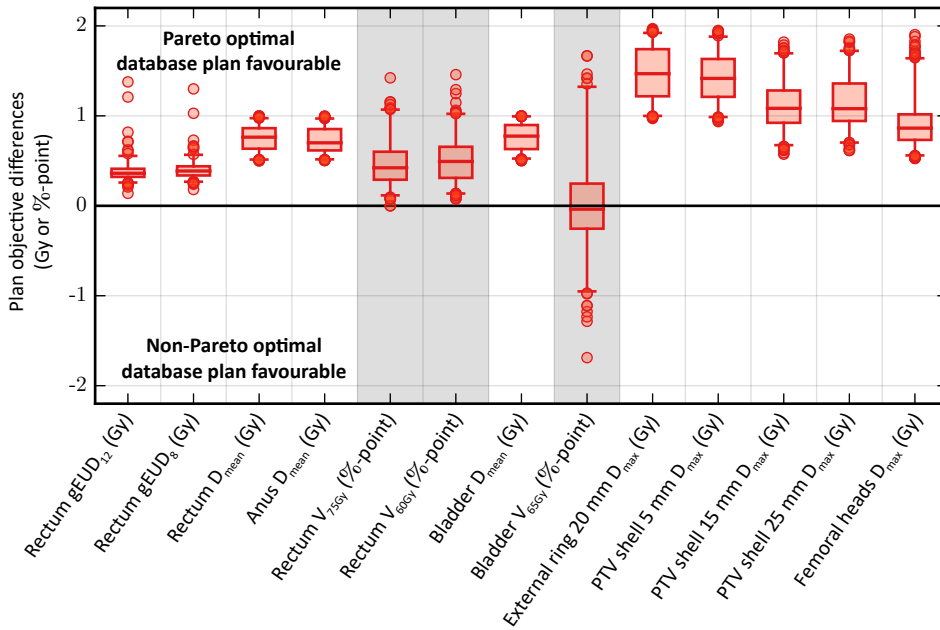


Figure 4.6: Boxplots of the differences in plan objective values (white background) and other clinically relevant evaluation criteria values (grey background) between plans in the non-Pareto-optimal database and the Pareto-optimal database. For all plan objectives, the differences are always positive. The lines in the boxes represent medians, the boxes are between the first and third quartile, the whiskers are between the 2.5th and 97.5th percentile, and the circles are outliers.

4.4 DISCUSSION

The purpose of this study was to develop and test a new automated procedure for configuring the RPM, a multi-objective method used for fast automated treatment planning. An appropriate RPM configuration is defined by user preferences: statistical population-based quantities together with desired lower bounds are specified for the differences observed in objective values between training set plans and RPM generated plans. Cross-validation was used to test the performance of the automated configuration procedure for the RPM. In practice, this procedure needs to be performed only once per treatment site. After the RPM configuration has been generated, a clinically favourable and Pareto-optimal treatment plan can be automatically generated for each patient in that treatment site.

In our previous studies (Van Haveren et al., 2017a,b), we demonstrated capability in generating high-quality plans for prostate cancer using the RPM, and for head and neck cancer patients using the LRPM. In these studies, however, the algorithm was configured using a labour-intensive and time-consuming interactive process. In the prostate study (Van Haveren et al., 2017a), the quality of the RPM generated plans was similar to that of the automatically generated plans using the 2péc method. In this paper, we showed that using

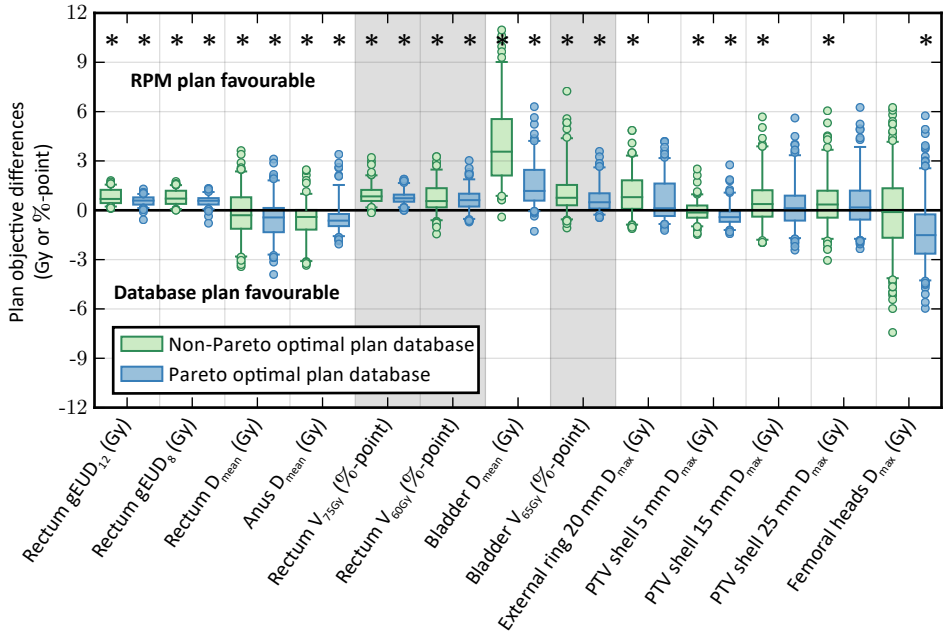


Figure 4.7: Boxplots of the differences in plan objective values (white background) and other clinically relevant evaluation criteria values (grey background) between the plan databases (non-Pareto-optimal and Pareto-optimal) and the RPM generated plans, evaluated on the test set. The lines in the boxes represent medians, the boxes are between the first and third quartile, the whiskers are between the 2.5th and 97.5th percentile, and the circles are outliers.

the RPM with an automated configuration resulted in similar or better plan quality with respect to population median parameters. Also, the severity of the outliers unfavourable for the RPM were within acceptable ranges. This indicates that automated configuration of the RPM should be preferred to manual configuration.

For both small and large training set sizes, we found that the population median parameters were mostly in favour of the RPM. To further avoid outliers unfavourable for the RPM of the rectum gEUD_{12} and gEUD_8 (figure 4.4), training on larger training set sizes and several training folds was required. On the other hand, differences between RPM generated plans and database plans for clinically relevant criteria that quantify delivery of high dose to the rectum, such as the rectum $V_{75\text{Gy}}$ and $V_{60\text{Gy}}$, performed well regardless of the training set size (figure 4.4). From this observation, the automated procedure for the RPM configuration may be improved by including clinically relevant criteria in the user preferences (table 4.2) for evaluation while retaining the convex plan objectives in the RPM fluence map optimisations.

The database plans were automatically generated using the clinically applied 2pc method, which ensures consistent plan quality of all plans. Possibly, this consistency in plan quality made it possible to generate appropriate RPM configuration for small training

set sizes. Without this consistency, larger training set sizes may be required.

Database plans that are not Pareto optimal can also be the starting point for applying the automated configuration procedure of the RPM. Relevant trade-off information, i.e. the Lagrange multipliers, can then be obtained by projecting the non-Pareto-optimal plan onto the Pareto front. It is likely that the boxplots in figure 4.7 are more favourable for the RPM when starting with a database of non-Pareto-optimal plans instead of Pareto-optimal plans. Since the database plans are not Pareto optimal, there is a “free” improvement for each plan objective without violating the constraints, see figure 4.6. For the boxplots in figure 4.7, this extra space resulted in a larger improvement for the bladder mean dose than was requested by the user preferences in table 4.2. In principle, this extra space could have also resulted in a larger improvement for another objective, since generating an RPM configuration according to the user preferences in table 4.2 for a training set of non-Pareto-optimal plans converges easier and faster than for Pareto-optimal plans. When observing that one objective performs better than requested, it thus means that the user preferences are too easily obeyed for the training set plans. To obtain a better result for another objective instead, the user preferences for that objective should be set more strict.

For a first time user of the automated configuration procedure, it may be difficult to select user preferences that yield clinically desired results. Since the second step of the automated configuration (section 4.2.3.2) is iterative, it could help to define the user preferences iteratively. In each iteration, boxplots of all plan objective differences, and median and individual DVHs can be presented to the user. Then, the user can adapt his/her preferences accordingly, and a new iteration starts. Evidently, this would make the procedure semi-automatic rather than automatic, but could improve the user-friendliness of the configuration procedure.

Although all computations for the automated configuration procedure of the RPM can be done offline, the time-efficiency can be improved. The largest improvement would be to generate the RPM plans in parallel instead of in serial. This is possible since the RPM configuration is fixed in each iteration of adjusting the trade-off parameters, see figure 4.1 and section 4.2.3.2. Another improvement is using a sparser bixel grid to decrease the computation time per plan. Finally, it could help to first train on a part of the training set for faster convergence, and then gradually add other parts of the training set. Ideally, this strategy reduces the total number of RPM plans that need to be generated.

Compared to previous studies (Van Haveren et al., 2017a,b), the RPM configuration procedure presented in this paper requires less information from the wish-list (table 4.1): only the constraints (volumes, types, and limits) and part of the objectives (volumes, types, and parameters) are required to initiate the procedure (priorities, and goal values of the objectives are not needed). Automating the selection of appropriate constraints and objectives is a topic for future research.

Another approach for automatically generating high-quality plans based on a training set is knowledge-based treatment planning, e.g. using the RapidPlan system (Varian Medical Systems, Palo Alto, USA). This commercially available planning tool (Yuan et al., 2012) first configures a DVH estimation model using information of at least 20 high-quality treatment plans. Then, the tool can be used to estimate OAR DVHs for new patients. Based on this estimation, a treatment plan can be automatically generated. Benefits of the approach presented in this paper are that Pareto-optimal plans are guaranteed, and that there is additional control in guiding the population-based plan quality by properly defining the user preferences.

The automated configuration procedure presented in this paper was applied to prostate cancer treatment with photon beams. There are no fundamental limitations for applying it to other (more challenging) tumour sites, or to treatment planning for particle therapy. These are topics for future studies.

4.5 CONCLUSIONS

This study presented a new fully automated configuration procedure for the reference point method (RPM), a multi-objective method used for fast automatic generation of Pareto-optimal treatment plans. After obtaining the RPM configuration, a clinically favourable plan can be generated for each patient. For prostate cancer patients, automatic RPM configurations were obtained for different-sized training plan sets. For each training set size, several training folds were used for cross-validation. All of the RPM configurations generated were of high quality, even when using only 9 plans in the training set. The proposed automated configuration procedure for the RPM has the potential to enhance plan quality and greatly reduce the manual configuration workload, thereby improving the efficiency and effectiveness of an automated clinical treatment planning workflow.

ACKNOWLEDGEMENTS

The authors would like to thank Yibing Wang of the Erasmus MC Cancer Institute for providing the CT scans and structure sets of the prostate cancer patients.

APPENDIX 4.A NON-PARETO-OPTIMAL DATABASE GENERATION

The starting point is a Pareto-optimal plan database. To automatically generate non-Pareto-optimal plans from these with slight but realistic deteriorations in objective values (i.e. the resulting non-Pareto-optimal plans are still of high quality, but slight improvements in objective values can be realised), our strategy was to divide the objectives f_i into three groups

1. I_{mean} , the indices corresponding to the mean dose objectives,
2. I_{max} , the indices corresponding to the maximum dose objectives,
3. $I_{\text{nonlinear}}$, the indices corresponding to the nonlinear objectives.

Then, the following convex optimisation problem was solved to generate a non-Pareto-optimal plan

$$\begin{aligned}
 & \underset{x \in X}{\text{minimise}} && \sum_{i \in I_{\text{mean}}} |f_i(x) - \bar{f}_i(x^*)| \\
 & \text{subject to} && \bar{d}_{\text{max}}^j(x^*) \leq d_{\text{max}}^j(x) && j \in I_{\text{max}}. \\
 & && \bar{\mathbf{d}}^j(x^*) \leq \mathbf{d}^j(x) && j \in I_{\text{nonlinear}},
 \end{aligned} \tag{4.9}$$

where $|\cdot|$ denotes the absolute norm, x^* is the corresponding Pareto-optimal fluence map, and

- $\bar{f}_i(x^*) = f_i(x^*) + u$, with $f_i(x^*)$ the objective value for fluence map x^* and u a sample from the $\text{Uniform}(0.5, 1)$ distribution,
- $\bar{d}_{\text{max}}^j(x^*) = d_{\text{max}}^j(x^*) + u$, with $d_{\text{max}}^j(x^*)$ the dose delivered to the voxel (one voxel per objective) corresponding with the maximum dose for fluence map x^* and u a sample from the $\text{Uniform}(0.5, 1)$ distribution,
- $\bar{\mathbf{d}}^j(x^*) = \mathbf{d}^j(x^*) + 10^{-4}$, with $\mathbf{d}^j(x^*)$ the dose vector of objective j for fluence map x^* . Dose values are thus slightly increased to ensure that the corresponding nonlinear objective value also slightly increases.

APPENDIX 4.B PLAN PROJECTION ONTO THE PARETO FRONT

For a non-Pareto-optimal plan, the RPM can be applied to project this plan onto the Pareto front. Suppose that r_i is the value achieved for objective f_i , $i \in [n]$. Then, the RPM plan parameters are set as follows

- $w_i = 1$, $i \in [n]$, i.e. the same gradient (of the reference path, see figure 4.2) for each plan objective is used,
- $c_i = -r_i$, $i \in [n]$, so that the reference path goes through the point representing the non-Pareto-optimal plan,

- $\rho_i = 10^{-16}$, $i \in [n]$, to guarantee Pareto optimality but so that the trade-off parameters have a negligible effect on the plan generated.

The optimisation problem to be solved is then given by (4.10).

$$\underset{x \in X}{\text{minimise}} \quad \left\{ \max_{i \in [n]} [f_i(x) - r_i] + 10^{-16} \sum_{i \in [n]} [f_i(x) - r_i] \right\}. \quad (4.10)$$

Automatic configuration of the reference point method for fully automated multi-objective treatment planning applied to oropharyngeal cancer

R. van Haveren, B.J.M. Heijmen, and S. Breedveld

Erasmus MC, University Medical Center Rotterdam, Department of Radiation Oncology, Rotterdam, The Netherlands

Medical Physics, Volume 47, Issue 4, 1499–1508, 4 February 2020

DOI: [10.1002/mp.14073](https://doi.org/10.1002/mp.14073)

ABSTRACT

Purpose: In automated treatment planning, configuration of the underlying algorithm to generate high-quality plans for all patients of a particular tumour type can be a major challenge. Often, a time-consuming trial-and-error tuning procedure is required. The purpose of this paper is to automatically configure an automated treatment planning algorithm for oropharyngeal cancer patients.

Methods: Recently, we proposed a new procedure to automatically configure the reference point method (RPM), a fast automatic multi-objective treatment planning algorithm. With a well-tuned configuration, the RPM generates a single Pareto-optimal treatment plan with clinically favourable trade-offs for each patient. The automatic configuration of the RPM requires a set of CT scans with corresponding dose distributions for training. Previously, we demonstrated for prostate cancer planning with 12 objectives that training with only 9 patients resulted in high-quality configurations. This paper further develops and explores the new automatic RPM configuration procedure for head and neck cancer planning with 22 objectives. Investigations were performed with planning CT scans of 105 previously treated unilateral or bilateral oropharyngeal cancer patients together with corresponding Pareto-optimal treatment plans. These plans were generated with our clinically applied 2-phase ϵ -constraint method (Erasmus-iCycle) for automated multi-objective treatment planning, ensuring consistent high quality and Pareto optimality of all plans. Clinically relevant, nonconvex criteria, such as dose-volume parameters and NTCPs, were included to steer the RPM configuration.

Results: Training sets with 20–50 patients were investigated. Even with 20 training plans, high-quality configurations of the RPM were feasible. Automated plan generation with the automatically configured RPM resulted in Pareto-optimal plans with overall similar or better quality than that of the Pareto-optimal database plans.

Conclusions: Automatic configuration of the RPM for automated treatment planning is feasible and drastically reduces the time and workload required when compared to manual tuning of an automated treatment planning algorithm.

5.1 INTRODUCTION

Generating high-quality IMRT or VMAT treatment plans for oropharyngeal cancer patients is challenging. A high dose is to be delivered to the planning target volume (PTV), which is in close proximity to many critical surrounding organs-at-risk (OARs) such as salivary glands, oral cavity, swallowing muscles, larynx, oesophagus, spinal cord, and brain-stem.

Several automated treatment planning approaches have been proposed in the literature (Hussein et al., 2018, Breedveld et al., 2019). This paper focusses on automated multi-objective fluence map optimisation to generate a single Pareto-optimal and clinically favourable treatment plan for each patient. Two algorithms for automated multi-objective optimisation of Pareto-optimal plans have been developed in our center: (1) the 2-phase ϵ -constraint (2pec) method (Breedveld et al., 2009) which is part of the clinically applied Erasmus-iCycle optimiser (Breedveld et al., 2012), and (2) the fast and fuzzy lexicographic reference point method (LRPM), see Van Haveren et al. (2017a). Previous studies (Voet et al., 2013a, Heijmen et al., 2018) have demonstrated that the quality of plans generated with the 2pec method is generally superior to that of manually generated plans. Main advantages of the LRPM over the 2pec method are faster plan generation with average relative speed-up factors of 12 for prostate (Van Haveren et al., 2017a) and 22 for head and neck cancer (Van Haveren et al., 2017b), and that trade-offs between all planning objectives are balanced simultaneously (LRPM) instead of pairwise (2pec method), allowing for large gains for some objectives at the cost of minor degradations for other objectives.

However, algorithms for automated planning have to be configured separately for all tumour sites. Interactive (manual) tuning of the configuration is a time-consuming and workload-intensive procedure for both the 2pec method (“wish-list” creation, see Breedveld et al. (2012), Heijmen et al. (2018)) and the LRPM (Van Haveren et al., 2017a,b). Recently, we proposed a new automatic procedure (Van Haveren et al., 2019) to configure the reference point method (RPM), see Wierzbicki (1982, 1986), Ogryczak and Kozłowski (2011), a special case of the LRPM (Van Haveren et al., 2017a). The procedure was successfully applied to prostate IMRT (Van Haveren et al., 2019), and adaptive prostate and cervix IMPT (Jagt et al., 2018, 2019).

This paper further develops and investigates the proposed automatic RPM configuration for a heterogeneous group of oropharyngeal cancer patients, with 22 objectives used in automatic plan generation. In previous work (Van Haveren et al., 2019), creation and evaluation of RPM configurations was based on convex plan criteria. This paper investigates the use of clinically more relevant nonconvex criteria such as dose-volume points or normal tissue complication probabilities (NTCPs). This allows for more flexible, intuitive, and clinically relevant automatic configurations. Dependency of the configuration quality on the (number of) selected training plans was included in the investigations.

Table 5.1: Wish-list used for generating the database plans with the 2p ϵ c method. The down-arrows (\downarrow) indicate that the objectives are to be minimised. Prescribed dose was $D_{\text{high}} = 70$ Gy for the PTV high, and $D_{\text{low}} = 54.25$ Gy for the PTV low.

Constraints					
	Volume	Type	Limit		
	PTV high	D_{max}	74.9 Gy	(= 107% of D_{high})	
	PTV high	D_{mean}	70.7 Gy	(= 101% of D_{high})	
	Spinal cord	D_{max}	42 Gy	(= 60% of D_{high})	
	Brainstem	D_{max}	49 Gy	(= 70% of D_{high})	
	PTV shell 0 mm	D_{max}	70 Gy	(= 100% of D_{high})	
	PTV shell 30 mm	D_{max}	35 Gy	(= 50% of D_{high})	
	Unspecified tissue	D_{max}	74.9 Gy	(= 107% of D_{high})	
Objectives					
Priority	Volume	Type	Goal	Sufficient	Parameters
1	PTV high	\downarrow LTCP	0.5	0.5	$D^p = D_{\text{high}},$ $\alpha = 0.8$
2	PTV low	\downarrow LTCP	0.5	0.5	$D^p = D_{\text{low}},$ $\alpha = 0.8$
3	Parotid glands	$\downarrow D_{\text{mean}}$	20 Gy		
4	SMGs	$\downarrow D_{\text{mean}}$	35 Gy		
5	MCS/MCP	$\downarrow D_{\text{mean}}$	25 Gy		
6	MCM/MCI	$\downarrow D_{\text{mean}}$	25 Gy		
7	PTV shell 5 mm	$\downarrow \text{gEUD}_{10}$	10 Gy		
	PTV shell 15 mm	$\downarrow \text{gEUD}_{10}$	10 Gy		
8	Oral cavity/Larynx	$\downarrow D_{\text{mean}}$	35 Gy		
9	Oesophagus	$\downarrow D_{\text{mean}}$	40 Gy		
10	Spinal cord/Brainstem	$\downarrow \text{gEUD}_{12}$	25 Gy		
11	PTV shell 40 mm	$\downarrow \text{gEUD}_8$	5 Gy		
	PTV shell 50 mm	$\downarrow \text{gEUD}_8$	5 Gy		
12	External ring 20 mm	$\downarrow D_{\text{max}}$	27.1 Gy		
13	Cochleas	$\downarrow D_{\text{mean}}$	35 Gy		

Abbreviations: PTV = planning target volume; SMG = submandibular gland; MCS = musculus constrictor superior; MCP = musculus constrictor cricopharyngeus; MCM = musculus constrictor medius; MCI = musculus constrictor inferior; LTCP = logarithmic tumour control probability; gEUD_r = generalised equivalent uniform dose with applied parameter r .

5.2 METHODS AND MATERIALS

5.2.1 Patient database

Planning CT scans of 105 previously treated unilateral and bilateral oropharyngeal cancer patients, together with a single corresponding Pareto-optimal treatment plan per scan, were included in a database. All patients were treated with a simultaneously integrated boost technique red similar to our clinical protocol (Wang et al., 2019). The high dose part of the PTV (PTV high) was prescribed 70 Gy, and the low dose part of the PTV (PTV low) was prescribed 54.25 Gy. A fixed coplanar equi-angular 23 beam setup was used for each patient to mimic VMAT-like dose distributions. The treatment was delivered in 35 fractions. In our clinical treatment planning workflow, the generated fleunce map is automatically converted to a VMAT plan using Monaco (Elekta AB, Sweden). In this study




Figure 5.1: Schematic overview of the automatic RPM configuration procedure.

however, plan comparisons are made with respect to the fluence maps so that the performance of both multi-objective methods are objectively compared (no bias due to VMAT segmentation).

Each Pareto-optimal plan in the database was generated with the 2p ϵ c method (Breedveld et al., 2009). The applied configuration (wish-list) for plan generation with 2p ϵ c method is presented in table 5.1. To achieve clinically acceptable coverage for both PTVs ($V_{95\%} \geq 98\%$), the logarithmic tumour control probability (Alber and Reemtsen, 2007) (LTCP) was used as objective function. For the OARs, the focus was either on minimising the mean dose (salivary glands, swallowing muscles, oral cavity, larynx, oesophagus, and cochleas) or on minimising the near maximum dose (spinal cord and brainstem) for which the generalised equivalent uniform dose (gEUD), see Niemierko (1997), with a high parameter value was used. Steering on the dose conformity was achieved by using maximum or near maximum doses to the PTV shells at 0, 5, 15, 30, 40, and 50 mm distance from the PTV. Entrance dose was controlled using the maximum dose to the external ring structure, which is the 20 mm ring inside the body contour. Hot spots were avoided by controlling the maximum dose in unspecified tissues.

5.2.2 Automatic RPM configuration

The automatic RPM configuration procedure (Van Haveren et al., 2019) applied in this paper is summarised in figure 5.1. For initialisation, a fraction of the patients in the database (section 5.2.1) was randomly selected for training (the remaining test patients were used to validate the configuration). Then, relevant data was acquired from the training plans (section 5.2.2.1) to create the final RPM configuration. (sections 5.2.2.2 and 5.2.2.3).

5.2.2.1 Data acquisition from training patients

The constraints and objectives used for plan generation with the 2p ϵ c method (table 5.1) were also the basis for plan generation with the RPM. For an RPM configuration, two 22-dimensional (or less if some OARs were not delineated) vectors were acquired from each training plan.

The first vector contained the values achieved for the up to 22 objectives used in the fluence map optimisation with the 2p ϵ c method (table 5.1).

The other vector contained, for each objective, a quantity related to the overall trade-offs made. More specifically, these were the Lagrange multipliers (one for each objective) resulting from the fluence map optimisation with the 2p ϵ c method. These Lagrange multipliers can be found as a byproduct of the optimisation (Breedveld et al., 2009, Van Haveren et al., 2019).

5.2.2.2 Automatic configuration procedure

The RPM automatically generates a fluence map by solving the minimisation problem

$$\underset{x \in X}{\text{minimise}} \left\{ \max_{i \in [n]} [w_i f_i(x) + c_i] + \sum_{i \in [n]} \rho_i (w_i f_i(x) + c_i) \right\}. \quad (5.1)$$

Here, x is the fluence map, X a constrained set, $f_1(x), \dots, f_n(x)$ the objectives, and the $w_1, c_1, \rho_1, \dots, w_n, c_n, \rho_n$ define an RPM configuration. The $w_1, c_1, \dots, w_n, c_n$ prioritise the objectives, and ρ_1, \dots, ρ_n quantify desired trade-offs between objectives. In the automatic procedure, each RPM configuration is iteratively generated. In the first iteration, the data acquired from the training database (section 5.2.2.1) was used to generate an initial RPM configuration (technical details in Van Haveren et al. (2019)). With this configuration, a single Pareto-optimal RPM plan can then be automatically generated for each training patient. Based on the differences observed between training and RPM generated plans for target coverage and other plan parameters, the RPM configuration was then either accepted or not (see section 5.2.2.3). If an RPM configuration was not accepted, the configuration was updated for the next iteration and the process was repeated. Updating the configuration is achieved by updating the trade-off parameters ρ_1, \dots, ρ_n . The general rule is that ρ_i is increased if its corresponding plan parameter scored worse than desired, but is decreased if the corresponding plan parameter scored better for the population than desired (details of the heuristic in Van Haveren et al. (2019)). If an RPM configuration was acceptable or if the RPM configuration is still not acceptable after 40 iterations (heuristic), the iterative process terminated and returned the final RPM configuration.

5.2.2.3 User-defined preferences for automatic RPM configuration

Each automatic RPM configuration is steered by a set of user-defined preferences. There are two types of preferences: (1) preferences regarding a minimum/maximum allowed value for a plan parameter in the RPM generated plans; (2) preferences regarding differences for a plan parameter between the training and RPM generated plans. An example for the first type is seen in the first row of table 5.2, which indicates that the minimum allowed value for $V_{95\%}$ of the PTV low and PTV high is 98% in all RPM generated plans. An example of the second type is seen in the second row of table 5.2, which indicates that the median value of all differences (database — RPM) in parotid gland NTCP (for both left and right)

Table 5.2: User preferences to create and evaluate an RPM configuration.

Plan parameter	Type	Lower bound	Planning objective
PTV low/PTV high $V_{95\%}$	minimum	98	PTVs LTCP
Parotids glands NTCP	median	0	Parotid glands D_{mean}
	5 th percentile	- 2.5	
SMGs/oral cavity NTCP	median	0	SMGs/oral cavity D_{mean}
	5 th percentile	- 4	
MCS D_{mean}	median	0	
	5 th percentile	- 2	
MCP D_{mean}	median	0	
	5 th percentile	- 2.5	
MCM/MCI D_{mean}	median	0	
	5 th percentile	- 3	
Larynx/oesophagus D_{mean}	median	0	
	5 th percentile	- 3	
Spinal cord/brainstem $gEUD_{12}$	median	- 1	
	5 th percentile	- 3	
Cochleas D_{mean}	1 st quartile	- 5	
PTV shell 5 mm $gEUD_8$	median	- 0.5	
	5 th percentile	- 3	
PTV shell 15 mm $gEUD_8$	median	- 0.75	
	5 th percentile	- 3.25	
PTV shell 40 mm $gEUD_8$	median	- 1.25	
	5 th percentile	- 3.75	
PTV shell 50 mm $gEUD_8$	median	- 1.5	
	5 th percentile	- 4	
External ring 20 mm D_{max}	median	- 1.5	
	5 th percentile	- 5	

Median = 50th percentile; 1st quartile = 25th percentile; PTV = planning target volume; SMG = submandibular gland; MCS = musculus constrictor superior; MCP = musculus constrictor cricopharyngeus; MCM = musculus constrictor medius; MCI = musculus constrictor inferior; LTCP = logarithmic tumour control probability; $gEUD_r$ = generalised equivalent uniform dose with applied parameter r .

is at least 0 (in %-point). Instead of the median value, other percentile values can be used as well. Multiple measures can be defined per plan parameter. If all measures are above the desired lower bounds, the RPM configuration is accepted.

For the automatic RPM configuration applied to prostate cancer (Van Haveren et al., 2019), only the convex constraints and objectives as applied in the wish-list were used for defining the user-preferences. A drawback of this approach is that clinically relevant plan quality criteria may involve nonconvex functions such as dose-volume points or models for predicting NTCPs. Therefore, we extended the previous methodology by allowing general nonconvex functions to be applied in the user preferences. The user preferences in table 5.2 were applied for creation and evaluation of all RPM configurations. The applied nonconvex functions were linked to convex surrogates, which were used in the plan optimisations (compare with table 5.1). The first row in table 5.2 specifies that for both the PTV high and PTV low, the $V_{95\%}$ should be at least 98% in all RPM generated plans. The following

Lyman NTCP model (Lyman, 1985) was applied for predicting xerostomia,

$$\text{NTCP}(D_{\text{mean}}) = (2\pi)^{-1/2} \int_{-\infty}^{(D_{\text{mean}} - 40)/16} \exp(-t^2/2) dt, \quad (5.2)$$

with D_{mean} the mean dose in a salivary gland or the oral cavity. The second row in table 5.2 specifies that the NTCP values in at least 50% of the RPM generated plans should be lower than those in the training plans, and that the NTCP values in at most 5% of the RPM generated plans can be 2.5 %-point higher than those in the training plans.

The aim of the user preferences in table 5.2 was to define an RPM configuration resulting in plans with: (1) sufficient target coverage for all patients ($V_{95\%} \geq 98\%$); (2) overall reduced NTCP values in salivary glands and oral cavity and reduced mean doses in the swallowing muscles. If needed to accomplish (1) and (2), moderate deteriorations were allowed for the spinal cord, brainstem, cochleas, and conformality measures (PTV shells and external ring). Both the median and 5th percentile were often used to both control the overall differences and to mostly avoid large unfavourable outliers for the RPM.

5.2.3 Variations in training sets

RPM configurations were established for various training sets: a variation of k -fold cross-validation was applied to training sets with 20 ($k = 5$) plans. Training sets with 35 ($k = 3$) and 50 ($k = 2$) plans were also tested. Selections of patients for the training folds was always random, with each patient only present in one fold. The quality of an RPM configuration was determined by comparing the RPM generated plans with the database plans for the test patients (patients not used for training) regarding the plan parameters defined in table 5.2. To visualise the heterogeneity of the training folds with 20 plans, the plan parameters for the corresponding database plans are shown in figure 5.2.

Paired two-sided Wilcoxon signed rank tests were applied to assess whether or not the differences in plan parameter values between database and RPM generated plans for the test patients were statistically significant ($p < 0.05$).

5.3 RESULTS

5.3.1 Target coverage

All database and RPM generated plans had clinically acceptable target coverage, i.e. the $V_{95\%}$ was at least 98% for both the PTV high and the PTV low.

Differences in target coverage between the database and corresponding RPM generated plans up to 1 %-point were observed. To focus on analysing differences in other plan parameters, all dose distributions were first scaled such that the $V_{95\%}$ for either the PTV low or PTV high was 98%.

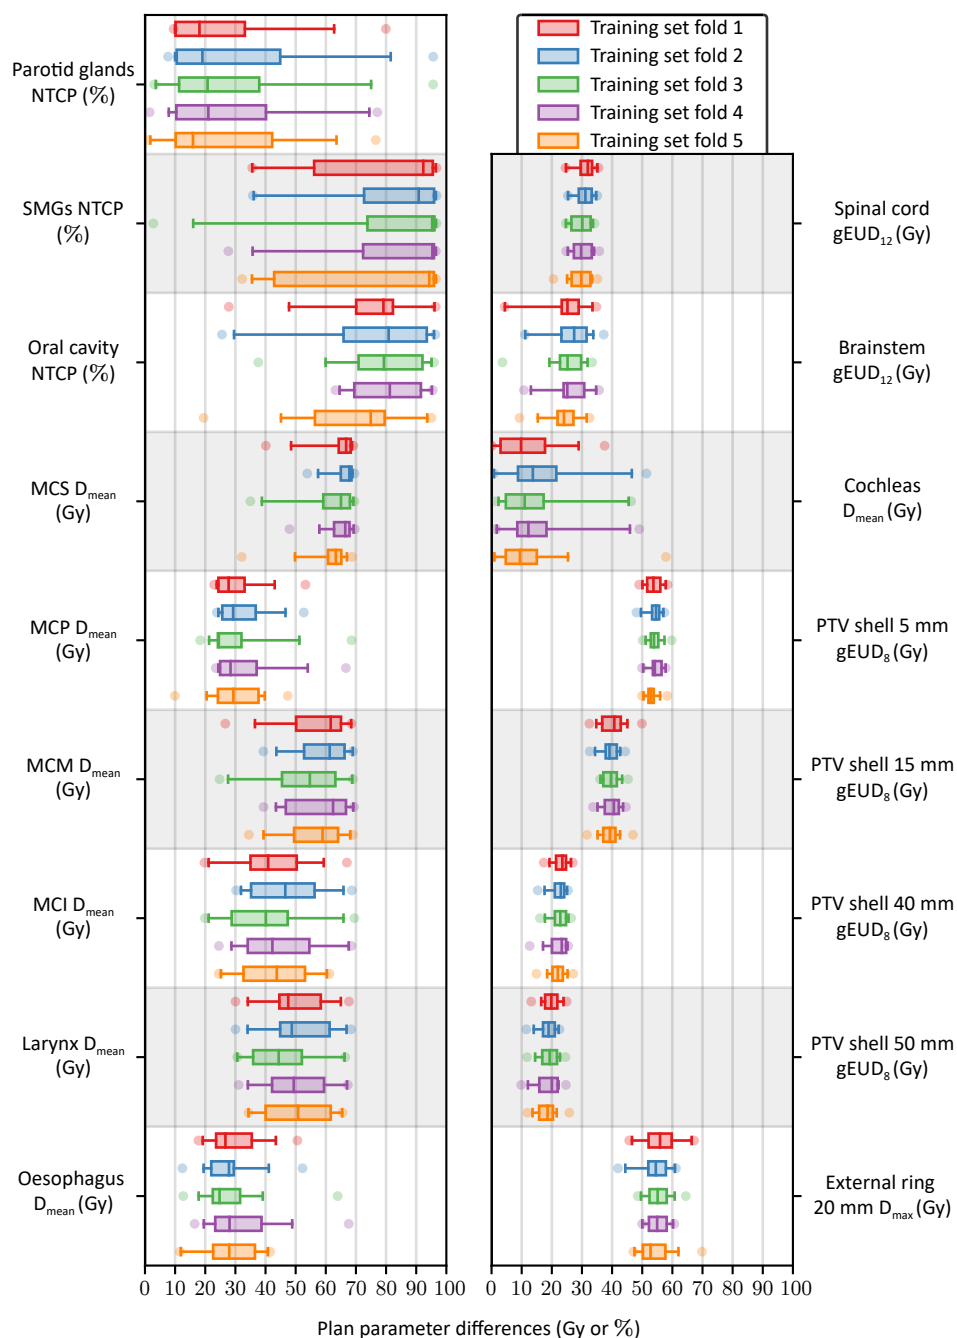


Figure 5.2: Boxplots of the plan parameter values (table 5.2) for the database plans corresponding to the five different training folds, each with 20 training patients. Vertical thick lines within the boxes are medians, boxes are between the first and third quartile, whiskers are between the 2.5th and 97.5th percentile, circles are outliers.

5.3.2 OAR sparing and conformality

For the five RPM configurations based on different sets of 20 training patients, the differences observed in plan parameters between database and RPM generated plans for the test patients are presented in figure 5.3.

For most plan parameters, the distribution of differences and the corresponding median difference for the five test folds were similar. The submandibular glands (SMGs), oral cavity, oesophagus, spinal cord, and brainstem showed overall better sparing for the RPM generated plans at the cost of some deterioration in conformality measures. The differences observed in plan parameters between RPM generated plans and database plans are in line with table 5.2, where preference is given to improve RPM generated plans regarding organ sparing by allowing some deterioration in conformality.

Figure 5.3 also shows outliers for some plan parameters, often in favour of the RPM. Since the performance of the RPM configuration (obtained with training set fold 4, see figure 5.2) on test fold 4 (see figure 5.3) is according to the user preferences (table 5.2), particularly for the oral cavity NTCP, this fold is analysed more in depth. For test fold 4, differences of the most important plan parameters are shown in figure 5.4 for the 15 plans with the most extreme outliers (both favourable and unfavourable for the RPM). As a reference, the last column in figure 5.4 shows the mean differences for *all* test fold 4 patients, clearly showing an overall gain for the RPM.

In another approach for comparing RPM generated plans with database plans, differences in all plan parameter values were summed for each patient in test fold 4. A histogram of the summed differences is presented in figure 5.5. The median of the summed differences was 13.1, indicating an advantage for the RPM ($p < 0.001$). This advantage was seen in 70 out of 85 patients in test fold 4.

In the [supporting information](#), results are presented for training with 35 and 50 patients. In general, it was found that increasing the number of training patients resulted in (1) slightly more consistent results for the test patients among the different folds with the same number of training patients, and (2) reduced severity of the outliers unfavourable for the RPM.

5.3.3 Computation times

All computations were performed on a dual Intel Xeon E5-2690 Linux server using an in-house developed solver tuned for radiotherapy treatment planning ([Breedveld et al., 2017](#)). On average, 5.6 min of computation time was required to generate a single RPM plan. Total computation times to automatically generate an RPM configuration ranged between 22.3 and 61.9 hours without any user interaction.




Figure 5.3: Boxplots of the differences in plan parameter values (table 5.2) between database plans and RPM generated plans for the five test folds corresponding to the five different RPM configurations with 20 training patients. Positive values are favourable for the RPM. Vertical thick lines within the boxes are medians, boxes are between the first and third quartile, whiskers are between the 2.5th and 97.5th percentile, circles are outliers, arrows indicate large outliers. Statistically significant differences ($p < 0.05$) in favour of database plans (\blacktriangleleft) or RPM plans (\blacktriangleright).




Figure 5.4: Differences in most important plan parameter values between database plans and RPM generated plans for the 15 most extreme outliers in test fold 4 (training with 20 patients), both favourable and unfavourable for the RPM. Positive values are favourable for the RPM. The last column shows the average results for all test patients.

5.4 DISCUSSION

The purpose of this study was to further develop and explore a recently introduced automatic configuration procedure for the RPM ([Van Haveren et al., 2019](#)), an algorithm for fast automated multi-objective treatment planning. The automatic configuration procedure requires a training set (delineated CT scans with corresponding treatment plans) as input.

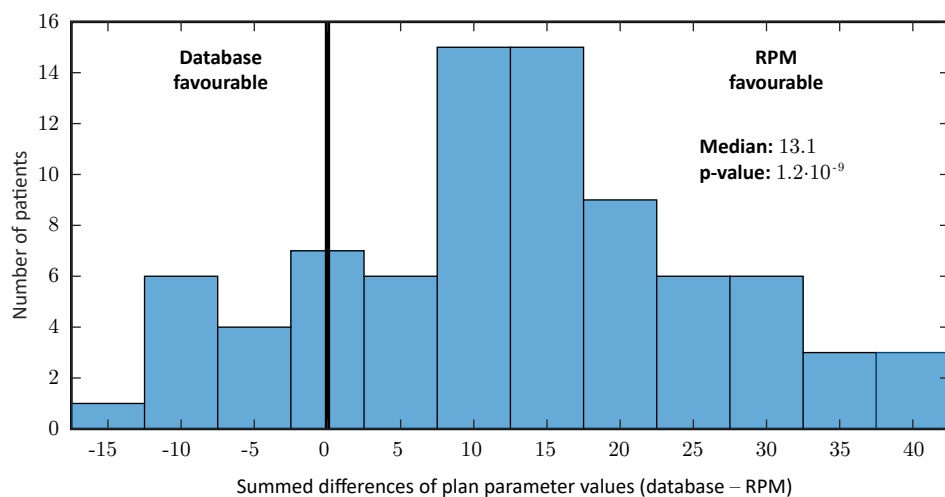


Figure 5.5: Histogram of the summed differences of plan parameter values (database – RPM) for all 85 patients in test fold 4 (training with 20 patients). The median gain of 13.1 indicates an advantage for the RPM ($p < 0.001$).

This study tested the automatic configuration for a heterogeneous group of unilateral and bilateral oropharyngeal cancer patients with planning based on 22 objectives, and demonstrated that high-quality configurations were obtained with only 20 training patients.

In previous work (Van Haveren et al., 2017b), the LRPM was used to automatically generate clinically favourable treatment plans for fifteen head and neck cancer patients. In that paper, part of the LRPM configuration (trade-off configuration) was established manually. This study improves on that work in several ways. First, we have shown that clinically favourable treatment plans for head and neck cancer patients can also be generated with the RPM (linear reference path) instead of the more complex LRPM (piecewise linear reference path). Secondly, it was shown that a single RPM configuration can generate clinically favourable plans for a larger patient database (105 patients instead of 15). Thirdly, in this work, the RPM configuration was automatically generated, removing the need for extensive manual tuning. Finally, a more heterogeneous patient database was included in this study, demonstrating flexibility of the RPM for automated treatment planning.

Whereas the user preferences for creating and evaluating RPM configurations in previous work (Van Haveren et al., 2019) were based exclusively on the convex planning objectives used in fluence map optimisation, this paper describes how nonconvex criteria, such as dose-volume parameters or NTCPs (e.g. see table 5.2), can be included by coupling them to correlated convex objectives. This made the automatic configuration more intuitive and clinically relevant, while the fluence map optimisation problem remained convex guaranteeing optimality of the plan generated.

Automatic RPM configurations were based on user preferences regarding population-

based differences between database and RPM generated plans (e.g. table 5.2). In practice, the lower bounds defined for the statistical population-based user preferences can be derived iteratively. For example, the first step can be to only define a median for each criterion, then perform a full configuration run, and then add or adjust the measures and lower bounds for criteria that showed undesired trade-offs. In this way the user iteratively gets a better understanding about which of the plan parameters are difficult to improve, and which are less difficult to improve. This procedure can then be repeated until a configuration is obtained that results in desirable trade-offs between all criteria. Tuning the entries in table 5.2 is easier for the user than tuning the RPM parameters directly, since the user is familiar with interpreting the plan parameters but not with the RPM parameters. Even with expert knowledge of the RPM, automatic configuration has shown to be superior (Van Haveren et al., 2017a, 2019) for prostate planning. Note that for any configuration, Pareto optimality of all RPM generated plans is guaranteed (Wierzbicki, 1986).

Compared to the automatic RPM configuration for automatic prostate planning (Van Haveren et al., 2019), we observed more variation in differences between database and RPM generated plans among the training folds (figures 5.3, 5.6, and 5.7). For example, for the different training sets of 20 patients (figure 5.3), slightly different trade-offs were observed among the different test folds: folds 1 and 2 showed better sparing of SMGs and oral cavity than folds 3 and 5 at the cost of degradations in the conformality measures. For training based on a larger training set of 50 patients (figure 5.7), the median differences were more consistent among the different test folds. However, differences in outliers were still present: fold 1 showed better sparing of SMGs and oral cavity than fold 2. This is likely due to the heterogeneous patient database (section 5.2.1). As can be seen in figure 5.7, the distribution of differences in plan parameter values for the test patients in fold 2 were slightly worse than desired (table 5.2) for the SMGs and oral cavity NTCP values. The recommendation for a heterogeneous group of patients is to generate various configurations, one for each different training fold, also with variation in training set sizes, in order to investigate variation in configuration quality related to the patient heterogeneity. Each of these configurations could include an iterative fine-tuning of the user preferences (see above and table 5.2). A single (large) test fold could ideally be the basis for all configurations (requiring many patients). Ideally, there is also a large evaluation fold with patients not used for training nor testing for final configuration selection and quality assessment.

Overall, the RPM generated plans showed a better OAR sparing at the cost of some decreased conformality. In figure 5.5, differences in OAR criteria values were added for each test plan and displayed in a histogram. The median improvement of 13.1 units is in favour of the RPM ($p < 0.001$). Technically, the maximum gain for this measure can be achieved by generating plans using the weighted sum method with equal weights (Miettinen, 1999). However, the RPM also ensures that the differences in criteria values corresponding to

OARs with high clinical priorities are within an acceptable range for each patient, which can be observed in figure 5.3.

A similar approach to automatic configuration of the RPM is knowledge-based planning (KBP), e.g. see [Ge and Wu \(2019\)](#). Both approaches rely on a set of training plans from previously treated patients. The main difference is that the training plans lead to explicit specification of the RPM parameters in the automatic configuration approach, while they are applied to create a model in KBP approach. This model is trained, using machine learning techniques such as deep-learning ([Chen et al., 2019](#)), support vector regression ([Ma et al., 2019](#)) or generative adversarial networks ([Babier et al., 2020](#)), to predict the DVHs or spatial dose distribution prediction. The predicted DVHs or dose distribution are then the basis for plan optimisation ([Babier and Boutilier, 2018](#), [Miguel-Chumacero et al., 2018](#), [Fan et al., 2019](#)). Both the automatic configuration and KBP approach report promising results.

The RPM automatically generates a Pareto-optimal fluence map plan and can thus not be directly delivered as the treatment device parameters are still unspecified. A recently developed automated segmentation algorithm ([Schipaanboord et al., 2019](#)) shows segmented plans are dosimetrically similar to the fluence map plans. The plan comparisons presented in this paper should therefore be an accurate representation of the plan comparisons after segmentation.

In this paper, the objectives and constraints in table 5.1 were used as a starting point for all automatic plan generations. A next step can be to eliminate the requirement to explicitly specify these objectives and constraints, which could possibly be achieved with inverse multi-objective optimisation techniques ([Chan et al., 2014](#)). This is a topic for further research.

5.5 CONCLUSIONS

A fully automated procedure for flexible and intuitive configuration of the reference point method (RPM), an algorithm for fast automated multi-objective plan generation, was tested for a heterogeneous group of oropharyngeal cancer patients. For each patient, the automatic RPM configuration allowed for fast automatic generation of a Pareto-optimal plan with clinically favourable trade-offs, even for configurations based on only 20 training patients. As requested, the configurations generally resulted in lower OAR doses than those in the database plans at the cost of slightly reduced conformality. The RPM also resulted in favourable outliers for doses in highly prioritised OARs. Automatic RPM configuration has great potential in replacing traditional time-consuming and labour-intensive treatment planning workflows relying on manual configuration.

APPENDIX 5.A SUPPORTING INFORMATION




Figure 5.6: Boxplots of the differences in plan parameter values (table 5.2) between database plans and RPM generated plans for the three test folds corresponding to the three different RPM configurations with 35 training patients. Positive values are favourable for the RPM. Vertical thick lines within the boxes are medians, boxes are between the first and third quartile, whiskers are between the 2.5th and 97.5th percentile, circles are outliers, arrows indicate large outliers. Statistically significant differences ($p < 0.05$) in favour of database plans (◀) or RPM plans (▶).

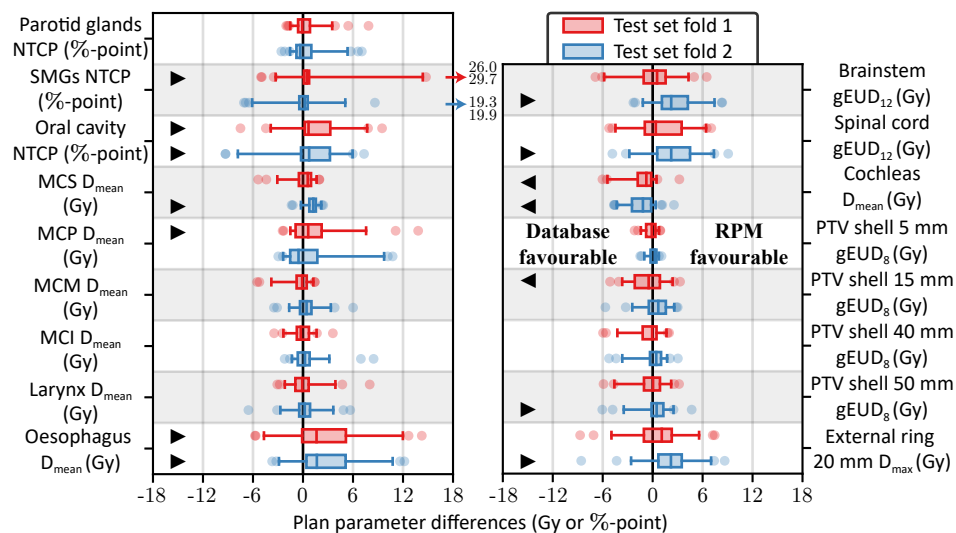


Figure 5.7: Boxplots of the differences in plan parameter values (table 5.2) between database plans and RPM generated plans for the two test folds corresponding to the two different RPM configurations with 50 training patients. Positive values are favourable for the RPM. Vertical thick lines within the boxes are medians, boxes are between the first and third quartile, whiskers are between the 2.5th and 97.5th percentile, circles are outliers, arrows indicate large outliers. Statistically significant differences ($p < 0.05$) in favour of database plans (\blacktriangleleft) or RPM plans (\blacktriangleright).

Plan-library supported automated replanning for online-adaptive intensity-modulated proton therapy of cervical cancer

T.Z. Jagt¹, S. Breedveld¹, R. van Haveren¹, R.A. Nout², E. Astreinidou², B.J.M. Heijmen¹, and M.S. Hoogeman¹

¹ Erasmus MC, University Medical Center Rotterdam, Department of Radiation Oncology, Rotterdam, The Netherlands

² Leiden University Medical Center, Department of Radiation Oncology, Leiden, The Netherlands

Acta Oncologica, Volume 58, Issue 10, 1440–1445, 23 October 2019

DOI: [10.1080/0284186X.2019.1627414](https://doi.org/10.1080/0284186X.2019.1627414)

ABSTRACT

Background: Intensity-modulated proton therapy is sensitive to inter-fraction variations, including density changes along the pencil-beam paths and variations in organ-shape and location. Large day-to-day variations are seen for cervical cancer patients. The purpose of this study was to develop and evaluate a novel method for online selection of a plan from a patient-specific library of prior plans for different anatomies, and adapt it for the daily anatomy.

Material and methods: The patient-specific library of prior plans accounting for altered target geometries was generated using a pretreatment established target motion model. Each fraction, the best fitting prior plan was selected. This prior plan was adapted using (1) a restoration of spot-positions (Bragg peaks) by adapting the energies to the new water equivalent path lengths; and (2) a spot addition to fully cover the target of the day, followed by a fast optimisation of the spot-weights with the reference point method (RPM) to obtain a Pareto-optimal plan for the daily anatomy. Spot addition and spot-weight optimisation could be repeated iteratively. The patient cohort consisted of six patients with in total 23 repeat-CT scans, with a prescribed dose of 45 Gy(RBE) to the primary tumour and the nodal CTV. Using a 1-plan-library (one prior plan based on all motion in the motion model) was compared to choosing from a 2-plan-library (two prior plans based on part of the motion).

Results: Applying the prior-plan adaptation method with one iteration of adding spots resulted in clinically acceptable target coverage ($V_{95\%} \geq 95\%$ and $V_{107\%} \leq 2\%$) for 37/46 plans using the 1-plan-library and 41/46 plans for the 2-plan-library. When adding spots twice, the 2-plan-library approach could obtain acceptable coverage for all scans, while the 1-plan-library approach showed $V_{107\%} > 2\%$ for 3/46 plans. Similar OAR results were obtained.

Conclusion: The automated prior-plan adaptation method can successfully adapt for the large day-to-day variations observed in cervical cancer patients.

6.1 INTRODUCTION

Highly localised dose deposition is possible in intensity-modulated proton therapy (IMPT) using the characteristic Bragg peak. At the same time, this treatment modality is sensitive to inter-fraction variations, including density changes along the pencil-beam paths and variations in organ-shape and location ([Lomax, 2008a,b](#)).

Large day-to-day variations can be seen in the shape and position of the cervix-uterus, mostly due to changes in filling of bladder, rectum and sigmoid. Displacements of the tip of the uterus of more than 3 cm between an empty-bladder and a full-bladder anatomy are common. In photon beam radiotherapy, a plan-of-the-day approach has been clinically implemented in several centres, in which a daily image is used to select the best fitting treatment plan from a plan-library ([Bondar et al., 2012](#), [Heijkoop et al., 2014](#)).

For cervical cancer IMPT, such an approach has been investigated by [Van de Schoot et al. \(2016\)](#). The cervix-uterus positions of a full and empty bladder CT scan were used to create an internal target volume (ITV) encompassing all possible positions. This ITV was divided into subITVs with which a patient specific plan-library was generated. All library plans were robustly optimised using 8 mm set-up errors and 3% range errors. For each simulated fraction, the library plan encompassing the daily CTV was selected, and recalculated on the daily anatomy without further (re-)optimisation. Despite the generous robustness settings, the selected plan resulted in inadequate CTV coverage in about 10% of the repeat-CT scans, due to “substantial deviating anatomy compared to the pretreatment derived full range ITV” ([Van de Schoot et al., 2016](#)). This shows that when the daily anatomy greatly deviates from the pretreatment observed motion, using a plan-library with robust treatment plans is insufficient to guarantee target coverage.

In this study we therefore propose to automatically adapt the treatment plan that is selected from the plan-library using our automated plan adaptation method developed for prostate cancer in previous work ([Jagt et al., 2018](#)). In this proposed prior plan strategy, the plan selected from a library of prior plans is adapted by an energy adaptation of the pencil-beams, followed by adding spots and a weight optimisation using the reference point method (RPM) using automatically tuned RPM-parameters. Outcomes were compared to forward dose calculation of the prior plans on the repeat-CT scans (no replanning), and to full, time-consuming multi-criteria optimisations for the daily scans (benchmark). To investigate the effect of using a prior plan as a warm-start for optimisation, outcomes were also compared to a time-constrained non-prior-plan strategy in which a new plan is generated from scratch for the daily scans.

6.2 METHODS AND MATERIALS

6.2.1 Patient data

This study included data of six patients with locally advanced cervical cancer selected from an institutional review board approved prospective study designed to investigate inter-fraction motion in cervical cancer patients. For every patient, a full and empty bladder CT scan was acquired pretreatment as well as four weekly repeat-CT scans, resulting in a total of 23 repeat-CT scans. More detailed background of the patient data can be found in the [Supplementary Materials](#).

6.2.2 Treatment planning volumes and prescription

The goal was to obtain clinically acceptable daily treatment plans for the repeat-CT scans. In the foreseen online-adaptive workflow, we assume that the structures are daily segmented automatically or with minimal user interaction. To account for intra-fraction uncertainties and inevitable uncertainties in the structure segmentation of the adaptive workflow, the daily targets were expanded with a margin: a PTV_{OAPT} (PTV Online-Adaptive Proton Therapy) was created by adding a 5 mm margin around the primal CTV and a 2 mm margin around the nodal CTV ([Heijkoop et al., 2015](#), [Van de Sande et al., 2016](#)). Prescribed dose to the PTV_{OAPT} was set to 45 Gy(RBE), using an RBE of 1.1, which was delivered by four beams (0°, 90°, 180°, and 270°).

6.2.3 The automated adaptive treatment planning method

The proposed strategy starts by selecting the best prior plan from a plan-library. The spot-positions are then restored by adapting the energy of each spot to the new water equivalent path length (WEPL). To adapt for changes in shape and location of the target, 3 000 new spots are added prior to the optimisation with the RPM. The combination of the spot addition and spot-weight optimisation can be repeated. In this study, we evaluated using the spot-position restoration in combination with zero (i.e. only optimise the restored spots), one and two iterations of adding spots and spot-weight optimisation. Figure 6.1 illustrates the workflow of the proposed strategy.

A detailed description of this approach is given in ([Jagt et al., 2018](#)). The two extensions of the existing method, namely the library of prior plans and RPM parameter tuning, are discussed below.

6.2.3.1 Plan-library generation for the prior-plan strategy

Due to potentially large day-to-day variations in the shape and location of the cervix-uterus, prior plans generated solely on either the full or empty bladder CT scan will likely result in insufficient spot coverage for the observed target deformations in the repeat-CT scans. For this reason, an in-house, non-rigid registration was used to derive for each patient a motion




Figure 6.1: The workflow of the automated prior-plan adaptive treatment planning method. A library of prior plans is generated for the patient. Each fraction, the best fitting prior plan is selected from the plan-library. A restoration of the spot-positions with a fixed number of energy layers follows, in which all Bragg peaks are restored to their new WEPL. After this the methods enters a loop in which new spots are added, the intensities are optimised and the non-contributing spots are deleted. The loop can be repeated; we evaluated running the loop zero, one or two times. The output is a Pareto-optimal treatment plan. Adapted from Jagt et al. (2018). Originally published in the [Supplementary Materials](#) (figure S1).

model describing the cervix-uterus shape for every possible bladder volume (Bondar et al., 2012). Using this pretreatment established motion model we created a “Complete ITV” including all observed motion, as well as a “Full” and “Empty” subITV, focusing on parts of the observed motion. The “Empty ITV” ranges from the cervix-uterus corresponding to the empty bladder to the cervix-uterus corresponding to a “half-full bladder” structure and the “Full ITV” ranges from the cervix-uterus corresponding to this half-full bladder to the cervix-uterus corresponding to the full bladder. Figure 6.2 shows an example of the three ITV structures in the sagittal view.

We investigated two library types for the prior-plan strategy:

- 1-plan-library: One prior plan, based on the “Complete ITV”.
- 2-plan-library: Two prior plans, based on the “Full” and “Empty” subITVs.

All prior treatment plans were generated based on a PTV_{Prior} which encompassed the Complete ITV or Full/Empty subITV enlarged with a 5 mm margin and the nodal CTV enlarged with a 4 mm margin. Anatomical differences not accounted for by the PTV_{Prior} are expected to be handled by adding new spots during replanning.

The prior treatment plans were generated using “Erasmus-iCycle”, our in-house developed treatment planning system for fully automated plan generation, combined with the “Astroid” dose engine. The optimisation iteratively adds and removes spots to the target, without time restrictions, see Breedveld et al. (2007), Kooy et al. (2010), Breedveld et al. (2012), Voet et al. (2013a), Van de Water et al. (2013, 2015) and the [Supplementary Materials](#) for more details. It is important to note that these prior plans were not intended as the definitive treatment plan, but serve as a warm-start for daily replanning.




Figure 6.2: An example of the ITV structures, shown in sagittal view. Top-left shows the cervix-uterus corresponding to the empty bladder CT scan in blue, and the cervix corresponding to the full bladder CT scan in orange. Top right shows the Complete ITV in green. Bottom left and right show respectively the Empty ITV and Full ITV in green. Originally published in the [Supplementary Materials](#) (figure S2).

6.2.3.2 Library plan selection strategies

In the case of a 2-plan-library, a selection had to be made between the two prior plans in the library. Traditionally this is done by comparing the bladder volume to a half-full bladder structure. In this work we selected based on the daily anatomy and the restored spot-positions, without assuming the cervix-uterus motion to be linked to bladder filling. The percentages of the total spots of the library plans that ended up in the daily target region after spot-position restoration were compared. If the difference was more than 1%-point, the plan with the highest percentage was selected as prior. If not, both plans fit the daily anatomy equally well. In that case, the plan with the most spots ending up in the daily target region after restoration was selected as prior.

6.2.3.3 RPM parameter tuning

The RPM is used in this study to automatically optimise the spot-weights in a single optimisation. The output is a Pareto-optimal solution, with objective trade-offs in line with the original (i.e. prior) plan. To get these trade-offs, the required RPM parameters were automatically tuned. As the results of the prior-plan adaptation method might depend on the RPM parameters, three-fold cross validation was applied. For each fold, two different patients were used for parameter tuning. The planning strategies using the found parameters were tested on the other four patients of each fold. Evaluation was done on all folds simultaneously: i.e. on 46 plans (two plans for each scan). More information on the RPM, the RPM parameter tuning and the individual folds can be found in the [Supplementary Materials](#) and Ogryczak (1997), Ogryczak and Kozłowski (2011), Van Haveren et al. (2017a,b), Jagt et al. (2018), Van Haveren et al. (2019).

6.2.4 Comparison and evaluation of the methods

In this study, we benchmarked the results of the prior-plan strategy against fully multi-criteria optimised plans. These benchmark plans were generated for each fraction on the PTV_{OAPT} with the same approach as was used for the prior plans (above).

Besides the time-consuming full multi-criteria benchmark optimisation, we investigated a replanning strategy that does not require a prior plan. New spots are placed in the target region, which are then optimised using the RPM. Two approaches for the non-prior-plan strategy were investigated:

- New-Spots-E3: New spots were positioned in a regular grid, using a 5 mm lateral spacing and an energy spacing three times the longitudinal width of the Bragg peak (at 80% of the peak height).
- Sampled-New-Spots-3x: New spots were iteratively added as was done for the benchmark and prior plans. To limit the calculation times, the optimisation was stopped after three iterations.

We compared the prior-plan strategy to the non-prior-plan strategy to see whether the use of a prior plan as a warm-start is beneficial for either plan quality or calculation time. More details on the non-prior-plan strategy approaches can be found in the [Supplementary Materials](#). Other strategies, energy spacings and number of iterations are also reported there.

Table 6.1 gives an overview of the different methods which were included in the evaluation.

For each repeat-CT scan, the dose distributions of all strategies (forward calculation of the prior, prior-plan strategy, non-prior-plan strategy and benchmark) were checked to see whether they fulfilled the planning criteria ($V_{95\%} \geq 95\%$ and $V_{107\%} \leq 2\%$) for the

Table 6.1: Overview of the different treatment plans that are compared.

Method	Explanation
No replanning: Forward dose calculation of prior plan on daily CT	Prior treatment plan selected from a plan-library, recalculated for each aligned repeat-CT scan as if it would have been delivered to that scan. Note that as these prior plans were not intended for treatment, the results are only shown to illustrate that replanning is required.
1-plan-library/2-plan-library: Prior-plan strategy	Prior treatment plan selected from a plan-library, adapted for each repeat-CT scan by an energy layer constrained WEPL correction followed by zero, one or two (-0x, -1x, -2x) iterations of spot addition (adding 3 000 spots per iteration) and RPM optimisation.
Sampled-New-Spots-3x/New-Spots-3x: Non-prior-plan strategy	Treatment plan generated by placing only new spots in the target region and using the RPM to optimise the spot intensities on the PTV_{OAPT} for each repeat-CT scan. Spots were either positioned in a regular grid, or randomly selected from a very fine regular grid using a limited number of iterations.
Benchmark	Treatment plan optimised from scratch using “Erasmus-iCycle” on the PTV_{OAPT} for each repeat-CT scan. Currently the best achievable plan if no time constraints would apply. This plan was included as a benchmark of obtainable plan quality.

PTV_{OAPT} . In addition, all dose distributions were visually checked for hot-spots inside and outside the target volumes.

For the PTV, we report the $V_{95\%}$ and $V_{107\%}$. For rectum, bladder and bowelbag, we report the $V_{30Gy(RBE)}$, D_{mean} and D_{max} and, for the sigmoid, femoral heads and whole body (patient) we report the D_{max} .

All calculations were performed on a dual Intel Xeon E5-2690 server.

6.3 RESULTS


6.3.1 Results for the targets

All prior plans achieved the $V_{95\%}$ and $V_{107\%}$ requirements for the respective PTV_{Prior} volumes. Table 6.2 shows for each treatment strategy the number of plans that met the target demands. It can be seen that forward calculation of the prior plans, i.e. without replanning, always resulted in inadequate target coverage. Replanning using the prior-plan approaches without the addition of new spots (0x) achieved sufficient $V_{95\%}$ values, but too high $V_{107\%}$ values. Adding spots once (1x) yielded acceptable target coverage for more than 80% of the plans. Acceptable coverage was only obtained for all plans with the 2-plan-library-2x approach.

For the non-prior-plan strategy, Table 6.2 shows that while using a fine regular grid (New-Spots-E3) always resulted in acceptable target coverage, iteratively sampling new spots (Sampled-New-Spots-3x) achieved the demands in only 84% of the plans.

Table 6.2: For each treatment strategy, the number of plans that meet the prescribed target demands.

	$V_{95\%} \geq 95\%$ & $V_{107\%} \leq 2\%$	Calculation times (min.) mean (min – max)
1-plan-library No replanning	0/46	–
2-plan-library No replanning	0/46	–
1-plan-library-0x	2/46	1.9 (1.6 – 2.4)
2-plan-library-0x	0/46	2.1 (1.7 – 2.5)
1-plan-library-1x	37/46	4.2 (3.2 – 5.2)
2-plan-library-1x	41/46	4.2 (3.4 – 5.4)
1-plan-library-2x	43/46	6.4 (5.1 – 8.3)
2-plan-library-2x	46/46	6.6 (5.3 – 8.4)
Sampled-New-Spots-3x	39/46	7.1 (5.7 – 8.4)
New-Spots-E3	46/46	40.7 (25.0 – 78.4)
Benchmark	46/46	56.4 (25.3 – 85.1)

**Figure 6.3:** Boxplots depicting the OAR differences between the prior-plan strategies with two iterations of adding spots and benchmark. In blue, the prior-plan strategy is shown using the 1-plan-library approach and in green the 2-plan-library approach. Negative deviations depict scans for which the OAR value is lower in the RPM plan than in benchmark. Statistically significant differences (Wilcoxon signed-rank test, 1% significance level, $p < 0.01$) are indicated by asterisks.

6.3.2 Results for the OARs

In Figure 6.3, the OAR results obtained using the 1-plan-library-2x approach and the 2-plan-library-2x approach are compared to the OAR results of the benchmark plans. The highest prioritised criteria (D_{\max}) deteriorated less than 5 Gy(RBE) compared to the benchmark plans, where some resulted in even lower doses (bowelbag, sigmoid D_{\max}). For the 2-plan-library-2x approach, the largest deviation (+14%-point) was seen for the rectum $V_{30\text{Gy(RBE)}}$, obtaining a value of 63%, where the benchmark plan had a value of 49%.

In Figure 6.4, the OAR results of the best approach of the prior-plan strategy (2-plan-library-2x) and the two approaches of the non-prior-plan strategy are compared to bench-

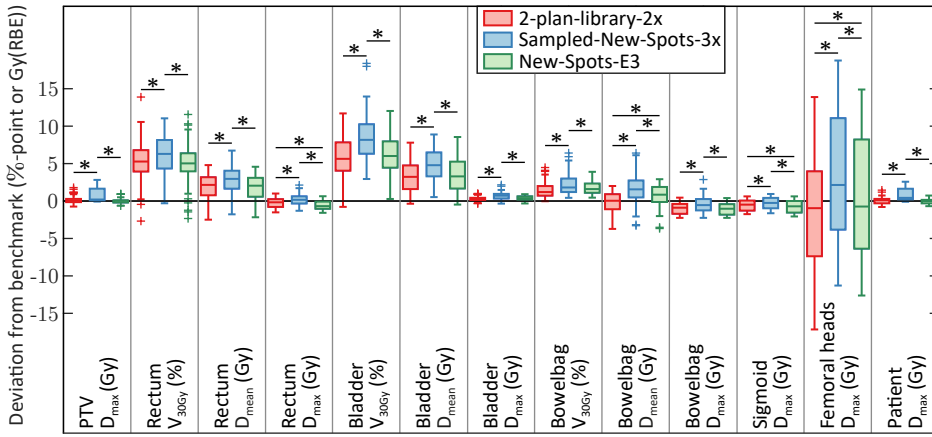


Figure 6.4: Boxplots depicting the OAR deviations from benchmark for the best prior-plan strategy approach and the non-prior-plan strategy approaches. In red the prior-plan strategy is shown using the 2-plan-library-2x approach, in blue the non-prior strategy using the Sampled-New-Spots-3x approach and in green the non-prior strategy using the New-Spots-E3 approach. Negative deviations depict scans for which the OAR value is lower in the RPM plan than in benchmark. Statistically significant differences (Wilcoxon signed-rank test, 1% significance level, $p < 0.01$) are indicated by asterisks.

mark. Similar OAR results were obtained for the 2-plan-library-2x and the New-Spots-E3 approaches, while the Sampled-New-Spots-3x approach showed slightly larger deviations from benchmark.

6.3.3 Calculation times

Generating the library of prior plans took on average 1.5 h per plan, including dose calculation. Table 6.2 shows the total calculation times required for all treatment strategies, excluding the final dose calculation. In the prior-plan strategy, the spot-position restoration step took on average 5.9 s (range 4.4 – 7.4) per restored plan. After restoration, the dose deposition matrix was recalculated in on average 1.2 min (1.0 – 1.4). Without the addition of spots, the RPM spot-weight optimisation took on average 28.0 s (19.7 – 54.8). Adding new spots and calculating their dose deposition matrices was completed in on average 1.4 min per iteration (0.9 – 2.5). With the addition of new spots, the average calculation time of the RPM spot-weight optimisation increased to 1 min (0.6 – 1.7) per iteration.

6.4 DISCUSSION

In this study, we combined a plan-library approach with a previously developed RPM adaptive method in a prior-plan strategy. The combination of selecting a prior plan and adding new pencil-beams could account for density changes along the pencil-beam paths and large inter-fraction shape changes of targets and OARs. Clinically acceptable treatment plans

were obtained for all plans when using the 2-plan-library-2x approach. One iteration of spot addition was already sufficient for more than 80% of the plans.

Plans were considered acceptable if they achieved $V_{95\%} \geq 95\%$ and $V_{107\%} \leq 2\%$. As all discussed strategies were completely automated, these demands were strictly checked, even though slight deviations might be clinically acceptable. The latter can be incorporated by automatically notifying the user when the plan is within a prescribed bandwidth of the demands.

Applying the prior treatment plans without replanning resulted in inadequate target coverage, while in [Van de Schoot et al. \(2016\)](#), most scans obtained acceptable target coverage. The differences can be explained by the fact that our prior plans were not intended for actual dose delivery, but only as a warm-start for daily replanning. For this reason, no robustness was used in the optimisation of the prior plans. In the [Supplementary Materials](#), we show that robustness against uncertainties in stopping power prediction can be added at the cost of a limited increase in optimisation time.

Comparing the 1-plan-library with the 2-plan-library approaches, table 6.2 shows that the 2-plan-library performed slightly better, but the differences in success rate in terms of target coverage were small. Using a 2-plan-library-2x approach always resulted in clinically acceptable target coverage. Similar OAR results were obtained with the 1-plan-library-2x and the 2-plan-library-2x approaches (figure 6.3).

We showed that a non-prior-plan strategy in which no warm-start is used can result in clinically acceptable treatment plans when sufficient spots are included. Figure 6.4 and table 6.2 show that the plans from New-Spots-E3 obtain similar OAR results as the 2-plan-library-2x approach, while requiring over six times longer calculation times. This makes this option infeasible for online replanning. Although the New-Spots-Sampling-3x approach requires similar calculation times as the proposed prior-plan method with two iterations of spot addition, only 39/46 plans fulfilled the target demands (table 6.2). To obtain good results for all plans would require more iterations of spot addition, again indicating that without a warm-start (i.e. a non-prior-plan approach) calculation times increase.

Several approaches of adaptive IMPT have been reported in the literature. An offline approach was proposed by [Kurz et al. \(2016\)](#), in which a new plan is generated to serve as an update for the next fraction. One hour was required for deformable image registration (DIR), optimisation and dose calculation. Adaptation reduced over-dosage in the targets and partially improved OAR sparing. [Bernatowicz et al. \(2018\)](#) compared dose restoration methods using new spots (no prior plan) aiming to restore a given reference dose distribution. Without restoration less than 45% of the repeat-CT scans achieved adequate target coverage; with restoration this improved to 100%. The difference is that our method optimises the dose distribution to the daily anatomy, while their restoration methods intend to only restore a prior dose distribution. [Botas et al. \(2018\)](#) developed online-adaptation

approaches based on cone beam CTs (CBCTs) in which only spots from the prior plan were used. A spot restoration was applied using DIR and, if necessary, this was followed by a weight tuning. Applying only a spot restoration was found to be insufficient; combined with a weight tuning acceptable results were obtained. Calculations were done using GPU-based Monte Carlo.

For conventional radiotherapy, the combination of a restoration followed by a weight optimisation was reported by [Ahunbay et al. \(2010\)](#). Segment aperture morphing combined with segment weight optimisation showed to improve target coverage and OAR sparing. Adaptation was possible in 10 min. Recently adaptive planning methods for the MR-linac have been described by [Winkel et al. \(2019\)](#). Two main categories of the adaptation methods are described as “adapt-to-position” and “adapt-to-shape”. Our replanning method would belong to the latter. Other centres have implemented a stereotactic MR-guided adaptive workflow (SMART) ([Bohoudi et al., 2017](#), [Henke et al., 2018](#), [Tyran et al., 2018](#), [Finazzi et al., 2019](#)). Generating online-adaptive plans when target and OAR aims were not met resulted in adequate target coverage and better OAR sparing.

In this study, we decided to ignore the simultaneously integrated boost that is recommended in the EMBRACE II protocol. Including the boost would not alter the workflow. If more spots are needed to cover the boost this could slightly increase the calculation times. For the prior plan selection approach, a threshold of 1%-point was used for the initial selection criterion. While this setting is admittedly ad hoc, it demonstrated good results. Also, the proposed replanning methods currently add 3 000 new spots to the optimisation per iteration, which is the same number as was used in the optimisation of the benchmark and prior plans. It is possible that using a different number in the replanning could result in acceptable results after fewer iterations. Fewer iterations might also be achieved by using a different beam set-up (i.e. not 0°, 90°, 180°, and 270°), as other beam set-ups could be more robust against the daily anatomical variations. Finally, this proof of principle study was conducted on a small dataset consisting of six patients. Further investigation based on more data is necessary.

In conclusion, large day-to-day variations such as seen in cervical cancer radiotherapy can be accounted for in IMPT by applying a fast and automated prior-plan adaptation method. Selecting a prior plan from a plan-library, adapting its pencil-beams to the new WEPL, adding new spots and optimising the spot-weights resulted in clinically acceptable treatment plans on daily anatomies. The use of a library of prior plans significantly reduced the optimisation times to obtain clinically acceptable treatment plans.

Fast and exact Hessian computation for a class of nonlinear functions used in radiation therapy treatment planning

R. van Haveren, and S. Breedveld

Erasmus MC, University Medical Center Rotterdam, Department of Radiation Oncology, Rotterdam, The Netherlands

Physics in Medicine and Biology, Volume 64, Issue 16, 16NT01, 14 August 2019

DOI: [10.1088/1361-6560/ab1e17](https://doi.org/10.1088/1361-6560/ab1e17)

ABSTRACT

In nonlinear optimisation, using exact Hessian computations (full-Newton) hold superior convergence properties over quasi-Newton methods or gradient-based methods. However, for medium-scale problems, computing the Hessian can be computationally expensive and thus time-consuming. For solvers dedicated to a specific problem type, it can be advantageous to hard-code optimised implementations to keep the computation time to a minimum.

In this paper we derive a computationally efficient canonical form for a class of additively and multiplicatively separable functions. The major computational cost is reduced to a single multiplication of the data matrix with itself, allowing simple parallelisation on modern-day multi-core processors. We present the approach in the practical application of radiation therapy treatment planning, where this form appears for many common functions. In this case, the data matrices are the dose-influence matrices. The method is compared against automatic differentiation.

7.1 INTRODUCTION

Nonlinear optimisation involves optimising a generally nonlinear *objective function* $F : X \rightarrow \mathbb{R}$, where $X \subseteq \mathbb{R}^n$ and n is an integer, possibly subject to a number of generally nonlinear constraints. Without loss of generality, the objective function needs to be minimised so that the corresponding optimisation problem is

$$\begin{aligned} & \underset{\mathbf{x} \in X}{\text{minimise}} && F(\mathbf{x}) \\ & \text{subject to} && G_i(\mathbf{x}) \leq 0 \quad i = 1, \dots, p, \end{aligned} \tag{7.1}$$

where the objective function and *constraint functions* $G_i : X \rightarrow \mathbb{R}$, for $i \in \{1, \dots, p\}$ with p an integer, are twice continuously differentiable.

Several algorithms attempting to solve a general constrained optimisation problem (7.1) have been developed, e.g. trust region methods (Moré and Sorensen, 1983, Conn et al., 2000, Nocedal and Wright, 2006) and interior-point methods (Mehrotra, 1992, Wright, 1997, Forsgren et al., 2002, Boyd and Vandenberghe, 2004, Nocedal and Wright, 2006, Aleman et al., 2010, Breedveld et al., 2017) such as HOPDM (Gondzio, 1996), LOQO (Vanderbei and Shanno, 1999) or IPOPT (Wächter and Biegler, 2006). These methods are based on solving a sequence of approximate minimisation problems and converge to the minimum of the objective function subject to the constraints. In each iteration, there is a procedure that determines a direction $\Delta \mathbf{x}$ used to update the current iterate. To determine that direction using full-Newton methods (Forsgren et al., 2002, Nocedal and Wright, 2006), evaluations at the current iterate of the gradients and Hessians of the objective and constraint functions are required.

One alternative is using quasi-Newton methods (Nocedal and Wright, 2006) which iteratively approximate the Hessian by using the information of the gradients of the previous iterations, e.g. with the BFGS formula (Nocedal and Wright, 2006). While the computation of the approximate Hessian may be faster, quasi-Newton methods require more iterations to converge which may result in longer optimisation times.

Several methods exist to obtain exact first and second order derivatives. One method, used in modern general solvers, is automatic differentiation (Bücker et al., 2006). However, these techniques are usually implemented using operator overloading which is only supported by a limited number of programming languages, e.g. C++ (Griewank et al., 1996, Walther and Griewank, 2012). Another method is hard-coding the derivatives. For *application-specific solvers*, only a limited number of functions are used so that only a limited number of first and second derivatives need to be hard-coded.

In the specific application of radiation therapy, the problem definition (i.e. treatment protocol) and corresponding functions remain identical across multiple data (i.e. patients). Also in the application of radiation therapy treatment planning, frequently used nonlinear functions share the same canonical form (section 7.2) including the generalised equi-

valent uniform dose (gEUD) (Niemierko, 1999), logarithmic tumour control probability (LTCP) (Alber and Reemtsen, 2007), a continuous approximation of the dose-volume criterion (DVH) (Alber and Reemtsen, 2007), and near-maximum approximations such as the log-sum-exp (LSE) (El Ghaoui, 2018). A disadvantage of hard-coding the derivatives is that this is error prone, although unit tests can be implemented to check the similarity between the hard-coded and numerical derivatives, e.g. using finite differences. Using the finite differences themselves to compute gradients or Hessians is undesirable, due to long computation times (especially for Hessian computations) and their inexactness.

Although computation times for automatic differentiation are promising (Gower and Mello, 2012, Walther and Griewank, 2012), we focus on the alternative for hard-coding the derivatives since it is easier to implement if a limited number of nonlinear functions are used which also share the same canonical form. For this canonical form, we propose an efficient and exact method for computing the Hessian by deriving the first and second order derivatives analytically. We also demonstrate that the method proposed is computationally more efficient and scalable than the modern automatic differentiation approach ADOL-C (Walther and Griewank, 2012).

It should be noted that full-Newton methods are limited to continuous and differentiable medium-scaled optimisation problems. In radiation therapy, full-Newton methods are efficient in for example brachytherapy (Guthier et al., 2017), intensity-modulated proton therapy (IMPT) (Van de Water et al., 2013, Jagt et al., 2018), pre-optimisation in clinical treatment planning (Voet et al., 2013a, 2014), fluence map segmentation (Balvert and Craft, 2017), direct aperture optimisation (Men et al., 2007), or any other (intermediate) problem that can be posed as medium-scale size.

Problem formulations which are large-scaled, such as integrated fluence map and beam orientation optimisation, high-resolution volumetric modulated arc therapy (VMAT), or IMPT using regular grids are in general not solvable using full-Newton methods, due to the problem size and solution times. Quasi-Newton methods such as limited memory BFGS or gradient methods (Nocedal and Wright, 2006) are feasible approaches. However, including general nonlinear constraints is not trivial. Nevertheless, gradient information is required, for which this note also provides a convenient canonical form.

The radiation therapy treatment planning problem can be modelled in various ways, depending on the application and approaches, to keep the problem computationally tractable while ensuring clinical deliverability. This may result in solving multiple medium-scale problem, a series of unconstrained medium-scale problems, or a combination of these. The intention of this note is not to discuss the plethora of possible modelling approaches, but to provide an efficient implementation for solving small to medium-scaled full-Newton problems.

We start by presenting the main result in section 7.2, of which the formal proofs can

be found in appendix 7.A. Section 7.3 demonstrates the application to radiation therapy, followed by the discussion and conclusions in sections 7.4 and 7.5. Appendix 7.B provides ready-to-use parameters for some commonly used functions in radiation therapy.

7.2 MAIN RESULT

7.2.1 Class of nonlinear functions

Constructing a function within the class of nonlinear functions used in radiation therapy treatment planning follows the next steps:

- (1) First, the linear relation between the fluence intensities and dose per voxel is established. In vector form, the formula is,

$$d(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b}, \quad (7.2)$$

which is an *affine* operation. Here, the vector $\mathbf{x} = (x_1, \dots, x_n)$ contains the fluence intensities (decision vector in optimisation problem (7.1)), and $\mathbf{d} = (d_1, \dots, d_m)$ is the dose vector. The $m \times n$ matrix \mathbf{A} is the dose-influence (data) matrix.

- (2) Secondly, a *separable* function of the dose vector, $u(\mathbf{d})$, is specified. We consider *additively separable* functions,

$$u(\mathbf{d}) = \sum_{i=1}^m u_i(d_i), \quad (7.3)$$

and *multiplicatively separable* functions,

$$u(\mathbf{d}) = \prod_{i=1}^m u_i(d_i). \quad (7.4)$$

The function $u_i(d_i)$ specifies the operation applied to the dose to voxel i .

- (3) Finally, another function v is composed with the separable function from step 2. The final function f is then calculated as

$$f(\mathbf{x}) = (v \circ u \circ d)(\mathbf{x}) = v(u(d(\mathbf{x}))). \quad (7.5)$$

An assumption is that the functions u_i and v are twice continuously differentiable for $i \in \{1, \dots, m\}$.

7.2.2 Canonical form for the derivatives

We are interested in the point evaluation of both the gradient column vector (7.6) and the Hessian matrix (7.7) of the nonlinear function (7.5),

$$\nabla_{\mathbf{x}} f := \left(\frac{\partial f}{\partial x_i} \right)_{i=1, \dots, n}^T, \quad (7.6)$$

$$\nabla_{\mathbf{x}}^2 f := \left(\frac{\partial^2 f}{\partial x_i \partial x_j} \right)_{i,j=1, \dots, n}. \quad (7.7)$$

In appendix 7.A, we prove that (7.6)-(7.7) can be written in the following canonical and computationally efficient form (7.8)-(7.9) for both the additive separable (7.3) and multiplicative separable (7.4) form,

$$\nabla_{\mathbf{x}} f = c_1 \mathbf{A}^T \mathbf{e}_1, \quad (7.8)$$

$$\nabla_{\mathbf{x}}^2 f = c_1 \mathbf{A}^T \text{diag}(\mathbf{e}_2) \mathbf{A} + c_2 (\mathbf{A}^T \mathbf{e}_1) (\mathbf{A}^T \mathbf{e}_1)^T, \quad (7.9)$$

for some scalars c_1, c_2 and vectors $\mathbf{e}_1, \mathbf{e}_2$ (see appendix 7.A). Here, we use the notation $\text{diag}(\mathbf{e}_2)$ for the sparse diagonal matrix having the entries of the vector \mathbf{e}_2 on the diagonal and zeros for all off-diagonal entries. The Hessian matrix can thus be written as a linear combination of a matrix product (main computational cost) and an outer product (negligible computational cost).

7.2.3 Computational efficiency

Equations (7.8)-(7.9) show analytical and canonical formulae for calculating the gradient vector and Hessian matrix for the class of nonlinear functions described in section 7.2.1.

A few additional steps can be taken to further reduce the computational load. The first step is to calculate an intermediate vector \mathbf{w} as follows

$$\mathbf{w} = \mathbf{A}^T \mathbf{e}_1. \quad (7.10)$$

Then the gradient of f with respect to \mathbf{x} can be calculated,

$$\nabla_{\mathbf{x}} f = c_1 \mathbf{w}.$$

The constant c_1 and vector \mathbf{w} can then be reused for the point evaluation of the Hessian matrix of f with respect to \mathbf{x} ,

$$\nabla_{\mathbf{x}}^2 f = c_1 \mathbf{A}^T \text{diag}(\mathbf{e}_2) \mathbf{A} + c_2 \mathbf{w} \mathbf{w}^T.$$

The computational effort of the Hessian matrix can be reduced further if all elements of \mathbf{e}_2 are nonnegative. The latter holds for additively separable functions if all $u_i(d_i)$ are convex, and for multiplicatively separable functions if all $u_i(d_i)$ are log-convex. Compute

$$\mathbf{Q} = \text{diag} \left(\mathbf{e}_2^{1/2} \right) \mathbf{A}, \quad (7.11)$$

by scaling the columns of \mathbf{A} with \mathbf{e}_2 . This operation is continuous in memory, and is practically negligible in time. The Hessian can then be computed by ¹

$$\nabla_{\mathbf{x}}^2 f = c_1 \mathbf{Q}^T \mathbf{Q} + c_2 \mathbf{w} \mathbf{w}^T, \quad (7.12)$$

i.e. the Hessian matrix is simply a linear combination of a symmetric matrix-matrix product and an outer product which can be efficiently computed using the BLAS (Basic Linear Algebra Subprograms) routines `_SYMM` and `_GER`. If the vector \mathbf{e}_2 does contain negative elements, then the diagonal scaling (7.11) is performed with \mathbf{e}_2 itself, and the operation (7.12) is then performed with $\mathbf{A}^T \mathbf{Q}$ using the `_GEMM` routine. For sparse \mathbf{A} , the techniques described in Gustavson (1978), Pissanetzky (1984), Breedveld et al. (2006) can be applied.

7.2.4 Automatic differentiation

An alternative to deriving explicit forms for the derivatives is to apply *automatic differentiation*. We compare the performance of our approach from section 7.2.3 against ADOL-C (Automatic Differentiation by OverLoading in C++), version 2.6.3. (Walther and Griewank, 2012). We found that automatic differentiation does not perform optimal when directly applied to \mathbf{x} , due to the large size of the matrix \mathbf{A} . Instead, we apply automatic differentiation on the dose \mathbf{d} and compute the derivatives to \mathbf{x} by

$$\nabla_{\mathbf{x}} f = \mathbf{A}^T \nabla_{\mathbf{d}} f, \quad (7.13)$$

$$\nabla_{\mathbf{x}}^2 f = \mathbf{A}^T \nabla_{\mathbf{d}}^2 f \mathbf{A}, \quad (7.14)$$

where the gradient vector $\nabla_{\mathbf{d}} f$ and Hessian matrix $\nabla_{\mathbf{d}}^2 f$ are obtained by automatic differentiation. For the Hessian, this introduces an extra matrix-matrix operation in the case that $\nabla_{\mathbf{d}}^2 f$ is not a diagonal matrix.

7.3 EXAMPLES IN RADIATION THERAPY

For patients suffering from cancer, radiation therapy is often part of the treatment. The aim of this therapy is to sufficiently irradiate the tumour while reducing the doses to the surrounding healthy tissues as much as possible to minimise the risk of treatment-induced side-effects. These goals are usually conflicting, making treatment planning a *multi-objective* optimisation problem (Miettinen, 1999). This generalised the notion of an optimal solution, e.g. *Pareto optimality*. Several techniques have been developed to deal with such problems, e.g. *interactive* methods (Granat and Makowski, 2000, Long et al., 2012) or *automated* methods (Wilkens et al., 2007, Breedveld et al., 2009, Van Haveren et al., 2017a).

¹In the original publication (Van Haveren and Breedveld, 2019), the constant c_1 was missing in equation (7.12). Details can be found in the corrigendum Van Haveren and Breedveld (2020).

Regardless of the approach used for treatment planning, (possibly several) regular non-linear medium-scale optimisation problems (7.1) need to be solved. Commonly, there are around 1 000–15 000 decision variables, representing small nonnegative photon or proton *intensities*, and the number of data points, i.e. number of rows of matrix \mathbf{A} , for computing the functions are in the order $O(10^5)$ (see section 7.3.2). More importantly, the problems are dense and thus require an efficient implementation of the solver.

In section 7.3.1, we elaborate on the actual implementation of the *generalised equivalent uniform dose* (gEUD) in *Erasmus-iCycle* (Breedveld et al., 2012) which is our in-house developed algorithm for multi-objective optimisation of beam angles and intensity profiles. Erasmus-iCycle is in clinical use for treatment of patients (Voet et al., 2013a, Sharfo et al., 2015). The nonlinear optimisation problems (7.1) in Erasmus-iCycle are solved using a custom primal-dual interior-point method (Breedveld et al., 2017). For pointers on the overall algorithm, such as number of iterations and time for other procedures during an iteration, we refer the reader to that paper. Besides the gEUD, other frequently used functions in radiation therapy are also separable, mainly in the additively separable form (7.3). Examples and parameters for commonly used functions are presented in appendix 7.B. In section 7.3.2, typical dimensions of radiation therapy treatment planning problems are given, and corresponding computation times of the gradients and Hessians and compared against those of ADOL-C.

7.3.1 Generalised equivalent uniform dose example

The generalised equivalent uniform dose (gEUD) is given by

$$\text{gEUD}(\mathbf{d}) = \left(\frac{1}{m} \sum_{i=1}^m d_i^a \right)^{1/a}, \quad (7.15)$$

with parameter $0 < |a| < \infty$. The gEUD (7.15) is bounded by the minimum ($a \rightarrow -\infty$) and maximum dose ($a \rightarrow \infty$), and equals the mean dose in the special case that $a = 1$. The gEUD is smooth and convex for $a \geq 1$, and thus fits within a convex optimisation framework.

The inner part of the gEUD (7.15) is additively separable: its decomposition is as follows

$$u_i(d_i) = d_i^a / m, \quad \text{for } i = 1, \dots, m \quad (7.16)$$

$$v(u) = u^{1/a}. \quad (7.17)$$

For $u(\mathbf{d}) = \sum_{i=1}^m u_i(d_i)$, equations (7.23) yield

$$\begin{aligned} c_1 &= \frac{1}{a} u(\mathbf{d})^{1/a-1}, \\ c_2 &= \frac{1}{a} \left(\frac{1}{a} - 1 \right) u(\mathbf{d})^{1/a-2}, \\ \mathbf{e}_1 &= \left(\frac{a}{m} d_i^{a-1} \right)_{i=1, \dots, m}, \\ \mathbf{e}_2 &= \left(\frac{a[a-1]}{m} d_i^{a-2} \right)_{i=1, \dots, m}. \end{aligned}$$

Following theorem 2 (appendix 7.A), the gradient and Hessian matrix of the gEUD can be computed by

$$\begin{aligned} \nabla_{\mathbf{x}} f &= c_1 \mathbf{A}^T \mathbf{e}_1, \\ \nabla_{\mathbf{x}}^2 f &= c_1 \mathbf{A}^T \text{diag}(\mathbf{e}_2) \mathbf{A} + c_2 (\mathbf{A}^T \mathbf{e}_1) (\mathbf{A}^T \mathbf{e}_1)^T. \end{aligned}$$

7.3.2 Computational effort

To demonstrate the efficiency of the canonical form in radiation therapy treatment planning, we measured the computation times for different steps in the process, and compared these times against those for automatic differentiation. We have performed the operations for the gEUD function (7.15).

The cases are derived from the TROTS open dataset (Breedveld and Heijmen, 2017). Overall properties are summarised in table 7.1 where the *density* is defined as the fraction of nonzero elements. For storage reasons, matrices were designated as sparse if this resulted in a more efficient storage than storing the matrix in dense format. The computations were run singlethreadedly on an Intel Xeon E5-2690, running at 2.9 GHz. We used the Intel Math Kernel Library version 10.3 for the dense operations. For sparse operations, we used custom code based on Gustavson (1978), Pissanetzky (1984), Breedveld et al. (2006). All computations were performed in double precision arithmetic. The resulting Hessian matrices are *a priori* assumed to be dense.

Table 7.2 shows the time required to compute the required constants c_1 , c_2 and vectors \mathbf{e}_1 , \mathbf{e}_2 , and subsequently the gradient and Hessian for the canonical form presented in this paper. The operations are performed as mentioned in section 7.2.3.

As a reference, we implemented the gEUD function in C++ using ADOL-C. We report the computation times in table 7.3 for the gradient vector $\nabla_{\mathbf{d}} f$ and Hessian $\nabla_{\mathbf{d}}^2 f$, and the subsequent times for computing $\nabla_{\mathbf{x}} f$ and $\nabla_{\mathbf{x}}^2 f$.

7.4 DISCUSSION

The main result of this note is twofold: (1) simple derivation of first and second order derivatives, and (2) an efficient canonical form for their computation. To keep the scope of the

Table 7.1: Typical data characteristics for a variety of optimisation problems in radiation therapy treatment planning. Each problem can have both sparse and dense matrices, for which the matrix sizes are given. The density level only applies to the sparse matrices. The number of rows is related to the number of voxels, the number of columns represents the number of intensities.

	Dense rows	Sparse rows	Density	Columns	Storage (MiB)
Protons 01	0	332 599	4.3%	1 080	239
Liver 01	31 409	75 180	7.4%	1 118	363
Prostate CK 01	27 087	42 636	7.2%	2 260	573
Prostate VMAT 101	65 524	11 787	15.0%	1 972	1 039
Head-and-Neck 01	0	82 403	13.5%	9 977	1 694
Head-and-Neck-Alt 01	68 463	13 940	18.5%	9 977	5 605

Table 7.2: Computation times in seconds to compute the gradient and Hessian of the gEUD function for several cases. The preparation consists of computing the constants c_1 , c_2 and vectors e_1 , e_2 .

	Preparation	Gradient	Hessian	Total
Protons 01	0.07	0.02	2.20	2.29
Liver 01	0.02	0.02	2.87	2.92
Prostate CK 01	0.02	0.04	8.27	8.33
Prostate VMAT 101	0.02	0.07	12.45	12.55
Head-and-Neck 01	0.02	0.14	146.67	146.84
Head-and-Neck-Alt 01	0.02	0.40	312.48	312.90

Table 7.3: Computation times in seconds for computing the gradient and Hessian for the gEUD for several cases. ADOL-C is the time required to compute the gradient and Hessian in the dose domain by automatic differentiation. The Matrix column represents the time required for the matrix multiplications.

	Gradient		Hessian		Total
	ADOL-C	Matrix	ADOL-C	Matrix	
Protons 01	0.31	0.02	4468.29	386.25	4854.88
Liver 01	0.14	0.02	537.42	48.55	586.13
Prostate CK 01	0.08	0.04	327.08	87.10	414.31
Prostate VMAT 101	0.10	0.07	396.94	134.59	531.71
Head-and-Neck 01	0.14	0.14	235.68	806.88	1042.84
Head-and-Neck-Alt 01	0.15	0.40	236.29	906.95	1143.80

note on the analytical form, we did not fully exploit the matrix-vector and matrix-matrix operations and used standard algorithms. There are many approaches to further improve the computational performance. A straightforward one is to use multithreading: especially for large problems, a speed-up factor of 6 can be achieved when using 8 threads. One advanced approach is presented in [Breedveld et al. \(2017\)](#), where each matrix undergoes a row and column permutation, and is stored in separate tiles. The matrix-matrix operation is successively performed using multiple precision arithmetic, i.e. using single precision floating point operations if possible. This resulted in a performance improvement of a factor 2.2 for this operation.

While automatic differentiation can be applied to more general functions, its disadvantage is that the computation is slower than using the canonical form presented in this note. For computation of the gradient, the canonical form is 2–15 times faster than with automatic differentiation, and is therefore also an advantage for quasi-Newton methods. For computation of the Hessian, automatic differentiation also requires the multiplication of three matrices instead of two as for the canonical form. Another problem for automatic differentiation is the affine part of the problem, i.e. the pencil-beam matrix \mathbf{A} . ADOL-C seems incapable of using this information directly in an efficient way, so a separation between the “dose domain” and “fluence domain” is necessary. As a result, problems where the number of rows in \mathbf{A} is much larger than the number of columns (much more voxels than intensities, e.g. the Protons case) result in excessive computation times. Finally, parallelisation of automatic differentiation is not straightforward (Bischof et al., 2008).

While the computational analysis in this note is based on the gEUD cost-function, the computational results hold for all common functions (table 7.4). As indicated by table 7.2, the main computational factor is the matrix algebra.

7.5 CONCLUSIONS

We derived a canonical and computationally efficient form for analytically computing the gradient and Hessian regarding functions of a general class of nonlinear functions. In particular, we focussed on compositions where part of the function was either additively or multiplicatively separable, and where affine maps of the decision variables were taken using a generally dense data matrix.

The application is mainly interesting for application-specific solvers, such as solvers for optimisation problems in radiation therapy treatment planning. For medium-scale problems, the results indicate that the canonical form efficiently computes the analytical Hessian for a general class of nonlinear functions. Especially large problems benefit from modern computer architecture using multithreading and standardised libraries.

APPENDIX 7.A PROOFS AND DERIVATION

7.A.1 Additively separable form

To obtain the gradient (7.6) and Hessian (7.7), we exploit the additive separability property of the function (7.3). For convenience, we first define the function g as

$$g(\mathbf{d}) = (v \circ u)(\mathbf{d}) = v(u(\mathbf{d})). \quad (7.18)$$

Lemma 1 implies a canonical form for gradient and the Hessian for the function (7.18). Then, theorem 2 shows the desired canonical forms (7.8)-(7.9).

Lemma 1. For the function g defined by (7.18) where $u(\mathbf{d})$ is additively separable, let v, u_1, \dots, u_m be twice continuously differentiable. Then, the gradient and Hessian of g with respect to \mathbf{d} can be written in the canonical form

$$\nabla_{\mathbf{d}} g = c_1 \mathbf{e}_1, \quad (7.19)$$

$$\nabla_{\mathbf{d}}^2 g = c_1 \text{diag}(\mathbf{e}_2) + c_2 (\mathbf{e}_1)(\mathbf{e}_1)^T. \quad (7.20)$$

Here, $c_1, c_2 \in \mathbb{R}$ are scalars, and $\mathbf{e}_1, \mathbf{e}_2 \in \mathbb{R}^m$ are vectors.

Proof. Since u is additively separable, we have $u(\mathbf{d}) = \sum_{i=1}^m u_i(d_i)$. So, the first and second order partial derivatives of g are

$$\frac{\partial g}{\partial d_i} = \frac{dv}{du} \frac{du_i}{dd_i}, \quad (7.21)$$

$$\frac{\partial^2 g}{\partial d_i \partial d_j} = \begin{cases} \frac{d^2 v}{du^2} \frac{du_i}{dd_i} \frac{du_j}{dd_j}, & i \neq j \\ \frac{dv}{du} \frac{d^2 u_i}{dd_i^2} + \frac{d^2 v}{du^2} \left(\frac{du_i}{dd_i} \right)^2, & i = j \end{cases} \quad (7.22)$$

where $i, j \in \{1, \dots, m\}$. The first order derivatives (7.21) contain the term $\frac{du_i}{dd_i}$ which reappears as an outer product in the second order derivatives (7.22). For the diagonal entries of the Hessian matrix (7.20) (i.e. the cases $i = j$ in (7.22)), the terms $\frac{dv}{du} \frac{d^2 u_i}{dd_i^2}$ need to be added. Thus, we need

$$\begin{aligned} c_1 &= \frac{dv}{du}, \\ c_2 &= \frac{d^2 v}{du^2}, \\ \mathbf{e}_1 &= \left(\frac{du_i}{dd_i} \right)_{i=1, \dots, m}, \\ \mathbf{e}_2 &= \left(\frac{d^2 u_i}{dd_i^2} \right)_{i=1, \dots, m}. \end{aligned} \quad (7.23)$$

□

The canonical forms (7.8)-(7.9) can now be proven in case the function u is additively separable (7.3).

Theorem 2. For $f(\mathbf{x}) = (g \circ d)(\mathbf{x})$ where $d(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b}$, and g defined by (7.18) where $u(\mathbf{d})$ is additively separable, let v, u_1, \dots, u_m be twice continuously differentiable. Then, the gradient and Hessian of f with respect to \mathbf{x} can be written in the canonical form

$$\nabla_{\mathbf{x}} f = c_1 \mathbf{A}^T \mathbf{e}_1, \quad (7.24)$$

$$\nabla_{\mathbf{x}}^2 f = c_1 \mathbf{A}^T \text{diag}(\mathbf{e}_2) \mathbf{A} + c_2 (\mathbf{A}^T \mathbf{e}_1) (\mathbf{A}^T \mathbf{e}_1)^T. \quad (7.25)$$

Here, $c_1, c_2 \in \mathbb{R}$ are scalars, and $\mathbf{e}_1, \mathbf{e}_2 \in \mathbb{R}^m$ are vectors.

Proof. Since $d(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b}$ is affine, standard calculus yields

$$\begin{aligned}\nabla_{\mathbf{x}}f &= \mathbf{A}^T \nabla_{\mathbf{d}}g, \\ \nabla_{\mathbf{x}}^2f &= \mathbf{A}^T \nabla_{\mathbf{d}}^2g \mathbf{A}.\end{aligned}\tag{7.26}$$

Substituting the result of lemma 1 into the equations (7.26) leads to the desired form

$$\begin{aligned}\nabla_{\mathbf{x}}f &= c_1 \mathbf{A}^T \mathbf{e}_1, \\ \nabla_{\mathbf{x}}^2f &= c_1 \mathbf{A}^T \text{diag}(\mathbf{e}_2) \mathbf{A} + c_2 (\mathbf{A}^T \mathbf{e}_1) (\mathbf{A}^T \mathbf{e}_1)^T,\end{aligned}$$

so the scalars c_1, c_2 , and vectors $\mathbf{e}_1, \mathbf{e}_2$ (7.23) in lemma 1 are exactly the unknowns needed. \square

If, in addition to the conditions stated in theorem 2, the point evaluation of $\frac{dv}{du}$ is nonzero (i.e. $c_1 \neq 0$), the Hessian matrix (7.9) can be written more elegantly as

$$\nabla_{\mathbf{x}}^2f = c_1 \mathbf{A}^T \text{diag}(\mathbf{e}_2) \mathbf{A} + c (\nabla_{\mathbf{x}}f)(\nabla_{\mathbf{x}}f)^T,$$

where $c := c_2/c_1^2$.

7.A.2 Multiplicatively separable form

A similar result can be obtained in case that the function $u(\mathbf{d})$ is multiplicatively separable (7.4). However, there is a degenerate case if the point evaluation of the product equals zero. Lemma 3 and theorem 4 show the canonical forms (7.8)-(7.9) for the nondegenerate case.

Lemma 3. *For the function g defined by (7.18) where $u(\mathbf{d})$ is multiplicatively separable, let v, u_1, \dots, u_m be twice continuously differentiable and $u(\mathbf{d}) \neq 0$. Then, the gradient and Hessian of g with respect to \mathbf{d} can be written in the canonical form*

$$\nabla_{\mathbf{d}}g = c_1 \mathbf{e}_1,\tag{7.27}$$

$$\nabla_{\mathbf{d}}^2g = c_1 \text{diag}(\mathbf{e}_2) + c_2 (\mathbf{e}_1)(\mathbf{e}_1)^T.\tag{7.28}$$

Here, $c_1, c_2 \in \mathbb{R}$ are scalars, and $\mathbf{e}_1, \mathbf{e}_2 \in \mathbb{R}^m$ are vectors.

Proof. Since u is multiplicatively separable, we have $u(\mathbf{d}) = \prod_{i=1}^m u_i(d_i)$. So, the first and second order partial derivatives of g are

$$\frac{\partial g}{\partial d_i} = u(\mathbf{d}) \frac{dv}{du} \left(\frac{1}{u_i(d_i)} \frac{du_i}{dd_i} \right),\tag{7.29}$$

$$\frac{\partial^2 g}{\partial d_i \partial d_j} = \begin{cases} \left[u(\mathbf{d}) \frac{dv}{du} + u(\mathbf{d})^2 \frac{d^2 v}{du^2} \right] \left(\frac{1}{u_i(d_i)} \frac{du_i}{dd_i} \right) \left(\frac{1}{u_j(d_j)} \frac{du_j}{dd_j} \right), & i \neq j \\ u(\mathbf{d}) \frac{dv}{du} \left(\frac{1}{u_i(d_i)} \frac{d^2 u_i}{dd_i^2} \right) + u(\mathbf{d})^2 \frac{d^2 v}{du^2} \left(\frac{1}{u_i(d_i)} \frac{du_i}{dd_i} \right)^2, & i = j \end{cases}\tag{7.30}$$

where $i, j \in \{1, \dots, m\}$. The canonical form of the gradient and Hessian (7.27)-(7.28) can be obtained by defining

$$\begin{aligned} c_1 &= u(\mathbf{d}) \frac{dv}{du}, \\ c_2 &= u(\mathbf{d}) \frac{dv}{du} + u(\mathbf{d})^2 \frac{d^2v}{du^2}, \\ \mathbf{e}_1 &= \left(\frac{1}{u_i(d_i)} \frac{du_i}{dd_i} \right)_{i=1, \dots, m}, \\ \mathbf{e}_2 &= \left(\frac{1}{u_i(d_i)} \frac{d^2u_i}{dd_i^2} - \left[\frac{1}{u_i(d_i)} \frac{du_i}{dd_i} \right]^2 \right)_{i=1, \dots, m}, \end{aligned} \quad (7.31)$$

which are well defined since $u(\mathbf{d}) \neq 0$. \square

In the degenerate case that the point evaluation leads to $u(\mathbf{d}) = 0$, i.e. at least one of the $u_i(d_i)$ equals zero, the vectors $\mathbf{e}_1, \mathbf{e}_2$ are not well defined. For those u_i , we can set $u_i(d_i) = \epsilon > 0$, implying well defined scalars and vectors (7.31). The gradient and Hessian (7.27)-(7.28) then become the limits

$$\nabla_{\mathbf{d}} g = \lim_{\epsilon \downarrow 0} [c_1 \mathbf{e}_1], \quad (7.32)$$

$$\nabla_{\mathbf{d}}^2 g = \lim_{\epsilon \downarrow 0} [c_1 \text{diag}(\mathbf{e}_2) + c_2 (\mathbf{e}_1)(\mathbf{e}_1)^T]. \quad (7.33)$$

For the nondegenerate case, the canonical forms (7.8)-(7.9) can now be proven in case the function u is multiplicatively separable (7.4).

Theorem 4. For $f(\mathbf{x}) = (g \circ d)(\mathbf{x})$ where $d(\mathbf{x}) = \mathbf{A}\mathbf{x} + \mathbf{b}$, and g defined by (7.18) where $u(\mathbf{d})$ is multiplicatively separable, let v, u_1, \dots, u_m be twice continuously differentiable and $u(\mathbf{d}) \neq 0$. Then, the gradient and Hessian of f with respect to \mathbf{x} can be written in the canonical form

$$\begin{aligned} \nabla_{\mathbf{x}} f &= c_1 \mathbf{A}^T \mathbf{e}_1, \\ \nabla_{\mathbf{x}^2} f &= c_1 \mathbf{A}^T \text{diag}(\mathbf{e}_2) \mathbf{A} + c_2 (\mathbf{A}^T \mathbf{e}_1) (\mathbf{A}^T \mathbf{e}_1)^T. \end{aligned}$$

Here, $c_1, c_2 \in \mathbb{R}$ are scalars, and $\mathbf{e}_1, \mathbf{e}_2 \in \mathbb{R}^m$ are vectors.

Proof. The proof is similar to the proof of theorem 2, only now it follows from lemma 3. The unknowns are thus given by equations (7.31). \square

For the degenerate case $u(\mathbf{d}) = 0$, the gradient and Hessian can be obtained by formulae (7.26) using the limits (7.32)-(7.33).

APPENDIX 7.B PARAMETERS FOR COMMONLY USED FUNCTIONS

In table 7.4, a concise overview of the decompositions (either additive or multiplicative, see section 7.2.1) and derivative parameters (for computing the gradient and Hessian, see section 7.2.2) of commonly used functions in radiation therapy treatment planning is provided.

Table 7.4: Decompositions and derivative parameters for commonly used functions in radiation therapy treatment planning. Parameter d_c represents the critical dose level for either an organ-at-risk, or the tumour (prescribed dose).

Function	Decomposition	Derivative parameters
$\text{gEUD}(\mathbf{d}; a)^*$ $= \left(\frac{1}{m} \sum_{i=1}^m d_i^a \right)^{1/a}$	$u_i(d_i) = d_i^a / m$ $u(\mathbf{d}) = \sum_{i=1}^m u_i(d_i)$ $v(u) = u^{1/a}$	$c_1 = \frac{1}{a} u(\mathbf{d})^{1/a-1}$ $c_2 = \frac{1}{a} \left(\frac{1}{a} - 1 \right) u(\mathbf{d})^{1/a-2}$ $\mathbf{e}_1 = \left(\frac{a}{m} d_i^{a-1} \right)_{i=1, \dots, m}$ $\mathbf{e}_2 = \left(\frac{a}{m} [a-1] d_i^{a-2} \right)_{i=1, \dots, m}$
$\text{LSE}(\mathbf{d})$ $= \ln \left(\sum_{i=1}^m e^{d_i} \right)$	$u_i(d_i) = e^{d_i}$ $u(\mathbf{d}) = \sum_{i=1}^m u_i(d_i)$ $v(u) = \ln(u)$	$c_1 = u(\mathbf{d})^{-1}$ $c_2 = -u(\mathbf{d})^{-2}$ $\mathbf{e}_1 = (e^{d_i})_{i=1, \dots, m}$ $\mathbf{e}_2 = (e^{d_i})_{i=1, \dots, m}$
$\text{QO}(\mathbf{d}; d_c)$ $= \frac{1}{m} \sum_{i=1}^m (d_i - d_c)_+^2$	$u_i(d_i) = \frac{1}{m} (d_i - d_c)_+^2$ $u(\mathbf{d}) = \sum_{i=1}^m u_i(d_i)$ $v(u) = u$	$c_1 = 1$ $c_2 = 0$ $\mathbf{e}_1 = \left(\frac{2}{m} (d_i - d_c)_+ \right)_{i=1, \dots, m}$ $\mathbf{e}_2 = \left(\frac{2}{m} \mathbf{1}_{\{d_i > d_c\}} \right)_{i=1, \dots, m}$
$\text{QU}(\mathbf{d}; d_c)$ $= \frac{1}{m} \sum_{i=1}^m (d_c - d_i)_+^2$	$u_i(d_i) = \frac{1}{m} (d_c - d_i)_+^2$ $u(\mathbf{d}) = \sum_{i=1}^m u_i(d_i)$ $v(u) = u$	$c_1 = 1$ $c_2 = 0$ $\mathbf{e}_1 = \left(\frac{2}{m} (d_c - d_i)_+ \right)_{i=1, \dots, m}$ $\mathbf{e}_2 = \left(\frac{2}{m} \mathbf{1}_{\{d_c > d_i\}} \right)_{i=1, \dots, m}$
$\text{DVH}_{\text{approx}}(\mathbf{d}; d_c, a)$ $= \frac{1}{m} \sum_{i=1}^m \frac{(d_i/d_c)^a}{1+(d_i/d_c)^a}$	$u_i(d_i) = \frac{1}{m} \frac{(d_i/d_c)^a}{1+(d_i/d_c)^a}$ $u(\mathbf{d}) = \sum_{i=1}^m u_i(d_i)$ $v(u) = u$ $K_i = d_i/d_c, \quad i = 1, \dots, m$	$c_1 = 1$ $c_2 = 0$ $\mathbf{e}_1 = \left(\frac{a}{m d_c} \frac{K_i^{a-1}}{(1+K_i^a)^2} \right)_{i=1, \dots, m}$ $\mathbf{e}_2 = \left(\frac{a K_i^{a-2}}{m d_c^2} \frac{(1-a) K_i^{a+a-1}}{(1+K_i^a)^3} \right)_{i=1, \dots, m}$
$\text{LTCP}(\mathbf{d}; d_c, a)$ $= \frac{1}{m} \sum_{i=1}^m e^{-a(d_i - d_c)}$	$u_i(d_i) = \frac{1}{m} e^{-a(d_i - d_c)}$ $u(\mathbf{d}) = \sum_{i=1}^m u_i(d_i)$ $v(u) = u$	$c_1 = 1$ $c_2 = 0$ $\mathbf{e}_1 = \left(-\frac{a}{m} e^{-a(d_i - d_c)} \right)_{i=1, \dots, m}$ $\mathbf{e}_2 = \left(\frac{a^2}{m} e^{-a(d_i - d_c)} \right)_{i=1, \dots, m}$
$\text{TCP}(\mathbf{d}; d_c, a)$ $= \prod_{i=1}^m \exp \left[-\frac{1}{m} e^{-a(d_i - d_c)} \right]$	$u_i(d_i) = \exp \left[-\frac{1}{m} e^{-a(d_i - d_c)} \right]$ $u(\mathbf{d}) = \prod_{i=1}^m u_i(d_i)$ $v(u) = u$	$c_1 = u(\mathbf{d})$ $c_2 = u(\mathbf{d})$ $\mathbf{e}_1 = \left(\frac{a}{m} e^{-a(d_i - d_c)} \right)_{i=1, \dots, m}$ $\mathbf{e}_2 = \left(-\frac{a^2}{m} e^{-a(d_i - d_c)} \right)_{i=1, \dots, m}$

* gEUD parameter $a \neq 0$.

Abbreviations: gEUD = generalised equivalent uniform dose; LSE = log-sum-exp; QO = quadratic overdose; QU = quadratic underdose; DVH_{approx} = approximate dose volume histogram objective; (L)TCP = (logarithmic) tumour control probability (Munro and Gilbert, 1961, Alber and Reemtsen, 2007).

Functions: $(\cdot)_+ = \max(\cdot, 0)$;

$$\mathbf{1}_{\{d_i > d_c\}} = \begin{cases} 1 & \text{for } d_i > d_c \\ 0 & \text{for } d_i \leq d_c \end{cases} \quad \mathbf{1}_{\{d_c > d_i\}} = \begin{cases} 1 & \text{for } d_c > d_i \\ 0 & \text{for } d_c \leq d_i. \end{cases}$$

Discussion

8.1 INTRODUCTION

This thesis focussed on the LRPM and RPM to improve on the current automated multi-objective optimization of beam intensity profiles as implemented in the 2p ϵ c method in Erasmus-iCycle, which is in clinical use for automated treatment planning. The aims were to decrease the computation time required for automatic plan generation while retaining or possibly improving the plan quality, and to develop an intuitive procedure for automatic configuration of the RPM based on training with previously obtained dose distributions. LRPM and RPM generated plans were compared to those generated with the 2p ϵ c method.

8.2 RPM VERSUS LRPM FOR AUTOMATIC PLANNING

For automated multi-objective treatment planning, the underlying algorithm should consider prioritised goals to balance all planning aims for the PTV and OARs, and all other plan parameters. This section describes differences between the RPM and LRPM, with implications on the applicability as a treatment planning algorithm.

Technically, the LRPM is a lexicographic extension of the RPM, i.e. the RPM is a special case of the LRPM. The configuration parameters of the LRPM and RPM consist of a strictly increasing reference path (parametric curve in the objective space) and a set of trade-off parameters. The fundamental difference between the RPM and LRPM is that a linear reference path is used in the RPM, while the LRPM uses a nonlinear path, e.g. a piecewise linear path, see figure 8.1. A key consequence, when using the trade-off parameters, is that the LRPM indifference curves have multiple bends (due to the nonlinearity of the reference path) while the RPM indifference curves do not, which is crucial for the Pareto-optimal solution generated, see figure 8.1.

Initially, in chapters 2 and 3, the LRPM was designed to closely follow the prioritised goals specified in a wish-list to mimic the strict lexicographic optimisation of the 2p ϵ c

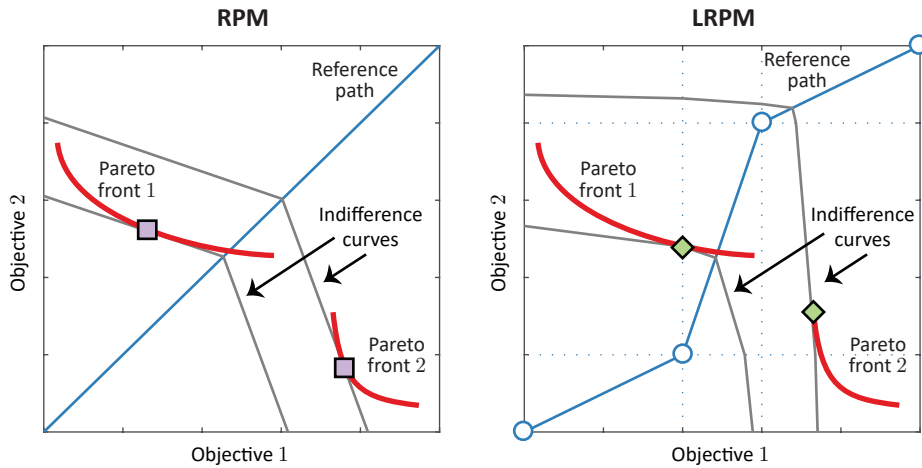


Figure 8.1: Solutions generated with the RPM (squares) and LRPM (diamonds) for the same two Pareto fronts. LRPM indifference curves bend each time they intersect with the dotted lines, causing a steeper trade-off for the solutions generated.

method. Achieving the goals in a prioritised manner was approximated (since strict lexicographic optimisation introduces numerical issues, see discussion chapter 2) by using the LRPM with a piecewise linear reference path. In chapter 2, a procedure (algorithm 2.5.1) was developed to automatically convert a wish-list to a piecewise linear reference path. Then, the other set of LRPM parameters, i.e. the trade-off parameters, were determined manually.

For automatic plan generation of all patients in a treatment group, only one reference path is used. Pareto fronts of patients differ both in shape and location due to the uniqueness of each patient's anatomy. Importantly, the Pareto front does not need to intersect with the reference path which becomes more likely when the objective space has a high dimensionality (many objectives). Therefore, the plan generated is often determined by how the indifference curves are bent, see figure refdiscussion:fig:(l)rpm. The quality of the plans generated are thus more determined by the trade-off parameters (which shape the indifference curves) than the reference path parameters. For the LRPM, the indifference curves bend multiple times so that the plan generated may perform well for the higher prioritised goals but can have a poor overall trade-off when the Pareto front is located far from the reference path. For the RPM, the overall trade-off is balanced similarly for both Pareto fronts close and far from the reference path.

Therefore, in chapters 4 and 5, a new strategy was developed for automatic configuration of the RPM. There are four major differences in this strategy compared to that in chapters 2 and 3. First, the trade-off parameters are automatically determined. Secondly, determining the reference path does no longer require specification of goal values and pri-

orities in a wish-list. Instead, the reference path and trade-off parameters are determined fully by data acquisition from Pareto-optimal training plans. Thirdly, in determining the reference path, the trade-offs made in the training plans are already considered. Lastly, the RPM configuration can be steered based on user-specified preferences for differences between training plans and RPM generated plans.

For the head and neck (HN) cancer treatment site, the LRPM was used for automatic planning in chapter 3, while the RPM was used in chapter 5, albeit with a different strategy. Differences against the dose distributions generated with the 2p ϵ c method were similar, indicating no preference for one planning algorithm over the other. In chapter 3, the first strategy (reference path automatically converted from wish-list and trade-off parameters manually tuned) was validated for automatic IMRT plan generation of 15 HN cancer patients receiving bilateral irradiation without a boost. Manual tuning of trade-off parameters was done for 5 out of 15 patients. In chapter 5, the automatic RPM configuration strategy was used for validating automatic IMRT plan generation for a heterogeneous group of 105 unilateral and bilateral oropharynx cancer patients including a boost technique, with training on 20 patients. In comparison with the partially manual strategy in chapter 3, the automatic strategy in chapter 5 was validated on a larger and more heterogeneous group of HN cancer patients. This is a clear indication that the RPM is suitable for configuring automated planning for complex sites with a relatively small number of training plans.

Automatic configuration was also implemented for the LRPM (not published). The LRPM and RPM resulted in similar performance for HN planning, indicating no preference between the LRPM and RPM for complex planning. After further inspection of the piecewise linear reference path of the LRPM that was generated by the automatic configuration algorithm, the path only bended outside of typical objective values observed in the plans generated. The automatic configuration algorithm thus generated a linear path in the range of objective values that are usually observed in treatment plans, and not a piecewise linear path, while the algorithm was given the freedom to do so. Possibly, this indicates that the LRPM is less suitable than the RPM for automatic configuration.

In summary, there seems to be no added benefit of using the more complex LRPM over the RPM for automatic plan generation. The configuration of the RPM already allows for sufficient model complexity, and shows more flexibility for a heterogeneous group of patients. In fact, the more complex LRPM comes with the drawback that indifference curves are bended multiple times which may lead to the generation of plans with a poor global trade-off between objectives. For implementation and interpretation purposes, the RPM is the preferred planning algorithm. Therefore, the remainder of this chapter focusses on the RPM.

8.3 RPM VERSUS 2p ϵ c METHOD REGARDING AUTOMATIC CONFIGURATION

Automatic configuration of an automatic treatment planning algorithm is a procedure for automatic creation of the configuration parameters that allow for automatic generation of high-quality and Pareto-optimal plans for all patients of a certain tumour type.

In chapters 4–6, automatic configuration of the RPM was developed and validated. The configuration procedure was based on a database of delineated CT scans together with corresponding dose distributions. Also, the objectives and constraints used in planning were known. The aim was to automatically obtain an RPM configuration that results, for each database patient, in an automatically generated RPM plan with a dose distribution similar to the corresponding dose distribution in the database, and if possible, with overall a clinical advantage over the database plans. This is known as group-based inverse multi-objective optimisation (Chan et al., 2014, Chan and Lee, 2017).

There are several reasons why automatic configuration of the 2p ϵ c method is more challenging than automatic RPM configuration: (1) all the configuration parameters of the RPM are combined in a single optimisation problem to generate a plan, while the 2p ϵ c method involves solving consecutive optimisation problems where the parameters are gradually added, (2) configuration of the 2p ϵ c method requires specification of the objectives' priorities together with aimed goal values which are not required for the automatic RPM configuration, and (3) in formulating the next ϵ -constraint optimisation problem, the 2p ϵ c method uses automatic, complex decision making to derive an upper bound for an objective which can either be the corresponding goal value itself (parameter of the 2p ϵ c method) if it was attainable, or a 3% relaxation of the minimum objective value if the goal value was not attainable.

The configuration of the 2p ϵ c method is mostly characterised by the prioritisation of the objectives and corresponding goal values. To apply inverse optimisation for the 2p ϵ c method, the starting point should be the last ϵ -constraint optimisation problem. In this problem, the upper bounds used for the other objectives can be extracted from the given dose distributions. However, it is unknown which objective was minimised since the prioritisation of the objectives is missing. Suppose that this prioritisation is known (or guessed), the upper bound for the lowest prioritised objective in the second-last ϵ -constraint optimisation is unknown. This upper bound comes from the first phase of the 2p ϵ c method, and requires backtracking of the complex decision making: it could either be an actual goal value in the wish-list or a 3% relaxation of the attained minimum. To properly backtrack this decision making, both the goal value itself (unknown parameter) and the attained minimum of the objective (depending on the solution of another ϵ -constraint optimisation problem) in the first phase of the 2p ϵ c method should be known. Applying inverse optimisation for the 2p ϵ c method is therefore highly complex. Possibly, heuristics for choosing the prioritisation and goal values of the objectives could be explored.

For automatic configuration of the RPM, a list of user-defined preferences regarding differences with the training plans is used to iteratively determine a desired configuration. For the 2p ϵ c method, it would not be clear how the parameters (priorities and goal values of the objectives) should be changed to improve certain objectives for a certain amount. Also, it may even not be possible to improve upon an objective because of the mandatory 3% relaxation (this prevents generation of plans on the extreme parts of a Pareto front). Possibly, this relaxation could be lowered although some relaxation is always necessary to prevent numerical issues in the optimisation problems (Klepikova, 1985, Breedveld et al., 2009).

Since the 2p ϵ c method is used clinically at Erasmus MC, wish-lists are available for several treatment sites. For the automatic RPM configuration, training dose distributions are required as input. This allows for a synergy between automatic RPM configuration and the current workflow for clinical treatment planning: for each treatment site that has a wish-list available, the input dose distributions for the RPM configuration can be automatically generated with the 2p ϵ c method. In this way, the starting point remains the wish-list, so that intuitive tuning of the automatically generated plans remains available.

In summary, the RPM involves solving a single optimisation problem in which all objectives are minimised simultaneously while the 2p ϵ c method minimises one objective after another. Although this makes the RPM less intuitive than the 2p ϵ c method, it also makes the RPM more suitable for automatic configuration. Backtracking a single optimisation problem, in which all parameters are present, is more straightforward than backtracking a sequence of optimisation problems that also includes complex decision making.

8.4 RPM FOR ADAPTIVE PLANNING

Besides the automatic generation of a high-quality and Pareto-optimal plan for each patient, the RPM also features short optimisation times for FMO-based planning, typically around 1 minute for prostate or 5 minutes for HN. Therefore, the RPM could become suitable for online-adaptive planning applications. On the other hand, it should be noted that (automatic) delineation (Elmahdy et al., 2019, Kosmin et al., 2019) times are not included here.

For adaptive planning with photons, the RPM can automatically generate a high-quality fluence-map-based plan so that clinicians can validate the resulting dose distribution within minutes after the delineation has finished. However, before the treatment can be delivered, the fluence-map-based plan needs to be segmented for delivery at the treatment unit.

For adaptive planning with protons, the daily changes in patient anatomy (differences in organ position and tissue density along the pencil-beam paths) have a much more severe effect on the dose distribution than in photon planning: a re-optimisation for each fraction

may be essential. In chapter 6, an automated strategy is described that restores spot positions of the planning CT, adds spots to sufficiently cover the tumour, and then applies an RPM optimisation of the spot weights to generate an IMPT plan for the daily anatomy. When using the same spot positions in the RPM generated plan and the Erasmus-iCycle plan, the resulting dose distributions were similar. In practice however, the selection of spot positions cannot be optimal due to time constraints. There is thus a trade-off between computation time and plan quality. Even for non-optimal spot positions, all RPM plans in chapter 6 showed sufficient PTV coverage with acceptable dose values for the OARs. In contrast to application of the RPM for adaptive photon planning, the IMPT plans are directly deliverable by a proton treatment device, i.e. there is no need for segmentation as for photons.

8.5 FUTURE WORK

8.5.1 Beam angle optimisation with the RPM

This thesis investigates the use of the RPM for automatic generation of IMRT plans with predefined fixed beam set-ups. In principle, the RPM can also be used for beam angle optimisation (BAO), as implemented in Erasmus-iCycle ([Breedveld et al., 2012](#)). The RPM can then be used, instead of the 2p ϵ c method, for fast solving of the corresponding IMRT optimisation problems with fixed beams. For the current selection of the best candidate beam (from a fixed list of beams) in Erasmus-iCycle, an equivalence formulation between the 2p ϵ c method and the weighted sum method is used to reduce computational times (the 2p ϵ c method requires solving multiple optimisation problems, while the weighted sum method requires solving a single optimisation problem). With the RPM, there is no need to switch to the weighted sum method for selecting the best candidate beam, since the RPM also requires solving a single optimisation problem. An added benefit is that, without switching methods, the exact same decision making is used (instead of an approximate) to find the next best beam direction.

Possibly, a fixed configuration of the RPM may be optimal for a certain number of beams being used in the final plan, but less favourable for less or more beams. This potential dependency of the RPM configuration on the number of beams in the final plan needs further investigation.

8.5.2 Local Pareto front generation

Both the clinically used 2p ϵ c method and the RPM automatically generate, per patient, a single high-quality Pareto-optimal treatment plan. These individual plans are generated using group-based configurations, i.e. a configuration is tuned to achieve a high average plan quality in the patient group.

Possibly, with personalised configurations (manual planning instead of automatic planning), higher individual plan quality could be obtained. A set of alternative Pareto-optimal plans representing a local Pareto front around the patient's standard plan (generated with the 2p ϵ c method or RPM) could lead to different plan selections. In particular when using a group-based configuration for planning of a heterogeneous or complex treatment site (e.g. in chapter 5), outliers are likely and the generation of a local Pareto front could lead to the selection of a clinically more optimal plan.

During this PhD project, opportunities were explored to use the RPM for generating local Pareto fronts around a given Pareto-optimal plan. We found that this could be done, but that many plans were needed to evenly distribute the plans on the local Pareto front. While this can be done for any number of objectives (dimensions), computing the local Pareto front becomes exponentially more expensive as the dimensionality increases (Messac and Mattson, 2004). This is known as the "curse of dimensionality". Sandwich approaches have been successful in generating the local Pareto front for up to 8 objectives (Craft et al., 2006).

Possibly, computationally more efficient sandwich approaches (Shao and Ehrgott, 2008, Rennen et al., 2011, Bokrantz, 2013) can be applied for generating local Pareto fronts to better deal with the dimensionality issue. These approaches generate the Pareto front iteratively, where in each iteration, the Pareto front is improved by generating plans that that best improve the inner and outer approximations of the Pareto front.

8.5.3 Automatic selection of objective functions

A key condition of the proposed procedure for automatic RPM configuration, was that the constraint and objective functions used in planning were known. Possible techniques for automatic selection of the objectives are described in (Chan et al., 2014). The starting point can be solving an optimisation problem in which the trade-offs encoded in the given dose distribution are preserved while many objective functions are used. From the solution, relative importance factors for each objective can be extracted. Then, the effect of removing an objective with a low importance factor on the dose distribution can be tested. If the dose distribution changes significantly for some patient, the objective should be included in the group-based configuration of the treatment algorithm. On the other hand, if removing an objective hardly has any effect on the dose distributions for (almost) all patients, the corresponding objective is not needed in planning.

8.6 TOWARDS CLINICAL INTRODUCTION OF THE RPM

8.6.1 Applicability of the automatic RPM configuration

In this thesis, the RPM and its automatic configuration procedure have been tested for prostate IMRT, HN IMRT, and online-adaptive IMPT of cervical cancer. Possibly, the system

could also be used for fast brachytherapy planning.

The automatic RPM configuration (chapters 4 and 5) requires input of Pareto-optimal training plans, and user preferences that specify acceptable differences between the training and RPM generated plans regarding plan parameters, e.g. for PTV coverage, OAR sparing, and conformality measures.

The first step is to collect a set of dose distributions for the treatment site to be configured. As a general guideline, around 10 to 20 of these dose distributions should be used for training, and the remaining dose distributions for testing the quality of the RPM configuration on a configuration-independent set of patients (cross-validation is recommended to test the algorithm's performance). The input dose distributions need to be Pareto optimal, and thus requires specification of the objectives used for planning all patients in the treatment site. Also, agreement about possible constraints for all patients within the treatment site is required. When a wish-list is already available for this treatment site, both the objectives and constraints in the wish-list can be used. Then, to automatically generate a Pareto-optimal plan that is representative for a given clinical plan, the latter could be projected on the Pareto front with the RPM (see the appendices of chapter 4).

Secondly, user preferences need to be specified to steer towards an acceptable RPM configuration. An initial RPM configuration is automatically generated from the training plans and does not yet require any user preferences. The user can then (gradually) define clinically relevant criteria, e.g. PTV $V_{95\%}$, NTCP or a planning objective, to evaluate the plan quality. Visually, the RPM generated plans (which are always Pareto optimal) can be compared to the training plans using a boxplot of the plan differences for each defined criterion. From the plan differences, user preferences can be gradually defined to steer the RPM configuration. For example, the initial goal could be to generate RPM plans that have overall similar median criteria values to those in the training plans. Then, possible outliers unfavourable for the RPM could be reduced by adding preferences for the fifth percentile of the corresponding criteria differences. Finally, the training plans could be improved upon by adjusting the user preferences so that the RPM generated plans improve on some of the criteria. To achieve this, it may be necessary to loosen user preferences for other criteria since all RPM generated plans are Pareto optimal. For example, to achieve better OAR sparing in the RPM generated plans, plan parameters for conformality measures are allowed to be somewhat worse than those in the training plans.

Possibly, a user interface (UI) could improve the user-friendliness for steering the RPM configuration. The UI should support addition of criteria for plan comparison, and showing the achieved and desired values for each criterion, e.g. with a boxplot. For each training patient, visual comparison of DVHs and dose distributions should be supported.

The RPM generated plans can either be compared to the Pareto-optimal training plans or the clinical training plans. It should be noted that the comparison with the clinical train-

ing plans is not entirely fair in IMRT planning since the RPM generated plans are fluence-based while the clinical plans are segmented. However, with a well-performing segmentation algorithm, the RPM generated plans can be similar before and after segmentation (Schipaanboord et al., 2019). When comparing Pareto-optimal RPM generated plans with non-Pareto-optimal clinical plans, it should be noted that it is generally easier to find an RPM configuration than when the RPM generated plans are compared with Pareto-optimal training plans. This is because for a non-Pareto-optimal plan, at least one of the objectives can be improved without deteriorating the others. For comparison with non-Pareto-optimal training plans, it is thus recommended to find an RPM configuration of which the resulting RPM plans overall improve the plan quality.

8.6.2 Deliverability of treatment plans

The automatic treatment plan generation with the RPM is based on multi-objective fluence map optimisation. The resulting plan cannot be directly delivered, as the parameters of the treatment device are still to be defined.

A straightforward solution is to apply an automatic segmentation phase after the fluence map optimisation to enable generation of plans that are deliverable with a multileaf collimator (Schipaanboord et al., 2019). This approach is currently used in the clinic at Erasmus MC for fluence maps generated with the 2p ϵ c method. It would thus be a small effort to adapt this approach for fluence maps generated with the RPM.

Another possible solution may be to apply direct aperture optimisation (DAO), which is a nonconvex problem (Bataar et al., 2005) and therefore difficult to solve to optimality. Heuristics such as column generation techniques (Romeijn et al., 2005, Cassioli and Unkelbach, 2012) have been proven effective in IMRT planning.

8.6.3 Quality assurance

The RPM is a complex, non-intuitive treatment planning algorithm. Plans with lower quality than expected may be difficult to detect. Therefore, it is advisable to check the correctness of the plans generated before introducing the RPM clinically.

Knowledge-based (KB) automatic quality assurance (QA) models have been developed for patient-specific QA of Erasmus-iCycle plans (Wang et al., 2016, 2017, 2019). A KB model can be trained with a set of RPM treatment plans for a specific treatment site to predict dose parameters for new patients. In this way, possible outliers can be detected when differences between planned dose and predicted dose is large.

References

- Ahunbay, E., Peng, C., Holmes, S., Godley, A., Lawton, C., and Li, X. Online adaptive re-planning method for prostate radiotherapy. *International Journal of Radiation Oncology Biology Physics*, 77(5):1561–1572, 2010. doi: [10.1016/j.ijrobp.2009.10.013](https://doi.org/10.1016/j.ijrobp.2009.10.013).
- Alber, M. and Reemtsen, R. Intensity modulated radiotherapy treatment planning by use of a barrier-penalty multiplier method. *Optimization Methods and Software*, 22(3):391–411, 2007. doi: [10.1080/10556780600604940](https://doi.org/10.1080/10556780600604940).
- Aleman, D., Glaser, D., Romeijn, H., and Dempsey, J. Interior point algorithms: guaranteed optimality for fluence map optimization in IMRT. *Physics in Medicine and Biology*, 55(18):5467–5482, 2010. doi: [10.1088/0031-9155/55/18/013](https://doi.org/10.1088/0031-9155/55/18/013).
- Babier, A. and Boutilier, J. Knowledge-based automated planning for oropharyngeal cancer. *Medical Physics*, 45(7):2875–2883, 2018. doi: [10.1002/mp.12930](https://doi.org/10.1002/mp.12930).
- Babier, A., Mahmood, R., McNiven, A., Diamant, A., and Chan, T. Knowledge-based automated planning with three-dimensional generative adversarial networks. *Medical Physics*, 47(2):297–306, 2020. doi: [10.1002/mp.13896](https://doi.org/10.1002/mp.13896).
- Balvert, M. and Craft, D. Fast approximate delivery of fluence maps for IMRT and VMAT. *Physics in Medicine and Biology*, 62(4):1225–1247, 2017. doi: [10.1088/1361-6560/aa56b6](https://doi.org/10.1088/1361-6560/aa56b6).
- Bataar, D., Hamacher, H., Ehrgott, M., and Woeginger, G. Decomposition of integer matrices and multileaf collimator sequencing. *Discrete Applied Mathematics*, 152(1–3): 6–34, 2005. doi: [10.1016/j.dam.2005.04.008](https://doi.org/10.1016/j.dam.2005.04.008).
- Bernatowicz, K., Geets, X., Barragan, A., Janssens, G., Souris, K., and Sterpin, E. Feasibility of online IMPT adaptation using fast, automatic and robust dose restoration. *Physics in Medicine and Biology*, 63(8):085018, 2018. doi: [10.1088/1361-6560/aaba8c](https://doi.org/10.1088/1361-6560/aaba8c).

- Bischof, C., Guertler, N., Kowarz, A., and Walther, A. *Parallel Reverse Mode Automatic Differentiation for OpenMP Programs with ADOL-C*, pages 163–173. Lecture Notes in Computational Science and Engineering. Springer, Berlin, Heidelberg, 2008. doi: [10.1007/978-3-540-68942-3_15](https://doi.org/10.1007/978-3-540-68942-3_15).
- Bohoudi, O., Bruynzeel, A., Senan, S., Cuijpers, J., Slotman, B., Lagerwaard, F., and Palacios, M. Fast and robust online adaptive planning in stereotactic MR-guided adaptive radiation therapy (SMART) for pancreatic cancer. *Radiotherapy and Oncology*, 125(3): 439–444, 2017. doi: [10.1016/j.radonc.2017.07.028](https://doi.org/10.1016/j.radonc.2017.07.028).
- Bokrantz, R. Distributed approximation of Pareto surfaces in multicriteria radiation therapy treatment planning. *Physics in Medicine and Biology*, 58(11):3501–3516, 2013. doi: [10.1088/0031-9155/58/11/3501](https://doi.org/10.1088/0031-9155/58/11/3501).
- Bondar, M., Hoogeman, M., Mens, J., Quint, S., Ahmad, R., Dhawtal, G., and Heijmen, B. Individualized nonadaptive and online-adaptive intensity-modulated radiotherapy treatment strategies for cervical cancer patients based on pretreatment acquired variable bladder filling computed tomography scans. *International Journal of Radiation Oncology Biology Physics*, 83(5):1617–1623, 2012. doi: [10.1016/j.ijrobp.2011.10.011](https://doi.org/10.1016/j.ijrobp.2011.10.011).
- Botas, P., Kim, J., Winey, B., and Paganetti, H. Online adaption approaches for intensity modulated proton therapy for head and neck patients based on cone beam CTs and Monte Carlo simulations. *Physics in Medicine and Biology*, 64(1):015004, 2018. doi: [10.1088/1361-6560/aaf30b](https://doi.org/10.1088/1361-6560/aaf30b).
- Boyd, S. and Vandenberghe, L. *Convex Optimization*. Cambridge University Press, 2004.
- Breedveld, S. and Heijmen, B. Data for TROTS – The Radiotherapy Optimisation Test Set. *Data in Brief*, 12:143–149, 2017. doi: [10.1016/j.dib.2017.03.037](https://doi.org/10.1016/j.dib.2017.03.037). <https://hdl.handle.net/1765/116520>.
- Breedveld, S., Storchi, P., M., K., and Heijmen, B. Fast, multiple optimizations of quadratic dose objective functions in IMRT. *Physics in Medicine and Biology*, 51(14):3569–3579, 2006. doi: [10.1088/0031-9155/51/14/019](https://doi.org/10.1088/0031-9155/51/14/019).
- Breedveld, S., Storchi, P., Keijzer, M., Heemink, A., and Heijmen, B. A novel approach to multi-criteria inverse planning for IMRT. *Physics in Medicine and Biology*, 52(20): 6339–6353, 2007. doi: [10.1088/0031-9155/52/20/016](https://doi.org/10.1088/0031-9155/52/20/016).
- Breedveld, S., Storchi, P., and Heijmen, B. The equivalence of multi-criteria methods for radiotherapy plan optimization. *Physics in Medicine and Biology*, 54(23):7199–7209, 2009. doi: [10.1088/0031-9155/54/23/011](https://doi.org/10.1088/0031-9155/54/23/011).

- Breedveld, S., Storchi, P., Voet, P., and Heijmen, B. iCycle: Integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans. *Medical Physics*, 39:951–963, 2012. doi: [10.1118/1.3676689](https://doi.org/10.1118/1.3676689).
- Breedveld, S., Van den Berg, B., and Heijmen, B. An interior-point implementation developed and tuned for radiation therapy treatment planning. *Computational Optimization and Applications*, 68(2):209–242, 2017. doi: [10.1007/s10589-017-9919-4](https://doi.org/10.1007/s10589-017-9919-4).
- Breedveld, S., Craft, D., Van Haveren, R., and Heijmen, B. Multi-criteria optimisation and decision-making in radiotherapy. *European Journal of Operational Research*, 277(1):1–19, 2019. doi: [10.1016/j.ejor.2018.08.019](https://doi.org/10.1016/j.ejor.2018.08.019).
- Bücker, M., Corliss, G., Hovland, P., Naumann, U., and Norris, B. *Automatic Differentiation: Applications, Theory, and Implementations (Lecture Notes in Computational Science and Engineering)*. Springer-Verlag, Berlin, Heidelberg, 2006.
- Buschmann, M., Sharfo, A., Penninkhof, J., Seppenwoolde, Y., Goldner, G., Georg, D., Breedveld, S., and Heijmen, B. Automated volumetric modulated arc therapy planning for whole pelvic prostate radiotherapy. *Strahlentherapie und Onkologie*, 194(4):333–342, 2018. doi: [10.1007/s00066-017-1246-2](https://doi.org/10.1007/s00066-017-1246-2).
- Cabrera, G., Ehrgott, M., Mason, A., and Raith, A. A matheuristic approach to solve the multiobjective beam angle optimization problem in intensity-modulated radiation therapy. *International Transactions in Operational Research*, 25(1):243–268, 2018. doi: [10.1111/itor.12241](https://doi.org/10.1111/itor.12241).
- Cassoli, A. and Unkelbach, J. Aperture shape optimization for IMRT treatment planning. *Physics in Medicine and Biology*, 58(2):301–318, 2012. doi: [10.1088/0031-9155/58/2/301](https://doi.org/10.1088/0031-9155/58/2/301).
- Chan, T., Craig, T., Lee, T., and Sharpe, M. Generalized inverse multiobjective optimization with application to cancer therapy. *Operations Research*, 62(3):680–695, 2014. doi: [10.1287/opre.2014.1267](https://doi.org/10.1287/opre.2014.1267).
- Chan, T. C. Y. and Lee, T. Trade-off preservation in inverse multi-objective convex optimization. *European Journal of Operational Research*, 270(1):25–39, 2017. doi: [10.1016/j.ejor.2018.02.045](https://doi.org/10.1016/j.ejor.2018.02.045).
- Chen, X., Men, K., Li, Y., Yi, J., and Dai, J. A feasibility study on an automated method to generate patient-specific dose distributions for radiotherapy using deep learning. *Medical Physics*, 46(1):56–64, 2019. doi: [10.1002/mp.13262](https://doi.org/10.1002/mp.13262).
- Conn, A., Gould, N., and Toint, P. *Trust Region Methods*. MOS-SIAM Series on Optimization, 2000.

- Craft, D. and Monz, M. Simultaneous navigation of multiple Pareto surfaces, with an application to multicriteria IMRT planning with multiple beam angle configurations. *Medical Physics*, 37(2):736–741, 2010. doi: [10.1118/1.3292636](https://doi.org/10.1118/1.3292636).
- Craft, D., Halabi, T., Shih, H., and Bortfeld, T. Approximating convex Pareto surfaces in multiobjective radiotherapy planning. *Medical Physics*, 33(9):3399–3407, 2006. doi: [10.1118/1.2335486](https://doi.org/10.1118/1.2335486).
- Della Gala, G., Dirkx, M., Hoekstra, N., Fransen, D., Van de Pol, M., Heijmen, B., and Petit, S. Automated treatment plan generation for advanced stage NSCLC patients. *Radiotherapy and Oncology*, 119:S124, 2016. Supplement 1.
- Dijkema, T., Raaijmakers, C., Ten Haken, R., Roesink, J., Braam, P., Houweling, A., Moerland, M., Eisbruch, A., and Terhaard, C. Parotid gland function after radiotherapy: The combined Michigan and Utrecht experience. *International Journal of Radiation Oncology Biology Physics*, 78(2):449–453, 2010. doi: [10.1016/j.ijrobp.2009.07.1708](https://doi.org/10.1016/j.ijrobp.2009.07.1708).
- El Ghaoui, L. Optimization models and applications. <http://livebooklabs.com/keepies/c5a5868ce26b8125>, 2018. Livebook visited Spring 2018.
- Elmahdy, M., Jagt, T., Zinkstok, R., Qiao, Y., Shahzad, R., Sokooti, H., Yousefi, S., Incrocci, L., Marijnen, C., Hoogeman, M., and Staring, M. Robust contour propagation using deep learning and image registration for online adaptive proton therapy of prostate cancer. *Medical Physics*, 46(8):3329–3343, 2019. doi: [10.1002/mp.13620](https://doi.org/10.1002/mp.13620).
- Fan, J., Wang, J., Chen, Z., Hu, C., Zhang, Z., and Hu, W. Automatic treatment planning based on three-dimensional dose distribution predicted from deep learning technique. *Medical Physics*, 46(1):370–381, 2019. doi: [10.1002/mp.13271](https://doi.org/10.1002/mp.13271).
- Finazzi, T., Palacios, M., Spoelstra, F., Haasbeek, C., Bruynzeel, A., Slotman, B., Lagerwaard, F., and Senan, S. Role of on-table plan adaptation in MR-guided ablative radiation therapy for central lung tumors. *International Journal of Radiation Oncology Biology Physics*, 104(4):933–941, 2019. doi: [10.1016/j.ijrobp.2019.03.035](https://doi.org/10.1016/j.ijrobp.2019.03.035).
- Forsgren, A., Gill, P., and Wright, M. Interior methods for nonlinear optimization. *SIAM Review*, 44(4):525–597, 2002. doi: [10.1137/S0036144502414942](https://doi.org/10.1137/S0036144502414942).
- Ge, Y. and Wu, Q. Knowledge-based planning for intensity-modulated radiation therapy: A review of data-driven approaches. *Medical Physics*, 46(6):2760–2775, 2019. doi: [10.1002/mp.13526](https://doi.org/10.1002/mp.13526).
- Gondzio, J. Multiple centrality corrections in a primal-dual method for linear programming. *Computational Optimization and Applications*, 6(2):137–156, 1996. doi: [10.1007/BF00249643](https://doi.org/10.1007/BF00249643).

- Gower, R. and Mello, M. A new framework for the computation of Hessians. *Optimization Methods and Software*, 27(2):251–273, 2012. doi: [10.1080/10556788.2011.580098](https://doi.org/10.1080/10556788.2011.580098).
- Granat, J. and Makowski, M. Interactive specification and analysis of aspiration-based preferences. *European journal of Operational Research*, 122(2):469–485, 2000. doi: [10.1016/S0377-2217\(99\)00248-9](https://doi.org/10.1016/S0377-2217(99)00248-9).
- Griewank, A., Juedes, D., and Utke, J. ADOL-C: A package for the automatic differentiation of algorithms written in C/C++. *ACM Transactions on Mathematical Software*, 22(2): 131–167, 1996. doi: [10.1145/229473.229474](https://doi.org/10.1145/229473.229474).
- Gustavson, F. Two fast algorithms for sparse matrices: multiplication and permuted transposition. *ACM Transactions on Mathematical Software*, 4(3):250–269, 1978. doi: [10.1145/355791.355796](https://doi.org/10.1145/355791.355796).
- Guthier, C., Damato, A., Viswanathan, A., Hesser, J., and Cormack, R. A fast multitarget inverse treatment planning strategy optimizing dosimetric measures for high-dose-rate (HDR) brachytherapy. *Medical Physics*, 44(9):4452–4462, 2017. doi: [10.1002/mp.12410](https://doi.org/10.1002/mp.12410).
- Van Haveren, R. and Breedveld, S. Fast and exact Hessian computation for a class of non-linear functions used in radiation therapy treatment planning. *Physics in Medicine and Biology*, 64(16):16NT01, 2019. doi: [10.1088/1361-6560/ab1e17](https://doi.org/10.1088/1361-6560/ab1e17).
- Van Haveren, R. and Breedveld, S. Corrigendum: Fast and exact Hessian computation for a class of non-linear functions used in radiation therapy treatment planning (2019 *Phys. Med. Biol.* 64 16NT01). *Physics in Medicine and Biology*, 65(7):079501, 2020. doi: [10.1088/1361-6560/ab7ac0](https://doi.org/10.1088/1361-6560/ab7ac0).
- Van Haveren, R., Breedveld, S., Keijzer, M., Voet, P., Heijmen, B., and Ogryczak, W. Lexicographic extension of the reference point method applied in radiation therapy treatment planning. *European Journal of Operational Research*, 263(1):247–257, 2017a. doi: [10.1016/j.ejor.2017.04.062](https://doi.org/10.1016/j.ejor.2017.04.062).
- Van Haveren, R., Ogryczak, W., Verduijn, G., Keijzer, M., Heijmen, B., and Breedveld, S. Fast and fuzzy multi-objective radiotherapy treatment plan generation for head and neck cancer patients with the lexicographic reference point method (LRPM). *Physics in Medicine and Biology*, 62(11):4318–4332, 2017b. doi: [10.1088/1361-6560/62/11/4318](https://doi.org/10.1088/1361-6560/62/11/4318).
- Van Haveren, R., Heijmen, B., and Breedveld, S. Automatically configuring the reference point method for automated multi-objective treatment planning. *Physics in Medicine and Biology*, 64(3):035002, 2019. doi: [10.1088/1361-6560/aaf9fe](https://doi.org/10.1088/1361-6560/aaf9fe).

- Heijkoop, S., Langerak, T., Quint, S., Bondar, L., Mens, J., Heijmen, B., and Hoogeman, M. Clinical implementation of an online adaptive plan-of-the-day protocol for nonrigid motion management in locally advanced cervical cancer IMRT. *International Journal of Radiation Oncology Biology Physics*, 90(3):673–679, 2014. doi: [10.1016/j.ijrobp.2014.06.046](https://doi.org/10.1016/j.ijrobp.2014.06.046).
- Heijkoop, S., Langerak, T., Quint, S., Mens, J., Zolnay, A., Heijmen, B., and Hoogeman, M. Quantification of intra-fraction changes during radiotherapy of cervical cancer assessed with pre- and post-fraction Cone Beam CT scans. *Radiotherapy and Oncology*, 117(3): 536–541, 2015. doi: [10.1016/j.radonc.2015.08.034](https://doi.org/10.1016/j.radonc.2015.08.034).
- Heijmen, B., Voet, P., Fransen, D., Penninkhof, J., Milder, M., Akhlat, H., Bonomo, P., Casati, M., Georg, D., Goldner, G., Henry, A., Lilley, J., Lohr, F., Marrazzo, L., Pallotta, S., Pellegrini, R., Seppenwoolde, Y., Simontacchi, G., Steil, V., Stieler, F., Wilson, S., and Breedveld, S. Fully automated, multi-criterial planning for Volumetric Modulated Arc Therapy – An international multi-center validation for prostate cancer. *Radiotherapy and Oncology*, 128(2):343–348, 2018. doi: [10.1016/j.radonc.2018.06.023](https://doi.org/10.1016/j.radonc.2018.06.023).
- Henke, L., Kashani, R., Robinson, C., Curcuru, A., DeWees, T., Bradley, J., Green, O., Michalski, J., Mutic, S., Parikh, P., and Olsen, J. Phase I trial of stereotactic MR-guided online adaptive radiation therapy (SMART) for the treatment of oligometastatic or unresectable primary malignancies of the abdomen. *Radiotherapy and Oncology*, 126(3): 519–526, 2018. doi: [10.1016/j.radonc.2017.11.032](https://doi.org/10.1016/j.radonc.2017.11.032).
- Hussein, M., Heijmen, B., Verellen, D., and Nisbet, A. Automation in intensity modulated radiotherapy treatment planning—a review of recent innovations. *British Journal of Radiology*, 91:20180270, 2018. doi: [10.1259/bjr.20180270](https://doi.org/10.1259/bjr.20180270).
- Jagt, T., Breedveld, S., Van Haveren, R., Heijmen, B., and Hoogeman, M. An automated planning strategy for near real-time adaptive proton therapy in prostate cancer. *Physics in Medicine and Biology*, 63(13):135017, 2018. doi: [10.1088/1361-6560/aacaa7](https://doi.org/10.1088/1361-6560/aacaa7).
- Jagt, T., Breedveld, S., Van Haveren, R., Nout, R., Astreinidou, E., Heijmen, B., and Hoogeman, M. Plan-library supported automated replanning for online-adaptive intensity-modulated proton therapy of cervical cancer. *Acta Oncologica*, 58(10):1440–1445, 2019. doi: [10.1080/0284186X.2019.1627414](https://doi.org/10.1080/0284186X.2019.1627414).
- Jee, K., McShan, D., and Fraass, B. Lexicographic ordering: intuitive multicriteria optimization for IMRT. *Physics in Medicine and Biology*, 52(7):1845–1861, 2007. doi: [10.1088/0031-9155/52/7/006](https://doi.org/10.1088/0031-9155/52/7/006).

- Kaliszewski, I. *Quantitative Pareto Analysis by Cone Separation Technique*. Kluwer Academic Publishers, Dordrecht, 1994.
- Kaliszewski, I. and Michalowski, W. Efficient solutions and bounds on tradeoffs. *Journal of Optimization Theory and Applications*, 94(2):381–394, 1997. doi: [10.1023/A:1022687729559](https://doi.org/10.1023/A:1022687729559).
- Klepikova, M. The stability of lexicographic optimization problems. *USSR Computational Mathematics and Mathematical Physics*, 25(1):21–29, 1985. doi: [10.1016/0041-5553\(85\)90037-0](https://doi.org/10.1016/0041-5553(85)90037-0).
- Kooy, H., Clasio, B., Lu, H., Madden, T., Bentefour, H., Depauw, N., Adams, J., Trofimov, A., Demaret, D., Delaney, T., and Flanz, J. A case study in proton pencil-beam scanning delivery. *International Journal of Radiation Oncology Biology Physics*, 76(2):624–630, 2010. doi: [10.1016/j.ijrobp.2009.06.065](https://doi.org/10.1016/j.ijrobp.2009.06.065).
- Korhonen, P. and Wallenius, J. A Pareto race. *Naval Research Logistics*, 35(6):615–623, 1988. doi: [10.1002/1520-6750\(198812\)35:6<615::AID-NAV3220350608>3.0.CO;2-K](https://doi.org/10.1002/1520-6750(198812)35:6<615::AID-NAV3220350608>3.0.CO;2-K).
- Kosmin, M., Ledsam, J., Romera-Paredes, B., Mendes, R., Moinuddin, S., de Souza, D., Gunn, L., Kelly, C., Hughes, C., Karthikesalingam, A., Nutting, C., and Sharma, R. Rapid advances in auto-segmentation of organs at risk and target volumes in head and neck cancer. *Radiotherapy and Oncology*, 135:130–140, 2019. doi: [10.1016/j.radonc.2019.03.004](https://doi.org/10.1016/j.radonc.2019.03.004).
- Kurz, C., Nijhuis, R., Reiner, M., Ganswindt, U., Thieke, C., Belka, C., Parodi, K., and Landry, G. Feasibility of automated proton therapy plan adaptation for head and neck tumors using cone beam CT images. *Radiation Oncology*, 11(1):64, 2016. doi: [10.1186/s13014-016-0641-7](https://doi.org/10.1186/s13014-016-0641-7).
- Leinders, S., Breedveld, B., Méndez Romero, A., Schaart, D., Seppenwoolde, Y., and Heijmen, B. Adaptive liver stereotactic body radiation therapy: Automated daily plan re-optimization prevents dose delivery degradation caused by anatomy deformations. *International Journal of Radiation Oncology Biology Physics*, 87(5):1016–1021, 2013. doi: [10.1016/j.ijrobp.2013.08.009](https://doi.org/10.1016/j.ijrobp.2013.08.009).
- Lomax, A. Intensity modulated proton therapy and its sensitivity to treatment uncertainties 1: the potential effects of calculational uncertainties. *Physics in Medicine and Biology*, 53(4):1027–1042, 2008a. doi: [10.1088/0031-9155/53/4/014](https://doi.org/10.1088/0031-9155/53/4/014).
- Lomax, A. Intensity modulated proton therapy and its sensitivity to treatment uncertainties 2: the potential effects of inter-fraction and inter-field motions. *Physics in Medicine and Biology*, 53(4):1043–1056, 2008b. doi: [10.1088/0031-9155/53/4/015](https://doi.org/10.1088/0031-9155/53/4/015).

- Long, T., Matuszak, M., Feng, M., Fraass, B., Ten Haken, R., and Romeijn, H. Sensitivity analysis for lexicographic ordering in radiation therapy treatment planning. *Medical Physics*, 39(6):3445–3455, 2012. doi: [10.1118/1.4720218](https://doi.org/10.1118/1.4720218).
- Lyman, J. Complication probability as assessed from dose-volume histograms. *Radiation Research Supplement*, 8(2s):S13–S19, 1985. doi: [10.2307/3583506](https://doi.org/10.2307/3583506).
- Ma, M., Kovalchuk, N., Buyyounouski, M., Xing, L., and Yang, Y. Dosimetric features-driven machine learning model for DVH prediction in VMAT treatment planning. *Medical Physics*, 46(2):857–867, 2019. doi: [10.1002/mp.13334](https://doi.org/10.1002/mp.13334).
- Mehrotra, S. On the implementation of a primal-dual interior point method. *SIAM Journal on Optimization*, 2(4):575–601, 1992. doi: [10.1137/0902028](https://doi.org/10.1137/0902028).
- Men, C., Romeijn, H., Taşkın, C., and Dempsey, J. An exact approach to direct aperture optimization in IMRT treatment planning. *Physics in Medicine and Biology*, 52(24):7333–7352, 2007. doi: [10.1088/0031-9155/52/24/009](https://doi.org/10.1088/0031-9155/52/24/009).
- Messac, A. and Mattson, C. Normal constraint method with guarantee of even representation of complete Pareto frontier. *AIAA Journal*, 42(10):2101–2111, 2004. doi: [10.2514/1.8977](https://doi.org/10.2514/1.8977).
- Miettinen, K. *Nonlinear Multiobjective Optimization*, volume 12 of *International Series in Operations Research and Management Science*. Kluwer Academic Publishers, Dordrecht, 1999.
- Miettinen, K., Ruiz, F., and Wierzbicki, A. Introduction to multiobjective optimization: Interactive approaches. In *Multiobjective Optimization*, Lecture Notes in Computer Science, pages 27–57. Springer Berlin Heidelberg, 2008. doi: [10.1007/978-3-540-88908-3_2](https://doi.org/10.1007/978-3-540-88908-3_2).
- Miguel-Chumacero, E., Currie, G., Johnston, A., and Currie, S. Effectiveness of multi-criteria optimization based trade-off exploration in combination with RapidPlan for head & neck radiotherapy planning. *Radiation Oncology*, 13(1):229, 2018. doi: [10.1186/s13014-018-1175-y](https://doi.org/10.1186/s13014-018-1175-y).
- Monz, M., Küfer, K., Bortfeld, T., and Thieke, C. Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning. *Physics in Medicine and Biology*, 53(4):985–998, 2008. doi: [10.1088/0031-9155/53/4/011](https://doi.org/10.1088/0031-9155/53/4/011).
- Moré, J. and Sorensen, D. Computing a trust region step. *SIAM Journal on Scientific and Statistical Computing*, 4(3):553–572, 1983. doi: [10.1137/0904038](https://doi.org/10.1137/0904038).

- Munro, T. and Gilbert, C. The relation between tumour lethal doses and the radiosensitivity of tumour cells. *The British Journal of Radiology*, 34(400):246–251, 1961. doi: [10.1259/0007-1285-34-400-246](https://doi.org/10.1259/0007-1285-34-400-246).
- Murdoch-Kinch, C., Kim, H., Vineberg, K., Ship, J., and Eisbruch, A. Dose-effect relationships for the submandibular salivary glands and implications for their sparing by intensity modulated radiotherapy. *International Journal of Radiation Oncology Biology Physics*, 72(2):373–382, 2008. doi: [10.1016/j.ijrobp.2007.12.033](https://doi.org/10.1016/j.ijrobp.2007.12.033).
- Niemierko, A. Reporting and analyzing dose distribution: A concept of equivalent uniform dose. *Medical Physics*, 24(1):103–110, 1997. doi: [10.1118/1.598063](https://doi.org/10.1118/1.598063).
- Niemierko, A. A generalized concept of equivalent uniform dose (EUD). *Medical Physics*, 26(1):1101, 1999.
- Nocedal, J. and Wright, S. *Numerical Optimization*. Springer, 2006.
- Ogryczak, W. Preemptive reference point method. In *Multicriteria Analysis*, pages 156–167. Springer Berlin Heidelberg, 1997. doi: [10.1007/978-3-642-60667-0_16](https://doi.org/10.1007/978-3-642-60667-0_16).
- Ogryczak, W. and Kozłowski, B. Reference point method with importance weighted ordered partial achievements. *TOP*, 19(2):380–401, 2011. doi: [10.1007/s11750-009-0121-4](https://doi.org/10.1007/s11750-009-0121-4).
- Petit, S., Seco, J., and Kooy, H. Increasing maximum tumor dose to manage range uncertainties in IMPT treatment planning. *Physics in Medicine and Biology*, 58(20):7329–7341, 2013. doi: [10.1088/0031-9155/58/20/7329](https://doi.org/10.1088/0031-9155/58/20/7329).
- Pissanetzky, S. *Sparse matrix technology*. Academic Press, 1984.
- Rennen, G., Van Dam, E., and Den Hertog, D. Enhancement of sandwich algorithms for approximating higher-dimensional convex Pareto sets. *INFORMS Journal on Computing*, 23(4):493–517, 2011. doi: [10.1287/ijoc.1100.0419](https://doi.org/10.1287/ijoc.1100.0419).
- Romeijn, E., Ahuja, R., Dempsey, J., and Kumar, A. A column generation approach to radiation therapy treatment planning using aperture modulation. *SIAM Journal on Optimization*, 15(3):838–862, 2005. doi: [10.1137/040606612](https://doi.org/10.1137/040606612).
- Rossi, L., Breedveld, S., Heijmen, B., Voet, P., Lanconelli, N., and Aluwini, S. On the beam direction search space in computerized non-coplanar beam angle optimization for IMRT–prostate SBRT. *Physics in Medicine and Biology*, 57(17):5441–5458, 2012. doi: [10.1088/0031-9155/57/17/5441](https://doi.org/10.1088/0031-9155/57/17/5441).

- Rossi, L., Breedveld, S., Aluwini, S., and Heijmen, B. Noncoplanar beam angle class solutions to replace time-consuming patient-specific beam angle optimization in robotic prostate stereotactic body radiation therapy. *International Journal of Radiation Oncology Biology Physics*, 92(4):762–770, 2015. doi: [10.1016/j.ijrobp.2015.03.013](https://doi.org/10.1016/j.ijrobp.2015.03.013).
- Van de Sande, M., Creutzberg, C., Van de Water, S., Sharfo, A., and Hoogeman, M. Which cervical and endometrial cancer patients will benefit most from intensity-modulated proton therapy? *Radiotherapy and Oncology*, 120(3):397–403, 2016. doi: [10.1016/j.radonc.2016.06.016](https://doi.org/10.1016/j.radonc.2016.06.016).
- Schipaanboord, B., Breedveld, S., Rossi, L., Keijzer, M., and Heijmen, B. Automated prioritised 3D dose-based MLC segment generation for step-and-shoot IMRT. *Physics in Medicine and Biology*, 64(16):165013, 2019. doi: [10.1088/1361-6560/ab1df9](https://doi.org/10.1088/1361-6560/ab1df9).
- Van de Schoot, A., De Boer, P., Crama, K., Visser, J., Stalpers, L., Rasch, C., and Bel, A. Dosimetric advantages of proton therapy compared with photon therapy using an adaptive strategy in cervical cancer. *Acta Oncologica*, 55(7):892–899, 2016. doi: [10.3109/0284186X.2016.1139179](https://doi.org/10.3109/0284186X.2016.1139179).
- Shao, L. and Ehrgott, M. Approximately solving multiobjective linear programmes in objective space and an application in radiotherapy treatment planning. *Mathematical Methods of Operations Research*, 68(2):257–276, 2008. doi: [10.1007/s00186-008-0220-2](https://doi.org/10.1007/s00186-008-0220-2).
- Sharfo, A., Voet, P., Breedveld, S., Mens, J., Hoogeman, M., and Heijmen, B. Comparison of VMAT and IMRT strategies for cervical cancer patients using automated planning. *Radiotherapy and Oncology*, 114(3):395–401, 2015. doi: [10.1016/j.radonc.2015.02.006](https://doi.org/10.1016/j.radonc.2015.02.006).
- Sharfo, A., Dirkx, M., Bijman, R., Schillemans, W., Breedveld, S., Aluwini, S., Pos, F., Incrocci, L., and Heijmen, B. Late toxicity in the randomized multicenter HYPRO trial for prostate cancer analyzed with automated treatment planning. *Radiotherapy and Oncology*, 128(2):349–356, 2018. doi: [10.1016/j.radonc.2018.05.028](https://doi.org/10.1016/j.radonc.2018.05.028).
- Storchi, P. and Woudstra, E. Calculation of the absorbed dose distribution due to irregularly shaped photon beams using pencil beam kernels derived from basic beam data. *Physics in Medicine and Biology*, 41(4):637–656, 1996. doi: [10.1088/0031-9155/41/4/005](https://doi.org/10.1088/0031-9155/41/4/005).
- Tyran, M., Jiang, N., Cao, M., Raldow, A., Lamb, J., Low, D., Luterstein, E., Steinberg, M., and Lee, P. Retrospective evaluation of decision-making for pancreatic stereotactic MR-guided adaptive radiotherapy. *Radiotherapy and Oncology*, 129(2):319–325, 2018. doi: [10.1016/j.radonc.2018.08.009](https://doi.org/10.1016/j.radonc.2018.08.009).

- Vanderbei, R. and Shanno, D. An interior-point algorithm for nonconvex nonlinear programming. *Computational Optimization and Applications*, 11(1-3):231–252, 1999. doi: [10.1023/A:1008677427361](https://doi.org/10.1023/A:1008677427361).
- Voet, P., Dirkx, M., Breedveld, S., Fransen, D., Levendag, P., and Heijmen, B. Toward fully automated multicriterial plan generation: A prospective clinical study. *International Journal of Radiation Oncology Biology Physics*, 85(3):866–872, 2013a. doi: [10.1016/j.ijrobp.2012.04.015](https://doi.org/10.1016/j.ijrobp.2012.04.015).
- Voet, P., Dirkx, M., Breedveld, S., and Heijmen, B. Automated generation of IMRT treatment plans for prostate cancer patients with metal hip prostheses: Comparison of different planning strategies. *Medical Physics*, 40(7):071704, 2013b. doi: [10.1118/1.4808117](https://doi.org/10.1118/1.4808117).
- Voet, P., Dirkx, M., Breedveld, S., Al-Mamgani, A., Incrocci, L., and Heijmen, B. Fully automated volumetric modulated arc therapy plan generation for prostate cancer patients. *International Journal of Radiation Oncology Biology Physics*, 88(5):1175–1179, 2014. doi: [10.1016/j.ijrobp.2013.12.046](https://doi.org/10.1016/j.ijrobp.2013.12.046).
- Wächter, A. and Biegler, L. On the implementation of an interior-point filter line-search algorithm for large-scale nonlinear programming. *Mathematical Programming Series A*, 106(1):25–57, 2006. doi: [10.1007/s10107-004-0559-y](https://doi.org/10.1007/s10107-004-0559-y).
- Walther, A. and Griewank, A. *Getting started with ADOL-C in Combinatorial Scientific Computing*. Chapman-Hall CRC Computational Science, ed. U Naumann and O Schenk, 2012.
- Wang, Y., Breedveld, S., Heijmen, B., and Petit, S. Evaluation of plan quality assurance models for prostate cancer patients based on fully automatically generated Pareto-optimal treatment plans. *Physics in Medicine and Biology*, 61(11):4268–4282, 2016. doi: [10.1088/0031-9155/61/11/4268](https://doi.org/10.1088/0031-9155/61/11/4268).
- Wang, Y., Heijmen, B., and Petit, S. Prospective clinical validation of independent DVH prediction for plan QA in automatic treatment planning for prostate cancer patients. *Radiotherapy and Oncology*, 125(3):500–506, 2017. doi: [10.1016/j.radonc.2017.09.021](https://doi.org/10.1016/j.radonc.2017.09.021).
- Wang, Y., Heijmen, B., and Petit, S. Knowledge-based dose prediction models for head and neck cancer are strongly affected by interorgan dependency and dataset inconsistency. *Medical Physics*, 46(2):934–943, 2019. doi: [10.1002/mp.13316](https://doi.org/10.1002/mp.13316).
- Van de Water, S., Kraan, A., Breedveld, S., Schillemans, W., Teguh, D., Kooy, H., Madden, T., Heijmen, B., and Hoogeman, M. Improved efficiency of multi-criteria IMPT treatment planning using iterative resampling of randomly placed pencil beams. *Physics in Medicine and Biology*, 58(19):6969–6983, 2013. doi: [10.1088/0031-9155/58/19/6969](https://doi.org/10.1088/0031-9155/58/19/6969).

- Van de Water, S., Kooy, H., Heijmen, B., and Hoogeman, M. Shortening delivery times of intensity modulated proton therapy by reducing proton energy layers during treatment plan optimization. *International Journal of Radiation Oncology Biology Physics*, 92(2): 460–468, 2015. doi: [10.1016/j.ijrobp.2015.01.031](https://doi.org/10.1016/j.ijrobp.2015.01.031).
- Wierzbicki, A. A mathematical basis for satisficing decision making. *Mathematical Modelling*, 3:391–405, 1982. doi: [10.1016/0270-0255\(82\)90038-0](https://doi.org/10.1016/0270-0255(82)90038-0).
- Wierzbicki, A. On the completeness and constructiveness of parametric characterizations to vector optimization problems. *OR Spectrum*, 8:73–87, 1986. doi: [10.1007/BF01719738](https://doi.org/10.1007/BF01719738).
- Wilkens, J., Alaly, J., Zakarian, K., Thorstad, W., and Deasy, J. IMRT treatment planning based on prioritizing prescription goals. *Physics in Medicine and Biology*, 52(6):1675–1692, 2007. doi: [10.1088/0031-9155/52/6/009](https://doi.org/10.1088/0031-9155/52/6/009).
- Winkel, D., Bol, G., Kroon, P., Van Asselen, B., Hackett, S., Werensteijn-Honingh, A., Intven, M., Eppinga, W., Tijssen, R., Kerkmeijer, L., De Boer, H., Mook, S., Meijer, G., Hes, J., Willemsen-Bosman, M., De Groot-Van Breugel, E., Jürgenliemk-Schulz, I., and Raaymakers, B. Adaptive radiotherapy: The Elekta Unity MR-linac concept. *Clinical and Translational Radiation Oncology*, 18:54–59, 2019. doi: [10.1016/j.ctro.2019.04.001](https://doi.org/10.1016/j.ctro.2019.04.001).
- Wright, S. *Primal-Dual Interior-Point Methods*. SIAM, 1997.
- Yuan, L., Ge, Y., Lee, W., Yin, F., Kirkpatrick, J., and Wu, Q. Quantitative analysis of the factors which affect the interpatient organ-at-risk dose sparing variation in IMRT plans. *Medical Physics*, 39(11):6868–6878, 2012. doi: [10.1118/1.4757927](https://doi.org/10.1118/1.4757927).

List of publications

FIRST AUTHOR PUBLICATIONS

Van Haveren, R., Breedveld, S., Keijzer, M., Voet, P., Heijmen, B., and Ogryczak, W. Lexicographic extension of the reference point method applied in radiation therapy treatment planning. *European Journal of Operational Research*, 263(1):247–257, 2017a. doi: [10.1016/j.ejor.2017.04.062](https://doi.org/10.1016/j.ejor.2017.04.062).

Van Haveren, R., Ogryczak, W., Verduijn, G., Keijzer, M., Heijmen, B., and Breedveld, S. Fast and fuzzy multi-objective radiotherapy treatment plan generation for head and neck cancer patients with the lexicographic reference point method (LRPM). *Physics in Medicine and Biology*, 62(11):4318–4332, 2017b. doi: [10.1088/1361-6560/62/11/4318](https://doi.org/10.1088/1361-6560/62/11/4318).

Van Haveren, R., Heijmen, B., and Breedveld, S. Automatically configuring the reference point method for automated multi-objective treatment planning. *Physics in Medicine and Biology*, 64(3):035002, 2019. doi: [10.1088/1361-6560/aaf9fe](https://doi.org/10.1088/1361-6560/aaf9fe).

Van Haveren, R., Heijmen, B., and Breedveld, S. Automatic configuration of the reference point method for fully automated multi-objective treatment planning applied to oropharyngeal cancer. *Medical Physics*, 47(4):1499–1508, 2020. doi: [10.1002/mp.14073](https://doi.org/10.1002/mp.14073).

Van Haveren, R. and Breedveld, S. Fast and exact Hessian computation for a class of non-linear functions used in radiation therapy treatment planning. *Physics in Medicine and Biology*, 64(16):16NT01, 2019. doi: [10.1088/1361-6560/ab1e17](https://doi.org/10.1088/1361-6560/ab1e17).

CO-AUTHOR PUBLICATIONS

Breedveld, S., Craft, D., Van Haveren, R., and Heijmen, B. Multi-criteria optimisation and decision-making in radiotherapy. *European Journal of Operational Research*, 277(1):1–19, 2019. doi: [10.1016/j.ejor.2018.08.019](https://doi.org/10.1016/j.ejor.2018.08.019).

- Jagt, T., Breedveld, S., Van Haveren, R., Heijmen, B., and Hoogeman, M. An automated planning strategy for near real-time adaptive proton therapy in prostate cancer. *Physics in Medicine and Biology*, 63(13):135017, 2018. doi: [10.1088/1361-6560/aacaa7](https://doi.org/10.1088/1361-6560/aacaa7).
- Jagt, T., Breedveld, S., Van Haveren, R., Nout, R., Astreinidou, E., Heijmen, B., and Hoogeman, M. Plan-library supported automated replanning for online-adaptive intensity-modulated proton therapy of cervical cancer. *Acta Oncologica*, 58(10):1440–1445, 2019. doi: [10.1080/0284186X.2019.1627414](https://doi.org/10.1080/0284186X.2019.1627414).
- Jagt, T., Breedveld, S., Van Haveren, R., Heijmen, B., and Hoogeman, M. Online-adaptive versus robust IMPT planning for prostate cancer: how much do we gain? *Radiotherapy and Oncology*, 151:228–233, 2020. doi: [10.1016/j.radonc.2020.07.054](https://doi.org/10.1016/j.radonc.2020.07.054).

Summary

This thesis focussed on introducing the *lexicographic reference point method* (LRPM) and *reference point method* (RPM) for fast automatic generation of Pareto-optimal and clinically favourable intensity-modulated radiotherapy (IMRT) plans, automatically configuring the RPM based on a set of training plans, and exploring the automatic RPM configuration for daily adaptive re-planning in intensity-modulated proton therapy (IMPT).

After a brief introduction to radiotherapy in chapter 1, the process of both manual trial-and-error and automated radiotherapy treatment planning is described. The focus lies on Erasmus-iCycle, which is the clinically applied treatment planning system for automatic multi-objective optimisation of beam angles and intensity profiles. The core of this system is the multi-objective 2-phase ϵ -constraint (2p ϵ c) method for automatic multi-objective optimisation of Pareto-optimal intensity profiles given a set of static beam angles. The LRPM and RPM aim to improve the 2p ϵ c method by (1) decreasing optimisation times, and (2) reducing the configuration workload of the planning algorithm, both without deteriorating the treatment plan quality.

Chapter 2 introduces the LRPM as a multi-objective optimisation method for fast automatic generation of Pareto-optimal treatment plans. Applying the LRPM requires solving a single optimisation problem, while a number of consecutive optimisation problems need to be solved for the clinically used 2p ϵ c method. This number of problems grows linearly with the number of objective functions used in the plan generation. The LRPM lexicographically extends the existing RPM which allows for mimicking the 2p ϵ c method, i.e. approximating the lexicographic optimisation of objectives towards predefined aspiration levels or goal values. In addition, it is explained how the LRPM can impose a priori bounds on partial trade-offs. The LRPM was tested for automatic plan generation of 30 randomly selected prostate cancer patients. For each patient, this resulted in a plan with similar dosimetric quality to that of the plan generated with the 2p ϵ c method. The LRPM achieved a relative speed-up factor of 10, decreasing the average plan computation time from 12.4 to

1.2 minutes.

In chapter 3, the LRPM is applied to automatically generate Pareto-optimal plans for 15 head and neck cancer patients who received bilateral neck irradiation. For each patient, both the LRPM and 2pec method were applied for automatic Pareto-optimal plan generation. All automatically generated plans were clinically relevant. Compared to the plans generated with the 2pec method, some patients had large favourable gains for certain objectives (large dose reductions for some organs-at-risk or OARs) in the LRPM generated plan at the cost of small degradations for other objectives. The LRPM was thus able to achieve better global trade-offs between objectives for some patients. For the other patients, the plans had similar quality. Because of the many objectives (up to 22) used in planning, the LRPM decreased average plan computation times from 209.2 to 9.5 minutes, a relative speed-up factor of 22.

Chapter 4 proposes an algorithm for automatic configuration of the RPM. Configuration of the 2pec method for automated planning currently involves a labour-intensive and time-consuming interactive tuning procedure that needs to be repeated for each treatment site/planning protocol. The aim of configuring an automatic planning system is to automatically generate a high-quality plan for each patient, without the need for further interactive manual fine-tuning. The automatic configuration algorithm developed for the RPM is based on plan characteristics in a set of training plans. In this study, a prostate cancer database was used that included 287 planning CT scans together with corresponding Pareto-optimal plans generated with the 2pec method. With the training plans, an RPM configuration is automatically obtained by iteratively steering the configuration towards user-defined preferences that specify differences allowed between the training and RPM generated plans. These preferences include achieving similar PTV coverage, improving rectum sparing, and allowing slightly worse conformality. To investigate the quality of the RPM configurations, plans were automatically generated for a group of test patients (patients not used for RPM training). These investigations were done for different-sized training sets and folds. All resulting configurations were of high quality and consistent with the user preferences, even for a group of only 9 training patients. Automating the configuration can greatly reduce the time required in manual configuration strategies.

Chapter 5 investigates the automatic RPM configuration for automated multi-objective planning of oropharynx cancer patients with 22 objectives. Clinically relevant nonconvex criteria, not used for plan generation, including dose-volume points and normal tissue complication probabilities (NTCPs), were introduced to improve the intuitiveness of iterative steering towards a desired RPM configuration. This was tested on a database including planning CT scans of 105 oropharynx cancer patients with corresponding Pareto-optimal plans generated with the 2pec method. Automatically obtained RPM configurations, based on training with only 20 patients, resulted in Pareto-optimal IMRT plans with overall sim-

ilar or better quality than that of the Pareto-optimal database plans. While configuration of the 2p ϵ c method is complex and labour intensive, configuring the RPM is automatic, intuitive, and requires much less hands-on time.

In chapter 6, an automatic re-planning strategy is described for online-adaptive IMPT of cervical cancer patients. IMPT is sensitive to daily variations in anatomy, which can be large for cervical cancer patients. Therefore, a strategy was developed for online plan selection from a patient-specific library of plans, generated using a target motion model. Each day, the best-fitting plan is selected and automatically adapted to the daily anatomy. The automatic plan adaptation restores spot positions of the planning CT by adapting the proton energies, adds spots to sufficiently cover the tumour, and then applies an RPM optimisation of the spot weights to generate a Pareto-optimal IMPT plan. This study included 23 repeat-CT scans of six patients, whose primary tumour and nodal CTV were prescribed a dose of 45 Gy(RBE, relative biological effectiveness). Plan adaptation resulted in clinically acceptable target coverage ($V_{95\%} \geq 95\%$ and $V_{107\%} \leq 2\%$) for 37/46 plans using a 1-plan-library (one prior plan based on all motion), and for 41/46 plans using a 2-plan-library (two prior plans based on part of the motion) in case of a single iteration of spot additions. With two iterations, only 3 plans showed a $V_{107\%} > 2\%$ for the 1-plan-library approach, while all 46 plans were clinically acceptable for the 2-plan-library approach. Both libraries showed similar OAR sparing.

Chapter 7 describes an analytical canonical form to efficiently compute first and second order derivatives for certain classes of additively and multiplicatively separable functions. These derivatives are used in the interior-point method, implemented in Erasmus-iCycle, for solving medium-scale nonlinear optimisation problems in radiotherapy treatment planning. The canonical form was hard-coded for commonly used nonlinear radiotherapy functions (e.g. the *logarithmic tumour control probability* or the *generalised equivalent uniform dose*). For computing the second order derivatives, the major cost is a symmetric matrix multiplication with itself which allows for parallel computing on multi-core processors. Computation times for the derivatives with the canonical form were shorter than those resulting from automatic differentiation.

Chapter 8 contains a general discussion on the LRPM and RPM, including directions for future research, and steps towards clinical introduction.

Samenvatting

De focus van dit proefschrift ligt op het introduceren van de LRPM (“lexicographic reference point method”) en de RPM (“reference point method”) voor het snel en automatisch genereren van Pareto-optimale en klinisch relevante IMRT (“intensity-modulated radiotherapy”) bestralingsplannen, het automatisch configureren van de RPM gebruikmakende van reeds bestaande bestralingsplannen, en het onderzoeken van de automatische configuratie van de RPM voor dagelijks adaptief herplannen in IMPT (“intensity-modulated proton therapy”).

Na een korte introductie van de radiotherapie in hoofdstuk 1 wordt het proces van zowel handmatige en automatische optimalisatie van bestralingsplannen in de radiotherapie toegelicht. De focus ligt op Erasmus-iCycle, wat het klinisch gebruikte systeem is voor het automatisch genereren van een behandelplan door bundelhoeken en intensiteitsprofielen te optimaliseren op basis van meerdere doelfuncties. De kern van dit systeem is de $2p\epsilon c$ (“2-phase ϵ -constraint”) methode voor het automatisch genereren van Pareto-optimale intensiteitsprofielen (gegeven een aantal bundelhoeken) op basis van meerdere doelfuncties. De beoogde verbeteringen van de LRPM en de RPM ten opzichte van de $2p\epsilon c$ methode zijn (1) het verminderen van de benodigde rekentijd, en (2) het verminderen van de werklast voor het configureren van het planningsalgoritme, beide onder de voorwaarde dat de plankwaliteit niet verslechtert.

Hoofdstuk 2 introduceert de LRPM als optimalisatiemethode op basis van meerdere doelfuncties voor het snel en automatisch genereren van Pareto-optimale bestralingsplannen. Er hoeft slechts één optimalisatieprobleem opgelost te worden voor het toepassen van de LRPM, terwijl er een aantal opeenvolgende optimalisatieproblemen opgelost moeten worden voor de klinisch gebruikte $2p\epsilon c$ methode. Dit aantal groeit lineair met het aantal gebruikte doelfuncties voor het genereren van het bestralingsplan. De LRPM is een lexicografische uitbreiding van de reeds bestaande RPM die gebruikt kan worden om de werking van de $2p\epsilon c$ methode na te bootsen, d.w.z. benadering van de lexicografische optimalisa-

tie van de doelfuncties met bijbehorende doelwaarden. Ook wordt uitgelegd hoe de LRPM gebruikt kan worden om a priori gedefinieerde grenzen te stellen op de partiële trade-offs. De LRPM is getest door bestralingsplannen automatisch te genereren voor 30 willekeurig geselecteerde prostaatkanker patiënten. Dit resulteerde voor iedere patiënt in een bestralingsplan met gelijkwaardige dosimetrische kwaliteit als die van het bestralingsplan dat gegenereerd was met de 2p ϵ c methode. Relatief gezien was de LRPM 10 keer sneller, zodat de gemiddelde rekentijd afnam van 12.4 naar 1.2 minuten.

In hoofdstuk 3 wordt de LRPM toegepast om automatisch bestralingsplannen te genereren voor 15 hoofd-halskanker patiënten (bilaterale bestraling). For iedere patiënt is zowel de LRPM als de 2p ϵ c methode gebruikt voor het automatisch genereren van een Pareto-optimaal bestralingsplan. Al deze automatisch gegenereerde plannen waren klinisch relevant. In vergelijking met de 2p ϵ c methode, gaf het plan gegenereerd met de LRPM voor enkele patiënten veel gunstigere doelfunctiewaarden (grote dosisreductie voor sommige risico-organen, dit zijn gezonde organen die zoveel mogelijk ontzien moeten worden van de straling) wat ten koste ging van kleine verslechtingen van andere doelfunctiewaarden. De LRPM was dus in staat betere globale trade-offs tussen doelfuncties te verkrijgen voor enkele patiënten. Voor de andere patiënten hadden de plannen een gelijkwaardige kwaliteit. Door de vele doelfuncties (maximaal 22) in de optimalisaties nam de gemiddelde rekentijd met de LRPM af van 209.2 naar 9.5 minuten, wat relatief gezien 22 keer sneller is.

Hoofdstuk 4 introduceert een algoritme voor het automatisch configureren van de RPM voor automatische optimalisatie van bestralingsplannen op basis van meerdere doelfuncties. De configuratie van de 2p ϵ c methode voor het automatisch genereren van bestralingsplannen met de huidige procedure is een arbeids- en tijdsintensieve taak die herhaald dient te worden voor iedere tumorgroep/planningsprotocol. De automatische configuratie heeft als doel om automatisch een hoogwaardig bestralingsplan te genereren voor iedere patiënt, zonder daarna het plan nog handmatig aan te passen. Het algoritme voor de automatische configuratie ontwikkeld voor de RPM is gebaseerd op karakteristieken van een aantal bestralingsplannen gebruikt voor training van het algoritme. In dit artikel is een database van prostaatkanker patiënten gebruikt bestaande uit 287 planning CT-scans met bijbehorende Pareto-optimale bestralingsplannen die gegenereerd zijn met de 2p ϵ c methode. Met de bestralingsplannen gebruikt voor training wordt er automatisch een RPM configuratie bepaald door middel van een iteratief proces waarin gebruiker-specifieke voorkeuren nageleefd worden die toegestane toleranties beschrijven voor de verschillen tussen de voor training gebruikte en RPM gegenereerde bestralingsplannen. Deze voorkeuren kunnen zijn om gelijke dosiscoverage van de PTV ("planning target volume", het tumorvolume dat bestraald dient te worden rekening houdend met onzekerheden, vaak in de vorm van marges rond de tumor, bij het genereren en/of afgeven van het bestralingsplan) te verkrijgen, minder dosis in het rectum te verkrijgen, en een iets slechtere conformaliteit van de

dosis toe te laten. Om de kwaliteit van een RPM-configuraties te onderzoeken, worden deze gebruikt om automatisch bestralingsplannen te genereren voor een groep test patiënten (deze groep patiënten is niet gebruikt om het algoritme te trainen). Dit werd onderzocht voor verschillende grootten trainingsgroepen en “folds” (andere trainingsgroep van eenzelfde grootte). Alle verkregen configuraties waren van hoogwaardige kwaliteit en consistent met de gebruiker-specifieke voorkeuren, zelfs voor een trainingsgrootte van slechts 9 patiënten. Het automatiseren van de configuratie kan de benodigde tijd voor handmatige configuratiestrategieën flink verminderen.

Hoofdstuk 5 onderzoekt of de automatische configuratie voor de RPM gebruikt kan worden voor automatische optimalisatie van bestralingsplannen voor orofarynxkanker patiënten op basis van 22 doelfuncties. Het gebruik van klinisch relevante niet-convexe evaluatiefuncties, zoals dosis-volume punten en NTCPs (“normal tissue complication probabilities”, functie voor waarschijnlijkheid op orgaan-gerelateerde complicaties), zijn toegevoegd om het iteratief sturen naar een gewenste configuratie voor de RPM intuïtiever te maken. Dit werd getest op een database bestaande uit planning CT-scans van 105 orofarynxkanker patiënten met bijbehorende Pareto-optimale bestralingsplannen die waren gegenereerd met de 2p ϵ c methode. Automatisch verkregen RPM-configuraties, gebaseerd op trainen met slechts 20 patiënten, resulteerde in Pareto-optimale IMRT-bestralingsplannen met in het algemeen gelijke of betere kwaliteit dan die van de Pareto-optimale bestralingsplannen in de database. Terwijl de configuratie van de 2p ϵ c methode complex en arbeidsintensief is, is het configureren van de RPM automatisch, intuïtief en vergt veel minder interactie met de gebruiker.

In hoofdstuk 6 is een automatische online-adaptieve herplanningsstrategie beschreven voor IMPT van cervixkanker patiënten. Dagelijkse variaties van de anatomie, die groot kunnen zijn voor cervixkanker patiënten, hebben beduidend effect op een IMPT-bestralingsplan. Om die reden is er een strategie ontwikkeld voor online-selectie van het bestralingsplan uit een patiënt-specifieke collectie van vooraf beschikbare bestralingsplannen, gegenereerd met een bewegingsmodel voor de tumor. Iedere dag wordt het best geschikte plan geselecteerd en aangepast aan de hand van de dagelijkse anatomie. De automatische adaptatie van het bestralingsplan herstelt de spotposities van de planning CT-scan over door de proton energieën aan te passen, voegt voldoende spots toe aan het tumorweefsel, en past een RPM-optimalisatie van de spotgewichten toe om een Pareto-optimaal IMPT-bestralingsplan te genereren. In dit onderzoek zijn een totaal van 23 herhaalde CT-scans gebruikt behorende bij zes patiënten voor wie de voorgeschreven dosis voor zowel de primaire tumor als de nodale CTV (“clinical target volume”, het daadwerkelijke tumorvolume plus een kleine marge om rekening te houden met onzekerheden afkomstig van de gemaakte beeldopnamen) 45 Gy(RBE, relatieve biologische effectiviteit) was. Adaptatie van de bestralingsplannen leidde tot klinisch acceptabele tumordosis ($V_{95\%} \geq 95\%$

en $V_{107\%} \leq 2\%$) voor 37 van de 46 bestralingsplannen waarbij 1 bestralingsplan uit een patiënt-specifieke collectie (1 bestralingsplan gebaseerd op alle bewegingen) gebruikt werd, en voor 41 van de 46 bestralingsplannen waarbij 2 bestralingsplannen uit een patiënt-specifieke collectie (2 bestralingsplannen gebaseerd op alle bewegingen) gebruikt werden, en wanneer er een enkele iteratie van spottoevoegingen werd toegepast. Met twee iteraties hadden slechts drie bestralingsplannen een $V_{107\%} > 2\%$ waarbij gebruik gemaakt werd van 1 bestralingsplan uit een patiënt-specifieke collectie terwijl alle 46 bestralingsplannen klinisch acceptabel waren waarbij gebruikt gemaakt werd van 2 bestralingsplannen uit een patiënt-specifieke collectie. Beide van deze strategieën leidden tot vergelijkbare dosissen voor de risico-organen.

Hoofdstuk 7 beschrijft een analytische canoniieke vorm voor het efficiënt berekenen van de eerste en tweede orde afgeleiden voor zekere klassen van additief en multiplicatief separabele functies. Deze afgeleiden worden gebruikt in de inwendige puntmethode geïmplementeerd in Erasmus-iCycle om middelgrote optimalisatieproblemen in de radiotherapie op te lossen. De canoniieke vorm is expliciet gecodeerd voor veelgebruikte niet-lineaire functies in de radiotherapie (bijvoorbeeld de *waarschijnlijkheid op tumor controle* of het *gegeneraliseerde gemiddelde*). Voor het uitrekenen van de tweede orde afgeleiden zit de voornaamste tijd in een symmetrische matrixvermenigvuldiging met zichzelf wat goed te paralleliseren is op processoren met meerdere kernen. Rekeningtijden voor de afgeleiden met de canoniieke vorm waren korter dan die behorende bij automatische differentiatie.

Hoofdstuk 8 bevat een algemene discussie over de LRPM en de RPM, inclusief suggesties voor verder onderzoek en benodigde stappen voor klinische introductie.

PhD portfolio

Name PhD student: Rens van Haveren
Erasmus MC Department: Radiation Oncology
Research School: Molecular Medicine
PhD period: 2014 – 2021
Promotor: Prof.dr. B.J.M. Heijmen
Copromotor: Dr.ir. S. Breedveld

1. PhD training	Year
General courses	
Biomedical scientific writing	2016
Scientific integrity	2018
International conferences	
ICCR, London, England (oral presentation)	2016
ESTRO, Vienna, Austria (poster presentation)	2017
AAPM, Nashville, USA (oral presentation)	2018
ESTRO, Milan, Italy (poster presentation)	2019
ICCR, Montreal, Canada (oral presentation)	2019
National conferences	
NVKF, Zeist (oral presentation)	2016
RKF Scientific Project Day, Utrecht (oral presentation)	2017
RKF Scientific Project Day, Rotterdam (oral presentation)	2018
Invited speaker	
MCDM, Hamburg, Germany (oral presentation)	2015
INFORMS Healthcare, Rotterdam, The Netherlands (oral presentation)	2017

In-house oral presentations

R&D Meeting	2014, 2016, 2019
R&D Work Discussion	2015 (2x), 2016, 2017, 2018
Journal Club	2016, 2017, 2018
Research Day	2017
Refereeravond	2016, 2017, 2018

Other

Peer reviewed 5 papers for 3 journals

2. Teaching	Year
Supervising Bachelor students	
Larissa Scholte, Delft University of Technology	2015
Ghiline van Furth, Delft University of Technology	2016
Cynthia Slotboom, Delft University of Technology	2016
3. Distinctions and awards	Year
Best poster in Physics track, phiRO award, ESTRO, Vienna, Austria	2017

Curriculum Vitae

Rens van Haveren was born on 27 January, 1990 in Korendijk, The Netherlands. He completed high school at Helinium, Hellevoetsluis, in 2008 after which he started his study Applied Mathematics at the Delft University of Technology. During the Master's, he did his internship followed by his master thesis graduation project at the Department of Radiation Oncology of the Erasmus Medical Center in Rotterdam on the subject of automated radiotherapy treatment planning. He graduated in 2014 with honours.

Rens started his PhD at the Erasmus Medical Center in Rotterdam in 2014, and continued his research on automated radiation treatment planning. Under supervision of Ben Heijmen and Sebastiaan Breedveld, an automatic procedure to configure a multi-objective optimisation algorithm for fast automated treatment planning was developed, leading to this thesis.

Since September 2019, Rens works as a scientific researcher at Netherlands Forensic Institute in The Hague.

