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

van Haveren, Rens; Ogryczak, Włodzimierz; Verduijn, Gerda M.; Keijzer, Marleen; Heijmen, Ben; Breedveld, Sebastiaan

DOI
10.1088/1361-6560/62/11/4318

Publication date
2017

Document Version
Final published version

Published in
Physics in Medicine and Biology

Citation (APA)
van Haveren, R., Ogryczak, W., Verduijn, G. M., Keijzer, M., Heijmen, B., & Breedveld, S. (2017). 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. https://doi.org/10.1088/1361-6560/62/11/4318

PAPER

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

To cite this article: Rens van Haveren *et al* 2017 *Phys. Med. Biol.* **62** 4318

IOP Publishing | Institute of Physics and Engineering in Medicine Physics in Medicine & Biology

Phys. Med. Biol. **62** (2017) 4318–4332 https://doi.org/10.1088/1361-6560/62/11/4318

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

Rens van Haveren[1], Włodzimierz Ogryczak[2], Gerda M Verduijn[1], Marleen Keijzer[3], Ben J M Heijmen[1] and Sebastiaan Breedveld[1]

[1] Department of Radiation Oncology, Erasmus MC—Cancer Institute, PO Box 2040, 3000 CA Rotterdam, The Netherlands

[2] Warsaw University of Technology, Institute of Control and Computation Engineering, 00-665 Warsaw, Poland

[3] Delft University of Technology, Delft Institute of Applied Mathematics, PO Box 5031, 2600 GA Delft, The Netherlands

E-mail: r.vanhaveren@erasmusmc.nl

Received 9 August 2016, revised 20 October 2016
Accepted for publication 16 November 2016
Published 5 May 2017

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 (2pϵc) method. For each patient, the output of Erasmus-iCycle is a clinically favourable, Pareto optimal plan. The 2pϵc 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 2pϵc 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 22 relative to the 2pϵc method.

Keywords: automated radiotherapy treatment planning, prioritised multi-objective optimisation, multi-criteria, lexicographic reference point method, Erasmus-iCycle, head and neck cancer, IMRT

(Some figures may appear in colour only in the online journal)

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, 2009, Breedveld *et al* 2012, Jee *et al* 2007, Wilkens *et al* 2007, Haveren *et al* 2017), 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* (2pεc) 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 orientations, e.g. for Cyberknife (Rossi *et al* 2012, 2015), and intensity modulated proton therapy (IMPT) plans (Water *et al* 2013).

The 2pϵc 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 2pϵc 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 (Haveren *et al* 2017), as an alternative to the 2pϵc method. Similar to the 2pϵc method, input parameters are uniform for all patients with the same tumour type. In contrast to the 2pϵc 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) and 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 2pϵc 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 2pϵc method will be compared both regarding quality and computation time.

In section 2, descriptions of both the 2pϵc method and LRPM are provided with their configurations for HN cancer patients. In section 3, we analyse automatically generated plans for 15 HN cancer patients. Sections 4 and 5 discuss our findings, and conclude the paper.

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 $i \in \{1, 2, \ldots, n\}$ need to be minimised while obeying the imposed constraints.

The 2pϵc method and LRPM, both in-house developed for automated prioritised optimisation of radiotherapy treatment plans are discussed in sections 2.1 and 2.2, respectively. Details on the study of generated HN plans are discussed in section 2.3.

2.1. 2-phase ϵ-constraint method

For this study, the automated treatment plan generation with the 2pϵc method is performed with the wish-list in table 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 1 shows that the first priority is to decrease the LTCP (Alber and Reemtsen 2007),

$$\mathrm{LTCP}(d;\alpha,D^p)=\frac{1}{M}\sum_{j=1}^{M}\exp\left[\alpha(D^p-d_j)\right]. \tag{1}$$

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 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 2pϵc 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 (indication 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 have 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.

Table 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 (1); gEUD = generalised equivalent uniform dose.

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

2.2. Lexicographic reference point method

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 2pϵc 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 2pϵc 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 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 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 Haveren *et al* (2017)), a strictly monotonic *reference path* through these reference points is made, see figure 1(a). This path may be nonlinear in general, but we consider the piecewise linear case as in figure 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 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 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 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 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 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 1(d): the square is the result of using the non-fuzzy indifference curves (no

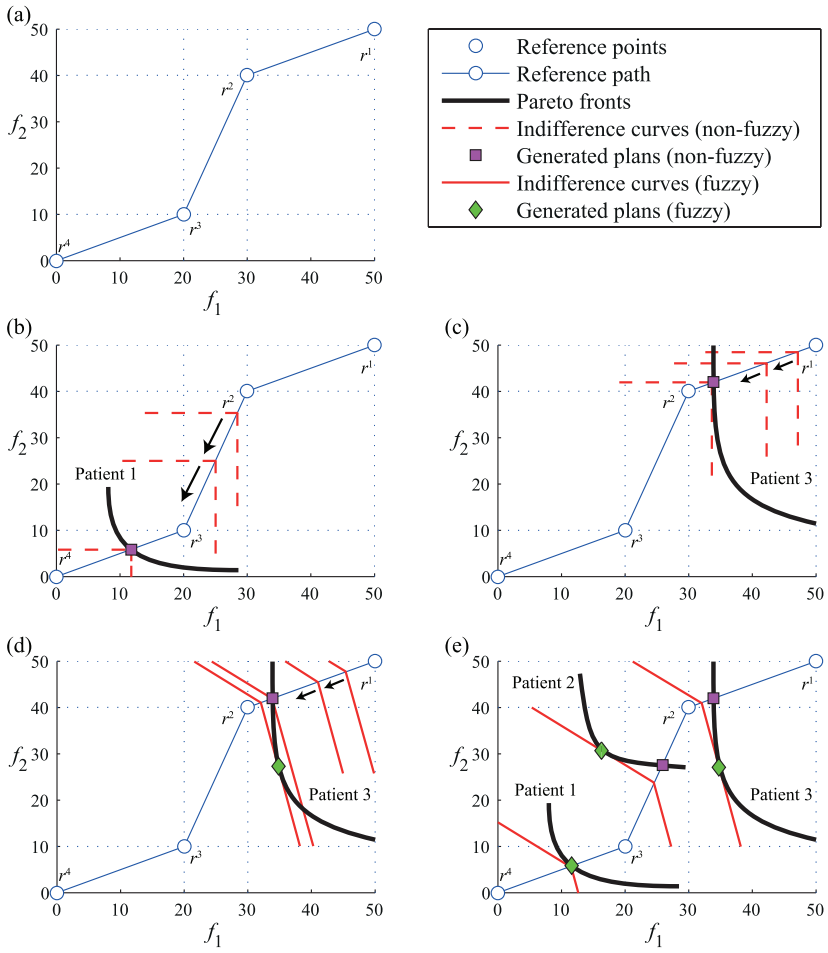

Figure 1. Principle LRPM for two objectives. (a) Reference points/path for lexicographic ordering, (b)–(c) plan selection with non-fuzzy indifference curves, (d) effect fuzzy indifference curve on plan selection, (e) plan selection for a group of patients using a uniformly configured LRPM.

trade-off parameters) and the diamond results from using fuzzy indifference curves. In some cases, e.g. for patient 1 in figure 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 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

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

Here, $r^j = (r_1^j, \ldots, r_n^j)$ 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 1(a). The ρ_i are the trade-off parameters used to bend the indifference curves, see figure 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 Haveren *et al* (2017).

2.2.2. Technical issues. There are two practical issues with applying the LRPM: the sufficient parameter values in the wish-list (table 1) and numerical issues for the LTCP (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

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

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

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

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. LLTCP = ln(LTCP): an equivalent convex objective which has a one-to-one correspondence with the (L)TCP.

2.3. Study setup

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 2pϵc method to automatically generate

a treatment plan. All plans were generated using a 23 equi-angular coplanar beam setup 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. 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 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 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 OARs that allow large favourable gains (without large degradations for other OARs) differ per patient (see figure 2) i.e. the LRPM 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, the distributions of differences in selected plan parameters are sketched using boxplots. Most medians are positive, and thus in favour of the LRPM. For the differences in maximum doses of the spinal cord and brainstem, the medians are even well above zero. The negative medians (in favour of the 2pϵc method) are only slightly below zero. From the 13 observed outliers in figure 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 2pϵc 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 2pϵc 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 2pϵc method, an observed speed-up factor of the mean computation times of 22 for the LRPM relative to the 2pϵc 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.

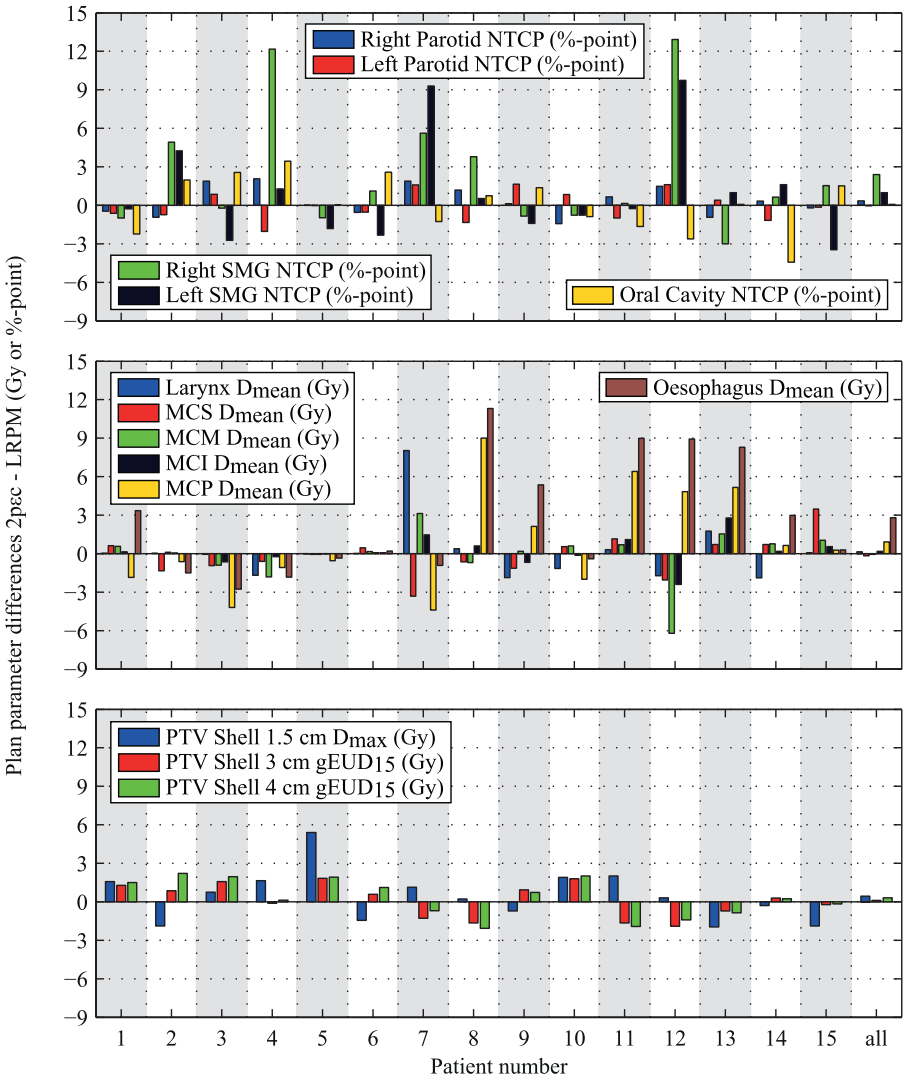

Figure 2. Plan differences (2pεc-LRPM) per patient for the most relevant treatment objectives and evaluation criteria. Positive values are in favour of the LRPM, whereas negative values indicate that the 2pεc method performed better. All generated plans are Pareto optimal. SMG = submandibular gland,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

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 2pϵc 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 2pϵc 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 2pϵc method can 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 (1) were replaced by equivalent LLTCP functions (section 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 2pϵc 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 2pϵc method needs to solve a sequence of optimisation problems to determine a clinically relevant Pareto optimal plan. On average, the 2pϵc method required solving 28 (range [23 32]) optimisation problems to generate a plan. The 2pϵc 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 2pϵc 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 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 2pϵc 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 2pϵc 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 (Haveren *et al* 2017), the latter were determined in an iterative manner, where differences with the 2pϵc 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 2.2.2. We have investigated this for generation of single boost prostate plans (Haveren *et al* 2017), and no difficulties were

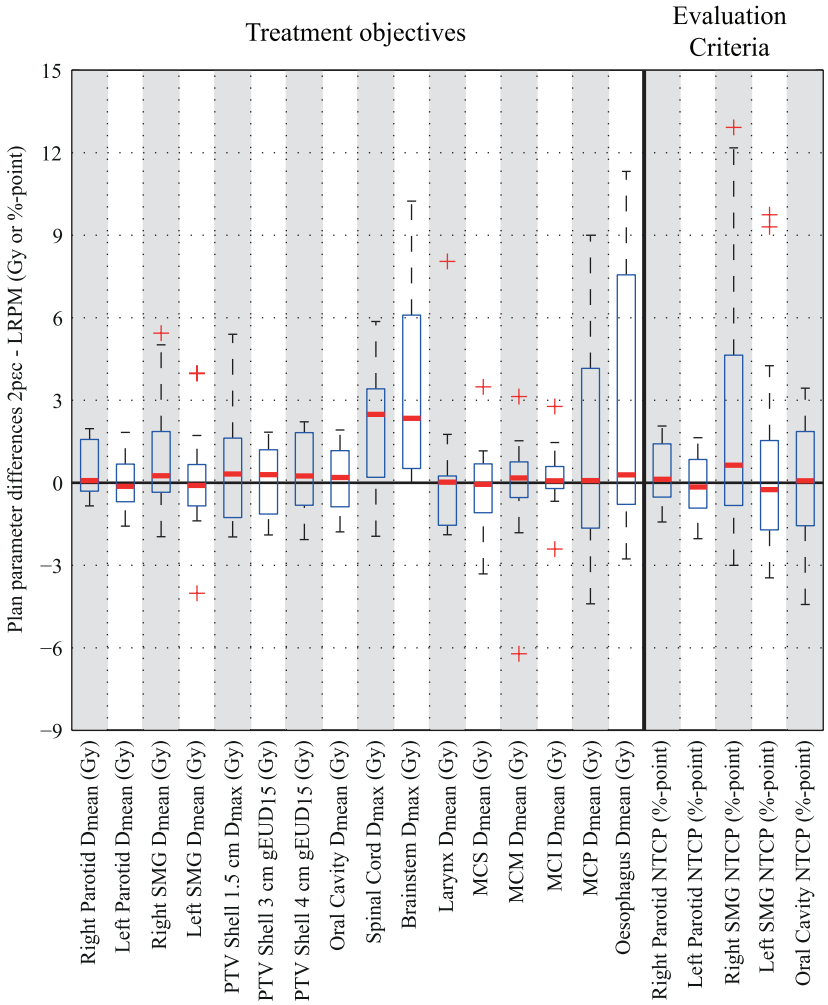

Figure 3. Boxplots of the plan differences (2pεc-LRPM) for treatment objectives and evaluation criteria. The bottom and top edges of the boxes are the first and third quartiles respectively, the horizontal solid lines represent the medians, the whiskers are at 1.5 times the interquartile range and the pluses represent outliers. SMG = submandibular gland, MCS = musculus constrictor superior, MCM = musculus constrictor medius, MCI = musculus constrictor inferior, MCP = musculus constrictor cricopharyngeus.

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 2pϵc method) for some patients. Both multi-objective methods have different unrelated mechanisms for the trade-offs. The 2pϵc 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), a clear distinction should be made between the plan comparison for individual cases (figure 2) and the distributions of plan differences for all patients (figure 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 (Haveren *et al* 2017). In this study, 30 randomly selected prostate cancer patients were considered. For each patient, treatment plans generated with the 2pϵc 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 2pϵc method were of high clinical quality. For these prostate cancer patients, both the 2pϵc 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 2pϵc method.

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 2pϵc 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 22 relative to the 2pϵc 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.

Acknowledgments

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. The Erasmus MC—Cancer Institute has research collaborations with Elekta AB, Stockholm, Sweden and Accuray Inc., Sunnyvale, USA.

References

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

Boyd S and Vandenberghe L 2004 *Convex Optimization* (Cambridge: Cambridge University Press)

Breedveld S, Van den Berg B and Heijmen B 2017 An interior-point implementation developed and tuned for radiation therapy treatment planning submitted

Breedveld S and Heijmen B 2017 Data for TROTS—The radiotherapy optimisation test set *Data in Brief* **17** 143–9

Breedveld S, Storchi P and Heijmen B 2009 The equivalence of multi-criteria methods for radiotherapy plan optimization *Phys. Med. Biol.* **54** 7199–209

Breedveld S, Storchi P, Keijzer M, Heemink A and Heijmen B 2007 A novel approach to multi-criteria inverse planning for IMRT *Phys. Med. Biol.* **52** 6339–53

Breedveld S, Storchi P, Voet P and Heijmen B 2012 iCycle: integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans *Med. Phys.* **39** 951–63

Craft D, Halabi T, Shih H and Bortfeld T 2006 Approximating convex Pareto surfaces in multiobjective radiotherapy planning *Med. Phys.* **33** 3399–407

Della Gala G, Dirkx M, Hoekstra N, Fransen D, Van de Pol M, Heijmen B and Petit S 2016 Automated treatment plan generation for advanced stage NSCLC patients *Radiother. Oncol.* **119** S124

Dijkema T, Raaijmakers C, Ten Haken R, Roesink J, Braam P, Houweling A, Moerland M, Eisbruch A and Terhaard C 2010 Parotid gland function after radiotherapy: the combined Michigan and Utrecht experience *Int. J. Radiat. Oncol. Biol. Phys.* **78** 449–53

Forsgren A, Gill P and Wright M 2002 Interior methods for nonlinear optimization *SIAM Rev.* **44** 525–97

Granat J and Makowski M 2000 Interactive specification and analysis of aspiration-based preferences *Eur. J. Oper. Res.* **122** 469–85

Haveren R V, Breedveld S, Keijzer M, Voet P, Heijmen B and Ogryczak W 2017 Lexicographic extension of the reference point method applied in radiation therapy treatment planning submitted

Heijkoop S, Langerak T, Quint S, Bondar L, Mens J, Heijmen B and Hoogeman M 2014 Clinical implementation of an online adaptive plan-of-the-day protocol for nonrigid motion management in locally advanced cervical cancer IMRT *Int. J. Radiat. Oncol. Biol. Phys.* **90** 673–9

Jee K, McShan D and Fraass B 2007 Lexicographic ordering: intuitive multicriteria optimization for IMRT *Phys. Med. Biol.* **52** 1845–61

Korhonen P and Wallenius J 1988 A Pareto race *Nav. Res. Log.* **35** 615–23

Leinders S, Breedveld B, Méndez Romero A, Schaart D, Seppenwoolde Y and Heijmen B 2013 Adaptive liver stereotactic body radiation therapy: automated daily plan reoptimization prevents dose delivery degradation caused by anatomy deformations *Int. J. Radiat. Oncol. Biol. Phys.* **87** 1016–21

Long T, Matuszak M, Feng M, Fraass B, Ten Haken R and Romeijn H 2012 Sensitivity analysis for lexicographic ordering in radiation therapy treatment planning *Med. Phys.* **39** 3445–55

Mehrotra S 1992 On the implementation of a primal-dual interior point method *SIAM J. Opt.* **2** 575–601

Miettinen K 1999 *Nonlinear Multiobjective Optimization* (Int. Series in Operations Research and Management Science vol 12) (Dordrecht: Kluwer)

Miettinen K, Ruiz F and Wierzbicki A 2008 *Multiobjective Optimization* (Lecture Notes in Computer Science) (Berlin: Springer) p 27

Murdoch-Kinch C, Kim H, Vineberg K, Ship J and Eisbruch A 2008 Dose-effect relationships for the submandibular salivary glands and implications for their sparing by intensity modulated radiotherapy *Int. J. Radiat. Oncol. Biol. Phys.* **72** 373–82

Niemierko A 1997 Reporting and analyzing dose distribution: a concept of equivalent uniform dose *Med. Phys.* **24** 103–10

Nocedal J and Wright S 2006 *Numerical Optimization* (Berlin: Springer)

Ogryczak W and Kozłowski B 2011 Reference point method with importance weighted ordered partial achievements *TOP* **19** 380–401

Petit S, Seco J and Kooy H 2013 Increasing maximum tumor dose to manage range uncertainties in IMPT treatment planning *Phys. Med. Biol.* **58** 7329–41

Rossi L, Breedveld S, Aluwini S and Heijmen B 2015 Noncoplanar beam angle class solutions to replace time-consuming patient-specific beam angle optimization in robotic prostate stereotactic body radiation therapy *Int. J. Radiat. Oncol. Biol. Phys.* **92** 762–70

Rossi L, Breedveld S, Heijmen B, Voet P, Lanconelli N and Aluwini S 2012 On the beam direction search space in computerized non-coplanar beam angle optimization for IMRT—prostate SBRT *Phys. Med. Biol.* **57** 5441–58

Sharfo A, Voet P, Breedveld S, Mens J, Hoogeman M and Heijmen B 2015 Comparison of VMAT and IMRT stategies for cervical cancer patients using automated planning *Radiother. Oncol.* **114** 395–401

Voet P, Dirkx M, Breedveld S, Al-Mamgani A, Incrocci L and Heijmen B 2014 Fully automated volumetric modulated arc therapy plan generation for prostate cancer patients *Int. J. Radiat. Oncol. Biol. Phys.* **88** 1175–9

Voet P, Dirkx M, Breedveld S, Fransen D, Levendag P and Heijmen B 2013a Toward fully automated multicriterial plan generation: a prospective clinical study *Int. J. Radiat. Oncol. Biol. Phys.* **85** 866–72

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

Water S V d, Kraan A, Breedveld S, Schillemans W, Teguh D, Kooy H, Madden T, Heijmen B and Hoogeman M 2013 Improved efficiency of multi-criteria IMPT treatment planning using iterative resampling of randomly placed pencil beams *Phys. Med. Biol.* **58** 6969–83

Wierzbicki A 1982 A mathematical basis for satisficing decision making *Math. Mod.* **3** 391–405

Wierzbicki A 1986 On the completeness and constructiveness of parametric characterizations to vector optimization problems *OR Spectrum* **8** 73–87

Wilkens J, Alaly J, Zakarian K, Thorstad W and Deasy J 2007 IMRT treatment planning based on prioritizing prescription goals *Phys. Med. Biol.* **52** 1675–92