



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

Rens van Haveren<sup>a)</sup>, Ben J. M. Heijmen, and Sebastiaan Breedveld

Department of Radiation Oncology, Erasmus MC, University Medical Center Rotterdam, 3015 GD, Rotterdam, The Netherlands

(Received 29 November 2019; revised 22 January 2020; accepted for publication 27 January 2020; published xx xxxx xxxx)

**Purpose:** In automated treatment planning, configuration of the underlying algorithm to generate high-quality plans for all patients of a particular tumor 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 favorable trade-offs for each patient. The automatic configuration of the RPM requires a set of computed tomography (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 two-phase  $\varepsilon$ -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. © 2020 The Authors. *Medical Physics* published by Wiley Periodicals, Inc. on behalf of American Association of Physicists in Medicine. [<https://doi.org/10.1002/mp.14073>]

Key words: automatic configuration, automated treatment planning, IMRT, oropharyngeal cancer, Pareto optimal, radiotherapy

## 1. INTRODUCTION

Generating high-quality intensity-modulated radiation therapy (IMRT) or volumetric modulated arc therapy (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, esophagus, spinal cord, and brainstem.

Several automated treatment planning approaches have been proposed in the literature.<sup>1,2</sup> This paper focuses on automated multi-objective fluence map optimization to generate a single Pareto optimal and clinically favorable treatment plan for each patient. Two algorithms for automated multi-objective optimization of Pareto optimal plans have been developed in our center: (a) the two-phase  $\varepsilon$ -constraint (2p $\varepsilon$ c) method<sup>3</sup> which is part of the clinically applied Erasmus-iCycle optimizer,<sup>4</sup> and (b) the fast and fuzzy lexicographic

reference point method<sup>5</sup> (LRPM). Previous studies<sup>6,7</sup> have demonstrated that the quality of plans generated with the 2p $\varepsilon$ c method is generally superior to that of manually generated plans. The main advantages of the LRPM over the 2p $\varepsilon$ c method are faster plan generation with average relative speed-up factors of 12 for prostate<sup>5</sup> and 22 for head and neck cancer,<sup>8</sup> and that trade-offs between all planning objectives are balanced simultaneously (LRPM) instead of pairwise (2p $\varepsilon$ c 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 tumor sites. Interactive (manual) tuning of the configuration is a time-consuming and workload-intensive procedure for both the 2p $\varepsilon$ c method (“wishlist” creation<sup>4,7</sup>) and the LRPM.<sup>5,8</sup> Recently, we proposed a new automatic procedure<sup>9</sup> to configure the reference point method<sup>10–12</sup> (RPM), a special case of the LRPM.<sup>5</sup> The procedure was successfully applied to prostate IMRT,<sup>9</sup> and adaptive prostate and cervix IMPT.<sup>13,14</sup>

© 2020 The Authors. *Medical Physics* published by Wiley Periodicals, Inc. on behalf of American Association of Physicists in Medicine. This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

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,<sup>9</sup> 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.

## 2. MATERIALS AND METHODS

### 2.A. 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 similar to our clinical protocol.<sup>15</sup> 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 equiangular 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 fluence map is automatically converted to a VMAT plan using Monaco (Elekta AB, Sweden). In this study 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<sub>ec</sub> method.<sup>3</sup> The applied configuration (wish-list) for plan generation with 2p<sub>ec</sub> method is presented in Table I. To achieve clinically acceptable coverage for both PTVs ( $V_{95\%} \geq 98\%$ ), the logarithmic tumor control probability<sup>16</sup> (LTCP) was used as the objective function. For the OARs, the focus was either on minimizing the mean dose (salivary glands, swallowing muscles, oral cavity, larynx, esophagus, and cochleas) or on minimizing the near maximum dose (spinal cord and brainstem) for which the generalized equivalent uniform dose<sup>17</sup> (gEUD) with a high parameter value was used. Steering on the dose conformality 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. The 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.

### 2.B. Automatic RPM configuration

The automatic RPM configuration procedure<sup>9</sup> applied in this paper is summarized in Fig. 1. For initialization, a

TABLE I. Wish-list used for generating the database plans with the 2p<sub>ec</sub> method. The down-arrows ( $\downarrow$ ) indicate that the objectives are to be minimized. Prescribed dose was  $D_{\text{high}} = 70$  Gy for the planning target volume (PTV) high, and  $D_{\text{low}} = 54.25$  Gy for the PTV low

| Volume             | Type                  | Limit (Gy)                      |         |            |   |
|--------------------|-----------------------|---------------------------------|---------|------------|---|
| <b>Constraints</b> |                       |                                 |         |            |   |
| PTV high           | $D_{\text{max}}$      | 74.9                            |         |            | (=107% of $D_{\text{high}}$ )               |
| PTV high           | $D_{\text{mean}}$     | 70.7                            |         |            | (=101% of $D_{\text{high}}$ )               |
| Spinal cord        | $D_{\text{max}}$      | 42                              |         |            | (=60% of $D_{\text{high}}$ )                |
| Brainstem          | $D_{\text{max}}$      | 49                              |         |            | (=70% of $D_{\text{high}}$ )                |
| PTV shell 0 mm     | $D_{\text{max}}$      | 70                              |         |            | (=100% of $D_{\text{high}}$ )               |
| PTV shell 30 mm    | $D_{\text{max}}$      | 35                              |         |            | (=50% of $D_{\text{high}}$ )                |
| Unspecified tissue | $D_{\text{max}}$      | 74.9                            |         |            | (=107% of $D_{\text{high}}$ )               |
| Priority           | Volume                | Type                            | Goal    | Sufficient | Parameters                                  |
| <b>Objectives</b>  |                       |                                 |         |            |   |
| 1                  | PTV high              | $\downarrow$ LTCP               | 0.5     | 0.5        | $D^p = D_{\text{high}}$ ,<br>$\alpha = 0.8$ |
| 2                  | PTV low               | $\downarrow$ LTCP               | 0.5     | 0.5        | $D^p = D_{\text{low}}$ ,<br>$\alpha = 0.8$  |
| 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        | $\downarrow$ gEUD <sub>10</sub> | 10 Gy   |            |   |
|                    | PTV shell 15 mm       | $\downarrow$ gEUD <sub>10</sub> | 10 Gy   |            |   |
| 8                  | Oral cavity/Larynx    | $\downarrow D_{\text{mean}}$    | 35 Gy   |            |   |
| 9                  | esophagus             | $\downarrow D_{\text{mean}}$    | 40 Gy   |            |   |
| 10                 | Spinal cord/brainstem | $\downarrow$ gEUD <sub>12</sub> | 25 Gy   |            |   |
| 11                 | PTV shell 40 mm       | $\downarrow$ gEUD <sub>8</sub>  | 5 Gy    |            |   |
|                    | PTV shell 50 mm       | $\downarrow$ gEUD <sub>8</sub>  | 5 Gy    |            |   |
| 12                 | External ring 20 mm   | $\downarrow D_{\text{max}}$     | 27.1 Gy |            |   |
| 13                 | Cochleas              | $\downarrow D_{\text{mean}}$    | 35 Gy   |            |   |

Abbreviations: gEUD<sub>r</sub> = generalized equivalent uniform dose with applied parameter  $r$ ; LTCP = logarithmic tumor control probability; MCI = musculus constrictor inferior; MCM = musculus constrictor medius; MCP = musculus constrictor cricopharyngeus; MCS = musculus constrictor superior; PTV = planning target volume; SMG = submandibular gland.

fraction of the patients in the database (Section 2.A) was randomly selected for training (the remaining test patients were used to validate the configuration). Then, relevant data were acquired from the training plans (Section 2.B.1) to create the final RPM configuration. (Sections 2.B.2 and 2.B.3).

#### 2.B.1. Data acquisition from training patients

The constraints and objectives used for plan generation with the 2p<sub>ec</sub> method (Table I) 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 up to 22 objectives used in the fluence map optimization with the 2p $\epsilon$ c method (Table I).

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 optimization with the 2p $\epsilon$ c method. These Lagrange multipliers can be found as a byproduct of the optimization.<sup>3,9</sup>

### 2.B.2. Automatic configuration procedure

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

$$\text{minimize}_{x \in X} \left\{ \max_{i \in [n]} [w_i f_i(x) + c_i] + \sum_{i \in [n]} \rho_i (w_i f_i(x) + c_i) \right\}. \quad (1)$$

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  $\rho_1, \dots, \rho_n$  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 2.B.1) was used to generate an initial RPM configuration (technical details<sup>9</sup>). 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 2.B.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  $\rho_1, \dots, \rho_n$ . The general rule is that  $\rho_i$  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<sup>9</sup>). 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.

### 2.B.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: (a) preferences regarding a minimum/maximum allowed value for a plan parameter in the RPM-generated plans; (b) 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 II, 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 II, which indicates that the median value of all differences (database – RPM) in the parotid gland NTCP (for both left and right) 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,<sup>9</sup> 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 II 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 optimizations (compare with Table I). The first row in Table II 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<sup>18</sup> was applied for predicting xerostomia,

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

with  $D_{\text{mean}}$  being the mean dose in a salivary gland or the oral cavity. The second row in Table II 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%-points higher than those in the training plans.

The aim of the user preferences in Table II was to define an RPM configuration resulting in plans with: (a) sufficient

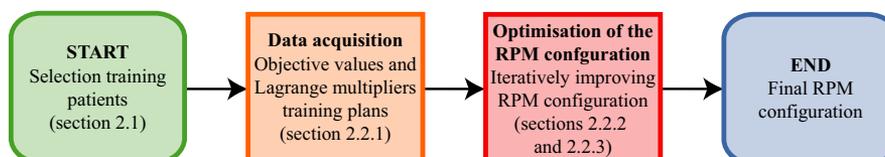


FIG. 1. Schematic overview of the automatic reference point method configuration procedure.

TABLE II. User preferences to create and evaluate an reference point method (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 $D_{\text{mean}}$   |
|  | 5th percentile | -2.5        |                                    |
| SMGs/oral cavity NTCP                        | Median         | 0           | SMGs/oral cavity $D_{\text{mean}}$ |
|  | 5th percentile | -4          |                                    |
| MCS $D_{\text{mean}}$                        | Median         | 0           |                                    |
|  | 5th percentile | -2          |                                    |
| MCP $D_{\text{mean}}$                        | Median         | 0           |                                    |
|  | 5th percentile | -2.5        |                                    |
| MCM/MCI $D_{\text{mean}}$                    | Median         | 0           |                                    |
|  | 5th percentile | -3          |                                    |
| Larynx/esophagus $D_{\text{mean}}$           | Median         | 0           |                                    |
|  | 5th percentile | -3          |                                    |
| Spinal cord/<br>brainstem gEUD <sub>12</sub> | Median         | -1          |                                    |
|  | 5th percentile | -3          |                                    |
| Cochleas $D_{\text{mean}}$                   | 1st quartile   | -5          |                                    |
| PTV shell 5 mm gEUD <sub>8</sub>             | Median         | -0.5        |                                    |
|  | 5th percentile | -3          |                                    |
| PTV shell 15 mm gEUD <sub>8</sub>            | Median         | -0.75       |                                    |
|  | 5th percentile | -3.25       |                                    |
| PTV shell 40 mm gEUD <sub>8</sub>            | Median         | -1.25       |                                    |
|  | 5th percentile | -3.75       |                                    |
| PTV shell 50 mm gEUD <sub>8</sub>            | Median         | -1.5        |                                    |
|  | 5th percentile | -4          |                                    |
| External ring 20 mm $D_{\text{max}}$         | Median         | -1.5        |                                    |
|  | 5th percentile | -5          |                                    |

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

target coverage for all patients ( $V_{95\%} \geq 98\%$ ); (b) overall reduced NTCP values in salivary glands and oral cavity and reduced mean doses in the swallowing muscles. If needed to accomplish (a) and (b), 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 unfavorable outliers for the RPM.

### 2.C. 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. Selection 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 II. To visualize the heterogeneity of the training folds with 20 plans, the plan parameters for the corresponding database plans are shown in Fig. 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 ( $P < 0.05$ ).

## 3. RESULTS

### 3.A. Target coverage

All database and RPM-generated plans had clinically acceptable target coverage, that is, 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 analyzing 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%.

### 3.B. 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 Fig. 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, esophagus, spinal cord, and brainstem showed overall better sparing for the RPM-generated plans at the cost of some deterioration in conformality measures. Differences observed in plan parameters between RPM-generated plans and database plans are in line with Table II, where preference is given to improve RPM-generated plans regarding organ sparing by allowing some deterioration in conformality.

Figure 3 also shows outliers for some plan parameters, often in favor of the RPM. Since the performance of the RPM configuration (obtained with training set fold 4, see Fig. 2) on test fold 4 (see Fig. 3) is according to the user preferences (Table II), particularly for the oral cavity NTCP, this fold is analyzed more in depth. For test fold 4, the differences in the most important plan parameters are shown in Fig. 4 for 15 plans with the most extreme outliers (both favorable and unfavorable for the RPM). As a reference, the last column in Fig. 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 Fig. 5. The median of

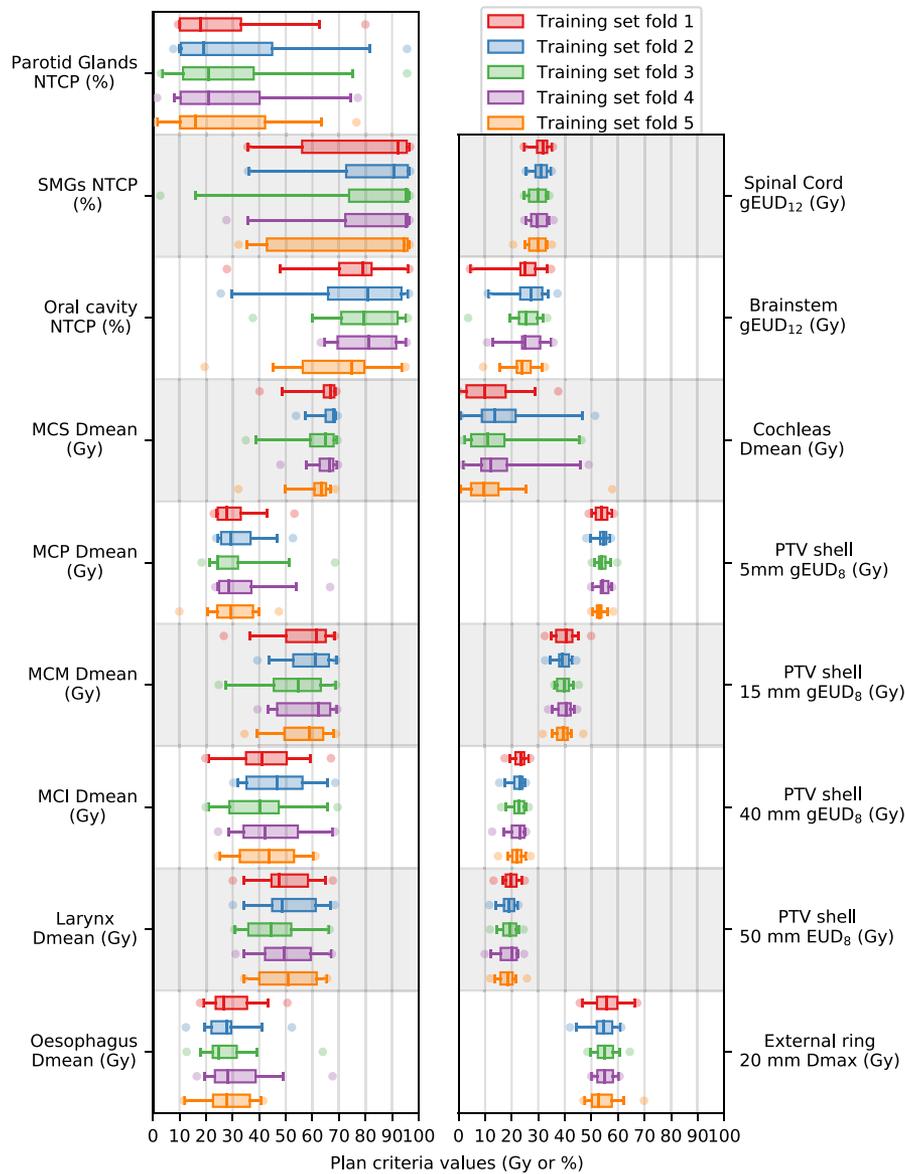


FIG. 2. Boxplots of the plan parameter values (Table II) for the database plans corresponding to the five different training folds, each with 20 training patients. Vertical thick lines within the boxes are medians, boxes are between the first and third quartile, whiskers are between the 2.5th and 97.5th percentile, circles are outliers.

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 supplementary material, results are presented for training with 35 and 50 patients. In general, it was found that increasing the number of training patients resulted in (a) slightly more consistent results for the test patients among the different folds with the same number of training patients, and (b) reduced severity of the outliers unfavorable for the RPM.

### 3.C. 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.<sup>19</sup> 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 h without any user interaction.

## 4. DISCUSSION

The purpose of this study was to further develop and explore a recently introduced automatic configuration procedure for the RPM,<sup>9</sup> 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. This study tested the automatic configuration for a heterogeneous group of

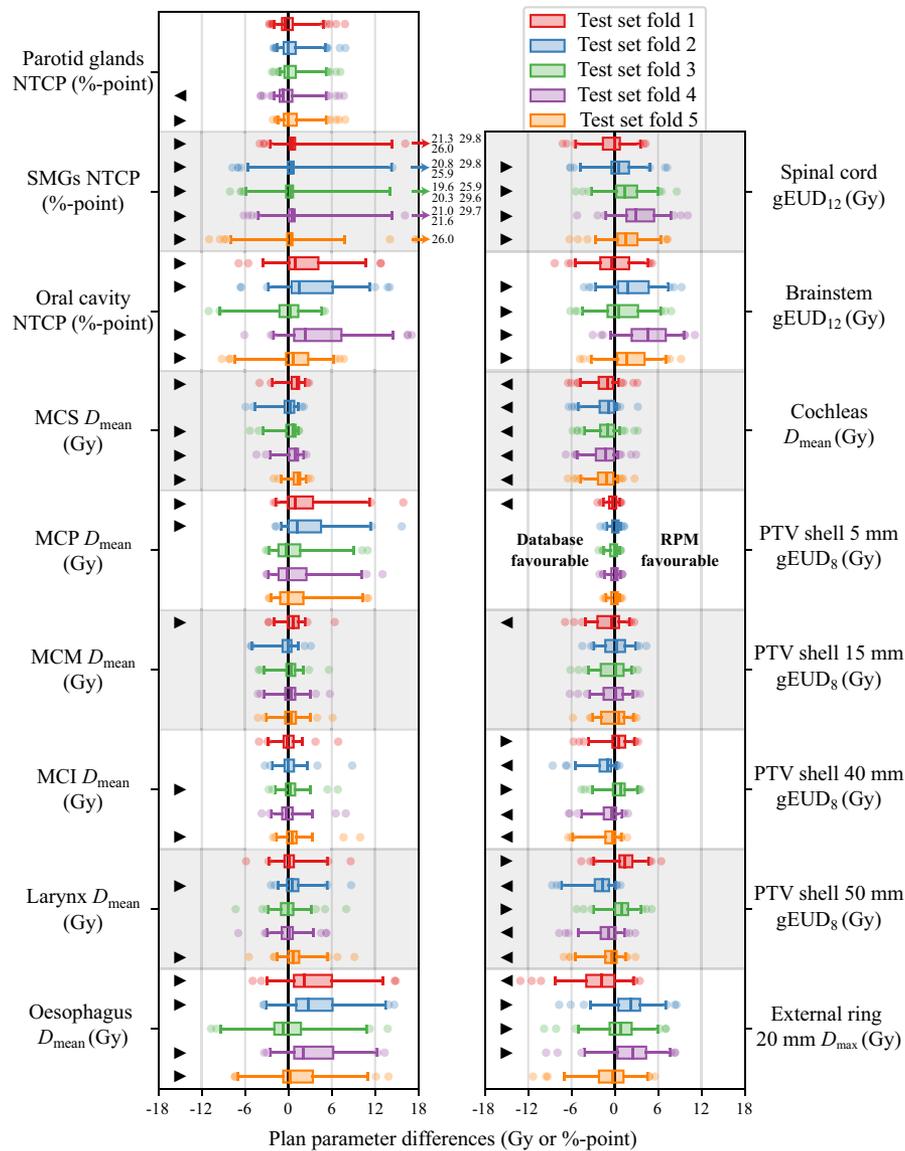


FIG. 3. Boxplots of the differences in plan parameter values (Table II) between database plans and reference point method (RPM)-generated plans for the five test folds corresponding to the five different RPM configurations with 20 training patients. Positive values are favorable for the RPM. Vertical 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 favor of database plans ( $\blacktriangleleft$ ) or RPM plans ( $\blacktriangleright$ ).

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,<sup>8</sup> the LRPM was used to automatically generate clinically favorable 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 favorable 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 favorable 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<sup>9</sup> were based exclusively on the convex planning objectives used in fluence map optimization, this paper describes how nonconvex criteria, such as dose-volume parameters or NTCPs (e.g., see Table II), can be included by coupling them to correlated convex objectives. This made the automatic configuration more intuitive and clinically relevant, while the fluence map optimization problem remained convex guaranteeing optimality of the plan generated.

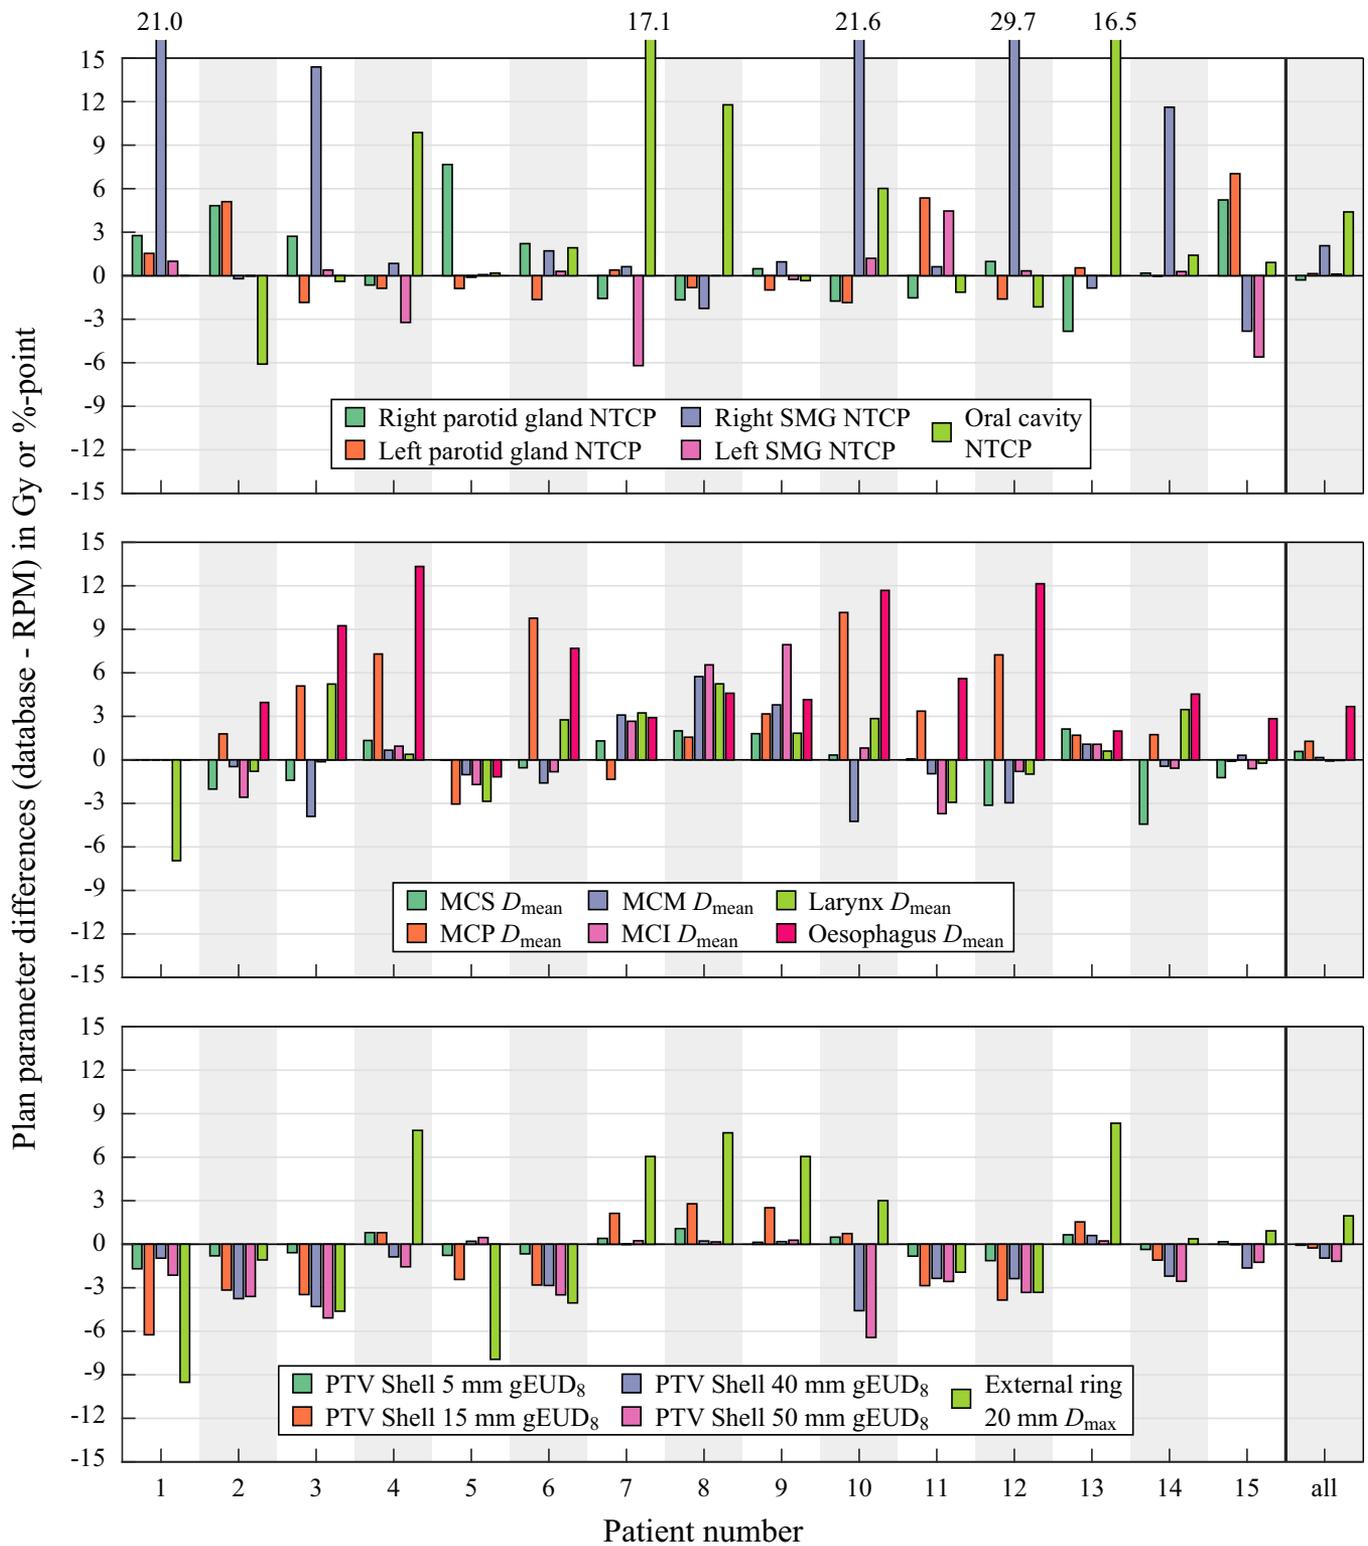


FIG. 4. Differences in most important plan parameter values between database plans and reference point method (RPM)-generated plans for the 15 most extreme outliers in test fold 4 (training with 20 patients), both favorable and unfavorable for the RPM. Positive values are favorable for the RPM. The last column shows the average results for all test patients.

Automatic RPM configurations were based on user preferences regarding population-based differences between database and RPM-generated plans (e.g., Table II). 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

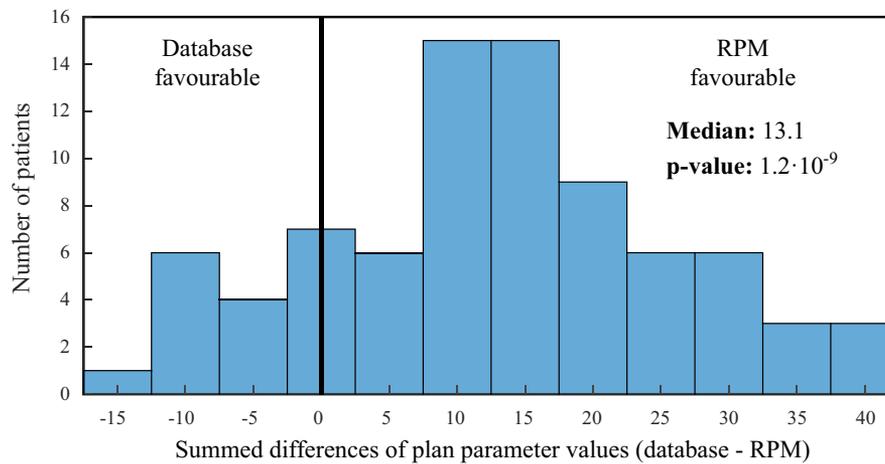


FIG. 5. Histogram of the summed differences of plan parameter values (database — reference point method) for all 85 patients in test fold 4 (training with 20 patients). The median gain of 13.1 indicates an advantage for the RPM ( $P < 0.001$ ).

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 II 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<sup>5,9</sup> for prostate planning. Note that for any configuration, Pareto optimality of all RPM-generated plans is guaranteed.<sup>11</sup>

Compared to the automatic RPM configuration for automatic prostate planning,<sup>9</sup> we observed more variation in differences between database and RPM-generated plans among the training folds (Fig. 3 and Figs. S1 and S2). For example, for the different training sets of 20 patients (Fig. 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 (Fig. S2), 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 2.A). As can be seen in Fig. S2, the distribution of differences in plan parameter values for the test patients in fold 2 were slightly worse than desired (Table II) 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 II). 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.

sense and clarity. 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 Fig. 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 favor 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.<sup>20</sup> 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 Fig. 3.

A similar approach to automatic configuration of the RPM is knowledge-based planning<sup>21</sup> (KBP). 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 the KBP approach. This model is trained, using machine learning techniques such as deep-learning,<sup>22</sup> support vector regression<sup>23</sup> or generative adversarial networks,<sup>24</sup> to predict the DVHs or spatial dose distribution prediction. The predicted DVHs or dose distribution are then the basis for plan optimization.<sup>25–27</sup> Both the automatic configuration and the 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<sup>28</sup> 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 I 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 optimization techniques.<sup>29</sup> This is a topic for further research.

## 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 favorable 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 favorable outliers for doses in highly prioritized OARs. Automatic RPM configuration has great potential in replacing traditional time-consuming and labor-intensive treatment planning workflows relying on manual configuration.

## CONFLICT OF INTEREST

The Erasmus MC Cancer Institute has research collaborations with Elekta AB, Stockholm, Sweden and Accuray Inc., Sunnyvale, USA. These companies were not involved in the work in this paper. The authors have no conflict to disclose.

<sup>a)</sup> Author to whom correspondence should be addressed. Electronic mail: r.vanhaveren@erasmusmc.nl.

## REFERENCES

- Hussein M, Heijmen B, Verellen D, Nisbet A. Automation in intensity modulated radiotherapy treatment planning – a review of recent innovations. *Br J Radiol.* 2018;91:20180270.
- Breedveld S, Craft D, Van Haveren R, Heijmen B. Multi-criteria optimisation and decision-making in radiotherapy. *Eur J Oper Res.* 2019;277:1-19.
- Breedveld S, Storchi P, Heijmen B. The equivalence of multi-criteria methods for radiotherapy plan optimization. *Phys Med Biol.* 2009;54:7199-7209.
- Breedveld S, Storchi P, Voet P, Heijmen B. iCycle: integrated, multi-criterial beam angle and profile optimization for generation of coplanar and non-coplanar IMRT plans. *Med Phys.* 2012;39:951-963.
- Van Haveren R, Breedveld S, Keijzer M, Voet P, Heijmen B, Ogryczak W. Lexicographic extension of the reference point method applied in radiation therapy treatment planning. *Eur J Oper Res.* 2017a;263:247-257.
- Voet P, Dirkx M, Breedveld S, Fransen D, Levendag P, Heijmen B. Toward fully automated multicriterial plan generation: a prospective clinical study. *Int J Radiat Oncol Biol Phys.* 2013;85:866-872.
- Heijmen B, Voet P, Fransen D, et al. Fully automated, multi-criterial planning for volumetric modulated arc therapy – an international multi-center validation for prostate cancer. *Radiother Oncol.* 2018;128:343-348.
- Van Haveren R, Ogryczak W, Keijzer M, Heijmen B, Breedveld S. Fast and fuzzy multi-objective radiotherapy treatment plan generation for head and neck cancer patients with the lexicographic reference point method. *Phys Med Biol.* 2017b;62:4318-4332.
- Van Haveren R, Heijmen B, Breedveld S. Automatically configuring the reference point method for automated multi-objective treatment planning. *Phys Med Biol.* 2019;64:035002.
- Wierzbicki A. A mathematical basis for satisficing decision making. *Math Mod.* 1982;3:391-405.
- Wierzbicki A. On the completeness and constructiveness of parametric characterizations to vector optimization problems. *OR Spectrum.* 1986;8:73-87.
- Ogryczak W, Kozłowski B. Reference point method with importance weighted ordered partial achievements. *TOP.* 2011;19:380-401.
- Jagt T, Breedveld S, Van Haveren R, Heijmen B, Hoogeman M. An automated planning strategy for near real-time adaptive proton therapy in prostate cancer. *Phys Med Biol.* 2018;63:135017.
- Jagt T, Breedveld S, Van Haveren R, et al. Plan-library supported automated replanning for online-adaptive intensity-modulated proton therapy of cervical cancer. *Acta Oncol.* 2019;58:1440-1445.
- Wang Y, Heijmen B, Petit S. Knowledge-based dose prediction models for head and neck cancer are strongly affected by interorgan dependency and dataset inconsistency. *Med Phys.* 2018;46:934-943.
- Alber M, Reemtsen R. Intensity modulated radiotherapy treatment planning by use of a barrier-penalty multiplier method. *Optim Methods Softw.* 2007;22:391-411.
- Niemierko A. Reporting and analyzing dose distribution: a concept of equivalent uniform dose. *Med Phys.* 1997;24:103-110.
- Lyman J. Complication probability as assessed from dose-volume histograms. *Radiother Res Suppl.* 1985;8:S13-S19.
- Breedveld S, Van Berg den Heijmen B. An interior-point implementation developed and tuned for radiation therapy treatment planning. *Comput Optim Appl.* 2017;68:209-242.
- Miettinen K. *Nonlinear Multiobjective Optimization.* Volume 12 of International Series in Operations Research and Management Science. Dordrecht: Kluwer Academic Publishers; 1999.
- Ge Y, Wu Q. Knowledge-based planning for intensity-modulated radiation therapy: a review of data-driven approaches. *Med Phys.* 2019;46:2760-2775.
- Chen X, Men K, Li Y, Yi J, Dai J. A feasibility study on an automated method to generate patient-specific dose distributions for radiotherapy using deep learning. *Med Phys.* 2019;46:56-64.
- Ma M, Kovalchuk N, Buyounouski M, Xing L, Yang Y. Dosimetric features-driven machine learning model for DVH prediction in VMAT treatment planning. *Med Phys.* 2019;46:857-867.
- Babier A, Mahmood R, McNiven A, Diamant A, Chan T. Knowledge-based automated planning with three-dimensional generative adversarial networks. *Med Phys.* 2020;47:297-306.
- Babier A, Boutilier J. Knowledge-based automated planning for oropharyngeal cancer. *Med Phys.* 2018;45:2875-2883.
- Miguel-Chumacero E, Currie G, Johnston A, Currie S. Effectiveness of multi-criteria optimization based trade-off exploration in combination with RapidPlan for head & neck radiotherapy planning. *Radiother Oncol.* 2018;13:229.
- Fan J, Wang J, Chen Z, Hu C, Zhang Z, Hu W. Automatic treatment planning based on three-dimensional dose distribution predicted from deep learning technique. *Med Phys.* 2019;46:370-381.
- Schipaanboord B, Breedveld S, Rossi L, Keijzer M, Heijmen B. Automated prioritised 3D dose-based MLC segment generation for step-and-shoot IMRT. *Phys Med Biol.* 2019;64:165013.
- Chan T, Craig T, Lee T, Sharpe M. Generalized inverse multiobjective optimization with application to cancer therapy. *Oper Res.* 2014;62:680-695.

## SUPPORTING INFORMATION

Additional supporting information may be found online in the Supporting Information section at the end of the article.

**Figure S1.** Boxplots of the differences in plan parameter values (Table II) 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 S2.** Boxplots of the differences in plan parameter values (Table II) 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 (◀) or RPM plans (▶).