Published in final edited form as:

Phys Med Biol. 2013 February 7; 58(3): 621–639. doi:10.1088/0031-9155/58/3/621.

A Column-generation-based Method for Multi-criteria Direct Aperture Optimization

Ehsan Salari and Jan Unkelbach

Department of Radiation Oncology, Massachusetts General Hospital and Harvard Medical School, Boston, MA 02114, USA

Ehsan Salari: salari.ehsan@mgh.harvard.edu

Abstract

Navigation-based multi-criteria optimization has been introduced to radiotherapy planning in order to allow the interactive exploration of trade-offs between conflicting clinical goals. However, this has been mainly applied to fluence map optimization. The subsequent leaf sequencing step may cause dose discrepancy, leading to human iteration loops in the treatment planning process that multi-criteria methods were meant to avoid. To circumvent this issue, this paper investigates the application of direct aperture optimization methods in the context of multi-criteria optimization. We develop a solution method to directly obtain a collection of apertures that can adequately span the entire Pareto surface. To that end, we extend the column generation method for direct aperture optimization to a multi-criteria setting in which apertures that can improve the entire Pareto surface are sequentially identified and added to the aperture collection. Our proposed solution method can be embedded in a navigation-based multi-criteria optimization framework, in which the treatment planner explores the trade-off between treatment objectives directly in the space of deliverable apertures. Our solution method is demonstrated for a paraspinal case where the trade-off between target coverage and spinal-cord sparing is studied. The computational results validate that our proposed method obtains a balanced approximation of the Pareto surface over a wide range of clinically relevant plans.

Keywords

Multi-criteria optimization; Direct aperture optimization; Pareto optimality; Column generation method

1. Introduction

Intensity-modulated radiation therapy (IMRT) is a radiation treatment modality capable of shaping complex dose distributions in the patient using a *multi-leaf collimator* (MLC) system ([29, 3]). Compared to the conventional 3D-conformal radiation therapy, IMRT allows for delivering the prescription dose to the target while sparing *organs at risk* (OARs) to a larger extent. IMRT planning has been traditionally performed in two sequential stages: (i) *fluence map optimization* (FMO) in which radiation beams are discretized into *beamlets* and mathematical optimization methods are used to determine the optimal intensity of all beamlets (for a review on the FMO problem see [26, 22]), and (ii) *leaf sequencing* (LS) where the obtained fluence maps are converted into *apertures* that can be delivered using the MLC [5, 1, 12].

IMRT planning involves finding a treatment plan that yields the desired trade-off between different treatment goals. Using the traditional two-stage approach to achieve this goal leads to iterative loops: The FMO problem is formulated and solved to optimize a weighted-sum of the treatment goals. Hence, to obtain the desired trade-off, the objective weights are manually changed in an iterative process. Additionally, there is dose discrepancy between FMO and LS stages since fluence maps are only approximately sequenced at the LS stage, particularly when delivery efficiency is taken into consideration. Moreover, in order to accurately measure the dose delivered by a treatment plan, the knowledge of the shape and intensity of the apertures employed is essential. Since this information is not available in the FMO stage, beamlet-based dose calculation methods are used during this stage, thereby contributing to the dose discrepancy (see, e.g., [14, 13]).

In recent years, the field of IMRT treatment plan optimization has evolved and has addressed the main shortcomings of the traditional approach. In particular, two major improvements, which have been implemented into current commercial treatment planning systems, are *multi-criteria optimization* (MCO) and *direct aperture optimization* (DAO).

1.1. Multi-criteria optimization

MCO methods [16, 6, 15] address the problem of finding the desired trade-off between different planning goals, which was traditionally obtained by manually changing objective weights in a time-consuming trial-and-error process. Two MCO solution approaches have been proposed and are currently in clinical use for IMRT treatment planning. The first approach is *lexicographic optimization* (LO) where the treatment objectives are prioritized based on their importance (see, e.g., [11, 30]). LO sequentially optimizes each individual objective while higher-priority objectives are constrained to their achieved values. The second approach is based on the paradigm of navigating on the Pareto surface ([20, 7]). In this approach a representative subset of Pareto-optimal treatment plans, the so-called *database* plans, are initially generated. Pareto-optimal plans are those for which we cannot improve any treatment objective unless some other objective deteriorates (for a review of Pareto optimality see [9]). The Pareto surface is then approximated using the convex combinations of the database plans. The user then navigates on the approximated Pareto surface to choose the Pareto-optimal plan that yields the desired trade-off between different treatment objectives.

1.2. Direct aperture optimization

DAO addresses the problem that the FMO dose distribution is degraded during the LS step. More specifically, DAO aims at directly optimizing the shape and intensity of deliverable MLC segments, i.e. apertures, and thus avoids the traditional two-step approach. Three types of DAO solution methods have been proposed: stochastic search methods [25, 17], gradient-based leaf refinement techniques [10], and column generation methods [21, 19]. Column generation is an iterative technique in which at each iteration, the aperture that promises the largest improvement to the objective function is identified and added to the treatment plan until a sufficient plan quality is achieved. This paper will consider the column-generation approach to DAO.

1.3. Multi-criteria direct aperture optimization

MCO and DAO have mostly been developed independently[‡]. In fact, combining both methodologies is related to an intrinsic problem: The MCO approach, and particularly navigation-based techniques, have been mainly developed in the context of the FMO

[‡]Recently, [18] has applied the LO approach to the DAO problem

problem, partly because they rely on forming linear combinations of treatment plans. Hence, the optimal fluence maps need to be sequenced into a collection of deliverable apertures, thereby causing dose discrepancy. This creates manual iteration loops in the treatment planning process that MCO was meant to avoid. To avoid this caveat, it is desirable to employ the navigation-based methods in the space of deliverable apertures. This yields the following two options:

1. Navigation-based techniques rely on forming convex combinations of the database plans. One can obtain each database plan independently using one of the proposed DAO techniques. In that case, different database plans may contain different sets of apertures. Hence, convex combinations of database plans will then use apertures from the union of all these collections, rendering a naive implementation of this approach clinically impractical due to the large number of apertures employed. Hence, a clinical application of this method requires limiting the number of apertures used in the plan during the navigation process. However, this is difficult to achieve since the navigation process is required to be interactive and all computations should be performed in real time.
2. The above problem can be avoided by ensuring all database plans use the same collection of apertures. In that case, the Pareto-surface navigation methods developed for the FMO problem are immediately applicable. However, this requires determining a single set of apertures that can adequately represent the entire Pareto surface within a clinically meaningful range.

This paper pursues the latter approach and suggests a solution method to obtain a set of apertures that is shared by all database plans and can span the entire Pareto surface. To that end, we generalize the DAO column generation method to a multi-criteria setting. The rest of this paper is organized as follows. In Section 2 we introduce the notation and formulate the multi-criteria DAO problem. In Section 3 we propose our solution approach to solve this problem and obtain a representation of the Pareto surface. In Section 4 the performance of our solution technique is demonstrated for a clinical case. In Section 5 we discuss some future research questions to be addressed, and in Section 6 we summarize and conclude the paper.

2. The Multi-criteria DAO (MCDAO) Problem

We let K denote the set of all deliverable apertures by the MLC. Moreover, we use the decision variable $y_k \in \mathbb{R}_+$ to represent the intensity of aperture $k \in K$. To measure the dose delivered to the patient, the relevant structures in the patient are conceptually discretized into a set of voxels, denoted by I . At unit intensity, aperture $k \in K$ deposits a dose of D_{ki} in voxel $i \in I$, which we call the *aperture dose deposition coefficient*. For beamlet-based dose calculation models, D_{ki} can be approximated via summing over dose deposition coefficients of beamlets that are exposed in aperture k . To measure the total dose deposited in voxel $i \in I$, we associate a variable d_i with that voxel. d_i can then be expressed as a linear combination of aperture intensities through the dose deposition coefficients as follows:

$$d_i = \sum_{k \in K} D_{ki} y_k.$$

Finally, we let $\mathbf{d} = (d_i : i \in I)^\top$ denote the vector of the dose distribution delivered to the patient.

We use a collection of treatment-plan evaluation criteria, indexed by $\ell \in L$, to measure the treatment plan quality. More specifically, we consider $F_\ell: \mathbb{R}_+^{|I|} \rightarrow \mathbb{R}$ for $\ell \in L$ as a convex function of the dose distribution $\mathbf{d}$ where, without loss of generality, we assume smaller values are preferred to larger ones. Many treatment-plan evaluation criteria, such as voxel-based piecewise quadratic penalties, equivalent uniform dose (EUD), and tail EUD (see [4]) are convex functions. Moreover, it has been shown that in a multi-criteria setting some non-convex criteria such as tumor control probability (TCP), normal-tissue complication probability (NTCP), and probability of uncomplicated tumor control (P_+) can be replaced by convex ones (see [23]).

The goal of the IMRT treatment planning problem is to find the collection of apertures and their associated intensities that yield the desired trade-off between different treatment-plan evaluation criteria. Hence, we formulate a multi-criteria problem in terms of all deliverable aperture intensities as follows:

$$\min \mathbf{F}(\mathbf{d})$$

subject to (M)

$$\begin{aligned} d_i &= \sum_{k \in K} D_{ki} y_k & i \in I \\ y_k &\geq 0 & k \in K. \end{aligned}$$

where $\mathbf{F}(\mathbf{d}) = (F_\ell(\mathbf{d}))_{\ell \in L}^\top$ denotes the vector of all treatment objectives. The notion of Pareto optimality can be used to quantify the trade-off between different treatment evaluation criteria. In particular, Pareto-optimal treatment plans are those for which one cannot improve a criterion value unless some other criterion value deteriorates. Other treatment plans are dominated by at least one Pareto-optimal plan and therefore do not need to be considered in the trade-off.

Since (M) is a convex problem, the weighted-sum method can be employed to generate all corresponding Pareto-optimal solutions (see, e.g., [9]). In this method, the multi-criteria problem is transformed into a single-criterion problem where the objective function is the weighted sum of all evaluation criteria. More specifically, a family of optimization problems is formulated as:

$$\min \mathbf{w}^\top \mathbf{F}(\mathbf{d})$$

subject to (M($\mathbf{w}$))

$$d_i = \sum_{k \in K} D_{ki} y_k \quad i \in I \quad (1)$$

$$y_k \geq 0 \quad k \in K. \quad (2)$$

where $\mathbf{w} \in \mathbb{R}_+^{|L|}$ are nonnegative weight vectors representing the relative importance of criteria $\ell \in L$. The optimal solution to (M($\mathbf{w}$)) for any nonnegative weight vector $\mathbf{w}$ yields a Pareto-optimal treatment plan. (M($\mathbf{w}$)) corresponds to the single-criterion DAO problem. In

the next section, we first briefly explain the column generation method used to solve $(M(\mathbf{w}))$ and then extend this method to characterize the set of Pareto-optimal solutions to (M) .

3. Solution Approach

3.1. The column generation method

Since there are an enormous number of apertures deliverable using the MLC, the cardinality of set K is extremely large, and thus, solving the full version of $(M(\mathbf{w}))$ is computationally prohibitive. Additionally, only those solutions that employ a limited number of apertures are clinically relevant. Therefore, column generation techniques have been proposed to find high-quality sparse solutions to $(M(\mathbf{w}))$ (see, e.g., [21, 19]). In particular, a column generation method to $(M(\mathbf{w}))$ is an iterative approach that starts with a subset of deliverable apertures $K' \in K$. At each iteration, the following two problems are solved:

1. In the *master problem*, a restricted version of $(M(\mathbf{w}))$, where the set of apertures is confined to K' , is solved. This yields the optimal aperture intensities y_k^* for $k \in K'$ and the corresponding dose distribution $\mathbf{d}^*$.
2. Given the solution to the master problem, the *pricing problem* is solved to identify a new promising aperture to be added to K' to improve the objective function. The pricing problem amounts to finding the aperture k that minimizes the gradient of the weighted-sum objective function $\mathbf{w}^T \mathbf{F}(\mathbf{d})$ with respect to the intensity of the new aperture.

The algorithm continues until a high-quality treatment plan is attained or no new promising aperture is found.

3.2. The pricing problem and its interpretation in the multi-criteria setting

In this section we formally define the pricing problem for $(M(\mathbf{w}))$ and provide an interpretation of the pricing problem in a multi-criteria setting. To formulate the pricing problem, we let

$$\nabla F_{\ell} = \left(\frac{\partial F_{\ell}}{\partial d_i} : i \in I \right)^T$$

denote the gradient of treatment evaluation criterion $\ell \in L$ with respect to the dose distribution. We let $\mathbf{D}_k = (D_{ki} : i \in I)$ be the row vector of dose deposition coefficients corresponding to aperture $k \in K$. We can then define

$$r_{k\ell}(\mathbf{d}) \equiv \mathbf{D}_k \nabla F_{\ell}(\mathbf{d})$$

to measure the rate of change in criterion $\ell \in L$ at dose distribution $\mathbf{d}$ as we increase the intensity of aperture $k \in K$. In other words, if we consider $\mathbf{d}$ as a function of the aperture intensities, $r_{k\ell}$ denotes the partial derivative of F_{ℓ} with respect to the intensity of aperture k , evaluated at the current dose distribution. Finally, we let $\mathbf{r}_k(\mathbf{d}) = (r_{k\ell}(\mathbf{d}) : \ell \in L)^T$ be the vector of rates of change in all treatment evaluation criteria.

Now consider the column generation method for a fixed weight vector $\mathbf{w}_p$ with unit length (i.e., the Euclidean norm $\|\mathbf{w}_p\|_2 = 1$). At each iteration, solving the restricted problem $(M(\mathbf{w}_p))$ using aperture set $K' \subseteq K$ yields the solution $(\mathbf{y}_p^*, \mathbf{d}_p^*)$. The pricing problem then

seeks for the aperture k that has the largest improvement rate in the current solution. This is measured by the partial derivative of the weighted-sum objective $\mathbf{w}_p^\top \mathbf{F}(\mathbf{d}_p^*)$ with respect to the intensity of aperture k . Thus, the pricing problem is formulated as:

$$\min_{k \in K} \mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*). \quad (3)$$

This has a geometric interpretation, which is illustrated in Figure 1: The aperture set K' spans a Pareto surface, containing all Pareto-optimal solutions to the restricted (M). The Pareto-optimal plan p in particular corresponds to a point on the Pareto surface with objective values $\mathbf{F}(\mathbf{d}_p^*)$. Moreover, $\mathbf{w}_p$ is the orthogonal vector to the Pareto surface at point $\mathbf{F}(\mathbf{d}_p^*)$. The vector $\mathbf{r}_k(\mathbf{d}_p^*)$ indicates the direction in the objective space along which the point $\mathbf{F}(\mathbf{d}_p^*)$ moves if aperture k is added to the plan with a small intensity y_k . One can observe that the objective function of the pricing problem, i.e. $\mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*)$, is the scalar projection of this vector onto the normal unit vector at point $\mathbf{F}(\mathbf{d}_p^*)$. Thus, as the intensity of aperture $k \in K$ increases, $\mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*)$ can be interpreted as the rate of change in the Pareto surface at point $\mathbf{F}(\mathbf{d}_p^*)$ along direction $\mathbf{w}_p$, which we call the *orthogonal rate of change*.

3.3. Extending the standard column generation to the multi-criteria setting

For the navigation-based MCO approach, we aim to approximate the Pareto surface within a clinically relevant region using multiple database plans. One may use the column generation method described in 3.1 to solve individual instances of $(M(\mathbf{w}))$ for different nonnegative weight vectors $\mathbf{w}$ to generate the database plans. However, that results in a distinct collection of apertures generated for each individual plan. The navigation process, performed in the convex hull of the database plans, will then use the union of all generated apertures. This results in clinically impractical treatment plans due to the large number of apertures used. To avoid this issue, we propose a column generation method to obtain a reasonably-sized collection of apertures that is shared among all database plans, and thus all corresponding convex-combination plans.

We consider the Pareto surface associated with the restricted problem (M) in which the set of apertures are confined to $K' \subseteq K$. Adding a new aperture $k \in K \setminus K'$ to K' does not worsen this Pareto surface, because the new aperture intensity can be set to zero if the added aperture is not beneficial for some plans along the current Pareto surface. In other words, the Pareto surface associated with $K' \cup \{k\}$ dominates the one corresponding to K' . Hence, by sequentially adding apertures to K' we can improve the current Pareto surface. This is schematically illustrated in Figure 2. Therefore, we propose a generalization of the column generation method to the multi-criteria setting, which can be applied to approximate the Pareto surface associated with the MCDAO problem.

1. At each iteration, the MCDAO master problem approximates the Pareto surface associated with the restricted set of apertures, i.e. K' . This is detailed in the following Section 3.4.
2. The MCDAO pricing problem then identifies a new aperture that promises the largest rate of improvement in the current Pareto surface. This new aperture is then added to K' . In Section 3.5, we discuss approaches to formulate the pricing problem.

The above procedure continues until a user-specified termination condition (e.g., limit on the number of generated apertures) is met.

3.4. The MCDAO master problem

Consider the restricted (M) with the aperture set K' . The master problem seeks to characterize the Pareto surface to the restricted (M). Several approximation techniques have been proposed to efficiently obtain a representation of the Pareto surface associated with a multi-criteria convex problem (see, e.g., [24]). For the purpose of our column generation method, it will be sufficient to obtain a simple approximation by generating a collection of Pareto-optimal plans, indexed by $p \in P$, using the weighted-sum method. More specifically, let $\mathbf{w}_p$ for $p \in P$ be nonnegative weight vectors. The solution to the restricted (M($\mathbf{w}_p$)), denoted by $(\mathbf{y}_p^*, \mathbf{d}_p^*)$, then yields a Pareto-optimal solution to the restricted (M). Thus, we characterize the Pareto surface associated with the restricted (M), using $|P|$ Pareto-optimal treatment plans. In this paper, we define a fixed set of weight vectors up front which are used during the course of the algorithm. However, generalizations to adaptively changing weight vectors can be considered in the future. One can observe that the MCDAO master problem is a simple extension where the master problem of a standard column generation method is solved $|P|$ times. Moreover, Pareto-optimal plans $p \in P$ share the same set of apertures (i.e., K') and provide a piecewise-linear approximation of the Pareto surface corresponding to the restricted (M).

3.5. The MCDAO pricing problem

In the pricing problem for the standard column generation approach described in Section 3.1, we seek for the aperture that promises the largest improvement in the weighted-sum of treatment evaluation criteria, i.e. $\mathbf{w}^\top \mathbf{F}$, corresponding to a single weight vector $\mathbf{w}$. In MCDAO, we seek for the aperture that yields the largest improvement of the entire Pareto surface. In other words, we aim at finding an aperture that is beneficial for all treatment plans along the Pareto surface. We have discussed in Section 3.2 that the standard pricing problem measures the orthogonal rate of change at a given point on the Pareto surface. We can generalize this idea and determine the aperture that minimizes the cumulative orthogonal rates of change over all sample points on the Pareto surface. In the following, we formulate a new pricing problem to find such an aperture. In Appendix A we derive an exact method to solve this formulation. Finally, we discuss some heuristic approaches to the pricing problem.

At each iteration of the column generation algorithm, solving the MCDAO master problem, which is described in Section 3.4, yields $|P|$ Pareto-optimal treatment plans that approximate the Pareto surface corresponding to aperture set K' . In order to measure the improvement rate of aperture k in the current Pareto surface, we propose to use the Pareto-optimal plans in P as sample points on the surface. In particular, we measure aperture k 's cumulative orthogonal rates of change at these sample points as a surrogate for its improvement rate in the entire Pareto surface. Hence, we formulate the MCDAO pricing problem as follows:

$$\min_{k \in K} \sum_{p \in P} \min \{0, \mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*)\}. \quad (4)$$

In order to motivate this formulation, we make the following consideration: A given aperture $k \in K$ may have either positive or negative orthogonal rate of change at different Pareto-optimal points $p \in P$. In particular, suppose aperture k has a positive orthogonal rate of change at point $\bar{p} \in P$ (i.e., $\mathbf{w}_{\bar{p}}^\top \mathbf{r}_k(\mathbf{d}_{\bar{p}}^*) > 0$). In that case, adding aperture k to K' and solving the restricted problem (M($\mathbf{w}_{\bar{p}}$)) yield the same Pareto-optimal point. In other words, aperture k neither improves nor deteriorates point $\bar{p}$. Thus, to measure the cumulative orthogonal

rates of change corresponding to aperture k we only consider those points $p \in P$ for which $\mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*) < 0$.

In Appendix A we reformulate (4) as a mixed-integer programming (MIP) problem, which can be solved using MIP solution techniques. Although the MIP solution approach described in Appendix A is computationally feasible for clinical instances, it is desirable to devise more efficient solution approaches for the MCDAO pricing problem. To that end, we also consider the following heuristic methods:

Individual orthogonal rate of change (IORC)—We consider a group of heuristic strategies in which at each iteration of the column generation algorithm, the aperture with the most negative orthogonal rate of change for an individual Pareto-optimal point $\bar{p} \in P$ is identified and added to the restricted (M). In this setting the pricing problem can be formulated as:

$$\min_{k \in K} \mathbf{w}_{\bar{p}}^\top \mathbf{r}_k(\mathbf{d}_{\bar{p}}^*). \quad (5)$$

This pricing problem can be solved using the solution techniques discussed in [21, 19]. There are many options to select a plan in a given iteration, and this is a potential topic of further research. For this paper, we consider the following two simple heuristics:

1. The weight vector $\mathbf{w}_{\bar{p}}$ is fixed during the course of the column generation algorithm. This essentially corresponds to solving a traditional DAO problem with a weighted-sum of evaluation criteria $\mathbf{w}_{\bar{p}}^\top \mathbf{F}$. We consider this approach for the sake of comparison in the computational results.
2. Alternatively, $\bar{p}$ may vary during the course of the algorithm. For instance, $\bar{p}$ can alternate within P such that at each iteration of the algorithm, the aperture with the most negative orthogonal rate of change measured at a different region of the Pareto surface is chosen.

Cumulative orthogonal rate of change (CORC)—Solving the MCDAO pricing problem as formulated in (4) requires employing MIP solution approaches owing to the use of the minimum operation in the objective function. This is necessary in order to account for the possibility that an aperture intensity can be set to zero, if that aperture does not improve the Pareto surface at a particular sample point. However, the pricing problem becomes much simpler if we drop the minimum operation. In this case, the pricing problem can be formulated as

$$\min_{k \in K} \sum_{p \in P} \mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*), \quad (6)$$

which has the same functional form as the IORC pricing problem and thus can be efficiently solved using solution approaches available for the standard pricing problem.

4. Computational Results

In this section we investigate the performance of our column generation method on a clinical paraspinal case depicted in Figure 3. We use nine equi-spaced 6 MV beams with a beamlet grid resolution of $0.5 \times 0.5 \text{ cm}^2$. To incorporate all beamlets that can contribute to treating the paraspinal target at each beam, a maximum of 20 MLC leaf pairs, with a maximum of 18

beamlets per MLC row, are considered in the treatment plan optimization. This yields a maximum of 360 beamlets per beam. Moreover, a voxel resolution of $2.96 \times 2.96 \times 2.5 \text{ mm}^3$ is used. The number of voxels at each structure is reported in Table 1. The matrix of beamlet dose deposition coefficients is computed using the pencil-beam dose calculation algorithm, Quadratic Infinite Beam (QIB), embedded in CERR 3.0 Beta 3 (see [8]). Voxel-based piecewise-quadratic penalties are used as treatment-plan evaluation criteria. More specifically, let S be the set of all structures and I_s be the set of voxels in structure $s \in S$. To measure the quality of the dose distribution in structure $s \in S$, we use evaluation criterion F_s as follows:

$$F_s(\mathbf{d}) = \frac{1}{|I_s|} \sum_{i \in I_s} \left[\lambda_s^+ (\max\{d_i - t_s, 0\})^2 + \lambda_s^- (\max\{t_s - d_i, 0\})^2 \right] \quad (7)$$

where t_s is the threshold dose and λ_s^- and λ_s^+ are the underdosing and overdosing penalties, respectively. We set $t_s = 0$ for all OARs and $t_s = 60$ for the target. Moreover, since there is no penalty associated with OAR underdosing, we set $\lambda_s^- = 0$ and $\lambda_s^+ = 1$. For the target volume, we set $\lambda_s^- = 7$ and $\lambda_s^+ = 1$.

4.1. Trade-off between spinal-cord sparing and target coverage

In a paraspinal case, the target typically wraps around the spinal cord, and thus, the trade-off between target coverage and spinal-cord sparing is of particular interest. We first consider a simplified case consisting of the planning target volume (PTV) and the spinal cord only and compare the two-dimensional Pareto surface obtained using different variants of the column generation method described in Section 3.5. We consider three weight vectors $\mathbf{w}_p$ for $p \in \{1, 2, 3\}$ corresponding to a balanced, OAR-, and target-emphasis plan, respectively (see Table 2). Figure 5(a) illustrates the DVH curves associated with the plans obtained by the traditional column generation method using these weight vectors.

Using different formulations of the MCDAO pricing problem in Section 3.5 and the three weight vectors, we run the column generation algorithm for 50 iterations starting with no initial apertures. This yields a total of 6 aperture sets, each containing 50 apertures, as follows:

MCDAO uses formulation (4) with $P = \{1, 2, 3\}$,

Balanced-IORC uses formulation (5) with $\bar{p} = 1$ (balanced weight vector),

OAR-IORC uses formulation (5) with $\bar{p} = 2$ (OAR-emphasis weight vector),

Target-IORC uses formulation (5) with $\bar{p} = 3$ (target-emphasis weight vector),

Alter-IORC uses formulation (5) with $\bar{p} \in \{1, 2, 3\}$ (alternating weight vector), and

CORC uses formulation (6) with $P = \{1, 2, 3\}$.

Given these aperture sets, we then obtain a representation of the complete Pareto surface associated with each set, i.e., in addition to the pareto-optimal plans that correspond to the weight vectors used, we generate supplementary points on the surface. This is performed using the ε -constraint method (see, e.g., [28]). Similarly, we obtain a Pareto-frontier approximation associated with the FMO problem. Details on the optimization solvers as well as the computational effort required to solve the pricing problem, are provided in Section 4.1.2. The obtained Pareto surfaces are shown in Figure 4. The Pareto-optimal plans associated with $\bar{p} = 1, 2$, and 3, obtained using the traditional column generation method, are marked on their corresponding Pareto surfaces. DVH curves of these plans are shown in Figure 5(a). In order to illustrate the difference between these Pareto surfaces, we constrain

the target objective value for OAR- and Target-IORC, MCDAO, and FMO to its optimal value in the balanced plan, as indicated in figure 4(a) with a dashed line, and obtain the corresponding Pareto-optimal plans. The DVH curves associated with these plans are shown in figure 5(b). Finally, for the MCDAO aperture set, we also obtained Pareto-optimal plans associated with balanced, OAR- and target-emphasized weights, which are marked on the corresponding Pareto surface.

We first compare the MCDAO Pareto surface with the Balanced-, OAR-, and Target-IORC Pareto surfaces (see Figure 4(a)). One can observe that for small target objective values, Target-IORC dominates other strategies. Similarly, for the Balanced- and OAR-IORC, it is observed that the corresponding Pareto surface dominates other strategies in a neighborhood of the plan that the apertures were originally generated for. On the other hand, the MCDAO Pareto surface considerably dominates Target-IORC in the region of good spinal-cord sparing and the OAR-IORC in the region of good target coverage. Additionally, the figure shows the FMO Pareto surface, which provides a lower bound on the MCDAO Pareto surface. We next compare MCDAO to the Alter-IORC and CORC in Figure 4(b). Clearly, MCDAO outperforms Alter-IORC, but the CORC method yields results very close to the MCDAO method.

4.1.1. Including additional structures—Above, we considered target and spinal cord as the only structures. Next, we consider the complete case incorporating the left and right lungs, esophagus, and unclassified tissue into the model, while still focusing on the trade-off between target and spinal cord. Voxels in the lungs and unclassified tissue are downsampled at a 1:4 and 1:8 ratio, respectively. Similar to the simplified case, we use three weight vectors, which are the balanced, target-, and OAR-emphasized weight vectors (see Table 2). Using these three weight vectors we first run the standard column generation method and obtain the corresponding balanced, target-, and OAR-emphasis plans. The DVH curves are shown in Figure 6(a). We next obtain a collection of 75 apertures using different variants of the column generation method. For each aperture set we then approximate the two-dimensional trade-off between the target coverage and spinal-cord sparing while constraining the objective values of other structures to the achieved values in the balanced plan. More specifically, we use ε -constraint method similar to the simplified case with the difference that the objective values corresponding to other structures are constrained to their optimal values of the balanced plan. Figure 6(b) compares the Pareto surface approximations associated with different strategies. Similar to the simplified case, the IORC Pareto surfaces are dominant within a small neighborhood around the balanced, OAR-, and target-emphasis plans. MCDAO is dominant in other regions. It should be noted that the trade-off curves end for different target objective values, which is due to constraining the objective values of other structures. In particular, the collection of apertures generated using OAR-IORC is not able to achieve low target objective values. The difference between trade-off curves associated with IORC and MCDAO in the complete case is larger than the simplified case.

4.1.2. Optimization solvers—At each iteration of the column generation algorithm, the master problem (i.e., the restricted $(M(\mathbf{w}))$) is solved using the primal-dual interior point method presented in [2]. In particular, we first substitute decision variables $\mathbf{d}$ using constraints in (1) and then solve the problem subject to only nonnegativity constraints in (2). The pricing problem for the MCDAO aperture set is solved by applying the IBM ILOG CPLEX Optimizer to the formulation presented in Appendix A. For the paraspinal case, this formulation contains 363 binary variables and 2729 constraints, which takes 2.5 seconds, on the average, to solve (see Appendix A for more information on the problem size). Moreover, the pricing problem corresponding to CORC and IORC aperture sets is solved using the polynomial-time algorithm provided by [21], whose computational time is negligible compared to the MCDAO pricing problem. As discussed above, for each aperture set, we

use the ϵ -constraint method to get a fine approximation of the Pareto surface. This requires solving a collection of single-criterion optimization problems in which the target objective is minimized while constraining the spinal-cord objective to different values. To solve these optimization problems we first transform (7) into a quadratic function using auxiliary variables and additional linear constraints. We then use CPLEX Barrier Optimizer to solve the resulting quadratically-constrained quadratic programming (QCQP) problems.

4.2. Discussion

It is observed that IORC methods dominate other strategies in a neighborhood of the plan that the apertures were generated for. This result is expected since all apertures in those sets were generated as to improve a particular plan. In contrast, the MCDAO strategy provides a well-balanced approximation of the Pareto surface across a wide range of clinically-relevant treatment plans.

For the paraspinal example, the approximate method CORC, which considers all three plans but solves an approximate pricing problem, shows similar performance compared to the MCDAO method for the simplified case. Using CORC has the advantage that the pricing problem can be solved efficiently without the use of integer programming methods. However, one can observe that MCDAO dominates CORC for small values of the target objective when other structures are considered in the model. More patient cases, and higher dimensional trade-offs need to be investigated in order to clarify if CORC generally yields similar performance as MCDAO. Moreover, for the case considered here, MCDAO and CORC outperform Alter-IORC. This suggests that there is a gain in considering multiple treatment plans simultaneously in the pricing problem, as opposed to considering only one plan at a time.

5. Remarks and Future Research

In this paper, we limited the computational results to the characterization of a two-dimensional trade-off between target coverage and spinal-cord sparing. This was achieved for the complete paraspinal case by constraining the objectives associated with lungs, esophagus, and unclassified tissue, to their corresponding values in the balanced plan. The motivation for that is to rigorously compare two-dimensional Pareto surfaces. Comparing Pareto surfaces in higher dimensions is more involved and difficult to visualize (see, e.g., [27]). For the two-dimensional trade-off example, consistent results verify that IORC methods that consider a single plan in the pricing problem outperform other approaches in the vicinity of the plan that apertures were designed for. In contrast, the MCDAO method provides a well-balanced approximation of the Pareto surface over a range of clinically-relevant treatment plans. The results obtained by the proposed MCDAO method motivates further work to characterize the performance of the algorithm in a clinical setting. In particular, it remains to be investigated whether this method is capable of approximating high-dimensional Pareto surfaces. Such future studies should also consider methodological extensions to the proposed method that are outlined in the rest of this section.

MCDAO versus multi-criteria LS

In the proposed MCDAO column generation method, all weight vectors are considered simultaneously over the course of the algorithm. Thus, this method can be viewed as a natural approach to enforcing the same set of apertures for different database plans. In contrast, traditional methods generate database plans independently by solving the FMO problem, which are then individually sequenced using LS methods. However, enforcing shared apertures for these plans at the LS stage requires new multi-criteria LS algorithms.

Although this can be an alternative direction of future research, such algorithms are, to our knowledge, currently not available.

Generalization to dose-constrained MCDAO framework

The current approach does not provide a mechanism for choosing the appropriate set of weight vectors. In particular, in this study, we have manually chosen three weight vectors for balanced, target-, and OAR-emphasized plans. These weight vectors are chosen to emphasize the clinically-relevant regions of the Pareto surface. Future work is needed to automate the selection of the sampling points. To address this issue, the current MCDAO framework can be generalized to a constrained formulation in which dose constraints limit the span of the Pareto surface to only clinically meaningful regions. This will then allow for using the anchor plans as sample points on the Pareto surface while ensuring the Pareto approximation remains clinically relevant, thereby eliminating the need to manually determine weight vectors.

Sampling the Pareto Surface

Currently we use $|P|$ weight vectors to sample the Pareto surface, which yields $|P|$ sample Pareto points. These weights are chosen such that we can sample different regions of the Pareto surface and are fixed during the course of the column generation algorithm. However, in principle, our column generation method allows for changing these weight vectors between different iterations. For instance, if there is no significant increase during the past iterations for a given weight vector, one can replace this vector with a new one. Thus, developing algorithms that adaptively determine sample points on the Pareto surface is an area of further research.

Hybrid approaches

It may not be always possible to approximate the entire Pareto surface adequately with a limited collection of apertures. This may, for example, be the case for a high-dimensional trade-off space. However, the proposed MCDAO approach can still be used to generate a basic set of apertures that is shared among all database plans. Subsequently, a limited number of plan-specific apertures can be added to improve individual regions of the Pareto surface.

6. Conclusion

In this paper we developed a navigation-based MCO approach in which the database plans share the same collection of apertures. This allows for conducting the navigation step in the space of deliverable apertures, thereby avoiding dose discrepancies caused by the LS stage. More specifically, we extended the traditional column generation method to a multi-criteria setting in which we sequentially add apertures to improve the entire Pareto surface. In particular, we approximated the improvement rate of an aperture in the Pareto surface by sampling Pareto-optimal points from different regions of the Pareto surface. We formulated a generalization of the standard pricing problem in which we seek for the aperture that yields the largest rate of improvement over all sample points on the Pareto surface, and developed several solution strategies for it. For a paraspinal case, where we consider the trade-off between target coverage and spinal-cord sparing, we performed a rigorous comparison of the algorithm to strategies that consider only individual plans in the pricing problem. The computational results suggest that the proposed method yields a balanced approximation of the Pareto surface over a wide range of clinically-relevant plans, which are considered in navigation-based planning. The results demonstrated here motivate further work to fully assess the role of the proposed column-generation-based method in deliverable Pareto-surface navigation.

Acknowledgments

The project described was supported by Award No R01 CA118200-01 from the National Cancer Institute. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Cancer Institute or the National Institutes of Health. The authors would like to thank David Craft for providing feedback on the manuscript.

References

1. Ahuja RK, Hamacher HW. A network flow algorithm to minimize beam-on time for unconstrained multileaf collimator problems in cancer radiation therapy. *Networks*. 2005; 45(1):36–41.
2. Aleman DM, Glaser D, Romeijn HE, Dempsey JF. Interior point algorithms: guaranteed optimality for fluence map optimization in IMRT. *Physics in Medicine and Biology*. 2010; 55(18):5467–5482. [PubMed: 20798458]
3. Bortfeld T. IMRT: a review and preview. *Physics in Medicine and Biology*. 2006; 51:R363–R379. [PubMed: 16790913]
4. Bortfeld T, Craft D, Dempsey JF, Halabi T, Romeijn HE. Evaluating target cold spots by the use of tail EUDs. *International Journal of Radiation Oncology* Biology* Physics*. 2008; 71(3):880–889.
5. Bortfeld TR, Kahler DL, Waldron TJ, Boyer AL. X-ray field compensation with multileaf collimators. *International Journal of Radiation Oncology* Biology* Physics*. 1994; 28(3):723–730.
6. Craft D, Halabi T, Shih HA, Bortfeld T. An approach for practical multiobjective IMRT treatment planning. *International Journal of Radiation Oncology* Biology* Physics*. 2007; 69(5):1600–1607.
7. Craft DL, Halabi TF, Shih HA, Bortfeld TR. Approximating convex Pareto surfaces in multiobjective radiotherapy planning. *Medical Physics*. 2006; 33:3399–3408. [PubMed: 17022236]
8. Deasy JO, Blanco AI, Clark VH. CERR: a computational environment for radiotherapy research. *Medical Physics*. 2003; 30:979–985. [PubMed: 12773007]
9. Ehrgott, M. *Multicriteria Optimization*. Springer Verlag; Berlin, Germany: 2005.
10. Hardemark, A.; Liander, H.; Rehbinder, H.; Löf, J. White Paper. Raysearch Laboratories; 2003. Direct machine parameter optimization with RayMachine[®] in Pinnacle[®].
11. Jee KW, McShan DL, Fraass BA. Lexicographic ordering: intuitive multicriteria optimization for IMRT. *Physics in Medicine and Biology*. 2007; 52:1845–1861. [PubMed: 17374915]
12. Kamath S, Sahni S, Li J, Palta J, Ranka S. Leaf sequencing algorithms for segmented multileaf collimation. *Physics in Medicine and Biology*. 2003; 48(3):307–324. [PubMed: 12608609]
13. Knoos T, Ahnesjö A, Nilsson P, Weber L. Limitations of a pencil beam approach to photon dose calculations in lung tissue. *Physics in Medicine and Biology*. 1995; 40:1411–1420. [PubMed: 8532755]
14. Knoos T, Ceberg C, Weber L, Nilsson P. The dosimetric verification of a pencil beam based treatment planning system. *Physics in Medicine and Biology*. 1994; 39:1609–1628. [PubMed: 15551534]
15. Küfer, KH.; Monz, M.; Scherrer, A.; Süß, P.; Alonso, F.; Sultan, ASA.; Bortfeld, T.; Thieke, C. Multicriteria optimization in intensity modulated radiotherapy planning. In: Pardalos, PM.; Romeijn, HE., editors. *Handbook of Optimization in Medicine*. Springer; 2009.
16. Küfer KH, Scherrer A, Monz M, Alonso F, Trinkaus H, Bortfeld T, Thieke C. Intensity-modulated radiotherapy—a large scale multi-criteria programming problem. *OR Spectrum*. 2003; 25(2):223–249.
17. Li Y, Yao J, Yao D. Genetic algorithm based deliverable segments optimization for static intensity-modulated radiotherapy. *Physics in Medicine and Biology*. 2003; 48:3353–3374. [PubMed: 14620063]
18. Long T, Fraass D, Matuszak M, Feng M, Ten Haken RK, Romeijn E. Sensitivity analysis for lexicographic ordering in radiation therapy treatment planning. *Medical Physics*. 2012; 39(6):3445–3455. [PubMed: 22755724]
19. Men C, Romeijn HE, Ta kin ZC, Dempsey JF. An exact approach to direct aperture optimization in IMRT treatment planning. *Physics in Medicine and Biology*. 2007; 52(24):7333–7352. [PubMed: 18065842]

20. Monz M, Küfer KH, Bortfeld TR, Thieke C. Pareto navigation - algorithmic foundation of interactive multi-criteria IMRT planning. *Physics in Medicine and Biology*. 2008; 53:985–998. [PubMed: 18263953]
21. Romeijn HE, Ahuja RK, Dempsey JF, Kumar A. A column generation approach to radiation therapy treatment planning using aperture modulation. *SIAM Journal on Optimization*. 2005; 15:838–862.
22. Romeijn HE, Dempsey JF. Intensity Modulated Radiation Therapy treatment plan optimization. *TOP*. 2008; 16(2):215–243.
23. Romeijn HE, Dempsey JF, Li JG. A unifying framework for multi-criteria fluence map optimization models. *Physics in Medicine and Biology*. 2004; 49:1991–2013. [PubMed: 15214537]
24. Ruzika S, Wiecek MM. Approximation methods in multiobjective programming. *Journal of Optimization Theory and Applications*. 2005; 126(3):473–501.
25. Shepard DM, Earl MA, Li XA, Naqvi S, Yu C. Direct aperture optimization: a turnkey solution for step-and-shoot IMRT. *Medical Physics*. 2002; 29:1007–1018. [PubMed: 12094970]
26. Shepard DM, Ferris MC, Olivera GH, Mackie TR. Optimizing the delivery of radiation therapy to cancer patients. *SIAM Review*. 1999; 41:721–744.
27. Teichert K, Süss P, Serna JJ, Monz M, Küfer KH, Thieke C. Comparative analysis of Pareto surfaces in multi-criteria IMRT planning. *Physics in Medicine and Biology*. 2011; 56(12):3669–3684. [PubMed: 21610294]
28. Vira, C.; Haimes, YY. *Multiobjective decision making: theory and methodology*. North Holland; New York: 1983.
29. Webb, S. *Intensity-Modulated Radiation Therapy*. Institute of Physics Publishing; Bristol, U.K: 2001.
30. Wilkens JJ, Alaly JR, Zakarian K, Thorstad WL, Deasy JO. IMRT treatment planning based on prioritizing prescription goals. *Physics in Medicine and Biology*. 2007; 52:1675–1692. [PubMed: 17327656]
31. Winston, WL. *Operations Research: applications and algorithms*. Brooks/Cole - Thomas Learning; Belmont, CA: 2004.

Appendix A. A mixed-integer programming formulation for the MCDAO pricing problem

In this section we discuss an exact method to solve the MCDAO pricing problem in (4) using mixed-integer programming (MIP) techniques. For a review of integer-programming modeling and solution techniques we refer to [31]. The objective function in formulation (4) has a mathematically inconvenient form and needs to be reformulated. In particular, consider aperture $k \in K$ and point $p \in P$: Their corresponding term in the objective function (i.e., $\min \{\mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*), 0\}$) can be substituted by variable $z_{pk} \geq 0$ via adding the following logical constraint to the problem:

$$\text{If } \mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*) < 0, \text{ then } z_{pk} \geq \mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*); \text{ else } z_{pk} \geq 0.$$

Moreover, we can replace the above logical constraint with the following set of constraints (see, e.g., [31]):

$$\begin{aligned} \mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*) - z_{pk} &\leq \Delta u_{pk} \\ -\mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*) &\leq \Delta(1 - u_{pk}) \\ -z_{pk} &\leq \Delta(1 - u_{pk}) \\ u_{pk} &\in \{0, 1\} \end{aligned}$$

where Δ is a sufficiently large user-specified parameter. Here, u_{pk} is a binary variable, which ultimately will take the value zero if $\mathbf{w}_p^\top \mathbf{r}_k(\mathbf{d}_p^*) < 0$ (i.e., if adding the aperture is guaranteed to locally improve plan p), and value one otherwise.

We next develop a binary representation of an aperture. For each beam angle, we consider a beamlet grid in which rows and columns are indexed by $n = 1, 2, \dots, N$ and $c = 1, 2, \dots, C$, respectively. For notational convenience, we let (n, c) represent the beamlet located at row n and column c . We then associate a binary decision variable $x_{(n,c)} \in \{0, 1\}$ with beamlet (n, c) indicating whether this beamlet is blocked ($x_{(n,c)} = 0$) in the aperture or not ($x_{(n,c)} = 1$). Thus, an aperture can be expressed as a binary vector $\mathbf{x} \in \{0, 1\}^{NC}$. To ensure that row convexity at each beamlet row n is met, $\mathbf{x}$ has to satisfy the following set of constraints:

$$x_{(n,c_1)} + x_{(n,c_2)} - x_{(n,c)} \leq 1 \quad \forall c, c_1, c_2: c_1 < c < c_2. \quad (\text{A.1})$$

The set of constraints in (A.1) ensure that if beamlets (n, c_1) and (n, c_2) are both exposed in the aperture, then any intermediate beamlet in that row (i.e., $\{(n, c) : c_1 < c < c_2\}$) should be exposed as well. Furthermore, we can aggregate the constraints in (A.1) corresponding to intermediate beamlets to obtain

$$x_{(n,c_1)} + x_{(n,c_2)} - \frac{1}{c_2 - c_1 - 1} \sum_{c=c_1+1}^{c_2-1} x_{(n,c)} \leq 1 \quad \forall c_1, c_2: c_1 + 1 < c_2.$$

Using this binary representation, we can express the orthogonal rate of change at plan p as

$$\sum_{n=1}^N \sum_{c=1}^C (\mathbf{w}_p^\top \mathbf{r}_{(n,c)}(\mathbf{d}_p^*)) x_{(n,c)},$$

where $\mathbf{r}_{(n,c)}(\mathbf{d}_p^*)$ is the vector of rates of change in treatment evaluation criteria at dose distribution $\mathbf{d}_p^*$ as the intensity of beamlet (n, c) increases. We can then reformulate the pricing problem in (4) as an MIP problem as follows:

$$\min \sum_{p \in P} z_p$$

subject to

$$\begin{aligned}
x_{(n,c_1)} + x_{(n,c_2)} - \frac{1}{c_2 - c_1 - 1} \sum_{c=c_1+1}^{c_2-1} x_{(n,c)} &\leq 1 & \forall n, c_1, c_2: c_1 + 1 < c_2 \\
\sum_{n=1}^N \sum_{c=1}^C (\mathbf{w}_p^\top \mathbf{r}_{(n,c)} (\mathbf{d}_p^*)) x_{(n,c)} - z_p &\leq \Delta u_p & p \in P \\
-\sum_{n=1}^N \sum_{c=1}^C (\mathbf{w}_p^\top \mathbf{r}_{(n,c)} (\mathbf{d}_p^*)) x_{(n,c)} &\leq \Delta (1 - u_p) & p \in P \\
-z_p &\leq \Delta (1 - u_p) & p \in P \\
z_p &\leq 0 & p \in P \\
u_p &\in \{0, 1\} & p \in P \\
x_{(n,c)} &\in \{0, 1\} & \forall n, c.
\end{aligned}$$

For limited number of sample points (i.e., $|P| < NC$), this formulation has $\mathcal{O}(NC)$ binary decision variables and $\mathcal{O}(NC^2)$ constraints. To determine the aperture with the largest improvement rate, the MIP formulation is solved for each beam direction considered in the problem. Finally, this formulation can be extended to incorporate MLC-hardware restrictions such as the leaf interdigitation constraint.

Figure 1.

The orthogonal rate of change in the Pareto surface at point $\mathbf{F}(\mathbf{d}_p^*)$ is defined as the scalar projection of $\mathbf{r}_k$ onto the normal unit vector $\mathbf{w}_p$. In this illustration, it is assumed that F_1 is the target objective. Thus, F_1 can be improved by adding the new aperture k so that the F_1 -component of $\mathbf{r}_k$ is negative. F_2 corresponds to an OAR objective. Thus, the F_2 -component of $\mathbf{r}_k$ is positive since adding aperture k with a small positive intensity locally increases F_2 .

Figure 2. Illustration of the solution approach: at each iteration, given the Pareto surface approximation associated with K' , a new aperture k is chosen and added to K' such that the Pareto surface is improved. The Pareto surface associated with the multi-criteria FMO problem dominates the MCDAO Pareto surface.

Figure 3. Representative CT slice for the paraspinal cancer case. The contours show the target volume (red), the spinal cord (orange) and the esophagus (green).

(a)

(b)

Figure 4. Comparing Pareto surfaces associated with (a) FMO, Balanced-, Target-, and OAR-IORC (b) FMO, Alter-IORC, and CORC, with MCDAO for the simplified case.

(a)

(b)

Figure 5. DVH curves associated with (a) balanced, OAR-, and target-emphasis plans marked on the corresponding Pareto surfaces (b) Pareto-optimal plans with a fixed target objective value, for the simplified case.

(a)

(b)

Figure 6. (a) DVH curves associated with the balanced, OAR-, and target-emphasis plans obtained for the complete case using the traditional column-generation method. (b) Comparing trade-off curves obtained for the complete case using IORC heuristics, CORC, and MCDAO.

Table 1

Number of voxels associated with each structure of the paraspinal case. In the optimization model, voxels in the lungs and skin are downsampled at a 1:4 and 1:8 ratio, respectively. Voxels that are not hit by any beamlets are not included.

Structure	Number of voxels	
	original	model
Target	7591	7591
Spinal cord	224	224
Esophagus	657	657
Left lung	39925	9976
Right lung	32755	8205
Skin	169202	21235
Total	250354	82376

Table 2

Normalized weight vectors, $\mathbf{w}_p / \|\mathbf{w}_p\|_\infty$ for $p \in P$, used in the simplified and complete cases. P contains three weight vectors.

Case	Weight vector	Target	Spinal cord	Lungs	Esophagus	Skin
Simplified	OAR emph.	1	0.14	0	0	0
	Balanced	1	0.07	0	0	0
	Target emph.	1	0.01	0	0	0
Complete	OAR emph.	1	0.20	0.02	0.02	0.02
	Balanced	1	0.10	0.02	0.02	0.02
	Target emph.	1	0.02	0.02	0.02	0.02