Phys. Med. Biol. **58** (2013) 7683–7697 doi:10.1088/0031-9155/58/21/7683

Deliverable navigation for multicriteria IMRT treatment planning by combining shared and individual apertures

,

Albin Fredriksson[1,2] **and Rasmus Bokrantz**[1,2]

[1] Optimization and Systems Theory, Department of Mathematics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden
[2] RaySearch Laboratories, Sveavägen 25, SE-111 34 Stockholm, Sweden

E-mail: albfre@kth.se and albin.fredriksson@raysearchlabs.com

Received 6 May 2013, in final form 19 August 2013
Published 14 October 2013
Online at stacks.iop.org/PMB/58/7683

Abstract
We consider the problem of deliverable Pareto surface navigation for step-and-shoot intensity-modulated radiation therapy. This problem amounts to calculation of a collection of treatment plans with the property that convex combinations of plans are directly deliverable. Previous methods for deliverable navigation impose restrictions on the number of apertures of the individual plans, or require that all treatment plans have identical apertures. We introduce simultaneous direct step-and-shoot optimization of multiple plans subject to constraints that some of the apertures must be identical across all plans. This method generalizes previous methods for deliverable navigation to allow for treatment plans with some apertures from a collective pool and some apertures that are individual. The method can also be used as a post-processing step to previous methods for deliverable navigation in order to improve upon their plans. By applying the method to subsets of plans in the collection representing the Pareto set, we show how it can enable convergence toward the unrestricted (non-navigable) Pareto set where all apertures are individual.

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

1. Introduction

A considerable fraction of the time required to generate intensity-modulated radiation therapy (IMRT) treatment plans is often spent on balancing target coverage against sparing of normal tissue. This time expense is related to the fact that manual tweaking of optimization parameters such as importance weights and dose–volume levels is often a blunt tool for controlling the shape of the planned dose distribution. Parameter-tuning moreover offers little overview of the full set of possible treatment plans, leading to a risk that the best options are not explored. As an aid in the decision making process, multicriteria optimization (MCO) methods that offer continuous exploration of the possible planning options have been developed (Craft *et al* 2007,

0031-9155/13/217683+15$33.00

Craft *et al* 2012, Monz *et al* 2008, Thieke *et al* 2007). Ultimately, the goal of these methods is to provide a comprehensive representation of all relevant treatment plans, along with the tools for a structured search over this set. To this end, the referenced methods, which we refer to as methods for Pareto surface navigation, typically use a database of precomputed treatment plans to approximate the Pareto optimal set. The Pareto set consists of the nondominated treatment plans, meaning the feasible plans such that there is no other feasible plan that is at least as good in all criteria and strictly better in one. By forming convex combinations in dose between the database plans, the treatment planner may continuously explore the tradeoffs of the considered patient case in real-time. In this paper, we consider application of such methods to treatment planning for step-and-shoot IMRT.

Pareto surface navigation was first developed with respect to treatment plans generated by optimization with respect to freely modulated fluence distributions, without consideration of the fact that an optimized fluence map requires conversion into multileaf collimator (MLC) apertures to be deliverable. Performing this conversion after the navigation need not degrade plan quality per se, but complicates decision making because the converted plan is not guaranteed to be acceptable only because the navigated plan is. At worst, fluence-based navigation may therefore result in an iterative process involving multiple rounds of adjustments and conversions, i.e., a form of treatment planning that MCO was intended to avoid. Pareto surface navigation where any navigated plan is deliverable would therefore greatly facilitate the practical use of MCO.

Bokrantz (2012) considered deliverable Pareto surface navigation for volumetric-modulated arc therapy. Following fluence-based navigation and conversion into a deliverable plan, a surface of plans with identical apertures was generated in a neighborhood of the deliverable plan, using segment weight optimization. Because all apertures are shared, navigation on this surface yields deliverable plans. This method thus allows for fine-tuning of the tradeoffs in the vicinity of the plan generated by the conversion algorithm.

The first step toward deliverable Pareto surface navigation for step-and-shoot IMRT was taken by Craft and Richter (2013). They used a database of segmented plans to represent the Pareto set. When convex combinations of such plans are taken, the resulting plan is deliverable, but with an increased number of apertures. They showed that the navigable approximation of the Pareto set does not deteriorate too much when the convex combinations of plans are restricted to just a few nonzero coefficients. When their convex combinations were restricted to two or three plans, they were still able to approximate a large number of points randomly sampled from a representation of a Pareto set (in an eight-dimensional objective space) with about respectively 10% and 5% deterioration in normalized objective value. Their approach to deliverable Pareto surface navigation was thus to combine few plans and accept a restriction of the navigation and an increase in the total number of apertures by a small integer factor ($\sim$3). It should be noted that there likely exists treatment plans of similar quality that can be delivered within fewer apertures than the plans produced by this form of navigation.

Salari and Unkelbach (2013) also considered the problem of deliverable Pareto surface navigation for step-and-shoot IMRT. Using column generation (Romeijn *et al* 2005), they created a single set of apertures well-suited for multiple plans, thereby creating a Pareto set representation where all treatment plans have identical apertures but different weightings of these. Navigation of this representation corresponds to a modification of segment weights and does therefore not increase the total number of apertures required for treatment plan delivery. Their approach to deliverable Pareto surface navigation was thus to maintain delivery time but sacrifice plan quality. The sacrifice in quality is likely to increase both with the number of plans in the Pareto set representation and with the number of objectives, because the restriction to identical apertures is more limiting if the apertures represent increasingly disparate planning goals.

The present paper concerns deliverable Pareto surface navigation by direct step-and-shoot (DSS) optimization (Hårdemark *et al* 2003) of multiple plans simultaneously, subject to constraints that enforce apertures to be shared among plans. The purpose of the simultaneous optimization of multiple plans is threefold:

(i) To improve upon the Pareto set representations generated by the previous approaches (Craft and Richter 2013, Salari and Unkelbach 2013) to deliverable Pareto surface navigation.

(ii) To generalize the previous methods for deliverable Pareto surface navigation to the situation where some apertures are shared between plans while others are individual.

(iii) To create Pareto set representations that converge to the unrestricted Pareto set as the number of plans increases.

2. Methods

We generate Pareto set representations consisting of plans for which a subset of the apertures is shared across the plans. If there are no shared apertures, then each plan in the Pareto set representation may be optimized independently. This property does not hold when there are shared apertures, because a change to the apertures of one plan may then affect the other plans. We therefore give an optimization problem formulation that explicitly incorporates a possible coupling between the plans. This optimization problem amounts to DSS optimization of all plans in the Pareto set representation simultaneously, subject to the constraint that the leaf position variables associated with some apertures must be identical across all plans.

2.1. Multicriteria direct step-and-shoot optimization

Direct step-and-shoot optimization is considered with respect to minimization of n conflicting criteria $f_1, \ldots, f_n$, with $n \geqslant 2$. All criteria are assumed to be functions of the dose distribution d, which in turn is a function of the MLC leaf positions λ and the nonnegative segment weights σ. The set of feasible leaf positions is denoted by Λ. Furthermore, the problem is subject to m constraints given by the functions of dose $c_1, \ldots, c_m$, which are assumed to be satisfied whenever they evaluate to zero or less. The MCO problem is then formulated

$$\begin{array}{lll} \underset{\lambda,\sigma}{\text{minimize}} & [f_1(d(\lambda,\sigma)) \cdots f_n(d(\lambda,\sigma))]^T & \\ \text{subject to} & c_i(d(\lambda,\sigma)) \leqslant 0, & i = 1, \ldots, m, \\ & \lambda \in \Lambda, & \\ & \sigma \geqslant 0. & \end{array} \tag{2.1}$$

The vector-valued objective function of (2.1) has no unambiguously defined unique minimizer in general. Instead, all elements in the entire set of Pareto optimal points to (2.1) are considered to be equally valid minimizers. The Pareto set is typically composed of an infinite number of points and can therefore not be easily determined. Instead, it may be approximated by a finite number of points, each obtained by minimization of a scalar-valued substitute for the vector-valued objective of problem (2.1) (Craft *et al* 2006, Rennen *et al* 2011, Bokrantz and Forsgren 2013). Different scalarization methods can be used to find such points. Here, we use the common technique of minimizing a weighted sum of the objective functions. Such scalarization amounts to the problem

$$\begin{array}{lll} \underset{\lambda,\sigma}{\text{minimize}} & \displaystyle\sum_{i=1}^{n} w_i f_i(d(\lambda,\sigma)) & \\ \text{subject to} & c_i(d(\lambda,\sigma)) \leqslant 0, & i = 1, \ldots, m, \\ & \lambda \in \Lambda, & \\ & \sigma \geqslant 0, & \end{array} \tag{2.2}$$

where w_i is the nonnegative weight for criterion i for $i = 1, \ldots, n$. The set of optimal solutions to a number of instances of problem (2.1) with different weights is then used to approximate the Pareto set.

In the present paper, DSS optimization with respect to the problem (2.2) is performed using gradient-based search with leaf positions and segment weights as variables, as described by Hårdemark *et al* (2003). The segment weights must be nonnegative while the leaf positions are subject to linear constraints that model the hardware limitations of the MLC system, e.g., bounds on the minimum leaf tip gap and interdigitation limitations. Alternative methods for optimization with respect to problem (2.2) include column generation methods where deliverable apertures are created iteratively during the optimization and then kept fixed (see, e.g., Romeijn *et al* (2005)), and stochastic search methods that change the shape of the apertures during optimization (Shepard *et al* 2002). Methods that combine column generation and gradient-based search have also been considered (Carlsson 2008, Cassioli and Unkelbach 2013). Because the dose distribution is a nonconvex function of the leaf positions, the referenced methods are not guaranteed to converge to the global optimum in finite time.

2.2. Direct step-and-shoot optimization using shared apertures

When some apertures are to be shared across the plans that represent the Pareto set, changing the apertures of one plan may affect the other plans. In order to take this coupling between the plans into account, we optimize all plans simultaneously. Assume that the Pareto set is approximated by n_p plans. For each plan p, let λ_p represent the leaf positions of all apertures and let σ_p be the vector of all segment weights. The shared apertures of plan p are denoted by $\lambda_{p,\mathcal{S}}$. Note that $\mathcal{S}$ is an index set over the leaves of the shared apertures and that $\lambda_{p,\mathcal{S}}$ is a subvector of λ_p. The multiple plan optimization problem may then be formulated as

$$\begin{array}{lll} \underset{\lambda,\sigma}{\text{minimize}} & \displaystyle\sum_{p=1}^{n_p}\sum_{i=1}^{n} w_{p,i} f_i(d(\lambda_p,\sigma_p)) & \\ \text{subject to} & c_i(d(\lambda_p,\sigma_p)) \leqslant 0, & p = 1,\ldots,n_p, \quad i = 1,\ldots,m, \\ & \lambda_p \in \Lambda, & p = 1,\ldots,n_p, \\ & \lambda_{p,\mathcal{S}} = \lambda_{p+1,\mathcal{S}}, & p = 1,\ldots,n_p - 1, \\ & \sigma_p \geqslant 0, & p = 1,\ldots,n_p. \end{array} \tag{2.3}$$

The coupling between the plans is comprised of the equality constraints, which enforce the leaf positions of the shared apertures to be equal over all plans.

Formulation (2.3) yields the best mean approximation of the Pareto set. When starting the optimization from a point in which plan p has objective value $\bar{z}_p$ for $p = 1, \ldots, n_p$, it is therefore possible that the optimization makes some of the plans worse than their respective values $\bar{z}_p$, provided the objective values of other plans improve more. A formulation that yields the best approximation of the Pareto set in the worst case (and that ensures that no plan deteriorates) is achieved by a change of the sum in (2.3) into a maximum operator applied to the difference between the plan objectives and $\bar{z}_p$. The new objective function is given by

$$\max_{p=1,\ldots,n_p}\left(\sum_{i=1}^{n} w_{p,i} f_i(d(\lambda_p,\sigma_p)) - \bar{z}_p\right).$$

Because the standard DSS optimization problem is nonconvex, the simultaneous optimization of multiple plans problem (2.3) is too. Since there may be multiple local optima, the solutions found by gradient-based methods depend on the initial values of the variables. Methods for initialization are elaborated in section 2.4.

2.3. Convergence toward the unrestricted Pareto set

The quality of segmented IMRT plans depends on the number of apertures that the plans are allowed to use. For each positive integer $n_{\mathcal{I}}$, there is a Pareto set consisting of plans with $n_{\mathcal{I}}$ apertures (which are individual to the plans). We call this set the $n_{\mathcal{I}}$-aperture Pareto set. Because an increased number of apertures yields an increased number of degrees of freedom, the $(n_{\mathcal{I}} + 1)$-aperture Pareto set dominates the $n_{\mathcal{I}}$-aperture Pareto set.

Navigation by the method of Craft and Richter (2013) uses convex combinations of k plans with $n_{\mathcal{I}}$ apertures each. When $k \geqslant 2$ plans from the $n_{\mathcal{I}}$-aperture Pareto set are combined, the quality of the resulting plan is at best as good as the plans on the $n_{\mathcal{I}}$-aperture Pareto set, but the plan requires $kn_{\mathcal{I}}$ apertures to be delivered. Because the $kn_{\mathcal{I}}$-aperture Pareto set dominates the $n_{\mathcal{I}}$-aperture Pareto set, there exists a plan on the $kn_{\mathcal{I}}$-aperture Pareto set that requires the same number of apertures to be delivered as the navigated plan, but that has as good or better objective values. For a fixed number of apertures, the navigated plans therefore never reach the unrestricted Pareto optimal set where all apertures are individual, no matter how many plans that are included in the representation of the Pareto set.

The method of Salari and Unkelbach (2013) constrains the apertures to be identical across all plans, which is a restriction of the feasible set to the treatment plan optimization problem. A Pareto set representation composed of plans with $n_{\mathcal{I}}$ shared apertures is therefore dominated by the $n_{\mathcal{I}}$-aperture Pareto set, in which all plans have $n_{\mathcal{I}}$ individual apertures. Also for this method, the navigated plans may therefore never reach the unrestricted Pareto set with all individual apertures, regardless of the number of plans in the representation. Because formulation (2.3) is a form of interpolation between the method of Craft and Richter and the method of Salari and Unkelbach, it also suffers from this shortcoming.

To obtain a method that yields representations that approach the unrestricted Pareto set, we propose an extension of the method introduced in section 2.1. Note that we do not claim convergence to the exact, globally optimal, Pareto set, but only to a set of locally optimal solutions to instances of problem (2.2) with the weights sufficiently densely sampled over the space of admissible weights. Convergence to the exact Pareto set requires a globally convergent method for DSS optimization and a scalarization other than weights, which can find all Pareto optimal solutions to a nonconvex problem. Nevertheless, we refer to the proposed extension as the 'convergent method.'

In the convergent method that we propose, we consider a representation where all apertures are individual, and navigation is performed by convex combinations of k plans at the time. So far, the convergent method is identical to the method of Craft and Richter, i.e., it gives a representation of the $n_{\mathcal{I}}$-aperture Pareto set where a navigated plan in general requires $kn_{\mathcal{I}}$ apertures to be delivered. Now suppose that formulation (2.3) is applied to each k-tuple of plans that can be formed during navigation (the k-tuples that correspond to $(k-1)$-dimensional faces of the Pareto set representation), and that all $kn_{\mathcal{I}}$ apertures of these plans are shared between the plans ($\mathcal{S}$ is set to the index set over the leaves of all $kn_{\mathcal{I}}$ apertures). The simultaneous optimization of multiple plans then yields a set of $kn_{\mathcal{I}}$-aperture plans with $kn_{\mathcal{I}}$ shared apertures that dominates the initial k-tuple of plans.

The convergent method is schematically illustrated in figure 1 with respect to a two-dimensional tradeoff. The figure shows the ideal non-navigable $2n_{\mathcal{I}}$-aperture Pareto set, a navigable approximation of this set using plans with $n_{\mathcal{I}}$ individual apertures, and an improved navigable approximation composed of plans with $2n_{\mathcal{I}}$ shared apertures, obtained by further optimization with respect to neighboring pairs of plans from the $n_{\mathcal{I}}$-aperture Pareto set. The interpretation of the figure is as follows. Plans p_1 and p_2 have $n_{\mathcal{I}}$ individual apertures each. These plans are optimal among plans with $n_{\mathcal{I}}$ apertures, and thus reside on the $n_{\mathcal{I}}$-aperture

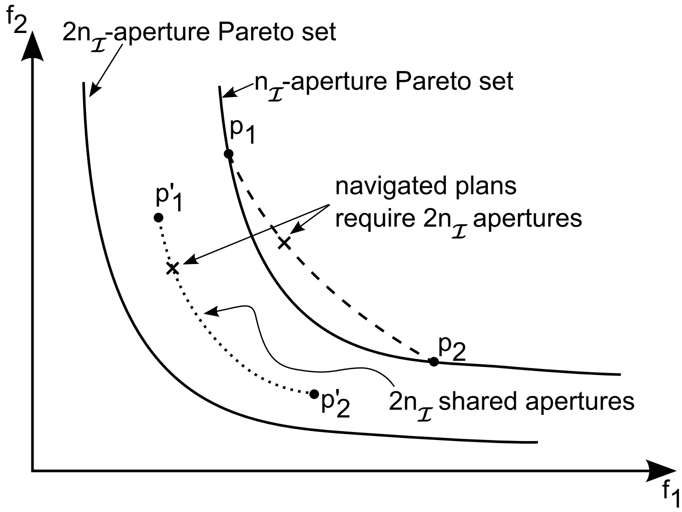

Figure 1. Dominance relations between ideal Pareto sets (solid), and navigable approximations of these sets produced by the method of Craft and Richter (2013) (dashed) and the convergent method (dotted).

Pareto set. A navigated plan that is a convex combination of p_1 and p_2 requires $2n_{\mathcal{I}}$ apertures to be delivered, but its quality is as best as good as that of a plan on the $n_{\mathcal{I}}$-aperture Pareto set. Further optimization of p_1 and p_2 according to (2.3), with $2n_{\mathcal{I}}$ shared apertures as variables and starting from the apertures of plans p_1 and p_2, leads to the improved $2n_{\mathcal{I}}$-aperture plans p'_1 and p'_2. These plans have a quality that is intermediate to plans on the $n_{\mathcal{I}}$-aperture Pareto set and plans on the $2n_{\mathcal{I}}$-aperture Pareto set. A convex combination of p'_1 and p'_2 requires $2n_{\mathcal{I}}$ apertures to be delivered.

As the number of plans in the initial representation of the $n_{\mathcal{I}}$-aperture set increases, the weights associated with the plans in each k-tuple tend toward each other. In figure 1, this corresponds to that the points p_1 and p_2 tend toward each other. This implies that the optimal apertures for these plans approach a single set of apertures that define a plan on the $kn_{\mathcal{I}}$-aperture Pareto set. In figure 1, this would mean that points p'_1 and p'_2 would approach a single point on the $2n_{\mathcal{I}}$-aperture Pareto set. The convergent method thus produces representations that tend toward the unrestricted Pareto set as the number of plans in the representations increases. The disadvantage of the method is that it contains no requirement that different k-tuples of plans should patch up smoothly. The optimization therefore transforms a continuous representation of the Pareto set into a set of non-connected patches. Navigation of a representation generated using the convergent method is therefore only piecewise continuous.

2.4. Computational study

We implemented multiple plan DSS optimization in a research version of RayStation 2.9 (RaySearch Laboratories, Stockholm, Sweden). The equality constraints in (2.3) were eliminated by variable reduction. The DSS optimization of RayStation optimizes leaf positions and segment weights (Hårdemark *et al* 2003), using a sequential quadratic programming algorithm similar to that described by Gill *et al* (2005). The optimization is performed with respect to dose calculated using a singular-value decomposition of pencil beam kernels, similar to Bortfeld *et al* (1993).

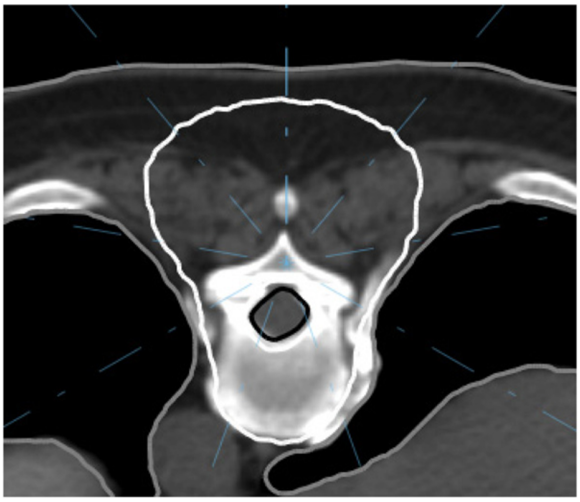

Figure 2. A transversal slice of the paraspinal case. The target is indicated by a white contour and the spinal cord by a black contour.

Table 1. Optimization problem formulation for the paraspinal case. The two-dimensional experiment considered the first two objectives only.

Objectives			Constraints		
Structure	Function	Dose level (Gy)	Structure	Function	Dose level (Gy)
Target	Uniform dose	60	Target	Minimum dose	53
Spinal cord	Maximum dose	0	Target	Maximum dose	69
External	Maximum dose	0	External	Maximum dose	69
			Spinal cord	Maximum dose	55

The implementation of the proposed method was evaluated with respect to a single patient case. Salari and Unkelbach (2013) applied their method to a paraspinal case with a bicriteria tradeoff between target homogeneity and sparing of the spinal cord. Similarly, we applied our method to a paraspinal case and planning criteria modeled by structure-specific least-squares penalties between the planned dose and reference dose levels. A transversal slice of the considered patient case is shown in figure 2. Treatment planning for this case was performed with respect to nine equispaced coplanar beams, a dosegrid resolution of $3 \times 3 \times 3$ mm^3, and a fluence grid resolution of 5×5 mm^2. We first considered a two-dimensional tradeoff, for which the results can be easily visualized. Second, we considered a three-dimensional tradeoff.

The optimization problem formulation used for the paraspinal case is presented in table 1. The optimization functions penalize quadratic (semi)deviation of dose. A maximum dose function with reference dose level $\hat{d}$ applied to a subset $\mathcal{V}$ of the patient volume is given by

$$\int_{\mathcal{V}} \max\{d_v - \hat{d}, 0\}^2 \, \mathrm{d}v, \tag{2.4}$$

where d_v denotes the dose to the point v in $\mathcal{V}$. Minimum dose and uniform dose functions have respectively the minimum and the identity operator substituted for the maximum in (2.4). Optimization functions posed as constraints are required to evaluate to zero or less. The vectors of weights w_p for the plans $p = 1, \ldots, n_p$ were throughout sampled from a grid of the set $\left\{w \in \mathbb{R}^n : \sum_{i=1}^n w_i = 1, w \geqslant 0\right\}$.

We characterize the effect of sharing different numbers of apertures qualitatively by plotting the resulting approximations of the Pareto sets. The objective function values of the plans with 60 individual segments are normalized to lie in the interval [0, 1], and the objective function values of the other plans are modified using the same scaling. We also give quantitative measures of how the Pareto sets deteriorate as the number of shared apertures increases. To this end, we measure how well each Pareto set representation approximates the representations consisting of plans with individual apertures only. As in Bokrantz and Forsgren (2013), how well a set Z approximates a set Z' is measured by the Hausdorff distance

$$d_H(Z, Z') = \max_{z' \in Z'} \min_{z \in Z} d(z, z') \tag{2.5}$$

with respect to the premetric

$$d(z, z') = \max_{i=1,\ldots,n} \max\{z_i - z_i', 0\}.$$

Since we apply a local search method to a nonconvex optimization problem, the solution will depend on the starting point. To facilitate a fair comparison of results, we therefore use identical starting points for a given number of segments regardless of how many of the apertures that are shared. The starting point used in the numerical experiments of this paper is constituted of the leaf positions and segment weights obtained by FMO with respect to the direct sum of all objectives, i.e., a plan representing all goals to some extent, followed by leaf-sequencing using the conversion algorithm of RayStation, described by Hårdemark *et al* (2003). We then separate the apertures into a set of shared apertures and a set of individual apertures. To select the apertures to be shared, we loop over the beams and select the first aperture for each considered beam until the desired number of shared apertures has been selected. If a beam has only one aperture left, it is not selected, since we want at least one individual aperture for each beam. The apertures resulting from the utilized leaf-sequencing algorithm are typically ordered roughly in size, with larger apertures first.

The initialization outlined above is arguably a crude method. It is therefore likely that better starting points can be obtained if a custom-made multicriteria leaf-sequencer is used that explicitly takes all planning criteria into account, and that also accounts for the fact that the plans have individual segment weights. The method for multicriteria column generation presented by Salari and Unkelbach (2013) is a viable such method. As they point out, their method can be used to generate a set of shared apertures, and in addition, some apertures that are individual to the plans. Formulation (2.3) can then be directly applied to the resulting Pareto set representation, as to further improve the quality of its constituent plans. For single-criterion IMRT planning, local refinement by gradient-based optimization of leaf positions and segment weights have been shown to improve the quality of plans generated by column generation (Carlsson 2008, Cassioli and Unkelbach 2013).

3. Results

3.1. Two-dimensional tradeoff

We first studied the two-dimensional tradeoff between target homogeneity and sparing of the spinal cord. Simultaneous optimization of multiple plans according to formulation (2.3) was performed with the number of plans n_p set to 5.

Figure 3 depicts the resulting Pareto set representations as a function of number of shared and number of individual apertures. Figure 4 depicts the dose–volume histograms (DVHs) of the resulting plans with a fixed objective value for the criterion on the target. Results are also shown for the convergent method. Note that delivery of the navigated plans

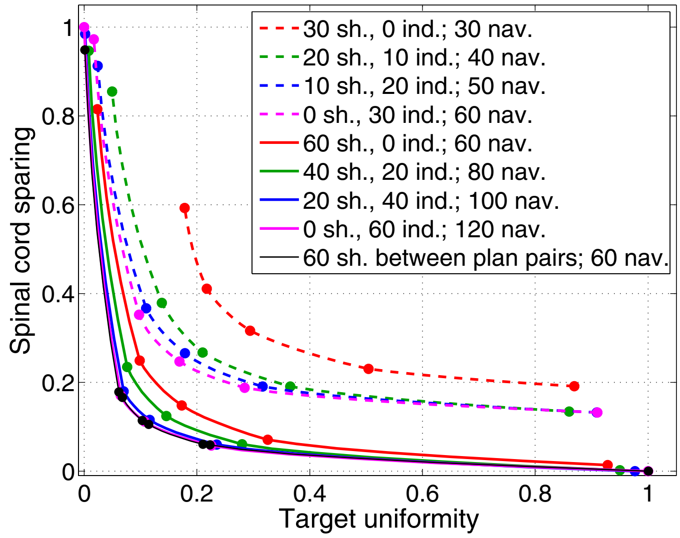

Figure 3. Pareto set representations for the two-dimensional tradeoff, as a function of the number of shared (sh.) and individual (ind.) apertures. The notation 'x nav.' means that x apertures are required for a navigated plan to be deliverable.

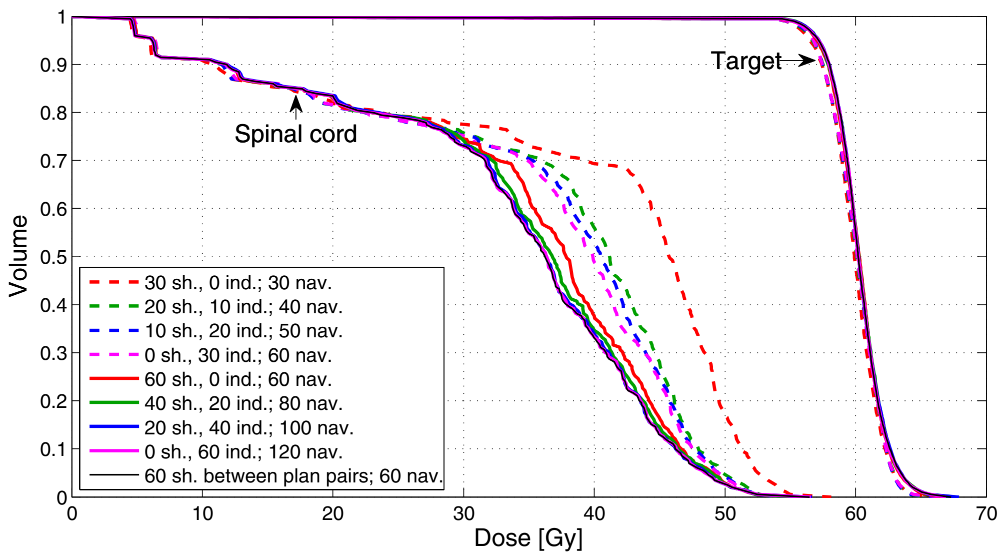

Figure 4. Dose–volume histograms for the two-dimensional tradeoff, of plans with a fixed objective value for the criterion on the target and a variable number of shared (sh.) and individual (ind.) apertures. The notation 'x nav.' means that x apertures are required for a navigated plan to be deliverable.

from these representations require $n_{\mathcal{S}} + 2n_{\mathcal{I}}$ apertures in general, where $n_{\mathcal{S}}$ is the number of shared apertures and $n_{\mathcal{I}}$ is the number of individual apertures. When the total number of segments is small (30 segments), constraints that all apertures must be equal leads to highly deteriorated plan quality and a small Pareto set representation. When a few individual apertures are allowed, most of this deterioration is counterbalanced. With a larger number of segments (60 segments), the difference between the Pareto set representations decreases. Moreover, the

Table 2. Approximation error according to (2.5) between Pareto set representations with variable number of shared (sh.) and individual (ind.) apertures and unrestricted Pareto set representations with 30 and 60 apertures for the two-dimensional tradeoff. The notation 'x nav.' means that x apertures are required for a navigated plan to be deliverable.

	Approximation error (%)	
Representation	0 sh., 30 ind.	0 sh., 60 ind.
30 sh., 0 ind.; 30 nav.	16.1	20.1
20 sh., 10 ind.; 40 nav.	3.9	13.8
10 sh., 20 ind.; 50 nav.	1.6	13.2
0 sh., 30 ind.; 60 nav.	0.0	13.2
60 sh., 0 ind.; 60 nav.	–	5.2
40 sh., 20 ind.; 80 nav.	–	2.9
20 sh., 40 ind.; 100 nav.	–	0.8
0 sh., 60 ind.; 120 nav.	–	0.0
60 sh.[a], 0 ind.; 60 nav.	–	0.5

[a] shared between pairs of plans.

60-segment representation with all apertures shared dominates the 30-segment representation with all apertures individual, which results in 60-segment plans as convex combinations are taken. The convergent method with 60 apertures shared between pairs of plans results in solutions very close to the Pareto set of plans of plans with 60 individual segments.

The approximation errors between the representations consisting of plans with different numbers of shared and individual apertures and those with only individual apertures are presented in table 2. The tabulated results show that when a few of the constraints enforcing apertures to be equal are removed, there is a large increase in approximation quality (at the cost of an increased number of segments), but as the number of individual apertures increases, the benefit of increasing the fraction of individual apertures decreases. The convergent method has only half a percent approximation error from the representation using only individual apertures.

The representations with 30 individual apertures, 60 shared apertures, and 60 shared apertures between plan pairs all result in navigated plans that are deliverable within 60 apertures. These representations approximate the unrestricted 60-segment Pareto set to within respectively 13%, 5%, and 0.5%.

3.2. Three-dimensional tradeoff

To see how the different methods scale with an increased number of criteria, we studied the three-dimensional tradeoff between target homogeneity, sparing of the spinal cord, and integral dose. The multiple plan optimization according to formulation (2.3) was performed with $n_p = 15$.

Since full three-dimensional Pareto sets cannot be easily visualized, we consider two-dimensional cuts of these sets to illustrate how the constraints on different numbers of shared apertures affect the solutions. Cuts are taken along the planes $\{x \in \mathbb{R}^3 : x_i = 0.4\}$ for $i = 1, 2, 3$. The resulting two-dimensional Pareto set representation slices are depicted in figure 5. The DVHs of the resulting plans with fixed objective values for the criteria on the target and external volume are shown in figure 6. Note that delivery of the navigated plans from these representations require $n_{\mathcal{S}} + 3n_{\mathcal{I}}$, where $n_{\mathcal{S}}$ is the number of shared apertures and $n_{\mathcal{I}}$ is the number of individual apertures. If the number of nonzero coefficients in the convex combinations of plans is restricted to 2, the coefficient 3 may be reduced to 2. When a small number of apertures is used (30 apertures), the constraints that all apertures must be equal

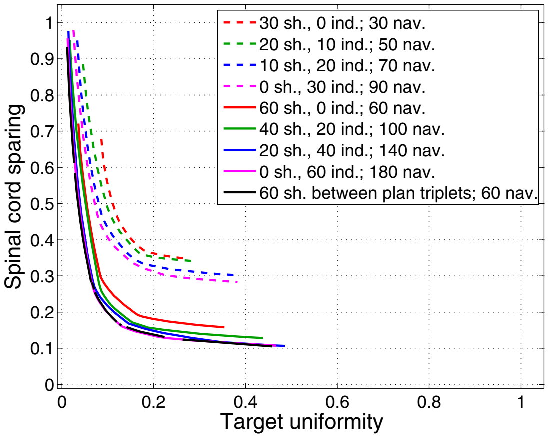

(a)

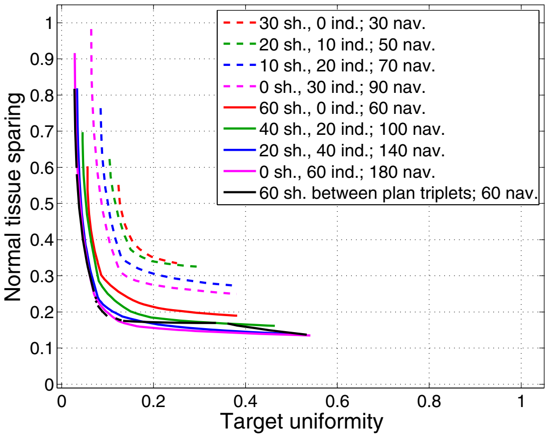

(b)

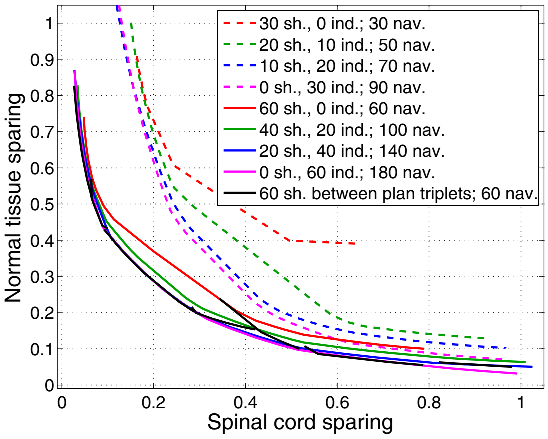

(c)

Figure 5. Two-dimensional cuts of three-dimensional Pareto set representations for the three-dimensional tradeoff, as a function of the number of shared (sh.) and individual (ind.) apertures. The notation 'x nav.' means that x apertures are required for a navigated plan to be deliverable.

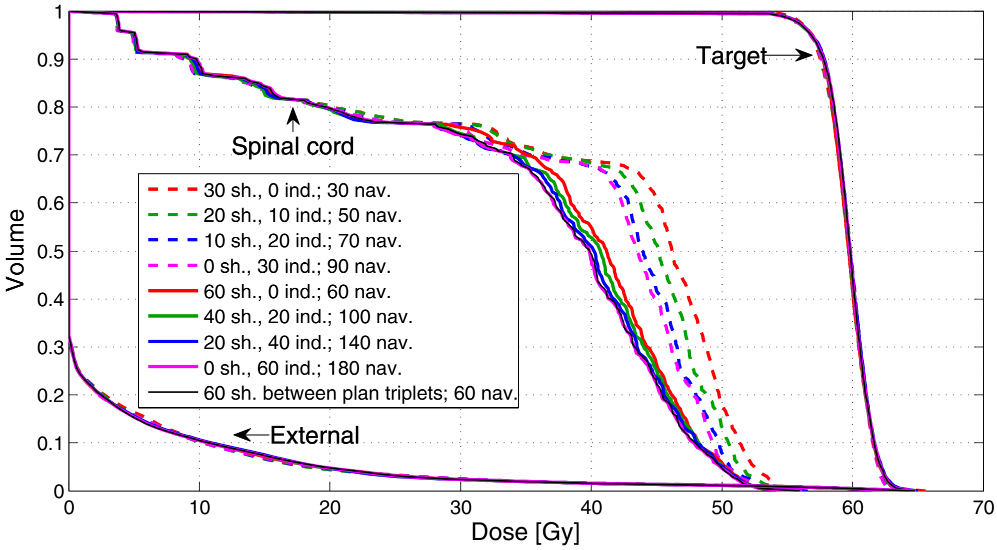

Figure 6. Dose–volume histograms for the three-dimensional tradeoff, of plans with fixed objective values for the criteria on the target and external volume and a variable number of shared (sh.) and individual (ind.) apertures. The notation 'x nav.' means that x apertures are required for a navigated plan to be deliverable.

Table 3. Approximation errors according to (2.5) between Pareto set representations with variable number of shared (sh.) and individual (ind.) apertures and unrestricted Pareto set representations with 30 and 60 apertures for the three-dimensional tradeoff. The notation 'x nav.' means that x apertures are required for a navigated plan to be deliverable.

	Approximation error (%)	
Representation	0 sh., 30 ind.	0 sh., 60 ind.
30 sh., 0 ind.; 30 nav.	14.1	17.9
20 sh., 10 ind.; 50 nav.	8.5	15.9
10 sh., 20 ind.; 70 nav.	4.3	15.3
0 sh., 30 ind.; 90 nav.	0.0	13.3
60 sh., 0 ind.; 60 nav.	–	8.8
40 sh., 20 ind.; 100 nav.	–	4.1
20 sh., 40 ind.; 140 nav.	–	2.9
0 sh., 60 ind.; 180 nav.	–	0.0
60 sh.[a], 0 ind.; 60 nav.	–	2.6

[a] shared between triplets of plans.

are highly restrictive. The benefit of allowing for a few individual apertures is large for the tradeoff between the two healthy structures (but results in an increased number of apertures), but not as evident for the tradeoff between target uniformity and healthy structure sparing as for the two-dimensional case. The differences are smaller when a larger number of segments is used (60 apertures). The convergent method with 60 apertures shared between triplets of plans results in solutions very close to the Pareto set representations of plans with 60 individual apertures.

The approximation errors between the representations consisting of plans with different numbers of shared and individual apertures and those with only individual apertures are presented in table 3. As for the two-dimensional case, these approximation errors quantify the comparatively large benefit of including some individual apertures and moreover show that as the fraction of individual apertures increases, the benefit of a further increase diminishes. The distance between the representation with 60 shared apertures and that with 60 individual apertures is larger than the corresponding distance for the two-dimensional case, and the benefit of adding a few individual apertures is smaller. The convergent method yields the best approximation of the representation consisting of plans with 60 individual apertures.

The representations with 30 individual apertures, 60 shared apertures, and 60 shared apertures between plan pairs all result in navigated plans that are deliverable within 60 apertures (if convex combinations from the representation with 30 individual apertures are restricted to two nonzero coefficients). These representations approximate the unrestricted 60-aperture Pareto set to within respectively 13% , 9%, and 3%.

4. Discussion

It has been hypothesized by Craft and Richter (2013) that the use of a few shared apertures can be beneficial because some characteristics are shared among the different Pareto optimal plans. Our numerical results support this hypothesis: there was a large benefit of including a few individual apertures (however at the cost of an increased number of apertures). As the fraction of individual apertures increased, the benefit of a further increase diminished; individual apertures were of diminishing marginal utility. Our data shows that for the considered low-dimensional case, it is better to share apertures than to use individual apertures, because the representations with 60 shared apertures dominated those with 30 individual apertures. Both of these representations result in deliverable plans with 60 apertures (if the latter method is restricted to convex combinations of two plans). It is important to note that the optimal partitioning is highly dependent on the optimization problem formulation. If the problem is formulated using narrow constraints, then all treatment plans will be rather similar and can therefore be well approximated by a single set of apertures. If the constraints are loose, then the Pareto optimal set will contain very disparate plans and therefore require individual apertures. Finally, as pointed out in section 2.4, to make the most out of any method that shares apertures, the shared apertures should be taken into account already in the leaf-sequencing or aperture generation. One method that accomplishes this is that of Salari and Unkelbach (2013).

Neither the method with only individual apertures nor that with only shared apertures came close to the unrestricted Pareto set. The method with only individual apertures resulted in about 13% approximation error compared to the unrestricted Pareto set for both the two-dimensional and the three-dimensional tradeoff, while the method with only shared apertures resulted in 5% and 9% approximation error for the two- and three-dimensional tradeoffs, respectively. The convergent method that we proposed cured this shortcoming by providing solutions that came very close to the unrestricted Pareto set. In some points, it even dominated the unrestricted Pareto set, a result that can be attributed to the fact that the optimization problems are nonconvex and that the different optimizations can therefore find different local minima. For the considered case, the convergent method yielded approximation errors of less than 1% and 3% for the two- and three-dimensional tradeoffs, respectively. The restriction that the apertures are shared between subsets of plans does not lead to any noticeable deterioration in objective values compared to the plans with 60 individual apertures.

In the first method that we proposed, there is a global pool of shared apertures, which are used by all plans. In the convergent method, there are local pools of shared apertures, which are shared across the plans in a neighborhood. The use of local pools of shared apertures is a significant drawback of the convergent method because it leads to discontinuities

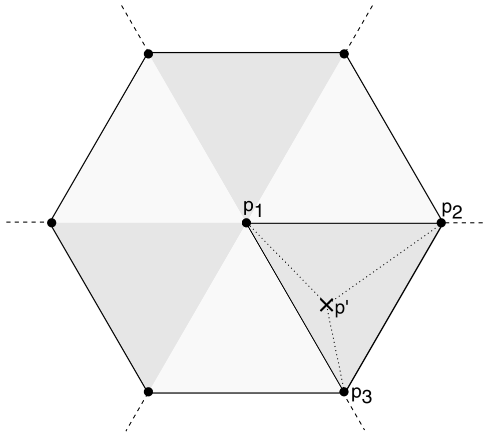

Figure 7. Navigation with respect to a Pareto set representation for a three-dimensional tradeoff. In general, each point has six neighbors. The navigated plan p' is a convex combination of p_1, p_2, and p_3, and requires 150 apertures to be delivered if each plan in the representation has 60 apertures and shares 10 apertures with each of its neighbors.

when the navigation shifts from one pool to another. A method that uses local pools of shared apertures but enforces continuity between the neighborhoods would potentially enable continuous navigation with high plan quality. The main obstacle of creating such a method is that the number of neighboring plans rapidly becomes large with increasing dimension. To exemplify, consider a three-dimensional tradeoff and a Pareto set representation that is isomorphic to a uniform tiling of the plane with equilateral triangles. Then, a general plan in the representation has six neighbors, as illustrated in figure 7. Suppose that each plan shares ten apertures with each of its neighbors and that it has no individual apertures in addition to these. Then, a general convex combination of three plans has $3 \times 60 = 180$ apertures, out of which $3 \times 10 = 30$ have been counted twice. A navigated plan thus requires 150 apertures to be delivered. Already in dimension 3, the saving in number of apertures is thus relatively small compared to the unrestricted case where each plan has 60 individual apertures and a navigated plan requires 180 apertures to be delivered.

5. Conclusion

The potential of deliverable Pareto surface navigation is clear: it not only allows the treatment planner to interactively explore the possible tradeoffs between competing planning criteria, but also keeps its promises in the sense that the plans that are seen during navigation can actually be delivered. To accomplish deliverable navigation, we generalized DSS optimization to a multicriteria setting in which all plans representing the Pareto set are simultaneously optimized under constraints that a subset of the apertures must be identical across plans. Such optimization can be used to improve upon the previous methods for deliverable Pareto surface navigation. Furthermore, we used the simultaneous DSS optimization of multiple plans to combine the benefits of the two previous approaches to deliverable navigation, which either combined a small number of plans or combined plans with the same apertures, and thereby created an improved approach with the possibility of converging to the unrestricted Pareto set, however with the drawback of discontinuities in the navigation.

Acknowledgments

We thank Anders Forsgren and Björn Hårdemark for constructive comments. The research of the first author was supported by the Swedish Research Council (VR).

References

Bokrantz R 2012 Multicriteria optimization for volumetric-modulated arc therapy by decomposition into a fluence-based relaxation and a segment weight-based restriction *Med. Phys.* **39** 6712–25

Bokrantz R and Forsgren A 2013 An algorithm for approximating convex Pareto surfaces based on dual techniques *INFORMS J. Comput.* **25** 377–93

Bortfeld T, Schlegel W and Rhein B 1993 Decomposition of pencil beam kernels for fast dose calculations in three-dimensional treatment planning *Med. Phys.* **20** 311–8

Carlsson F 2008 Combining segment generation with direct step-and-shoot optimization in intensity-modulated radiation therapy *Med. Phys.* **35** 3828

Cassioli A and Unkelbach J 2013 Aperture shape optimization for IMRT treatment planning *Phys. Med. Biol.* **58** 301

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

Craft D, Halabi T, Shih H and Bortfeld T 2007 An approach for practical multiobjective IMRT treatment planning *Int. J. Radiat. Oncol. Biol. Phys.* **69** 1600–7

Craft D, Hong T, Shih H and Bortfeld T 2012 Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy *Int. J. Radiat. Oncol. Biol. Phys.* **82** e83–90

Craft D and Richter C 2013 Deliverable navigation for multicriteria step and shoot IMRT treatment planning *Phys. Med. Biol.* **58** 87–103

Gill P, Murray W and Saunders M 2005 SNOPT: an SQP algorithm for large-scale constrained optimization *SIAM Rev.* **47** 99–131

Hårdemark B, Liander A, Rehbinder H and Löf J 2003 Direct machine parameter optimization with RayMachine® in Pinnacle³® White Paper, RaySearch Laboratories, Stockholm, Sweden

Monz M, Küfer K-H, Bortfeld T and Thieke C 2008 Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning *Phys. Med. Biol.* **53** 985–98

Rennen G, van Dam E and den Hertog D 2011 Enhancement of sandwich algorithms for approximating higher-dimensional convex Pareto sets *INFORMS J. Comput.* **23** 493–517

Romeijn E, Ahuja R, Dempsey J and Kumar A 2005 A column generation approach to radiation therapy treatment planning using aperture modulation *SIAM J. Optim.* **15** 838–62

Salari E and Unkelbach J 2013 A column-generation-based method for multi-criteria direct aperture optimization *Phys. Med. Biol.* **58** 621–39

Shepard D, Earl M, Li X, Naqvi S and Yu C 2002 Direct aperture optimization: A turnkey solution for step-and-shoot IMRT *Med. Phys.* **29** 1007–18

Thieke C, Küfer K-H, Monz M, Scherrer A, Alonso F, Oelfke U, Huber P, Debus J and Bortfeld T 2007 A new concept for interactive radiotherapy planning with multicriteria optimization: first clinical evaluation *Radiother. Oncol.* **85** 292–8