Published in final edited form as:

Phys Med Biol. 2013 January 7; 58(1): 87–103. doi:10.1088/0031-9155/58/1/87.

Deliverable navigation for multicriteria step and shoot IMRT treatment planning

David Craft and Christian Richter

Department of Radiation Oncology, Massachusetts General Hospital

Abstract

We consider Pareto surface based multi-criteria optimization for step and shoot IMRT planning. By analyzing two navigation algorithms, we show both theoretically and in practice that the number of plans needed to form convex combinations of plans during navigation can be kept small (much less than the theoretical maximum number needed in general, which is equal to the number of objectives for on-surface Pareto navigation). Therefore a workable approach for directly deliverable navigation in this setting is to segment the underlying Pareto surface plans and then enforce the mild restriction that only a small number of these plans are active at any time during plan navigation, thus limiting the total number of segments used in the final plan.

1 Introduction

Although intensity modulated radiation therapy (IMRT) has become a clinical standard for treating many cancers, including prostate, head and neck, pancreas, brain, and lung, the planning process – which relies heavily on automated computer optimization – remains a time consuming and iterative task. This is due to the fact that IMRT planning is inherently a search for a plan which balances conflicting desires, namely: prescribed dose to the target(s) and minimal dose to the organs at risk (OARs). Multi-criteria optimization (MCO) has been introduced as a technique to allow the planner to explore the patient-specific tradeoffs, and has been shown to reduce treatment planning time and achieve superior plans for selected clinical scenarios [1, 2].

Under the general title of MCO fall two distinct approaches to treatment planning: prioritized optimization, also referred to as lexicographic optimization, and Pareto surface based MCO. In prioritized optimization, a sequence of optimizations are run, tackling the most important objectives first and then moving down the list of lower objectives [3, 4, 5, 6]. For each optimization run, the results of the previous higher priority optimizations are used as constraints to ensure that the higher priorities are respected. This approach results in a single plan and if all goes well, this plan is used as the clinical treatment plan. In Pareto surface based MCO the planner interacts with the planning system to explore dosimetric tradeoffs – different Pareto optimal plans – in real time. For a given set of objectives and constraints, an IMRT plan is Pareto optimal if it is feasible (satisfies the constraints) and it is not possible to improve on any objective without worsening at least one other objective. The full collection of Pareto optimal plans, typically an infinite set, is referred to as the Pareto surface. A practical technique for exploring this infinite set involves precomputing a modest number (<50) of Pareto optimal plans scattered across the Pareto surface and approximating the continuous Pareto surface by forming convex combinations of these base plans in response to user requests. For example, by way of a set of graphical user interface slider bars, a user might make a request to reduce spinal cord dose, at which point the system computes a convex combination of the base plans which has an improved spinal cord dose. Computing these convex combinations can be done fast enough to give the user the experience of continuous real time navigation.

Exploring a Pareto surface of treatment plan possibilities has moved from concept to clinical deployment in recent years. The first commercial product to implement the Pareto surface approach to MCO (from here onward, MCO will be taken to mean Pareto surface based MCO, which is the version of MCO studied in this report) is RayStation 2.0 (RaySearch Laboratories, Stockholm). In this implementation and in prototype implementations [7, 8, 9], the Pareto surface plans, also called database plans, are fluence map optimized treatment plans. Fluence map plans are idealized plans: to make a plan clinically deliverable, the plan must be segmented into step and shoot segments or a dynamic sliding window plan. Segmentation causes the idealized fluence map plan to change, both because the fluence map ends up being approximated and because the dose calculation using multi-leaf collimator aperture shapes is a fundamentally different algorithm than the fluence map based dose calculation.

This degradation of the treatment plan during segmentation can negatively impact the MCO treatment planning process by adding an iteration loop to the treatment planning process that the MCO technology was designed to avoid. Anecdotally, this plan degradation happens more for complicated plans – complicated here being synonymous with optimal isodose contours being highly non convex, i.e. the dose distribution being deeply carved away from OARs – for example, bi-lateral target head and neck cancers and brain tumors situated near optic structures.

The purpose of this work is to describe a technique for IMRT Pareto surface plan navigation where the plans explored during navigation are already segmented, thus, ready-to-go. In this sense, this work could also be called what you see is what you get (WYSIWYG) navigation. We assume step and shoot IMRT. The challenges of WYSIWYG navigation for step and shoot Pareto surface MCO can be understood by considering the following:

1. The Pareto surface should in general be expansive enough to cover a range of treatment planning possibilities.
2. Plans far apart from one another on the Pareto surface will have different optimal segmentation, particularly when the number of segments is restricted.
3. A general point found during Pareto surface navigation is formed by a convex combination of up to N of the pre-computed Pareto surface plans, where N is the number of objective functions, i.e. the dimension of the Pareto surface. If each database Pareto optimal plan is pre-segmented into S segments, then a navigated-to convex combination plan may have up to $N \times S$ segments.
4. It is highly desirable from a treatment delivery time perspective to have a plan which consists of small number of segments.

Together, these items point to the challenge. If we strive for a minimal set of segments for the final deliverable plan and we also strive to have a Pareto surface which encompasses an expansive tradeoff space, it is necessary to use a different segmentation for each of the database plans. During navigation it is theoretically possible to combine up to N plans (in a clinical setting, N can be 10 or more), depending on where you are on the Pareto surface. Plans along lower-dimensional edges of the Pareto surface use fewer than N plans, see Figure 1 as well as section 2.2.

In this work we demonstrate that it is possible to continuously navigate a high dimensional discretely sampled Pareto surface while restricting the number of active plans used to represent the current location, without the restriction significantly impacting the plan quality. We demonstrate this in two ways. The first is a new navigation method based on visualizing 2D tradeoffs (cuts of the higher dimensional Pareto surface). We show analytically that this method naturally encourages the use of a small number of plans, and

hence a small number of total segments, to form the navigation plans. The second is a statistical analysis of the more traditional ND navigation technique. For this we employ an automatic random navigation simulation and analyze the resulting statistics. Both approaches, one theoretical and one numerical, support the idea that the combination of only a few plans ($\ll N$) is sufficient to closely represent navigated positions on an N -dimensional Pareto surface, and this in turn supports the a concept of WYSIWYG step and shoot IMRT Pareto surface based navigation using pre-segmented database plans.

2 Methods

We consider the following type of IMRT formulation:

$$\begin{aligned} &\text{optimize} && \{f_1(d), f_2(d), \dots, f_N(d)\} \\ &\text{subject to} && d = Dx \\ & && d \in C \\ & && x \geq 0 \end{aligned} \quad (1)$$

where f_i are the N objective functions describing the dose to a particular structure (e.g. mean dose to stomach), d is the vector of voxel doses, D is the dose-influence matrix, and x is a concatenation of all the fluence maps into a single beamlet fluence vector. The set C is a convex set of dose constraints (e.g. minimum dose to the target) that are designed to be met by all plans.

2.1 Pareto surface generation and base plan segmentation

For a given problem formulation, we use RayStation 2.0 to compute a set of points on the Pareto surface as follows. First the software computes the N anchor plans. These are single objective optimizations of the individual objectives, i.e. the best possible plan for each of the objectives. The $N+1$ th plan is termed the balance plan, and is found by optimizing an equi-weighted sum of the N objectives. Objectives are similarly scaled so that an equi-weighted sum does indeed give a balanced plan. For additional plans, RayStation optimizes equi-weighted sums of pairs of objective functions which are most highly anti-correlated. To determine level of correlation, the first $N+1$ plans are used to compute correlation coefficients between every pair of objectives. More advanced techniques for generating additional Pareto surface points, involving sandwiching the Pareto surface between lower and upper bounds and computing additional points to close the large gaps in these bounds, are described in [10, 11, 12] but are not used in this work. Here we compute a total of $2N$ plans using the anti-correlation method.

The key mathematical entity needed for navigation is the $M \times N$ matrix of objective function values, where M is the number of pre-computed Pareto plans. In this report $M = 2N$ as described above, but M can be any number larger than N in general. To access RayStation's $M \times N$ matrix for navigation simulation and evaluation outside of RayStation, it needs to be extracted manually. This was done for the "deliverable plan stage". Thus, before navigation simulation, each base plan was segmented and for each plan, the values of the N objective functions were read off, resulting in the $M \times N$ matrix.

With the Pareto matrix P in hand, we analyze two styles of navigation. In both forms of navigation, the main idea is to form in real time convex combinations of the M base plans (equivalently the M rows of the Pareto surface matrix P) which respond to the planner's request to improve some objective. If the base plans are already segmented, such that all plans during navigation are ready-to-deliver, there is a strong delivery efficiency incentive to keep the number of base plans used in the convex combinations to a minimum. The first navigation method we present and analyze theoretically, called *2D-cut navigation*, is new to

the field of Pareto surface exploration to our knowledge. It offers insight into why the number of active plans needed for convex combinations during navigation will often be small. The next method, *ND navigation*, is the traditional *N* sliders navigation method implemented in RayStation and prototype systems. 2D-cut navigation is presented because it offers a potentially more intuitive way to navigate high-dimensional surfaces, and also because one can analyze the plan-combining properties of it quite easily. However, given the dominance of ND navigation in the field, we numerically simulated the ND navigation style to show that it also has favorable plan-combining properties (low number of plans needed to approximate the user's navigation position on the Pareto surface).

2.2 2D-cut navigation

The method described here is based on the display of a 2D cut of the Pareto surface. The user selects which two objectives to view a 2D tradeoff curve of, and all of the other objectives are free, or upper bounded by a user-chosen value. The user then navigates by either 1) moving along the 2D tradeoff curve or 2) changing the bounds of one of the other objectives, which alters the 2D tradeoff curve when those constraints affect the currently viewed tradeoff. At any point during navigation the user may switch which two objectives are currently active.

We assume that P is ordered so that the anchor plans, in order, are the first rows of the matrix. We normalize the Pareto surface via:

$$P_{k,j}^{\text{norm}} = \frac{P_{k,j} - \min_m(P_{m,n=j})}{\max_m(P_{m,n=j}) - \min_m(P_{m,n=j})} \quad (2)$$

By expressing each objective in normalized form, between 0 and 1, the normalization assures equal influence of different objectives (or dimensions) for the multi dimensional distance calculation between different points on the Pareto surface that is introduced later. All calculations will be done in this normalized setting, and so without ambiguity we let P from here onward refer to the normalized Pareto surface.

The Normal Constraint Method (NCM) can be used to generate an even distribution of points along the tradeoff curve for the two chosen objectives [13, 14]. Let x and y denote the two dimensions selected to view the tradeoff curve of (e.g. $x = 7$ and $y = 3$ would be a possibility for say an 8D tradeoff). Let B denote the vector of constraints the user has selected. This is the current position of the constraint sliders, and can include the objectives being traded off (in which case it is just the extent of the 2D curve that will be affected).

The NCM method starts by finding the anchor plans for the 2D tradeoff selected:

$$\begin{aligned} &\text{minimize} && (\lambda' P)_k \\ &\text{subject to} && (\lambda' P)_j \leq B_j, \forall j \\ &&& \sum_{i=1}^M \lambda_i = 1 \\ &&& \lambda \geq 0. \end{aligned} \quad (3)$$

λ is the convex combination vector, and k is the index of the objective to minimize (will run for $k = x$ and y). This optimization finds the convex combination of plans in the database that minimizes the objective in question, subject to constraints B .

With those two anchors in hand, the next step is to form a moderate number of points (20 is likely to be more than enough) evenly spaced along the line connecting them. These are points in the 2D tradeoff space labeled $g_1, g_2, \dots, g_k, \dots$ in Figure 2. From each of these

points – not including the anchor points – we project down to find the intersection of the 2D tradeoff surface in the direction perpendicular to the line connecting the two anchor points.

To do this, we form the vector from anchor point A_x to anchor point A_y , and call this v :

$$v = A_y - A_x.$$

v is a $[2 \times 1]$ vector.

Let P_{xy} be the two user selected columns of the Pareto surface. Since there are M plans in the Pareto database, P_{xy} is an $m \times 2$ matrix. Let g_k be the point we are running. The following optimization yields the k th point on the 2D Pareto curve:

$$\begin{aligned} & \text{minimize} && (\lambda' P)_{xy} \\ & \text{subject to} && (\lambda' P)_j \leq B_j, \forall j \\ & && \lambda' P_{xy} v = g_k' v \\ & && \sum_i \lambda_i = 1 \\ & && \lambda \geq 0. \end{aligned} \quad (4)$$

The constraint $\lambda' P_{xy} v = g_k' v$ enforces that the solution is on the line perpendicular to the v and passing through g_k . Writing the constraint as $(\lambda' P_{xy} - g_k') v = 0$ makes this easier to see. Note, the simpler method to generate points on the 2D curve would be to replace the

$\lambda' (P_{xy} - g_k') v = 0$ constraint with $(\lambda' P)_x \leq b$ for various values of b (the so-called epsilon constraint method of Pareto surface generation) but this does not have the useful property of spreading the points evenly across the curve. As for using a weighted sum method, we opt not for that because, particularly for sparse Pareto surface representations (i.e. small number of database plans), the 2D cuts will be very “polyhedral” (i.e. Pareto surface surface will consist of large facets) and so many weights will yield the same vertex.

Moving around on the 2D curve is then a matter of tracking which two points you are between and updating the convex combination used.

2.3 On the number of active convex combination plans during 2D-cut navigation

Let $|B|$ denote the number of objectives that are constrained during 2D-cut navigation.

Proposition: The number of active plans needed to form the current navigation point is at most $2 + |B|$.

Proof. : Consider the linear programming formulation (4), which computes points on the 2D cut Pareto surface. Linear programming theory states that if a linear program is feasible and the optimal solution is bounded, there exists an optimal solution with DIM tight constraints (a constraint is tight if it is satisfied at equality), where DIM refers to the underlying dimension of the linear program, i.e. the number of optimization variables. For linear program (4), there are M optimization variables, the components of the vector λ . Therefore, there exists an optimal solution to formulation (4) with M tight constraints. Counting the number of constraints satisfied at feasible solution, we have the two equality constraints, we have (up to) $|B|$ navigation constraints. This means we need (at least) $M - (2 + |B|)$ additional constraints to be tight. The only other constraints in this linear program are the 0 constraints, so we need at least $M - (2 + |B|)$ components of λ set to zero, or stated another way, at most $2 + |B|$ components of λ positive.

One caveat: linear interpolation between two of the points along the surface as computed by the NCM method could in theory double the number of underlying plans used. This is not in

reality a problem since solving formulation 4 is very fast and therefore we can represent the 2D tradeoff curve easily with a high density of points and therefore avoid needing linear interpolation.

The above proof is sufficient to show that if the user invokes only a small number of constraints during navigation then the number of active plans used to form the current location on the Pareto surface is small. But beyond this, we hypothesize that even more general Pareto surface navigation can also be done with maintaining a small set of active plans. To that end we turn to ND navigation.

2.4 Traditional ND sliding navigation, and navigation simulation for generating statistics

An automatic navigation approach was established to generate data to confirm the hypothesis that even in general navigation, the number of active convex combination plans needed is not large. In order to simulate navigation to a random spot on the Pareto surface, random constraint levels were chosen for each normalized objective based on a uniform probability density function. For each constraint level combination, points (plans) on the Pareto surface were calculated by minimizing each of the objectives individually and also by minimizing the equi-weighted sum of all objectives using optimization formulation 3. The tighter the randomly chosen constraint levels, the higher the probability that it is not possible to find any remaining plan on the Pareto surface. In those cases the constraint level combination is discarded, resulting in a modified distribution of the used constraint levels. These optimizations lead to a large set of λ vectors, each representing a point of the Pareto surface.

In a last step the random points on the Pareto surface $\lambda'P$ were approximated by plans with a restricted number of non-zero λ_i values, so called “small number of plans approximation” resulting in $\sim \lambda'P$. Two different methods have been used:

- Direct restriction of the number of active plans: Allow only a fixed number of r nonzero plans. We chose the plans with the highest λ_i values, set all other plan $\lambda_i = 0$, and scale such that $\sum_{i=1}^M \tilde{\lambda}_i = 1$.
- Restriction of small λ_i values: Set all $\lambda_i \in \epsilon$ to zero and again scale remaining plans so that $\sum_{i=1}^M \tilde{\lambda}_i = 1$.

To evaluate the difference between the original and the restricted plans, we use the average of the absolute objective function differences:

$$d(\lambda, \tilde{\lambda}) = \frac{1}{N} \sum_{j=1}^N |(\lambda'P)_j - (\tilde{\lambda}'P)_j| \quad (5)$$

Plan quality after a $\tilde{\lambda}$ restriction is then defined as $Q(\lambda; \tilde{\lambda}) = 1 - d(\lambda; \tilde{\lambda})$. Furthermore, the mean number of non-zero λ_i and the distribution of λ_i -values were studied for every objective minimized.

3 Results

We analyze the following IMRT instance:

$$\begin{aligned}
&\text{minimize} && \{\text{mean left eye dose, mean right eye dose, mean chiasm dose} \\
& && \text{,mean brainstem dose, mean left cochlea dose,} \\
& && \text{mean right cochlea dose, mean right optical nerve dose,} \\
& && \text{mean left optical nerve dose, max dose to PTV,} \\
& && - \text{min dose to PTV, dose falloff penalty: 59.4 Gy to 0 Gy in 1 cm}\} \\
& && (6) \\
&\text{subject to} && d = Dx \\
& && d_i \geq 59.4 \text{ Gy}, \forall_i \in \text{GTV} \\
& && d_i \geq 40 \text{ Gy}, \forall_i \in \text{PTV} \\
& && d_i \leq 70 \text{ Gy}, \forall_i \in \text{skin} \\
& && d_i \leq 60 \text{ Gy}, \forall_i \in \text{brainstem} \\
& && x \geq 0
\end{aligned}$$

The main result for the 2D cut navigation is the proof given in section 2.3. However, to illustrate the 2D navigator and highlight the proof result, we display a snapshot of the navigation process in Figure 3. In this case we navigated to a plan that was a convex combination of 5 database plans. Using a cutoff value of 0.1, we reduce the number of active plans to 3, and the plan quality $Q = 99.2\%$. In this case, as viewed in the 2D plot in Figure 3, it appears the restricted plan (gray circle) is actually better than the original plan (black circle) but this is only for the two chosen objectives: the plan will have degraded in some other objective(s).

To study ND navigation, 12000 random points on the Pareto surface approximation were generated by minimizing each of the 11 objectives and a linear combination of all objectives (1000 points for each objective minimized) as described in the Methods section. The number of segments S for each database plans was chosen to be 50. Figure 4 shows the distribution of the (normalized) constraint levels that led to feasible plans. Each of the 12000 randomly chosen Pareto combinations was approximated with both methods described in section 2.4.

Figures 5 show histograms of the number of active plans and also of the λ_j values for the calculated 1000 Pareto surface points from minimizing the (negative) PTV-minimal-dose objective. Three different cases are shown: (a) before restriction, (b) after restriction to a total number of 3 or fewer active plans and (c) after restriction to $\lambda_j \leq 0.1$. The chosen PTV-minimal-dose objective was the hardest to be fulfilled, meaning that the mean number of plans used was the greatest compared to the other objectives. In Figure 6 the same analysis is shown for the 1000 points on the Pareto surface obtained when minimizing the linear combination of all 11 objectives. The linear combination objective required the fewest active λ_j on average. Table 1 summarizes the mean number of active plans and the plan quality for both, the PTV-minimal-dose objective and the linear combination objective.

Comparing both cases it can be seen that the λ_j distribution is tilted more to the left side with nearly no $\lambda_j \geq 0.9$ for the hard case compared to the easy case. Comparing the λ_j distribution for both restriction methods, it becomes clear that the restriction of the maximal number of plans still leads to a significant amount of very small λ_j values although their number is reduced compared to the unrestricted case. Obviously, this is not the case with the second restriction method. The fact that both restriction methods lead to approximately the same mean number of active plans is not an intrinsic characteristic of the restriction method, but only caused by the chosen restriction levels for both methods.

In fact, both restriction methods have been applied for different restriction levels, meaning (a) different numbers of active plans allowed or (b) different minimal values of λ_j . For both

cases Figure 7 shows the plan quality for different restriction levels and the two selected objectives from above. For the same restriction applied, the approximation of the Pareto surface points generated by minimizing the linear combination objective lead to better plan quality than the approximation of the Pareto surface points generated by minimization the PTV dose objective. Restricting the Pareto surface points to a high number of plans (5) or to a low minimal λ_i value ($\lambda_i = 0.05$) leads to nearly no degradation of the plan quality.

To evaluate how many plans are actually active for a given plan quality to be fulfilled, in Figure 8 the number of active plans is plotted over the plan quality achieved. The plot includes the data of all different restriction levels shown in the previous Figure to all 12000 Pareto surface points. From the presented data it can be seen that for a given plan quality on average fewer plans are used with the restriction of the λ values compared to the restriction of the maximal number of plans. This finding is in agreement with our understanding of the different restriction methods. The restriction to a maximal number of plans eliminates plans if the total number of non-zero plans is higher than allowed – even if they have a high contribution (a high λ_i). On the other hand, it does not remove plans with small contributions if the total number of non-zero plans is equal or below the allowed number. Both this is not the case for the second method – the restriction in the λ value space. Moreover this method is more flexible when going to different cases with different number of objectives. Then the same λ_i -restriction can still be applied, but a restriction to a maximal number of plans obviously needs to be adjusted. Therefore, we conclude that restricting the λ -space is the method of choice for reducing the number of plans.

Figure 9 shows the plan quality for all plans with a specific number of plans before and after application of the restriction. It can be seen that in general three plans are enough to achieve a plan quality of at least 0.95. Only in the very rare cases (0.1% of the initial random plans) when there are more than 7 plans used in the initial plan, the plan quality will be reduced by more than 5% on average. An interesting result is also that already 46% of the initial random plans use no more than 3 plans.

In Figure 10 the dose distribution of a random case that uses 6 plans and that was restricted to 3 plans is shown. For selection of this case, we ordered all cases that used 6 plans before and 3 plans after the application of the restriction according to their plan quality. The plan corresponding to the median plan quality of these cases is the one presented. After restriction, this plan showed a plan quality of 96.1 %. The dose difference after restriction is below 4% in all voxels. The dose volume histograms (DVH) before and after restriction are shown in Figure 11, which confirms that only minor deviations occur as a result of the restriction.

4 Discussion and conclusions

This work grew out of the desire to avoid the “create deliverable plan” step in the current MCO navigation technique. Imagining what an ideal system would look like, where one is sliding on a Pareto surface and observing dose volume histograms and isodose contours update continuously and in real time, it becomes apparent that the MLC apertures used should depend on where one is on the Pareto surface. Initial considerations of this had individual MLC segment shapes coming in and out of the active set of apertures. Using this technique however would require a much larger real time computation, since one would need to maintain in memory the dose to voxel contributions of the individual apertures and make sure the composite dose from the current set of apertures is feasible and optimal. Instead, we have relied on the following observation in this work: *a particular aperture makes sense in conjunction with the other apertures it was created with*. Stated another way, a particular aperture is useful because it is part of a full collection of apertures which

together make a good treatment plan. Therefore, we can simplify the navigation considerably by swapping apertures in and out not individually, but collectively as the entire plan that they were created for, and scaled as a group. This allows us to consider simply the dose cubes (the term dose cube means the entire dose distribution), or even more minimally, the functions of the dose cubes that serve as the plan descriptors (e.g. minimum target dose, mean brainstem dose) during navigation.

In order for this idea of navigating by mixing pre-segmented plans to be feasible, it needs to be the case that not too many plans need to be combined during navigation. In this work we have demonstrated that a limited number of plans suffices to approximate a general location on a high dimensional Pareto surface. We demonstrate for a brain case with an 11 dimensional Pareto surface, generally 5 or fewer plans are needed to form the navigated plan, and that even for plans requiring 6 base plans, 3 plans is enough for very close approximation, see Figure 9. We also show that instead of requiring a combinatorial search for “what 3 plans” to use, it is sufficient to use the plans of the original convex combination that have the highest contributions (the highest λ_i values). Assuming each pre-computed Pareto surface plan is segmented into say 50 segments, as was done in this work, the total number of segments will be $3 \times 50 = 150$ or fewer (fewer if the same segments were used in more than one of the plans).

A problem arises when one considers that if 50 plans is enough to generate a good plan, than 150 plans is overkill and therefore not a desirable solution. On the other hand if around 150 segments are needed, then the base plans, with only 50 segments, might be far from Pareto optimal. We next present two different lines of attack to deal with this issue.

As the number of plans on the original surface is increased, the error in forming a navigated plan with a restricted number of active plans will decrease. Indeed in the limit of a very fine grid of plans pre-computed on the surface, one could use a single plan – the closest one – to represent a given convex combination plan, see Figure 1. Populating a Pareto surface with a fine grid can be done in parallel and offline, and real-time navigating a surface with hundreds of plans is possible due to the simplicity of the navigation optimization problems that are solved, which do not involve the dose cube or the beamlets. In this mode of operation one would segment each database plan with as many segments are needed or clinically allowed for delivery time considerations and the navigation system would display the plan with the highest λ_i value.

If the number of pre-computed plans cannot be made large enough to allow for single plan navigation, or mixing is desired for smooth exploration of the planning options, then shared segments can also alleviate the problem of the pre-computed plans needing to be too coarsely segmented and thus not Pareto optimal. In this case one would find a pool of apertures that are used by all the plans and then additional specialized apertures for each plan. Direct aperture optimization (e.g. [15, 16, 17]) will be immediately applicable here.

The approach presented is a practical compromise, conceived to allow smooth MCO exploration of step and shoot plans without a large computational overhead. Ideally as the planner navigates the Pareto surface the number of segments would be changing to reflect the varying complexity of plans in different regions of the Pareto surface. Additionally, one would like to be able to explore the benefit of adding a few more apertures during navigation. In the current scheme, apertures come as a full plan bundle, so these features are not handled smoothly. Despite the drawbacks, we see the presented approach as a practical alternative to real time aperture optimization while sliding over the Pareto surface, which at the current state is not computationally feasible. It should also serve as groundwork for future developments on directly deliverable step and shoot MCO, which in its more mature

state will likely employ a hybrid approach of a pool of apertures shared by all of the database plans, specialized apertures for different plans on this Pareto surface, limiting the number of plans combined during navigation, and possibly real-time aperture weight optimization and the fast creation of helpful additional apertures. Future research will also be directed towards directly deliverable navigation for sliding window IMRT planning, which will require a different line of attack.

In terms of why so few plans are needed during Pareto navigation, we have provided an explanation in reference to the newly introduced 2D-cut navigation, where we show by linear programming theory that the number of active variables in an optimal solution is related to the number of function constraints that are binding during navigation. Another view of why the number of active plans will typically be much less than the theoretical maximum number of N is related to the fact that for high dimensional Pareto surfaces in IMRT planning, many of the objectives will be strongly correlated, leading to effectively lower dimensional Pareto surfaces, i.e. an effectively smaller N . For example, in a brain case, one might expect that minimizing the dose to the left eye also minimizes the dose to the left optic nerve. The high dimensional surfaces analyzed by Spalke *et al* [18] were shown to have an effective dimension of 3 or 4 when the dimensions were reduced using principal components analysis or a non-linear technique called the isomap method [19].

Both 2D and ND navigation analysis lead us to the same conclusion: relatively few plans are needed to approximate an arbitrary navigated-to location on an IMRT Pareto surface. This finding justifies our proposal for directly deliverable step and shoot IMRT MCO: pre-segment the base Pareto surface plans and restrict the number of plans used during navigation to form the convex combination plans. Direct deliverability is a crucial step to bringing the full power of MCO into clinical IMRT treatment planning.

Acknowledgments

This work was supported by NCI Grant 1 R01 CA103904-01A1: Multi-criteria IMRT Optimization and by a research collaboration with RaySearch Laboratories. Thanks to Ehsan Salari for helpful feedback on the manuscript.

References

1. Craft D, Hong T, Shih H, Bortfeld T. Improved planning time and plan quality through multicriteria optimization for intensity-modulated radiotherapy. *Int. J. Radiation Oncology Biol. Phys.* 2012; 82(1):e83–e90.
2. Hong T, Craft D, Carlsson F, Bortfeld T. Multicriteria optimization in intensity-modulated radiation therapy treatment planning for locally advanced cancer of the pancreatic head. *Int. J. Radiation Oncology Biol. Phys.* 2008; 72(4):1208–1214.
3. Breedveld S, Storchi PRM, Keijzer M, Heemink AW, Heijmen BJM. A novel approach to multi-criteria inverse planning for IMRT. *Physics in Medicine and Biology*. 2007; 52:6339. [PubMed: 17921588]
4. Jee KW, McShan D, Fraass B. Lexicographic ordering: intuitive multicriteria optimization for IMRT. *Physics in Medicine and Biology*. 2007; 52:1845. [PubMed: 17374915]
5. Wilkens J, Alaly J, Zakarian K, Thorstad W, Deasy J. IMRT treatment planning based on prioritizing prescription goals. *Physics in Medicine and Biology*. 2007; 52(6):1675–1692. [PubMed: 17327656]
6. Falkinger M, Schell S, Müller J, Wilkens J. Prioritized optimization in intensity modulated proton therapy. *Zeitschrift für Medizinische Physik*. 2011
7. Craft D, Monz M. Simultaneous navigation of multiple pareto surfaces, with an application to multicriteria IMRT planning with multiple beam angle configurations. *Medical physics*. 2010; 37:736. [PubMed: 20229883]

8. Monz M, Küfer K-H, Bortfeld T, Thieke C. Pareto navigation - algorithmic foundation of interactive multi-criteria IMRT planning. *Physics in Medicine and Biology*. 2008; 53(4):985–998. [PubMed: 18263953]
9. Thieke C, Küfer K-H, Monz M, Scherrer A, Alonso F, Nill S, Thilmann C, Bortfeld TR. Beyond weight factors: new concepts for defining and analysing dose optimisation. *Radiotherapy and Oncology*. 2004; 73
10. Bokrantz R, Forsgren A. An algorithm for approximating convex pareto surfaces based on dual techniques. *Inform Journal of Computing*. 2012
11. Craft D, Halabi T, Shih H, Bortfeld T. Approximating convex Pareto surfaces in multi-objective radiotherapy planning. *Medical Physics*. 2006; 33(9):3399–3407. [PubMed: 17022236]
12. Rennen G, Van Dam E, Den Hertog D. Enhancement of sandwich algorithms for approximating higher dimensional convex Pareto sets. *INFORMS Journal on Computing*. 2011; 23:493–517.
13. Messac A, Ismail-Yahaya A, Mattson CA. The normalized normal constraint method for generating the Pareto frontier. *Journal of the International Society of Structural and Multidisciplinary Optimization (ISSMO)*. 2003; 25:86–98.
14. Messac A, Mattson CA. Normal constraint method with guarantee of even representation of complete Pareto frontier. *AIAA Journal*. 2004 Oct.42:2101–2111.
15. Shepard D, Earl M, Li X, Naqvi S, Yu C. Direct aperture optimization: A turnkey solution for step-and-shoot IMRT. *Medical Physics*. 2002; 29:1007–1018. [PubMed: 12094970]
16. Hårdemark, B.; Liander, A.; Reh binder, H.; Löf, J. RaySearch White Paper. Stockholm, Sweden: RaySearch Laboratories AB; 2003. Direct machine parameter optimization with RayMachine® in Pinnacle®.
17. Men C, Romeijn H, Ta kın Z, Dempsey J. An exact approach to direct aperture optimization in IMRT treatment planning. *Physics in Medicine and Biology*. 2007; 52:7333. [PubMed: 18065842]
18. Spalke T, Craft D, Bortfeld T. Analyzing the main trade-offs in multiobjective radiation therapy treatment planning databases. *Physics in Medicine and Biology*. 2009; 54(12):3741–3754. [PubMed: 19478382]
19. Tenenbaum J, De Silva V, Langford J. A global geometric framework for nonlinear dimensionality reduction. *Science*. 2000; 290(5500):2319–2323. [PubMed: 11125149]

Figure 1.

A 3D Pareto surface. The black circles indicate computed Pareto optimal plans, and the shaded triangles represent the convex combinations of those possible that give rise to undominated plans, and thus serve as an approximation of the Pareto surface. Points *a* and *b* represent two convex combination plans that could occur during navigation. Since *a* is on an edge, it is formed with fewer than $N = 3$ plans. On the other hand, point *b* is in the center of the Pareto surface and thus requires N plans to form it. However, depending on the number of plans used to represent the Pareto surface, it may be possible to reduce the number of plans needed to represent these central plans by snapping the convex combinations to nearby edge points, as shown by the two dotted line projections from point *b*. The shorter projection shows that if the number of computed plans on the surface is greater – which we represent by adding another point, the gray point, to the surface and thus breaking that central facet into three smaller facets – then the deviation will be less.

Figure 2.

Diagram showing key idea of the normal constraint method (NCM) of computing points of a 2D tradeoff curve. $\lambda^* P_{xy}$, the optimal solution of a single run of formulation 4, is displayed.

Figure 3.

Here, the planner has chosen to view the tradeoff between brainstem mean dose and PTV maximum dose. The light gray Pareto curves display navigation history: the different curves arise from different constraint levels chosen for the other objectives. The black circle corresponds to the current navigation position without any λ restrictions, and the gray circle below it corresponds to the plan after restricting $\lambda_i = 0.1$. The navigation interface is also displayed. The colors above each slider indicate the sensitivity of the Pareto curve to a change in that particular bound. The details of this computation are beyond the scope of this report, but briefly, linear programming duality theory is used to compute the sensitivity to changes in the constraint levels.

Figure 4. Distribution of the actual normalized constraint levels. Starting with a uniform distribution of the normalized constraint levels, small constraint levels chosen simultaneously have a higher probability of an infeasible solution and are therefore more often discarded, leading to the triangle shaped distribution shown.

Figure 5.

Histograms of the number of active plans and also of the λ_i values for the calculated 1000 Pareto surface points by minimizing the **PTV-minimal dose objective**. Three different cases are shown: (a) unrestricted plans, (b) restriction to a total number of 3 or fewer active plans and (c) restriction to $\lambda_i \geq 0.1$.

Figure 6.

Histograms of the number of active plans and also of the λ_i values for the calculated 1000 Pareto surface points by minimizing the **Linear combination objective**. Three different cases are shown: (a) unrestricted plans, (b) restriction to a total number of 3 or fewer active plans and (c) restriction to $\lambda_i \geq 0.1$.

Figure 7. Plan quality for different restriction levels for both restriction methods applied: (a) number of active plans restriction and (b) λ_i value restriction. Data points are shown for two selected objectives minimized. Each data point corresponds to the mean value of 1000 calculated Pareto surface points. Uncertainty bars represent the standard deviation.

Figure 8.

Number of active plans plotted as a function of plan quality for both restriction methods. The plot contains all different restrictions shown in the previous figure for all 11+1 objectives minimized. The resulting plan quality was binned for each of both restrictions methods into bins with a bin size of 0.02. Therefore each data point corresponds to all restrictions resulting in a plan quality of ± 0.01 from the x-coordinate of the data point. Uncertainty bars represent the standard deviation.

		Percentage	Number of plans allowed							
			1	2	3	4	5	6	7	8
Number of plans used	1	0.3	1	-	-	-	-	-	-	-
	2	10.2	0.92±0.06	1	-	-	-	-	-	-
	3	35.4	0.87±0.06	0.97±0.03	1	-	-	-	-	-
	4	31.5	0.84±0.06	0.95±0.03	0.98±0.01	1	-	-	-	-
	5	16.9	0.83±0.06	0.93±0.03	0.97±0.02	0.99±0.01	1	-	-	-
	6	5.1	0.81±0.06	0.92±0.03	0.96±0.02	0.98±0.01	0.99±0.01	1	-	-
	7	0.6	0.79±0.07	0.9±0.04	0.95±0.02	0.97±0.02	0.99±0.01	0.99±0.01	1	-
	8	0.1	0.78±0.08	0.89±0.06	0.94±0.04	0.96±0.02	0.98±0.01	0.99±0.01	1.00±0.00	1

Figure 9.

Plan quality matrix: Plan quality degrades as you allow fewer plans for plan representation. Number of plans used refers to the unrestricted number of plans during standard navigation. When this number is restricted to a smaller number, shown across the top of the matrix, the plan quality drops. Still, restricting to 3 plans retains a plan quality of over 95% except in the rare instance, 1 out of 1000 in this case, where 8 plans are used in the unrestricted setting.

Figure 10.

Dose wash of a random plan with 6 plans before (left) and after restricting the number of plans to 3 (middle). On the right the dose difference is shown. Contours mark PTV (black), GTV (white) and organs at risk (green, light blue, dark and light yellow, dark and light orange).

Figure 11. Dose volume histogram for the case shown in Figure 10. The continuous and dashed lines represent different structures before and after the application of the restriction, respectively.

Table 1

Statistics of the number of active plans in the unrestricted case and after applying both restriction methods. The plan quality Q for two examples of both restriction methods is shown. Given uncertainties represent the 1σ confidence interval of the variation over the 1000 single Pareto surface points.

		Objective minimized	
		PTV minimal dose	Linear combination
Mean number of active plans	unrestricted	4.77 ± 0.95	3.04 ± 1.00
	restricted to $n_{\lambda_i} = 3$	3.00 ± 0.00	2.67 ± 0.53
	restricted to $\lambda_i = 0.1$	3.06 ± 0.89	2.15 ± 0.82
Plan quality Q	restriction to $n_{\lambda_i} = 3$	0.973 ± 0.022	0.994 ± 0.013
	restriction to $\lambda_i = 0.1$	0.978 ± 0.019	0.984 ± 0.019