Multiobjective navigation of external radiotherapy plans based on clinical criteria

Kuan-Min Lin [*][1] and Matthias Ehrgott[2]

[1]km.lin@auckland.ac.nz, University of Auckland
[2]m.ehrgott@lancaster.ac.uk, Lancaster University

August 7, 2017

Abstract

This study considers a navigation method for finding the most preferable radiotherapy plan from a discrete set using planner-defined clinical criteria. The method is based on repeatedly solving an optimisation model to identify a plan that best satisfies the aspiration values set by the planner. During navigation, the planner iteratively adjusts the aspiration values to match the preference information learned from previous plans until the most preferable plan is identified. The use of soft constraints to model aspiration values enables navigation among a discrete set and allows the planner to freely specify the aspiration values without producing an infeasible model. We demonstrate the use of the model by applying it to a prostate cancer case. This illustrates that improvements in optimisation criteria do not necessarily lead to improvements in clinical criteria. Hence the method obviates the need to simultaneously monitor both optimisation and clinical criteria in current navigation systems. Instead, the direct use of clinical criteria for navigation aids the planner to quickly identify the most preferable plan.

KEY WORDS: Multiobjective decision making, interactive multiobjective optimisation, navigation, data envelopment analysis

---

*Corresponding author; Address: Level 3, Uniservices House 70 Symonds Street, Auckland; Phone: +64 9 373 7599; Fax: +64 9 373 7468

1 Introduction

Radiotherapy treatment planning involves managing conflicting objectives regarding tumour control and the sparing of surrounding healthy structures. Managing these conflicting objectives in radiotherapy treatment planning is not trivial since the achievable trade-offs between the objectives are not clear. In the conventional iterative planning practice, a planner experiments with the planning parameters in order to generate a plan that satisfies a clinic specific protocol. As the effect of changing the planning parameters is unknown a priori, this planning practice is time consuming and without guarantee of finding the best plan available.

Multiobjective optimisation (MO) has been introduced to handle this planning difficulty. Within the context of MO, three distinct approaches have been applied: *a priori* approaches, *a posteriori* approaches and *interactive* approaches. An a priori approach requires the planner to provide the preference information so that the information can be incorporated in the solution process to generate a preferable plan. One example of an a priori approach is multiobjective goal programming, in which a sequence of optimisations is conducted to achieve prioritised treatment goals predefined by the planner in a step-wise manner (Falkinger et al., 2012; Breedveld et al., 2009; Wilkens et al., 2007; Jee et al., 2007). An a posteriori approach aims at producing a representative set of *efficient* plans that capture potential trade-offs in optimisation objectives (Lin et al., 2016; Bokrantz and Forsgren, 2013; Shao and Ehrgott, 2007; Craft et al., 2006). A plan is efficient if it cannot be improved in any optimisation objective without deteriorating at least one other optimisation objective. Given a representative efficient set, the planner can then explore possible trade-offs in optimisation objectives and select the most preferable plan from the set. In an interactive approach, the planner iteratively adjusts preferences based on the knowledge learned from the generated plans. One rare example of an interactive method in radiotherapy treatment planning is proposed by Ruotsalainen (2009), who uses a classification based method for the planning problem.

Given a set of treatment plans generated by an a posteriori approach, it can be time consuming to examine the quality of every plan in the representative efficient set. Instead, a navigation method that effectively guides the planner toward the most preferable plan is needed (Allmendinger et al., 2016). Essentially, a navigation method should allow the planner to move from one plan to another with more desirable planning trade-offs until the most preferable plan is found. Desirable planning trade-offs can be expressed as a navigation query which specifies

desirable improvements in certain planning criteria while allowing other criteria to deteriorate. The existing navigation methods proposed for radiotherapy treatment planning so far are based on solving multi-objective optimisation problems on the convex hull of a set of efficient solutions generated by a sandwiching method. In the method of Monz et al. (2008), the navigation query is specified as a constraint in a single-objective optimisation problem which finds a plan that minimises the maximum difference in objective function value relative to the current plan. Craft and Monz (2010) propose a reference-point based navigation method designed for multiple efficient sets arising from different beam orientations. The navigation query is specified by changing the values of a reference point. The reference point is projected to the representative non-dominated set(s) to identify a plan that best satisfies the query. Craft and Richter (2013) propose a "2D-cut" navigation which features the 2D trade-off curve of two chosen objectives while other objectives are either relaxed or bounded by certain values. In all of these methods, navigation queries are specified in terms of *optimisation criteria*, i.e., the objective functions used in the optimisation to generate the representative efficient set. These criteria are convex functions so that optimisation algorithms can be used to find an optimal solution efficiently. Convex criteria along with a convex feasible set allow continuous navigation of treatment plans. Note that in these navigation methods, while optimisation criteria are used to steer the navigation process in the desired direction, *clinical criteria*, which are criteria used to evaluate plan quality in clinical practice, are also continuously monitored.

However, optimisation criteria are nevertheless only surrogates used to generate treatment plans and are generally not relevant to plan evaluation (Li et al., 2012). In this study, we show that the optimisation criterion values may not correctly reflect the plan quality and hence make navigation toward desired clinical improvement non-trivial. Instead, we propose a navigation method that uses planner-specified clinical criteria to specify navigation queries. While many clinical criteria are convex or can be reformulated as convex functions (see, e.g., Romeijn et al. (2004); Hoffmann et al. (2008)), some clinical criteria are non-convex and non-continuous, for example, the dose-volume (DV) criteria (Deasy, 1997; Llacer et al., 2003) and treatment time (where changes in beam orientation and/or the number of segments result in discrete changes in treatment time, see, e.g., Bortfeld and Schlegel (1993). For a non-continuous set, continuous navigation, as implemented in existing navigation models, is not applicable.

Instead, in this study we consider navigation on a discrete set of deliverable

plans of which the beam orientation, segment shapes and segment weights have been computed in a prior optimisation run. The clinical criterion values are extracted from the plans and form the underlying data for the navigation model. A few studies propose navigation/decision support methods with such a set of plans. Rosen et al. (2005) propose to fit the DV parameters of the target with the DV parameters of the critical organs using linear functions. Navigation is conducted by changing the DV parameters for the critical organs among some control points in a trial-and-error manner. Using this approach, the planner needs to explore all the control points in order to find the most preferable plan.

Ehrgott and Winz (2008) propose a decision support system based on filtering out plans with inferior criteria values. To conduct navigation in this method, the planner needs to iteratively impose, relax or tighten filters. Ripsman et al. (2015) propose a ranking based navigation system which ranks the plans according to the sum of weighted criteria values where the weights of the clinical criteria are assigned by the planner. However, the change of weight vector required to move from one efficient plan to another is unknown and can only be discovered through a trial-and-error process.

Our proposed method is inspired by data envelopment analysis (DEA), which is a well known management science method for assessing the performance of a set of decision-making units (DMUs) that convert inputs into outputs (Charnes et al., 1978; Cooper et al., 2011). In an economic interpretation, the inputs represent the cost we pay for producing the outputs. DEA has been applied to radiotherapy treatment planning as a quality assessment tool (Lin et al., 2013). The idea is to consider dose delivered to healthy organs (inputs) as the cost for delivering dose to the cancerous cells (outputs). To utilise the DEA concept and allow intuitive navigation, clinical criteria are categorised into inputs or outputs where decreases in inputs and increases in outputs are considered favourable. In our navigation method, the planner sets aspiration values for the clinical criteria and an optimisation model is solved to identify a plan that best satisfies the aspiration values. These aspiration values are specified as soft constraints for the optimisation model, hence the planner can freely set the ideal clinical criterion values, without the risk of generating an infeasible model. The planner can then inspect the plan identified by the navigation method and iteratively adjust the aspiration values until the most preferable plan is identified.

The contribution of this study is two-fold. Firstly, we propose a DEA based navigation model that allows a decision maker to freely express his/her preferences

in criterion values. As a consequence, the model can be applied to non-convex and discrete sets without leading to infeasibility. Secondly, we propose a clinically oriented navigation practice for radiotherapy treatment planning that allows planners to intuitively identify the most preferable plan from a set, instead of identifying the most preferable plan using optimisation criteria. In Section 2, we introduce the proposed navigation method. In Section 3, we present a prototype implementation of the proposed model with regards to a prostate treatment planning problem. Discussion and conclusion then follow in Section 4.

2   Method

We first explain our method using terminologies from DEA and MO, followed by interpretation and illustration of our method.

2.1   Efficiency in navigation

Consider a set of *n* DMUs which are evaluated by the decision maker based on a set of criteria. The criteria are categorised into *I* inputs and *O* outputs where inputs are considered as costs to produce outputs. The performance of a DMU is evaluated by comparing its inputs and outputs to the production possibility set, i.e., a set of input-output combinations that is potentially attainable under a certain production process.

Typically, the decision maker would prefer DMUs that are characterised as efficient. A DMU is efficient if there is no indication in the production possibility set that some of the DMU's inputs or outputs can be improved without worsening some of its other inputs or outputs.

The concept of efficiency is illustrated in Figure 1 in which the production possibility set is shown in grey. In the figure, DMUs *A*, *B*, *C* and *D* are considered efficient since none of the other potentially attainable input-output combinations show a lower or equal input value and simultaneously a higher or equal output value than these DMUs. Other DMUs are considered inefficient since there are potentially attainable input-output combinations (e.g. DMUs *A*, *B*, *C* and *D*) that empirically suggest the inefficient DMUs can be improved in the input/output without worsening the output/input.

In this study, we mainly consider a production possibility set referred to as the *free disposal hull*, which is formed by a set of DMUs with the assumption of free disposability, i.e., each DMU can consume extra inputs while producing the same
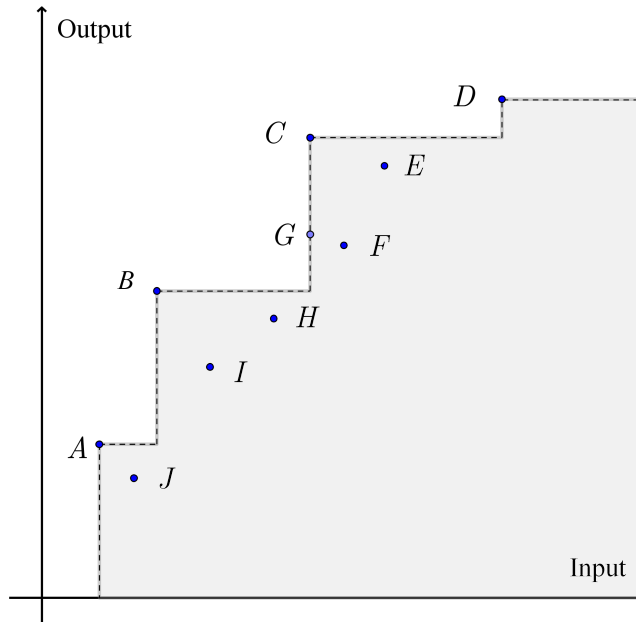
Figure 1: Illustrating the concept of efficiency. DMUs A, B, C and D are considered efficient while other DMUs are inefficient

levels of outputs or produce less outputs with the same levels of inputs. In Figure 1, the free disposal hull formed by the DMUs is shown in grey. The boundary of the production possibility set, illustrated by the dashed lines in Figure 1, is referred to as the *production frontier*. The production frontier consists of input-output combinations that cannot be further improved in at least one of the inputs or outputs. For example, DMU G cannot be improved further in input while improvement in output is achievable.

2.2 Identifying efficient DMUs using aspiration values

Given a large set, it can be impractical for the decision maker to examine each candidate DMU. Instead, the decision maker can propose a set of aspired criterion values (or aspiration values) and identify a DMU that best satisfies the aspiration values. In this subsection, we propose a model that allows the decision maker to perform such a task while at the same time ensuring efficiency of the DMU identified by the model.

Let $x \in \mathbb{R}^I_{>}$ and $y \in \mathbb{R}^O_{\leq}$ represent the aspiration values for the inputs and outputs, respectively. In addition, let the criterion values of the n DMUs be grouped into the input matrix $X \in \mathbb{R}^{I \times n}_{\geq}$ and the output matrix $Y \in \mathbb{R}^{O \times n}_{\geq}$ in which the ith col-

umn of X and Y contains the criterion values of the ith DMU. Note that $(x^T, y^T)^T$ is an arbitrary point specified by the decision maker and DMUs are actual points from the dataset.

Once the decision maker specifies the aspiration values, the following model is solved to identify a DMU that best satisfies $(x^T, y^T)^T$:

$$\max \quad \beta + \varepsilon \left(e^T s^- + e^T s^+ \right) \tag{1a}$$

$$\text{s.t.} \quad X\lambda \leq (1-\beta)x - s^- \tag{1b}$$

$$Y\lambda \geq (1+\beta)y + s^+ \tag{1c}$$

$$e^T \lambda = 1 \tag{1d}$$

$$\lambda \in \{0,1\}^n \tag{1e}$$

$$s^+, s^- \geq 0, \tag{1f}$$

where $\beta \in \mathbb{R}$, $\lambda \in \{0,1\}^n$, $s^- \in \mathbb{R}^I$, $s^+ \in \mathbb{R}^O$ are the decision variables. Vector e represents a vector of ones of appropriate dimension and $\varepsilon > 0$ is a very small number. In practice, there is no need to determine the value of ε. Instead, model (1) can be implemented using a 2-step lexicographic optimisation, as described in Section 2.3.

Proposition 1. *Model* (1) *is always feasible.*

Proof. Let $j^* \in \{1, \ldots, n\}$ and let $\lambda^{j^*} \in \mathbb{R}^n$ be a vector with a one in the j^*th entry and zeros in the other entries. Clearly λ^{j^*} satisfy constraints (1d) and (1e). By setting $\lambda = \lambda^{j^*}$ and re-arranging the variables, constraints (1b) and (1c) can be equivalently expressed as

$$\begin{aligned}
\beta &\leq \frac{X_{i,j^*} - x_i + s_i^-}{-x_i}, \quad i = 1, \ldots, I \\
\beta &\leq \frac{Y_{o,j^*} - y_o - s_o^+}{y_o}, \quad o = 1, \ldots, O.
\end{aligned} \tag{2}$$

Let s^{-*} and s^{+*} be zero vectors of dimension I and dimension O, respectively and let $\beta_{j^*} := \min \left\{ \frac{X_{i,j^*} - x_i}{-x_i}, \frac{Y_{o,j^*} - y_o}{y_o} : i = 1, \ldots, I, o = 1, \ldots, O \right\}$. It is obvious that $\beta = \beta_{j^*}, s^- = s^{-*}, s^+ = s^{+*}$ satisfy (2) and thus $(\beta_{j^*}, \lambda^{j^*}, s^{-*}, s^{-*})$ is a feasible solution of (1). $\square$

Model (1) can be interpreted as a mechanism that adjusts the aspiration values

7

to identify a best matching DMU. The adjusted aspiration values are expressed as the right-hand side of constraint (1b) and (1c), i.e., $(1-\beta)x - s^-$ and $(1+\beta)y + s^+$. Constraints (1d) and (1e) specify that only one element in vector λ can have a value of 1 with the rest of the elements having a value of zero. Thus, essentially, the model specifies that there must be one DMU with inputs and outputs better than or equal to the adjusted aspiration values. Since the production possibility set of model 1 is defined by the free disposal hull of the set of DMUs, the efficient DMU identified by the model must be one of the existing DMUs.

Two types of adjustments are used in the model: directional and additive. Directional adjustment is inspired by Chambers et al. (1996, 1998). The directional distance constraints $X\lambda \leq (1-\beta)\delta x$ and $Y\lambda \geq (1+\beta)\delta y$ allow the decision maker to specify a direction $(-\delta x^T, \delta y^T)^T$ for the adjustment of the aspiration values. In model (1) the direction is specified as $(-x^T, y^T)^T$ and β represents the magnitude of adjustment in the specified direction. As model 1 is always feasible (Proposition 1), the decision maker can choose the aspiration values $(-x^T, y^T)^T$ freely even if the values are outside of the production possibility set.

The other mechanism, additive adjustment, adjusts the individual aspiration values through s^- and s^+. In contrast to directional adjustment, which adjusts all criteria simultaneously, additive adjustment applies to each input/output independently and the adjustment only improves the aspiration values. Since the cost coefficient for additive adjustment, ε, is a very small value, additive adjustment is applied to the aspiration values with a lower priority than directional adjustment.

Hence, the adjustment of aspiration values in model (1) can be considered as a two step process where directional adjustment is conducted first, followed by additive adjustment. Directional adjustment moves a point to the production frontier where at least one input or one output cannot be further improved. The point on the production frontier is then adjusted by individual input/output through additive adjustment of s^+ and s^-. Additive adjustment ensures that the adjusted aspiration values equal the criterion values of one of the DMUs (since $\varepsilon(e^T s^- + e^T s^+)$ is maximised in the objective function, making constraints (1b) and (1c) binding). In addition, the DMU corresponding to the adjusted aspiration values must be efficient (Proposition 2).

Proposition 2. *The DMU selected by the optimal solution of model* (1) *must be efficient.*

Proof. If the selected DMU is not efficient, there must be another DMU from the
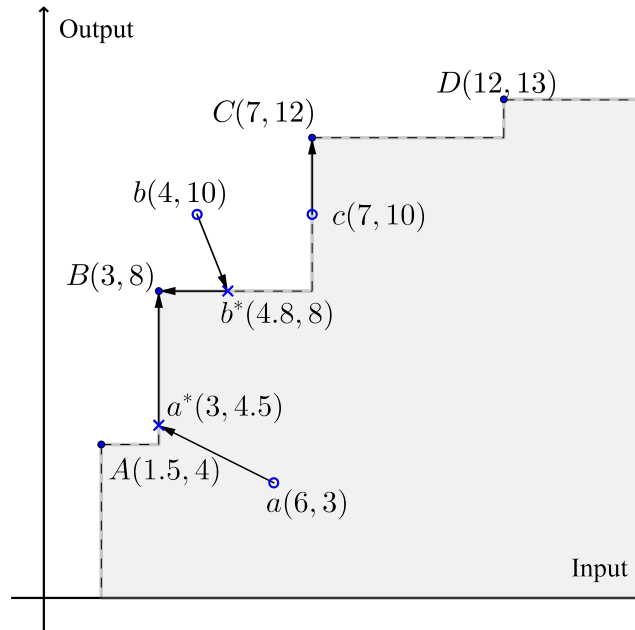
Figure 2: Illustration of model (1) with points *A*, *B*, *C* and *D* representing DMUs and points *a*, *b* and *c* representing aspiration values

set with better input/output values. Thus the adjusted aspiration values can be further improved, resulting in a better objective function value, which is contrary to the optimality of the solution. □

Figure 2 illustrates the use of model (1), in which a two-criterion problem with one input and one output is considered. In the figure, points *A*, *B*, *C* and *D* represent a set of efficient DMUs and points *a*, *b* and *c* each represent a set of aspiration values specified by the decision maker. Note that aspiration points are arbitrary points specified by the decision makers and may not be within the production possibility set, as shown by point *b*. Directional adjustment of aspiration values moves points *a* and *b* to points *a** and *b**. Notice that directional adjustment is determined by the corresponding aspiration values. Directional adjustment cannot be applied to point *c* as the point is already on the production frontier. Additive adjustment is then applied to points *a**, *b** and *c* to identify the corresponding efficient DMUs. The values of β, s^- and s^+ for points *a*, *b* and *c* are shown in Table 1.

Directional adjustment, which can be specified as $\beta(-x^T, y^T)^T$, allows the multiple inputs and the multiple outputs of the aspiration values to remain in the same proportion after adjustment. This is illustrated in Figure 3a where a two-criterion problem with only inputs is considered. In the figure, directional adjustment only

Table 1: Optimal β, s^- and s^+ values for the adjustment of points a, b and c in Figure 2

Point	$a(6,3)$	$b(4,10)$	$c(7,10)$
β	0.5	-0.2	0
(s^-, s^+)	(0, 3.5)	(1.8, 0)	(0, 2)
Identified DMU	B	B	C

moves aspiration points d and e along the corresponding ray emanating from the origin (shown in the dotted line) where any point on the ray has the same proportion in inputs as the corresponding aspiration values. Hence model (1) can be interpreted as a two step mechanism. The first step finds a point on the production frontier such that the point has criterion values in the same proportion among inputs and among outputs as the aspiration point. The second step then improves the adjusted aspiration point such that it equals one of the efficient DMUs.

By replacing the binary constraint of (1e) on λ, with convexity constraints, $\sum_{i=1}^{n} \lambda_i = 1$ and $0 \le \lambda_i \le 1 \; \forall i = 1, \ldots, n$, model (1) can be applied to a convex set where a convex combinations of inputs and outputs of existing DMUs are considered attainable. This convexity assumption is used in previous navigation methods where a convex combination of plans (or precisely, bixel intensities) is also a valid plan (Monz et al., 2008; Craft and Monz, 2010; Craft and Richter, 2013). Model (1) under convexity assumption is illustrated in Figure 3b where aspiration values represented by points d and e are adjusted to efficient points d^* and e^*. The convexity constraints correspond to the variable returns to scale assumption in the DEA context. Similarly, the constant returns to scale assumption in DEA can be facilitated by replacing (1e) with non-negativity constraints, $\lambda_i \ge 0 \; \forall i = 1, \ldots, n$.

2.3 Implementation

Determining the value of ε can be non-trivial in practice. Instead of solving model (1) directly, one can solve model (3) and (4) sequentially to get the same optimal solution without defining ε. Model (3) maximises β solely, without considering s^- and s^+, to obtain the optimal β value β^*. Model (4) maximises $e^T s^- + e^T s^+$ with the β value in the constraint set equal to β^*.
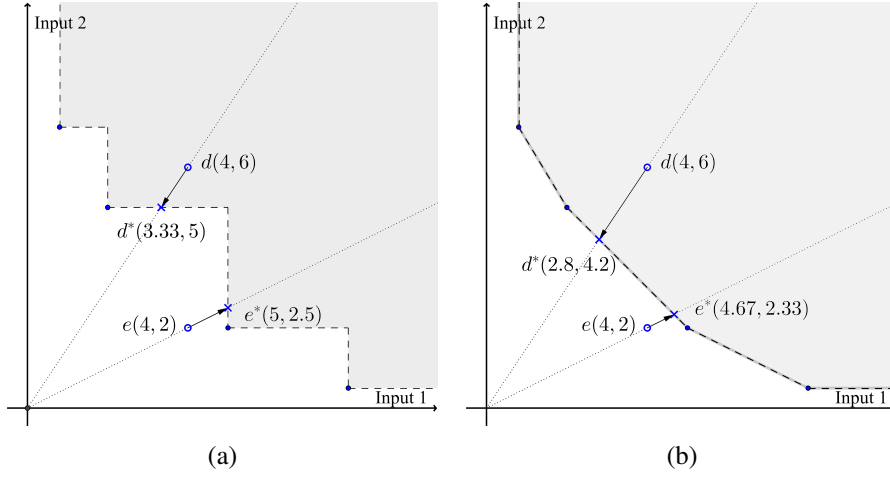
Figure 3: Illustration of directional adjustment for a two-input problem (a) with the assumption of free disposal hull (b) with the assumption of convexity.

$$
\begin{aligned}
\beta^* := \max \, & \beta \\
\text{s.t.} \quad & X\lambda \leq (1-\beta)x \\
& Y\lambda \geq (1+\beta)y \\
& e^T\lambda = 1 \\
& \lambda \in \{0,1\}^n.
\end{aligned}
\tag{3}
$$

$$
\begin{aligned}
\max \quad & e^T s^- + e^T s^+ \\
\text{s.t.} \quad & X\lambda \leq (1-\beta^*)x - s^- \\
& Y\lambda \geq (1+\beta^*)y + s^+ \\
& e^T\lambda = 1 \\
& \lambda \in \{0,1\}^n \\
& s^+, s^- \geq 0.
\end{aligned}
\tag{4}
$$

Proposition 3. *Models* (3) *and* (4) *are always feasible.*

Proof. Let $j^* \in \{1, \ldots, n\}$ and let $\lambda^{j^*} \in \mathbb{R}^n$ be a vector with one in the j^*th entry and zeros in the other entries and let

$$
\beta_{j^*} := \min \left\{ \frac{X_{i,j^*} - x_i}{-x_i}, \frac{Y_{o,j^*} - y_o}{y_o} : i = 1, \ldots, I, o = 1, \ldots, O \right\}.
$$

Following the proof of Proposition 1, it is easy to see that λ^{j^*} and β_{j^*} constitute a feasible solution for model (3). Consequently, given β^* and λ^* that satisfy constraints in model (3), constraints in model (4) will always be satisfied by s^+ and s^- equal to zero vectors. $\qquad\square$

2.4 Navigation mechanism with hard constraints

While the use of soft constraints in model (1) allows a decision maker to freely express the aspiration values, soft constraints do not allow exploration of alternative DMUs. In practice, changing the aspiration values may result in identifying the same DMU and the changes in (one or more) aspiration values required to identify an alternative DMU are not trivial. Instead, alternative DMUs can be accessed through the use of hard constraints. In this subsection, we illustrate the use of hard constraints in navigation. For convenience, we express the hard constraints for inputs only since extending the constraints to outputs is trivial.

Let the ith DMU be the DMU currently identified by the navigation system and let $e_j \in \mathbb{R}^I$ be a vector with one in the jth entry and zeros in all other entries. Constraint (5a) or (5b) can be added to model (1) in order to move from the current DMU to an alternative DMU where the kth input criterion is improved by ε_k or worsened by ε_k, respectively:

$$e_k^T X \lambda \leq X_{k,i} - \varepsilon_k, \tag{5a}$$
$$e_k^T X \lambda \geq X_{k,i} + \varepsilon_k. \tag{5b}$$

In our implementation (see Subsection 3.3), ε_k is the step size of the criterion slider bar defined as 1% of the value range of the kth input criterion, i.e., $\varepsilon_k = 0.01 \times \big(\max\{X_{k,j} : j \in I\} - \min\{X_{k,j} : j \in I\}\big)$.

While navigating, the decision maker may only be interested in DMUs with a certain set of input criteria, $\mathscr{L}$, better than or equal to certain values. This can be achieved by constraints of the form

$$e_l^T X \lambda \leq b_l, \quad \forall l \in \mathscr{L}, \tag{6}$$

where vector b records the upper bounding values for input criteria $l \in \mathscr{L}$.

Model (1) with constraints (5) and (6) forms the complete navigation model for this study. Notice that when hard constraints are imposed, soft constraints (1b) and (1c) still apply. That is, the navigation model still relies on the directional ad-

justment and the additive adjustment to identify an efficient DMU, although now the production possibility set is restricted by the hard constraints. In other words, the navigation mechanism still finds a DMU that best matches the aspiration values among DMUs that satisfy the hard constraints. Notice that the navigation model can be infeasible when one or more hard constraints exclude all DMUs from the feasible set. When infeasibility occurs, the decision maker will have to accept that the hard constraint(s) cannot be satisfied and should consider weakening the constraint(s) in order to find a solution. In our implementation, the hard constraints can only be imposed one at a time (by pressing a slider bar arrow button or by clicking a constraint check box, see Subsection 3.3). When a new hard constraint is added, model (1) with user-specified constraints of (5) and (6) is solved. Hence if a new constraint results in an infeasible model, the decision maker will be immediately notified.

3   Application to radiotherapy treatment planning

We apply the proposed navigation method to a prostate radiotherapy treatment planning problem. In section 3.1, we describe how a database of plans is generated. We present the clinical criteria used in this case study in Section 3.2. Section 3.3 presents the graphical user interface of the proposed navigation system and demonstrates the functionality of the system.

3.1   Database plan generation

The formulation used to generate treatment plans is based on the model of Holder (2003), as shown in (7). The dose deposition matrix $A \in \mathbb{R}^{m \times n}$, which represents the dose delivered to the m voxels from the n bixels under unit intensity, is partitioned and re-ordered into sub-matrices according to the structure type of the voxel, i.e., $A_T \in \mathbb{R}^{m_T \times n}$, $A_C \in \mathbb{R}^{m_C \times n}$ and $A_N \in \mathbb{R}^{m_N \times n}$ for planning target volume (PTV) T with m_T voxels, for critical organs C with m_C voxels and for normal tissue N with m_N voxels, respectively. Variables $x \in \mathbb{R}^n$ represent the radiation intensity of the n bixels. Dose lower bounds for the tumour are denoted as $LB_T \in \mathbb{R}^{m_T}$ and upper bounds for the tumour, critical organs and normal tissue are denoted as $UB_T \in \mathbb{R}^{m_T}$, $UB_C \in \mathbb{R}^{m_C}$ and $UB_N \in \mathbb{R}^{m_N}$, respectively. Variables $\alpha \in \mathbb{R}^{m_T}$, $\beta \in \mathbb{R}^{m_C}$ and $\gamma \in \mathbb{R}^{m_N}$ are the one-sided dose deviations from tumour lower bounds, critical organ upper bounds and normal tissue upper bounds, respectively.

13

$$\min \quad \left(\frac{1}{m_T}e^T\alpha, \frac{1}{m_C}e^T\beta, \frac{1}{m_N}e^T\gamma\right)$$
$$\text{s.t.} \quad LB_T - \alpha \le A_T x \le UB_T$$
$$A_C x \le UB_C + \beta \tag{7}$$
$$A_N x \le UB_N + \gamma$$
$$0 \le x, \alpha, \beta, \gamma$$

We consider three optimisation criteria: average one-sided dose deviation from the tumour dose lower bounds (criterion 1), average one-sided dose deviation from the rectal dose upper bounds (criterion 2) and average one-sided dose deviation from the bladder upper bounds (criterion 3). Other clinically relevant structures are involved in the formulation as constraints, e.g., voxels of the prostate are given a lower bound and an upper bound and voxels of the femur heads are given structure specific upper bounds. Model (7) can be equivalently reformulated with bixel intensity variables x replaced by segment intensity variables $\bar{x}$ and dose deposition matrix for bixels A replaced by dose deposition matrix for segments $\bar{A}$ (where each column of $\bar{A}$ represents the dose deposited to the voxels for a particular segment under unit intensity). The reformulated model is then solved with column generation integrated in the revised normal boundary intersection (RNBI) procedure (Shao and Ehrgott, 2007, 2016; Lin et al., 2016).

The RNBI procedure computes a representative set of non-dominated points for multiobjective linear programmes (MOLPs) (Shao and Ehrgott, 2007, 2016). The procedure constructs a set of equidistant reference points in the criterion space. For each reference point, a RNBI subproblem is solved to obtain the intersection point between the half-line emanating from the associated reference point and the feasible set, as illustrated in Figure 4. A RNBI subproblem of a reference point q corresponding to the reformulation of model (7) is shown in model (8). Model (8) is solved with the (single-objective) column generation technique, which iteratively generates columns that improve the current solution. The column generation technique essentially generates segments (represented by columns of $\bar{A}$) that potentially improve the current plan (Preciado-Walters et al., 2004; Romeijn et al., 2004). As a result, plans generated by column generation are readily deliverable, without the need to go through segmentation, which deteriorates plan quality (Rocha et al., 2012; Craft and Richter, 2013). In addition, one can also indirectly control the number of segments used in a treatment by terminating the column generation pro-
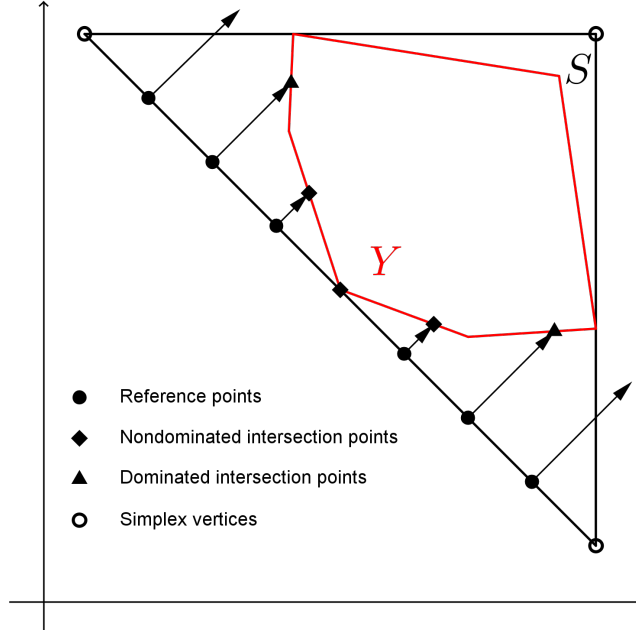
Figure 4: Illustration of the RNBI method. A simplex (*S*) that encloses the feasible set (*Y*) in the objective space is first created and a set of equidistant reference points is placed on the non-dominated set of *S*. For each reference point, a RNBI subproblem is solved to obtain the intersection point between the half-line emanating from the reference point and *Y*

cess when appropriate. For further details on column generation RNBI and its application to radiotherapy treatment planning, readers are referred to Lin et al. (2016).

$$
\begin{aligned}
\min \quad & t \\
\text{s.t.} \quad & q_1 + t = \frac{1}{m_T} e^T \alpha \\
& q_2 + t = \frac{1}{m_C} e^T \beta \\
& q_3 + t = \frac{1}{m_N} e^T \gamma \\
& LT_T - \alpha \leq \bar{A}_T \bar{x} \leq UB_T \\
& \bar{A}_C \bar{x} \leq UB_C + \beta \\
& \bar{A}_N \bar{x} \leq UB_N + \gamma \\
& 0 \leq \bar{x}, \alpha, \beta, \gamma, t
\end{aligned}
\tag{8}
$$

In the case study, 17 reference points (excluding reference points leading to an infeasible RNBI subproblem) are used in the RNBI method. During the column generation process, a solution is separately recorded in the database when the number of segments used in a treatment plan exceeds 40, 50, 60, 70, 80, 90 and 100. The column generation process terminates when a treatment plan consists of more than 100 segments. The plan generation step results in 119 radiotherapy treatment plans for the prostate treatment planning problem.

3.2 Navigation criteria as inputs and outputs

Once the plan database is ready, the planner specifies the navigation criteria to be used in evaluating plan quality. The navigation criteria, in terms of inputs and outputs, used for this planning problem are shown in Table 2. Left femur head and right femur head are abbreviated as LFH and RFH, respectively. Generalised equivalent uniform dose is abbreviated as gEUD. We denote the dose of which the associated iso-dose volume contains $v\%$ of the volume as D_v.

For dose conformity, we use the conformity index (CI) proposed by Lomax and Scheib (2003)

$$\text{CI} = \frac{V_{(T,d)}}{V_d}, \tag{9}$$

where $V_{(T,d)}$ is the d Gy isodose volume of the target and V_d is the d Gy isodose volume of the total patient volume. The CI value is considered as an output as higher values of CI are preferable (with a maximum value of 1). For target dose homogeneity, we use the homogeneity index (HI) proposed by Yoon (2007)

$$\text{HI} = \sqrt{\sum_{i \in I_T} (D_i - D_{mean})^2 \times \frac{v_i}{V}}, \tag{10}$$

where I_{PTV} is the index set for the target voxels and D_i and v_i are the dose and volume of voxel $i \in I$, respectively. D_{mean} is the mean dose and V is the volume of the target. Essentially, (10) measures the standard deviation of dose in the target volume. Hence HI is considered as an input in DEA terminology, because smaller dose deviation in the target is considered preferable.

It can be seen from Figure 5 that optimisation criteria may not correctly reflect the quality of clinical criteria. Ideally, a clinical criterion value, if considered as a function of optimisation criterion values, should be either strictly increasing or strictly decreasing. That is, an improvement in the optimisation criterion value should always lead to an improvement in the clinical criterion value. However,

Table 2: Navigation criteria used in the planning problem in terms of inputs and outputs.

| Inputs	Outputs
PTV HI	PTV D_{95}
Rectum gEUD	PTV CI
Rectum D_5	
Bladder D_{50}	
Bladder D_{25}	
LFH D_{10}	
RFH D_{10}	
Number of segments	

Figures 5a to 5c show that this is not the case. Thus when conducting navigation with optimisation criteria, it is necessary to also monitor the values of the clinical criteria in order to ensure attainment of a clinically acceptable plan.

3.3 Graphical user interface

A screen shot of the navigation system is shown in Figure 6. Three main components are used to conduct navigation: a criterion slider, an aspiration slider and a constraint check box. The value range of each slider is limited by the maximum and minimum of the corresponding criterion over all plans in the database. The minimum and maximum criterion values are also shown in the left and right text boxes under the corresponding aspiration slider. The user can use either the slider or the middle text box below each slider to change the corresponding criterion or aspiration value. When any of the aspiration values change, model (1) is updated accordingly and solved to obtain a plan that best satisfies the aspiration values. When the value of a criterion slider is changed, a hard constraint of the form of equation (5a) or (5b) is imposed on model (1) and the model is solved to obtain a plan that best satisfies the aspiration values under the hard constraint. If a current plan is replaced by an alternative plan, the values of the criterion sliders are changed to the values of the alternative plan. When a constraint check box is activated, a hard constraint of the form of (6) is added into model (1). The left and right text boxes below each criterion slider show the minimum and maximum feasible values for the corresponding criterion subject to existing hard constraints. When a criterion value of the current plan equals to the criterion's minimum/maximum feasible values with respect to the hard constraints, the background colour of the
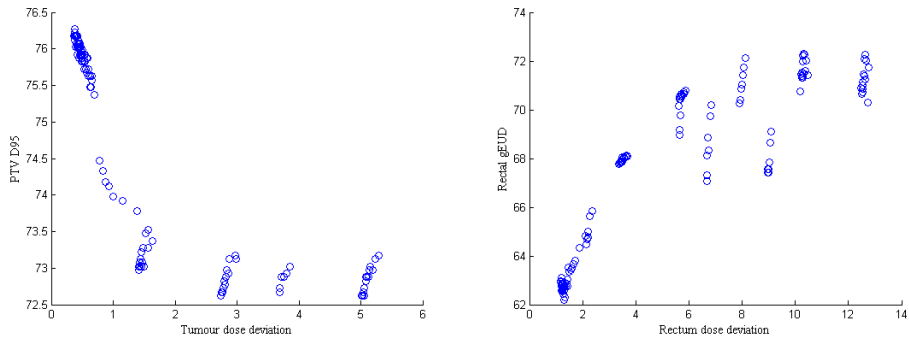
(a) PTV D_{95} against tumour dose deviation

(b) Rectum gEUD against rectum dose deviation
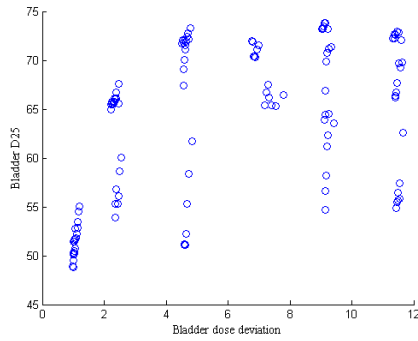
(c) Bladder D_{25} against bladder dose deviation

Figure 5: Plots showing that optimisation criterion values may not be appropriate quality indicators for clinical criteria.

corresponding text box will change to red, indicating further decrease/increase in the criterion value is not possible (unless one or more constraints are relaxed). If an aspiration value is/is not satisfied by the corresponding criterion value, the background colour of the criterion text box appears green/red. Similarly, the "Feasibility" text box shown on the top-left corner indicates if a navigation step is feasible (with green background colour) or not (with red background colour).

The navigation system is implemented in MATLAB and the optimisation models are solved with Gurobi. Clinical criterion values are calculated with CERR (Deasy et al., 2003). Experimenting with the prototype shows that the optimisation problem can be solved in real time for a dataset consists of 10000 plans. Additionally, we note that the optimisation model can also be solved analytically by enumerating over all possible λ vectors and find the β and slack variables that re-

18

Figure 6: Screen shot of the navigation system showing plan 5 with a set of aspiration values and constraints.

3.4   Demonstration of the navigation system

To demonstrate the use of the navigation system, we show a sequence of navigation steps. The navigation steps are composed of the following actions:

a Enter aspiration values

b Click on the right arrow of PTV D_{95} criterion value slider bar

c Click on rectal D_5 constraint check box

d Click on bladder D_{25} constraint check box

The action(s) performed and the plan identified in each step are shown in Table 3. Action(s) in brackets represent action(s) imposed from previous steps. The aspiration values used and criterion values of plans identified by the navigation system from the steps are shown in Table 4.

Table 3: Action(s) in each of the navigation step and the identified plan

Step	Action(s)	Plan ID
1	a	5
2	(a), b	66
3	(a), c, d	5
4	(a, c, d), b	infeasible
5	(a, c, d), d, b	9
6	(a, c, b), b	26
7	(a, c, b, b), b	60

In step 1, aspiration values for each criterion are set and the navigation system identifies plan 5 as the most preferable plan. All the criterion values of plan 5 achieve the corresponding aspiration values except for PTV D_{95}. Thus in step 2, action b is used to find a plan with an improved PTV D_{95} value, which results in plan 66 being identified. Although plan 66 achieves the PTV D_{95} aspiration value, many of the criterion values become worse than the corresponding aspiration values. In particular, we want to achieve the aspiration values for rectum D_5 and

Table 4: Aspiration values used and plans identified by the navigation steps in Table 3. Criterion values worse than the corresponding aspiration values are shown in bold.

Criteria	Aspiration	Plan 5	Plan 66	Plan 9	Plan 26	Plan 60
PTV D_{95}	74	**73.18**	75.83	**73.78**	**73.93**	74.13
PTV CI	0.6	0.77	**0.59**	0.67	0.67	0.65
PTV HI	1.7	1.62	1.32	**1.80**	**1.80**	**1.81**
rectum gEUD	67	63.54	**68.05**	64.83	64.35	63.68
rectum D_5	74	73.03	**76.18**	73.63	73.28	73.13
bladder D_{50}	45	36.13	35.58	41.93	41.93	44.43
bladder D_{25}	65	63.53	**67.68**	**66.43**	**65.33**	**66.18**
LFH D_{10}	35	15.78	17.58	7.98	8.33	9.23
RFH D_{10}	35	28.63	8.23	15.68	18.68	19.73
nSegment	70	40	**72**	42	50	70

bladder D_{25}. This can be specified with actions c and d. It is not surprising that plan 5 is identified by step 3 since from step 1 we know plan 5 is the most preferable plan for the given aspiration values while at the same time satisfying the rectum D_5 aspiration value and bladder D_{25} values. A screen shot of step 3 is shown in Figure 6. In step 4, one may attempt to find a plan with a better PTV D_{95} by conducting action b. However, it can be seen on Figure 6 that the PTV D_{95} of the current plan has reached the maximum feasible value under current constraints and action b would only lead to an infeasible model unless at least one of the constraints on rectum D_5 or bladder D_{25} is relaxed. In step 5, we relax the constraint on bladder D_{25} and try to improve PTV D_{95} through actions d and b, respectively. An improvement in PTV D_{95} is achieved, as shown by plan 9, although the aspired PTV D_{95} value is still not achieved. Further improvement in the PTV D_{95} value is attempted in steps 6 and 7, resulting in plans 26 and 60, respectively. Plan 60 shows that the aspired PTV D_{95} value is achieve at the cost of sacrificing the aspiration values for the inhomogeneity index and bladder D_{25}.

4   Discussion and conclusion

The benefit of MO in radiotherapy treatment planning has been demonstrated in several studies (Thieke et al., 2007; Hong et al., 2008; Craft et al., 2012; Wala et al., 2013; Chen et al., 2015). The MO based planning approach adopted in clinical practice uses linear interpolation of fluence maps as an approximation of the

efficient set. By doing so, one can quickly generate a large number of plans with modest computational expense. However, limitations are also encountered by this planning approach. Firstly, plans generated through interpolation are not truly efficient and may be further improved (Bokrantz and Miettinen, 2015). The quality of these plans is sacrificed in order to reduce the computational expense. Secondly, linear interpolation of fluence maps does not consider plan delivery. Thus, a plan selected from the efficient set needs to go through the segmentation step, which can significantly deteriorate the plan quality (McGarry et al., 2014; Rocha et al., 2012). To address this issue, Salari and Unkelbach (2013), Craft and Richter (2013) and Fredriksson and Bokrantz (2013) have proposed navigation methods for deliverable plans. These methods rely on the interpolation of segmented plans and/or segments to find deliverable plans, which neglect plan improvement potential (as stated in the first limitation). Thirdly, linear interpolation cannot be applied to plans with different beam angle configurations, thus plans generated from interpolation are subjected to the same beam angle configuration. However, the best beam angle configuration to achieve different treatment trade-offs can be different (see, e.g., Cabrera et al. (2016)). Excluding the search of different beam angle configurations in the plan generation process may prevent the best-quality plans to be found. Overall the effect of these limitations may lead to the generation of plans with considerable improvement potential.

An alternative MO practice is to generate a finely-sampled discrete representation of the efficient set in which each plan is given freedom in beam angle configuration, segment sets and segment intensities in order to attain a specific trade-off in treatment goals. Plan generation with this approach can be computationally expensive, especially for problems with many optimisation criteria. However, recent advancements of high performance computing, such as Jia et al. (2014), Tian et al. (2015) and Ziegenhein et al. (2013), alleviate this potential drawback. To further reduce the computational expense, one can also consider a two-stage planning practice where in the first stage, a coarser sample of the efficient set is generated and navigation is used to identify a close-to-ideal plan, followed by fine-tuning (see, e.g., Otto (2014); Ziegenhein et al. (2014)) of the dose distribution in the second stage.

The navigation method proposed in this study best facilitates the latter MO practice where navigation is conducted on a discrete set of plans. A distinct feature of the proposed navigation model, compared to previous navigation models, is its ability to navigate among a discrete set through the use of soft constraints. This

allows a planner to use any criteria, including the commonly used dose-volume criteria, which are known to be non-smooth, for navigation. Instead of using optimisation criteria, which may not perfectly correlate with clinical evaluation criteria, the use of clinical criteria enables the planner to assess plan quality and conduct navigation (including setting constraints) in a more intuitive manner. The knowledge derived from the navigation steps can then be used to form new preferences and the iterative process of navigation and forming new preference continues until the most preferable plan is identified.

References

R. Allmendinger, M. Ehrgott, X. Gandibleux, M. J. Geiger, K. Klamroth, and M. Luque. Navigation in multiobjective optimization methods. *Journal of Multi-Criteria Decision Analysis*, 2016. doi: 10.1002/mcda.1599.

R. Bokrantz and A. Forsgren. An algorithm for approximating convex Pareto surfaces based on dual techniques. *INFORMS Journal on Computing*, 25(2):377–393, 2013.

R. Bokrantz and K. Miettinen. Projections onto the Pareto surface in multicriteria radiation therapy optimization. *Medical Physics*, 42(10):5862–5870, 2015.

T. Bortfeld and W. Schlegel. Optimization of beam orientations in radiation therapy: Some theoretical considerations. *Physics in Medicine and Biology*, 38(2): 291–304, 1993.

S. Breedveld, P. R. Storchi, and B. J. Heijmen. The equivalence of multi-criteria methods for radiotherapy plan optimization. *Physics in Medicine and Biology*, 54(23):7199, 2009.

G. Cabrera, M. Ehrgott, A. J. Mason, and A. Raith. A matheuristic approach to solve the multiobjective beam angle optimization problem in intensity-modulated radiation therapy. *International Transactions in Operational Research*, pages 1–26, 2016.

R. G. Chambers, Y. Chung, and R. Färe. Benefit and distance functions. *Journal of Economic Theory*, 70(2):407–419, 1996.

R. G. Chambers, Y. Chung, and R. Färe. Profit, directional distance functions, and Nerlovian efficiency. *Journal of Optimization Theory and Applications*, 98(2): 351–364, 1998.

A. Charnes, W. W. Cooper, and E. Rhodes. Measuring the efficiency of decision making units. *European Journal of Operational Research*, 2(6):429–444, 1978.

H. Chen, B. A. Winey, J. Daartz, K. S. Oh, J. H. Shin, and D. P. Gierga. Efficiency gains for spinal radiosurgery using multicriteria optimization intensity modulated radiation therapy guided volumetric modulated arc therapy planning. *Practical Radiation Oncology*, 5(1):49–55, 2015.

W. W. Cooper, L. M. Seiford, and J. Zhu. Data envelopment analysis: History, models, and interpretations. In W. W. Cooper, L. M. Seiford, and J. Zhu, editors, *Handbook on Data Envelopment Analysis*, pages 1–39. Springer, New York, 2011.

D. Craft and M. Monz. Simultaneous navigation of multiple Pareto surfaces, with an application to multicriteria IMRT planning with multiple beam angle configurations. *Medical Physics*, 37(2):736–741, 2010.

D. Craft and C. Richter. Deliverable navigation for multicriteria step and shoot IMRT treatment planning. *Physics in Medicine and Biology*, 58(1):87–103, 2013.

D. L. Craft, T. F. Halabi, H. A. Shih, and T. R. Bortfeld. Approximating convex Pareto surfaces in multiobjective radiotherapy planning. *Medical Physics*, 33 (9):3399–3407, 2006.

D. L. Craft, T. S. Hong, H. A. Shih, and T. R. Bortfeld. Improved planning time and plan quality through multicriteria optimization for intensity–modulated radiotherapy. *International Journal of Radiation Oncology Biology Physics*, 82 (1):e83–e90, 2012.

J. O. Deasy. Multiple local minima in radiotherapy optimization problems with dose–volume constraints. *Medical Physics*, 24(7):1157–1161, 1997.

J. O. Deasy, A. I. Blanco, and V. H. Clark. CERR: A computational environment for radiotherapy research. *Medical Physics*, 30(5):979–985, 2003.

M. Ehrgott and I. Winz. Interactive decision support in radiation therapy treatment planning. *OR Spectrum*, 30(2):311–329, 2008.

M. Falkinger, S. Schell, J. Müller, and J. J. Wilkens. Prioritized optimization in intensity modulated proton therapy. *Zeitschrift fur Medizinische Physik*, 22(1): 21–28, 2012.

A. Fredriksson and R. Bokrantz. Deliverable navigation for multicriteria IMRT treatment planning by combining shared and individual apertures. *Physics in Medicine and Biology*, 58(21):7683–7697, 2013.

A. L. Hoffmann, D. den Hertog, A. Y. Siem, J. H. Kaanders, and H. Huizenga. Convex reformulation of biologically-based multi-criteria intensity-modulated radiation therapy optimization including fractionation effects. *Physics in Medicine and Biology*, 53(22):6345–6362, 2008.

A. Holder. Designing radiotherapy plans with elastic constraints and interior point methods. *Health Care Management Science*, 6(1):5–16, 2003.

T. S. Hong, D. L. Craft, F. Carlsson, and T. R. Bortfeld. Multicriteria optimization in intensity–modulated radiation therapy treatment planning for locally advanced cancer of the pancreatic head. *International Journal of Radiation Oncology Biology Physics*, 72(4):1208–1214, 2008.

K. W. Jee, D. L. McShan, and B. A. Fraass. Lexicographic ordering: Intuitive multicriteria optimization for IMRT. *Physics in Medicine and Biology*, 52(7): 1845–1861, 2007.

X. Jia, P. Ziegenhein, and S. B. Jiang. GPU-based high-performance computing for radiation therapy. *Physics in Medicine and Biology*, 59(4):R151–R182, 2014.

X. A. Li, M. Alber, J. O. Deasy, A. Jackson, K.-W. K. Jee, L. B. Marks, M. K. Martel, C. Mayo, V. Moiseenko, A. E. Nahum, A. Niemierko, V. A. Semenenko, and E. D. Yorke. The use and QA of biologically related models for treatment planning: Short report of the TG-166 of the therapy physics committee of the AAPM. *Medical Physics*, 39(3):1386–1409, 2012.

K.-M. Lin, J. Simpson, G. Sasso, A. Raith, and M. Ehrgott. Quality assessment for VMAT prostate radiotherapy planning based on data envelopment analysis. *Physics in Medicine and Biology*, 58(16):5753–5769, 2013.

K.-M. Lin, M. Ehrgott, and A. Raith. Integrating column generation in a method to compute a discrete representation of the non-dominated set of multi-objective linear programmes. *4OR*, pages 1–27, 2016.

J. Llacer, J. O. Deasy, T. R. Bortfeld, T. D. Solberg, and C. Promberger. Absence of multiple local minima effects in intensity modulated optimization with dose–volume constraints. *Physics in Medicine and Biology*, 48(2):183, 2003.

N. J. Lomax and S. G. Scheib. Quantifying the degree of conformity in radio-surgery treatment planning. *International Journal of Radiation Oncology Biology Physics*, 55(5):1409–1419, 2003.

C. K. McGarry, R. Bokrantz, J. M. O'Sullivan, and A. R. Hounsell. Advantages and limitations of navigation-based multicriteria optimization (MCO) for localized prostate cancer IMRT planning. *Medical Dosimetry*, 39(3):205–211, 2014.

M. Monz, K. Küfer, T. Bortfeld, and C. Thieke. Pareto navigation-algorithmic foundation of interactive multi-criteria IMRT planning. *Physics in Medicine and Biology*, 53(4):985, 2008.

K. Otto. Real-time interactive treatment planning. *Physics in Medicine and Biology*, 59(17):4845, 2014.

F. Preciado-Walters, R. Rardin, M. Langer, and V. Thai. A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer. *Mathematical Programming*, 101(2):319–338, 2004.

D. A. Ripsman, D. M. Aleman, and K. Ghobadi. Interactive visual guidance for automated stereotactic radiosurgery treatment planning. *Expert Systems with Applications*, 42(21):8337–8348, 2015.

H. Rocha, J. M. Dias, B. C. Ferreira, and M. C. Lopes. Discretization of optimal beamlet intensities in IMRT: A binary integer programming approach. *Mathematical and Computer Modelling*, 55(7-8):1969–1980, 2012.

H. E. Romeijn, J. F. Dempsey, and J. G. Li. A unifying framework for multi-criteria fluence map optimization models. *Physics in Medicine and Biology*, 49 (10):1991–2013, 2004.

I. Rosen, H. H. Liu, N. Childress, and Z. Liao. Interactively exploring optimized treatment plans. *International Journal of Radiation Oncology Biology Physics*, 61(2):570–582, 2005.

H. Ruotsalainen. *Interactive multiobjective optimization in model-based decision making with applications*. PhD thesis, Doctoral Thesis, University of Kuopio, Kuopio, Finland, 2009.

E. Salari and J. Unkelbach. A column-generation-based method for multi-criteria direct aperture optimization. *Physics in Medicine and Biology*, 58(3):621–639, 2013.

L. Shao and M. Ehrgott. Finding representative nondominated points in multiobjective linear programming. In *IEEE Symposium on Computational Intelligence in Multicriteria Decision Making*, pages 245–252. IEEE, 2007.

L. Shao and M. Ehrgott. Discrete representation of non-dominated sets in multi-objective linear programming. *European Journal of Operational Research*, 255 (3):687–698, 2016.

C. Thieke, K.-H. Küfer, M. Monz, A. Scherrer, F. Alonso, U. Oelfke, P. E. Huber, J. Debus, and T. Bortfeld. A new concept for interactive radiotherapy planning with multicriteria optimization: First clinical evaluation. *Radiotherapy and Oncology*, 85(2):292–298, 2007.

Z. Tian, F. Peng, M. Folkerts, J. Tan, X. Jia, and S. B. Jiang. Multi-GPU implementation of a VMAT treatment plan optimization algorithm. *Medical Physics*, 42(6):2841–2852, 2015.

J. Wala, D. Craft, J. Paly, A. Zietman, and J. Efstathiou. Maximizing dosimetric benefits of IMRT in the treatment of localized prostate cancer through multicriteria optimization planning. *Medical Dosimetry*, 38(3):298–303, 2013.

J. J. Wilkens, J. R. Alaly, K. Zakarian, W. L. Thorstad, and J. O. Deasy. IMRT treatment planning based on prioritizing prescription goals. *Physics in Medicine and Biology*, 52(6):1675–1692, 2007.

M. Yoon. A new homogeneity index based on the statistical analysis of dose volume histogram. *Journal of Applied Clinical Medical Physics*, 8(2):9–17, 2007.

P. Ziegenhein, C. P. Kamerling, M. Bangert, J. Kunkel, and U. Oelfke. Performance-optimized clinical IMRT planning on modern CPUs. *Physics in Medicine and Biology*, 58(11):3705–3715, 2013.

P. Ziegenhein, C. P. Kamerling, and U. Oelfke. Interactive dose shaping - efficient strategies for CPU-based real-time treatment planning. *Journal of Physics: Conference Series*, 489(1):012066–1–012066–6, 2014.