

Integrating column generation in a method to compute a discrete representation of the non-dominated set of multi-objective linear programmes

**Kuan-Min Lin · Matthias Ehrgott ·
Andrea Raith**

Received: date / Accepted: date

Abstract In this paper we propose the integration of column generation in the revised normal boundary intersection (RNBI) approach to compute a representative set of non-dominated points for multi-objective linear programmes (MOLPs). The RNBI approach solves single objective linear programmes, the RNBI subproblems, to project a set of evenly distributed reference points to the non-dominated set of an MOLP. We solve each RNBI subproblem using column generation, which moves the current point in objective space of the MOLP towards the non-dominated set. Since RNBI subproblems may be infeasible, we attempt to detect this infeasibility early. First, a reference point bounding method is proposed to eliminate reference points that lead to infeasible RNBI subproblems. Furthermore, different initialisation approaches for column generation are implemented, including Farkas pricing. We investigate the quality of the representation obtained.

To demonstrate the efficacy of the proposed approach, we apply it to an MOLP arising in radiotherapy treatment design. In contrast to conventional optimisation approaches, treatment design using column generation provides deliverable treatment plans, avoiding a segmentation step which deteriorates treatment quality. As a result total monitor units is considerably reduced. We also note that reference point bounding dramatically reduces the number of RNBI subproblems that need to be solved.

Keywords Multi-objective linear programming · Column generation · Revised normal boundary intersection method · Radiotherapy treatment design

Mathematics Subject Classification (2000) 90C29 · 90C05 · 90C90

KM. Lin

Department of Management Science, Lancaster University, Bailrigg, Lancaster LA1 4YX, United Kingdom. E-mail: k.lin1@lancaster.ac.uk, Tel.: +44 1524 510752

M. Ehrgott

Department of Management Science, Lancaster University, Bailrigg, Lancaster LA1 4YW, United Kingdom. E-mail: m.ehrgott@lancaster.ac.uk, Tel.: +44 1524 593888

A. Raith

Department of Engineering Science, The University of Auckland, 70 Symonds Street, Auckland 1010, New Zealand. E-mail: a.raith@auckland.ac.nz, Tel.: +64 9 373 7599 ext 81977

1 Introduction

Multi-objective optimisation (MOO) deals with optimisation problems involving several conflicting objectives. In MOO, a single solution that simultaneously optimises all objectives generally does not exist. Instead, MOO seeks for solutions that cannot improve in any single objective without deteriorating at least one other objective. Solutions with this property are referred to as efficient solutions. The points obtained by mapping the efficient solutions to the objective space are referred to as non-dominated points. The purpose of MOO is to obtain the non-dominated set and one efficient solution in the pre-image of every non-dominated point. A decision maker then has the task to select the most preferred non-dominated point and a corresponding efficient solution for the problem at hand. In multi-objective continuous optimisation, the non-dominated set consists of infinitely many non-dominated points. It is therefore impractical for a decision maker to examine all non-dominated points. Instead, a practical approach is to obtain a discrete representation of the non-dominated set satisfying some quality requirements (Sayin, 2000; Faulkenberg and Wiecek, 2010). Many methods that follow this approach have been proposed in the last two decades, as the paper by Faulkenberg and Wiecek (2010) shows. Given this representative non-dominated set, the decision maker can navigate through the non-dominated points and decide on the most preferred point. In this study we propose to integrate column generation in an approach to find a representative non-dominated set for multi-objective linear programmes (MOLPs).

Column generation is a technique that solves linear programmes by considering only a subset of the decision variables. The technique is particularly beneficial when the number of variables is much greater than the number of constraints. The idea is based on the fact that, typically, only a subset of variables is required in the basis to reach optimality; other variables are non-basic and have a value of zero. Column generation exploits this fact by only considering variables that have the potential to improve the objective function value, indicated by negative reduced costs. In each iteration of a column generation method, two problems need to be solved successively: the restricted master problem (RMP) and the subproblem (SP). RMP is the original problem with only a subset of variables. By solving the RMP, a vector of dual values associated with the constraints of the RMP is obtained. The dual information is passed on to the SP. The goal of the SP is to identify a new variable and an associated coefficient column with negative reduced cost, which can potentially improve the objective function value of the original problem. If such a variable and column can be identified, then they are added to RMP, which is re-optimised, and the next iteration begins. Otherwise, an optimal solution of RMP is also an optimal solution of the original problem.

Column generation methods in multi-objective optimisation are rare. Moradi et al (2015) present a column generation approach for the (linear) bi-objective multi-commodity minimum cost flow problem. Their algorithm incorporates column generation within a bi-objective simplex algorithm, which requires a modification of the objective function of the SP to a linear fractional function. The study of Salari and Unkelbach (2013) falls into the domain of non-linear programming, thus the subproblem is based on partial derivatives of individual objective functions. The aim of

[Salari and Unkelbach \(2013\)](#) is to approximate the entire non-dominated set using a limited number of variables. The basic idea is to use column generation to identify variables that potentially improve the non-dominated set approximation as a whole. To find such variables, multiple weighted-sum RMPs, where each RMP is associated with a unique non-negative weight vector, are solved. The partial derivatives obtained from solving each RMP are passed to a subproblem, which aggregates the individual subproblems corresponding to each RMP. A column obtained from solving the aggregated subproblem therefore potentially improves the majority of individual RMPs, thus improving the non-dominated set approximation as a whole. However, due to the use of weight vectors for the RMPs and the use of an aggregated subproblem, their method cannot guarantee that the whole non-dominated set is well approximated.

In this study, we propose to use column generation within a procedure that constructs an evenly distributed finite representative set of non-dominated points of an MOLP, i.e. the revised normal boundary intersection (RNBI) method of [Shao and Ehrgott \(2007, 2016\)](#). The RNBI method combines aspects of the global shooting method ([Benson and Sayin, 1997](#)) and the normal boundary intersection method ([Das and Dennis, 1998](#)) and has been proven to generate evenly distributed non-dominated points for MOLPs ([Shao and Ehrgott, 2007](#)). Unlike the method of [Salari and Unkelbach \(2013\)](#) in which a subproblem identifies a variable that improves the non-dominated set approximation in general, each of the column generation subproblems in our approach identifies a variable and an associated column to move a point in objective space in a direction that leads to non-dominance. In fact, if column generation is run to termination, i.e. optimality of the master problem, the resulting point will be on the boundary of the feasible set of the MOLP in objective space.

We apply our method to a multi-objective optimisation problem in radiotherapy treatment design. The goal of this problem is to identify a treatment plan (in the form of so-called fluence maps for several radiation beams) in order to deliver a tumouricidal dose of radiation to a planning target volume, while sparing healthy tissue. These conflicting goals naturally lead to formulations as multi-objective optimisation problems. We refer the reader to [Ehrgott et al \(2008a\)](#) for more details on optimisation methods in radiation oncology. By applying our column generation RNBI method to the treatment design problem, a set of representative treatment plans, each with a unique trade-off between objective function values, are generated. Given these plans, the oncologist can then decide on the plan that best benefits the patient.

Conventional and multi-objective approaches in radiotherapy treatment design generate treatment plans that often cannot be practically delivered by existing radiotherapy equipment. In order to make them deliverable, one needs to modify the treatment plans to incorporate physical delivery constraints and to reduce the total time a patient is exposed to radiation. This modification, which is referred to as segmentation, deteriorates treatment plan quality ([Rocha et al, 2012](#); [Craft and Richter, 2013](#)). Thus, should a treatment plan become unsatisfactory after segmentation, the treatment planner will have to re-optimize and find another plan. This iterative process makes the treatment design process inefficient. However, we shall see that the fluence map optimisation problem can be reformulated via decomposition to include the physical delivery constraints. It can then be solved by column generation to obtain treatment plans that are directly deliverable. In fact, we shall see that column gener-

ation produces plans that are close to optimality with a reduced delivery complexity, which are often preferable to optimal (but complex) plans in practice (Carlsson and Forsgren, 2014; Broderick et al, 2009).

In Section 2, we provide background and formulations of single objective column generation and the RNBI method. In Section 3, we introduce the column generation RNBI formulation and discuss implementation issues associated with the method, i.e. the detection of infeasibility through a reference point bounding method and initialisation of the process. The quality of the representative set obtained by our column generation based RNBI method is also discussed. In Section 4, we apply the method to a prostate radiotherapy treatment planning problem, followed by results and discussion in Section 5.

2 Column generation and the RNBI method

In this section we provide necessary details for column generation and the RNBI method. For further details of these two topics, we refer the readers to Lübbecke (2010) and Shao and Ehrgott (2007), respectively.

2.1 Column generation

Consider a single objective linear programme referred to as the master problem MP,

$$\begin{aligned} v_{MP} := \min & \sum_{j \in J} c_j x_j \\ \text{s.t.} & \sum_{j \in J} a^j x_j \geq b, \\ & x_j \geq 0, \quad j \in J, \end{aligned} \tag{MP}$$

with $|J| = n$ variables and m constraints. Each variable x_j is associated with a cost coefficient c_j and a constraint column $a^j \in \mathbb{R}^m$. The right-hand side constraint coefficients are specified by column $b \in \mathbb{R}^m$. The column generation technique considers a restricted master problem (RMP) which uses only a subset $J' \subseteq J$ of all variables. Because of this, the optimal solution x^* of RMP is worse than or equal to the optimal solution of MP in terms of the objective function. By solving the RMP, we obtain a dual solution π^* associated with the constraints of the RMP. In the Simplex method, the dual solution is used to calculate the reduced cost of each non-basic variable which indicates the unit change of the objective function value if the variable were to enter the basis. If the reduced cost of all non-basic variables is non-negative, the current basic feasible solution of RMP is an optimal solution to MP. Otherwise, a non-basic variable with negative reduced cost enters the basis, which improves the objective function if the entering variable takes a value greater than zero.

Column generation works in an analogous way. The vector of dual values π^* obtained from solving the RMP is passed into a subproblem,

$$\begin{aligned} \min \bar{c}_j &= c_j - (\pi^*)^T a^j \\ \text{s.t. } j &\in J. \end{aligned} \quad (\text{SP})$$

SP finds a variable x_{j^*} with lowest reduced cost \bar{c}_{j^*} . If \bar{c}_{j^*} is negative, the non-basic variable x_{j^*} and the coefficient column $(c_{j^*}, a^{j^*})^T$ are added to RMP and RMP is re-solved. Otherwise, an optimal solution of RMP is also an optimal solution of MP. Note that SP can be solved as an optimisation problem if the set J can be described by the feasible set X_J of an optimisation problem

$$\begin{aligned} \min c(\lambda) - (\pi^*)^T a(\lambda) \\ \text{s.t. } \lambda \in X_J, \end{aligned} \quad (1)$$

in which $c_j = c(\lambda)$ and $a^j = a(\lambda)$ and which has variable vector $\lambda \in X_J$.

The lowest reduced cost \bar{c}_{j^*} can be used to derive a lower bound on the optimal value v_{MP} of MP. Denote the optimal value of the current RMP as v_{RMP}^* . If there exists a constant κ with $\sum_{j \in J} x_j \leq \kappa$ for any optimal solution of MP, then we have a lower bound

$$v_{RMP}^* + \kappa \bar{c}_{j^*} \leq v_{MP}, \quad (2)$$

since we cannot improve the objective function value v_{RMP}^* by more than κ times the lowest reduced cost \bar{c}_{j^*} (Lübbecke, 2010).

2.2 The RNBI method

Consider an MOLP

$$\min\{Cx : x \in X\}, \quad (\text{MOLP})$$

where $C \in \mathbb{R}^{p \times n}$ is the cost coefficient matrix consisting of row vectors $c^k \in \mathbb{R}^n$ for $k = 1, \dots, p$. Throughout this paper we will assume that $X \subseteq \mathbb{R}^n$ is a non-empty compact polyhedral set (a polytope) of feasible solutions. The feasible set in objective space Y defined by

$$Y = \{Cx : x \in X\} \quad (3)$$

is also a polytope since it is the image under a linear mapping of the polytope X .

Definition 1 A feasible solution $\hat{x} \in X$ of MOLP is called *efficient* if there does not exist any $x \in X$ such that $c^k x \leq c^k \hat{x}$ for $k = 1, \dots, p$, and $Cx \neq C\hat{x}$. The set of all efficient solutions of the MOLP is called the efficient set X_E in decision space. The image in objective space $\hat{y} = C\hat{x}$ of efficient solution \hat{x} is called *non-dominated* point. The set of all non-dominated points is referred to as the non-dominated set Y_N in objective space.

We first explain the general idea of RNBI. A simplex S is constructed such that it contains Y and such that the non-dominated set S_N of S is a subsimplex of S . We denote by $\hat{S} := S_N$ the reference subsimplex. Reference points are positioned on \hat{S}

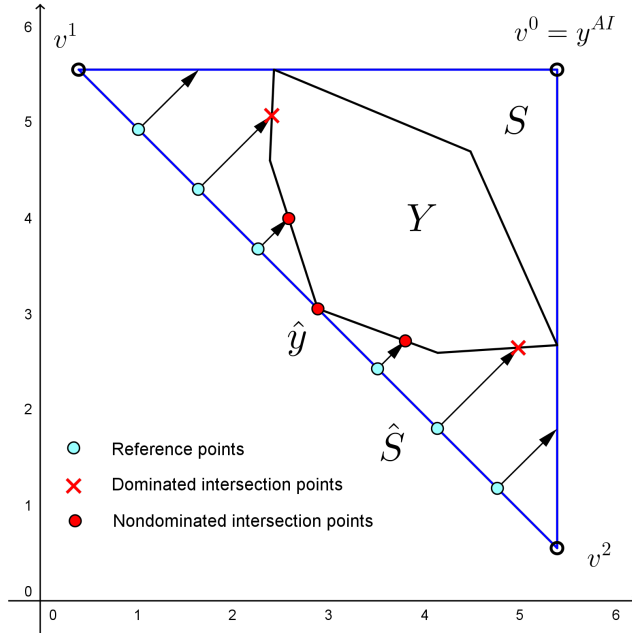


Fig. 1 The RNBI method: Illustration of the simplex S containing the feasible set in objective space Y , the reference subsimplex \hat{S} and the half-lines emanating from the reference points.

and for each reference point q , a half-line emanating from q in direction e is generated, where e is a vector of all ones. The RNBI subproblem then searches for the intersection point between the half-line and (the boundary of) Y closest to the reference point. As illustrated in Figure 1, not all half-lines intersect with Y . In this case the RNBI subproblem will be infeasible. In addition, some intersection points may be dominated. Hence in the last step, the algorithm checks the non-dominance of intersection points by solving one LP for each intersection point. The following subsections outline the RNBI method in more detail.

2.2.1 Constructing the reference subsimplex and choosing reference points

To construct the reference subsimplex \hat{S} , we first obtain scalar μ as $\mu := \min \{e^T y : y \in Y\}$. μ is attained at a non-dominated point \hat{y} of Y , as illustrated in Figure 1. We then derive the anti-ideal point y^{AI} of the MOLP, where $y_k^{AI} := \max \{y_k : y \in Y\}$ for $k = 1, \dots, p$ (Ehrgott, 2005). Based on μ and y^{AI} , we can define the $p + 1$ vertices $v^k \in \mathbb{R}^p, k = 0, 1, \dots, p$ of the simplex S that contains Y . Let $v^0 := y^{AI}$. For $k = 1, \dots, p$ and $l = 1, \dots, p$ let

$$v_l^k := \begin{cases} y_l^{AI} & \text{if } l \neq k, \\ \mu + y_k^{AI} - e^T v^0 & \text{if } l = k. \end{cases} \quad (4)$$

The convex hull of vertices $\{v^k : k = 0, 1, \dots, p\}$ is a p -dimensional simplex S that contains Y , as shown by Benson and Sayin (1997). The reference subsimplex \hat{S} , which is the non-dominated set of S , is defined by the convex hull of vertices

$\{v^k : k = 1, 2, \dots, p\}$. Reference points on \hat{S} can now be chosen as particular convex combinations of the extreme points of \hat{S} , i.e. a reference point q is given by

$$q = \sum_{k=1}^p \alpha_k v^k, \quad (5)$$

where α_k is the weighting of vertex k for $k = 1, \dots, p$ with $0 \leq \alpha_k \leq 1$ and $\sum_{k=1}^p \alpha_k = 1$. By varying the weighting for each vertex with a fixed increment η , an evenly distributed discrete set of points on the reference subsimplex \hat{S} can be generated (Benson and Sayin, 1997). Let the set of reference points be denoted Q .

2.2.2 Computing the intersection points and checking non-dominance

For each reference point $q \in Q$, RNBI computes the intersection point y of the half-line $\{q + te : t \geq 0\}$ and the boundary of Y by solving the RNBI subproblem

$$\min\{t : q + te \in Y; t \geq 0\}. \quad (\text{RNBIsub})$$

Notice that by construction, the all-ones vector e is the normal of the reference subsimplex \hat{S} . As illustrated in Figure 1, there are three scenarios for RNBIsub:

- RNBIsub is infeasible if and only if the half-line $\{q + te : t \geq 0\}$ does not intersect Y .
- RNBIsub has an optimal solution t^* , but the intersection point $q + t^*e$ of the half-line $\{q + te : t \geq 0\}$ and Y is dominated.
- RNBIsub has an optimal solution t^* and $q + t^*e$ is a non-dominated point of Y .

The first case is detected by infeasibility of RNBIsub. Because an intersection point may be a dominated point, it is necessary to check every intersection point for non-dominance. To do so, after obtaining all intersection points, a non-domination filter can be used to exclude some of the dominated points (Messac et al, 2003). This method allows fast elimination of some dominated intersection points but cannot guarantee the remaining points are non-dominated. Hence the non-dominance of the remaining intersection points \bar{y} must be verified, e.g. by solving the linear programme

$$\min\{\omega^T y : y \leq \bar{y}; y \in Y\}, \quad (6)$$

where $0 < \omega \in \mathbb{R}^p$ is an arbitrary strictly positive weight vector, for instance $\omega = e$. Then \bar{y} is non-dominated if and only if the optimal value of (6) is equal to $\omega^T \bar{y}$ (Ehrgott, 2005).

3 The RNBI method using column generation

To integrate column generation in the RNBI framework we solve RNBIsub using column generation. To do so, we adopt RNBIsub as the master problem. Following the background definitions for column generation in Section 2.1, we formulate the restricted master problem with a subset $J' \cup \{t\}$ of variables and with feasible set

defined by constraints $\sum_{j \in J'} a^j x_j \geq b$, $x_j \geq 0$ for all $j \in J'$ and $t \geq 0$. The condition $q + te \in Y$ of RNBISub is rewritten as constraints $q_k + t = \sum_{j \in J'} c_j^k x_j$ for $k = 1, \dots, p$. In this way, the objective functions of the original MOLP are incorporated in the restricted master problem as constraints of RNBISub, i.e. (7b), in addition to the original constraints of MP, i.e. (7c). The corresponding RMP is referred to as RMP-RNBISub, and is shown as follows.

$$\min \quad t \tag{7a}$$

$$\text{s.t.} \quad q_k + t = \sum_{j \in J'} c_j^k x_j, \quad k = 1, \dots, p, \tag{7b}$$

$$\sum_{j \in J'} a^j x_j \geq b, \tag{7c}$$

$$x_j \geq 0, \quad j \in J' \tag{7d}$$

$$t \geq 0. \tag{7e}$$

Notice that RMP-RNBISub is essentially the same as the RNBI subproblem but with only a subset of variables $j \in J'$. To conduct column generation on this RNBI subproblem, we solve RMP-RNBISub and the corresponding SP sequentially and iteratively. We remark that, in case column generation is terminated early, i.e. an optimal solution of RNBISub is not yet confirmed, the intersection point may be dominated. In contrast to the original RNBI method, non-dominance of the intersection points will not be checked, because only an optimal solution of RNBISub can define a non-dominated point.

As indicated in Section 2.2.2 RNBISub may be infeasible even in the presence of all variables. Hence, if we solve RMP-RNBISub with a subset of variables, it may be infeasible because either the constraints (7c) are not satisfied with a subset of variables or because the master problem RNBISub is infeasible, i.e. $\{q + te : t \geq 0\}$ does not intersect Y . The former case can be dealt with by the use of artificial variables to satisfy constraints (7c), see also Section 3.1. But in the latter case, many iterations of column generation may be wasted to detect the infeasibility. In fact, infeasibility of RNBISub could only be determined once all artificial variables are eliminated from the solution.

It will therefore be beneficial to identify reference points for which this is the case early to avoid attempts to solve RNBISub for such reference points. For convenience, we will from now on refer to reference points for which RNBISub is infeasible as infeasible reference points. In Section 3.1 we present a method, which we call reference point bounding, to identify infeasible reference points. To deal with infeasibility due to the restricted number of variables in RMP-RNBISub, we present three methods of initialisation in Section 3.2. Finally, we discuss the quality of the representation generated by column generation RNBI in Section 3.3.

3.1 Reference point bounding

One issue with the RNBI method, which stems from the use of the anti-ideal point in the definition of the covering simplex S and the reference subsimplex \hat{S} , is that there can be infeasible reference points, i.e. reference points for which RNBIsub is infeasible such that $\{q + te : t \geq 0\} \cap Y = \emptyset$. Because the components of y^{AI} may be far larger than the objective values of any non-dominated point, there can potentially be many reference points for which this is also the case, as shown in Fig. 1. Obviously, any effort invested in solving RNBIsub for infeasible reference points is wasted in the sense that it does not contribute to the computation of a representative set of non-dominated points. Therefore, solving RNBIsub using column generation if RNBIsub is in fact infeasible, can dramatically increase the computational time (see Section 4.2). In order to identify infeasible reference points we provide Theorem 1 characterising infeasible reference points and therefore defining the subset of feasible reference points of \hat{S} . We first state a lemma concerning the set of all feasible reference points.

Lemma 1 *The subset $\hat{Q} \subset \hat{S}$ of points q such that $\{q + te : t \geq 0\} \cap Y \neq \emptyset$ is a polytope.*

Proof The result follows, because \hat{Q} is the projection of polytope Y onto \hat{S} , which is a simplex on the hyperplane $e^T y = \mu$.

Theorem 1 *Let $q \in \hat{S}$ be a reference point. Then q is infeasible if and only if there is some $d \in \mathbb{R}^p \setminus \{0\}$ such that $d^T q < \min\{d^T z : z + te \in Y, z \in \hat{S}, t \geq 0\}$.*

Proof We first observe that the LP $\min\{d^T z : z + te \in Y, z \in \hat{S}, t \geq 0\}$ is always feasible, because $z^* = \hat{y}, t^* = 0$ with \hat{y} as defined in Section 2.2.1 is a feasible solution. It is also bounded, because Y is a compact set by assumption. Then $d^T q < \min\{d^T z : z + te \in Y, z \in \hat{S}, t \geq 0\}$ implies that q does not satisfy $q + te \in Y$ for any $t \geq 0$. Now let q be an infeasible reference point. Then $q \notin \hat{Q}$ as defined in Lemma 1. Hence there exists a hyperplane strictly separating q from \hat{Q} , i.e. there is $d \in \mathbb{R}^p \setminus \{0\}$ such that $d^T q < \min\{d^T z : z \in \hat{Q}\} = \min\{d^T z : z + ty \in Y, z \in \hat{S}, t \geq 0\}$.

Although Theorem 1 provides a theoretical characterisation of all feasible reference points, it is clearly impractical for implementation. Hence, we restrict ourselves to finding minimum and maximum values of each individual co-ordinate z_k of points on the reference subsimplex \hat{S} that are feasible reference points, i.e. we use the sufficient condition of Theorem 1 and apply it to vectors $d = e^k$ and $d = -e^k$ for $k = 1, \dots, p$, where e^k is the k th unit vector. We call this method reference point bounding.

The linear programme $\min\{d^T z : z + te \in Y, z \in \hat{S}, t \geq 0\}$ is solved for $d = e^k$ and $d = -e^k$ for $k = 1, \dots, p$. Let the optimal values be z_k^{\min} and z_k^{\max} , respectively. Then according to Theorem 1, reference points q with $q_k < z_k^{\min}$ or $q_k > z_k^{\max}$ for any $k \in \{1, \dots, p\}$ will be infeasible. Corollary 1 summarises the above argument.

Corollary 1 *If q is a reference point with $q_k < z_k^{\min}$ or $q_k > z_k^{\max}$ for some $k \in \{1, \dots, p\}$, then $\{q + te : t \geq 0\} \cap Y = \emptyset$.*




Fig. 2 Reference point bounding illustration. The red points indicate the bounding points for objective 1 and the green points are the eliminated reference points. Hence in this case, only the RNBI subproblems corresponding to the cyan reference points are solved.

Reference points that satisfy the condition of Corollary 1 are eliminated from the set Q of reference points and the corresponding RNBI subproblems are not solved. Figure 2 illustrates the bounds obtained by Corollary 1 for $k = 1$ in the same example used in Figure 1. In addition, we show the bounds obtained in the application of Section 4 in Figure 5.

3.2 Initialisation of RMP-RNBISub

Constraints (7b) may not be feasible given a limited set of variables. In addition, even after the reference point bounding procedure is applied, infeasible reference points may remain due to $\{q + te : t \geq 0\} \cap Y = \emptyset$. In this section we discuss how the infeasibility of RMP-RNBISub can be managed.

One way to handle the infeasibility is the Phase-1 approach, see e.g. Chvátal (1983), which adds non-negative artificial variables to satisfy constraints (7b) and (7c) while changing the objective function of the problem to minimise the sum of the artificial variables. The Big- M approach assigns large costs M to the artificial variables and minimises the sum of the original objective function plus the sum of the costed artificial variables. Using artificial variables, feasibility of RMP-RNBISub is assured. As soon as any of the artificial variables has a value of zero in a solution, the artificial variable can be removed. If any of the artificial variables remain positive

when the optimality condition is satisfied, we can conclude that RMP-RNBISub is infeasible because $\{q + te : t \geq 0\} \cap Y = \emptyset$.

We notice that in practice, column generation is rarely used to solve a (single objective) linear programme to optimality. In this situation, a possible approach is to perform column generation iterations on RMP-RNBISub until a specified termination condition, such as a pre-specified number of columns, is reached. One can, for example, conclude that a reference point is (approximately) feasible, if the solution satisfies constraints (7c) and the remaining total infeasibility in constraints (7b) is small enough, i.e. below a certain pre-determined threshold.

An alternative approach to manage infeasibility is to generate coefficient columns that show that the RMP is feasible (Andersen, 2001). The method is based on Farkas' lemma, which states that either $Ax = b, x \geq 0$ is feasible or there is a vector π with $\pi^T A \geq 0$ and $\pi^T b < 0$. The vector π corresponds to the dual vector of a linear programme. A linear programme is proved to be infeasible by finding a dual vector such that the condition $\pi^T Ax = \pi^T b$ can never be met due to opposite signs on the right-hand side and the left-hand side of the equation. Thus to prove that the restricted master problem is feasible, we can add a column a to A with $\pi^T a \leq 0$. Such a column can be found by solving $\min \{\pi^T a(\lambda) : \pi^T a(\lambda) \leq 0, \lambda \in X_J\}$. If no such column exists, we can conclude the corresponding master problem is infeasible. We will refer to this approach as Farkas pricing.

3.3 Quality of the representative set computed by the column generation RNBI method

Sayin (2000) defines three measures, coverage, uniformity and cardinality, to quantify the quality of a discrete representation of a set. A good representation of the non-dominated set should not contain an excessive number of points (low cardinality), should have points significantly different from one another (as indicated by high uniformity level) and should not neglect large portions of the non-dominated set (low coverage error).

Let $G \subset Y_N$ be a finite set of non-dominated points generated by the standard RNBI method using reference points $q \in Q$. Let H be the representative set generated by the RNBI method using column generation based on the same set of reference points. We shall write $g(q)$ and $h(q)$, respectively, to indicate the dependence of representative points on reference point q . The distance between two adjacent reference points is denoted as dq . Cardinality represents the number of points contained in the representation. It is clear that the number of points contained in H depends on the distance between adjacent reference points. In the rest of this section we discuss the quality of H in terms of uniformity level and coverage error.

The uniformity level δ of a representative set is measured by the distance between a pair of closest points in the set. The uniformity level of H can therefore be expressed by

$$\delta := \min_{h^i, h^j \in H, h^i \neq h^j} d(h^i, h^j) \quad (8)$$

with d being a metric. We shall use the Euclidean distance as metric in this paper. Assume h^k and h^l are the two closest points in H and let q^k and q^l be the corresponding reference points, as illustrated in Figure 3. By definition of the RNBI method, we know that vector $v^N = h^l - q^l$ must be perpendicular to vector $v^q = q^l - q^k$. Hence we have $\cos \theta = \|v^q\|/\|v^h\|$ where vector $v^h = h^l - h^k$ and θ is the angle between vectors v^h and v^q . To satisfy $h^l = q^l + t_l e$ with $t_l \geq 0$ being the absolute difference between q^l and h^l in all objectives, we must have $0 \leq \theta < \pi/2$, which corresponds to $0 < \cos \theta \leq 1$. Therefore, minimal $\|v^h\|$ occurs when $\cos \theta = 1$ and in that case the distance between h^l and h^k is $\|v^h\| = \|v^q\| = dq$. Therefore the lower bound on the uniformity level of H is dq , which is the same as that of G (Shao and Ehrgott, 2007).




Fig. 3 An illustration of the uniformity level for a representation produced by the RNBI method using column generation. The two diamonds represent a pair of closest representative points h^l and h^k and the circles represent the corresponding reference points q^l and q^k .

The coverage error ε indicates how accurately set H represents Y_N and can be expressed as

$$\varepsilon := \max_{y \in Y_N} \min_{h \in H} d(y, h). \quad (9)$$

Essentially, the coverage error ε is the maximum distance between a point in the non-dominated set and its closest point in the representation H . Notice that if RMP-RNBISub is not solved to optimality by column generation, $h \in H$ can be an intersection point of $\{q + te : t \geq 0\}$ with Y that is dominated even though $\{q + te : t \geq 0\}$ intersects Y in a non-dominated point.

Shao and Ehrgott (2007) show that the coverage error of G is at most $(\sqrt{pd}q)/2$. Hence the coverage error of H is bounded by the maximum distance between points $g(q)$ and $h(q)$ of H and G generated for reference points $q \in Q$ plus the coverage error of G which can be expressed as

$$\varepsilon \leq \max_{q \in Q} d(g(q), h(q)) + \sqrt{pd}q/2. \quad (10)$$

The term $d(g(q), h(q))$ can be derived from the difference between the objective function values of RNBISub and RMP-RNBISub for reference point q . If RNBISub is not solved to optimality, one can use a lower bound on the optimal value of RNBISub, e.g. equation (2), to estimate the coverage error.

Based on the above discussion, we can see that the quality of a representation generated by column generation RNBI depends on the distance between adjacent reference points. As the distance decreases, cardinality increases, the uniformity level decreases and the coverage error decreases. In addition, the coverage error also depends on the maximum distance between representative points $g(q)$ and $h(q)$ for reference points $q \in Q$, which depends on the termination condition of the column generation process. Consequently, given a problem at hand, one should select a dq and a column generation termination condition that results in appropriate uniformity level and coverage level for the representative non-dominated set.

4 Application of the column generation RNBI method in radiotherapy treatment design

In this section, we consider the so called fluence map optimisation problem of radiotherapy treatment design and describe how the column generation RNBI method can be applied to find a set of fluence maps that are deliverable without solving an additional segmentation problem and that define a representation of the entire set of non-dominated points of the multi-objective fluence map optimisation problem. The design of a radiotherapy treatment for cancer using optimisation methods has become an important application of optimisation with the introduction of intensity modulated radiotherapy treatment. It involves the determination of beam angles, beam intensities and a delivery schedule for the radiation using a gantry equipped with a linear accelerator. Its goal is to deliver a high and uniform radiation dose to the treatment planning target volume (PTV) while sparing surrounding healthy organs at risk (OARs) as much as possible. In the delivery of radiotherapy treatments, radiation fields pass through a device called multileaf collimator (MLC), which consists of a number of pairs of metal leaves that can move into and out of the path of the radiation independently (as illustrated in Figure 4). The MLC leaves block part of the radiation field which results in a radiation field of an irregular shape. An MLC leaf opening that is applied in the delivery of radiation is referred to as an aperture or a segment. For any beam direction, the application of multiple segments, each applied for a certain time, allows the delivery of modulated radiation intensity that results in the desired dose distribution. For further reading on radiotherapy treatment design, we refer readers to specialised textbooks e.g. [Webb \(2001\)](#); [Schlegel and Mahr \(2002\)](#).

From now on, we assume that beam directions are given and refer to [Ehrgott et al \(2008c\)](#) for an overview of the problem of determining beam directions. Hence, given a set of beam directions, we are interested in finding a design, consisting of a set of segments and the associated radiation intensities, which best benefits the patient. Conventionally, the radiotherapy treatment design problem is split into two sequential optimisation problems, the fluence map optimisation (FMO) problem and the segmentation problem. FMO is the problem of finding the optimal modulated radiation intensity for each beam direction. To model the FMO problem mathematically, the radiation field at a beam direction is discretised into small sized rectangular sub-fields called bixels. These correspond to the smallest openings of the MLC, i.e. are of the width of one of the leaves of the MLC and of the length that corresponds to the




Fig. 4 Illustration of an MLC device.

distance between two stop positions of the leaf. FMO finds the radiation intensity for each bixel such that a desirable dose distribution that meets the goals of the treatment can be delivered to the patient. The intensities for the bixels are referred to as the intensity pattern. In principle, one can deliver the intensity pattern bixel-by-bixel using bixel-sized segments. However, doing so would lead to an unrealistically long treatment time. In practice, the intensity pattern is realised by stacking a limited number of shaped radiation fields, each passing through an associated segment. Therefore, after obtaining the intensity pattern as output from solving the FMO problem, it is necessary to solve a so-called segmentation problem, which finds a set of segments that best realise the intensity pattern by, for instance, minimising the total beam-on time required to deliver the intensity pattern or by minimising the required number of segments (see Baatar et al (2005)).

The segmentation problem is an optimisation problem that needs to incorporate physical constraints of the MLC leaves. The elementary ones are collision constraints, that prevent opposing leaves to overlap and constraints that ensure the opening in any MLC row is continuous, i.e. all open bixels in a row are consecutive. Other constraints are specific to particular brands of MLCs, which is why we concentrate on the basic ones in this study. To avoid generating overly complex treatment plans that cannot be practically delivered, the bixel intensities are discretised into a range of intensity levels at the beginning of the segmentation process. As a result, the intensity pattern is realised approximately and the quality of the treatment plan deteriorates after segmentation. A survey of the literature on segmentation problems can be found in Ehrgott et al (2008b).

FMO needs to deal with several conflicting objectives associated with the PTV and the surrounding structures. The conflicting objectives in FMO have conventionally been handled by scalarisation (see Ehrgott et al (2008a) for a review). However, using this approach, if the generated intensity pattern is not satisfactory, the planner will have to iteratively adjust the plan optimisation parameters and re-optimize

until a satisfactory intensity pattern is found. This process is time consuming and without guarantee of finding the best possible intensity pattern under the patient specific conditions. Instead, multi-objective optimisation has been introduced to solve the FMO problem. By generating a representative set of non-dominated plans, the planner can browse the plans and choose the best one available without the iterative process. Several approaches have been proposed to solve multi-objective FMO problems, including goal programming methods (Falkinger et al, 2012; Breedveld et al, 2009; Wilkens et al, 2007; Jee et al, 2007), constraint methods (Hoffmann et al, 2006; Craft et al, 2005; Küfer et al, 2003; Hamacher and Küfer, 2002) and approximation methods (Shao and Ehrgott, 2008; Craft et al, 2006). In addition, RNBI has also been applied to multi-objective FMO problems by Shao and Ehrgott (2007).

Recently, a multi-objective FMO optimisation approach has been deployed in clinical practice (Craft and Richter, 2013). The approach approximates the non-dominated set using convex combinations of efficient solutions. However, since FMO does not consider plan delivery, the treatment plan generated from FMO needs to go through the segmentation process, which transforms an optimal intensity pattern into a limited number of segment intensities, thereby deteriorating plan quality, as demonstrated by Rocha et al (2012). If the deliverable plan is not satisfactory, the planner will have to re-optimize and find another plan. To avoid this drawback, Craft and Richter (2013), Salari and Unkelbach (2013) and Fredriksson and Bokrantz (2013) have proposed multi-objective approaches to find deliverable plans. These approaches use convex combinations of the segmented plans or conical combinations of the segments to approximate the feasible set of the FMO problem and then use multi-objective interactive optimisation methods to navigate among the non-dominated set of the approximated feasible set.

The clinically adopted MOO method uses a sandwiching method (see, e.g., Renen et al (2011) and Bokrantz and Forsgren (2013)) to generate an approximation of the non-dominated set, followed by plan navigation on the convex hull of a set of plans (Monz et al, 2008). While approximating the non-dominated set based on interpolation of a set of existing plans can reduce the computational expense, compared to generating a discrete representative non-dominated set, we note that the interpolated solution may be subject to further improvement potential due to approximation error (see, e.g., Bokrantz and Miettinen (2015)). In contrast, the RNBI procedure produces a discrete set of efficient plans that captures all potential treatment trade-offs. Moreover, it is guaranteed that each non-dominated point is no further than a known coverage error from the objective vector of one of the computed plans. Since it is also guaranteed that the objective vectors of the computed plans cover the entire non-dominated set (in terms of a guaranteed minimum distance, called the uniformity level, between any two of them) there is no need to consider convex combinations of plans. As interpolation is not used to form plans, each plan is given freedom in beam angle configuration, segment shapes and segment intensities and hence allows one to achieve the best-quality plans for different treatment trade-offs. In addition, since a set of discrete plans that captures different treatment trade-offs are generated, navigation can be conducted by examining the existing set of plans. Consequently, one can extract relevant (non-convex) clinical evaluation criteria, e.g., the dose-volume parameters and treatment delivery time, from the plans and use these criteria to find

the most preferable plan from the representative non-dominated set (Lin and Ehrgott, 2016).

Column generation has been used to generate deliverable plans for single objective radiotherapy plan optimisation (Preciado-Walters et al, 2004; Romeijn et al, 2005). Here, the physical delivery constraints in the segmentation process are considered in the column generation subproblem. Essentially, each column generated from the subproblem represents a segment that is likely to improve the objective function value. As a result, the solutions produced from column generation can be delivered without additional segmentation. As will be demonstrated by our results, column generation produces plans that are near-optimal and can be delivered with dramatically lower total monitor units than the corresponding optimal ones followed by segmentation. Such plans are desirable as they can be delivered with a shorter treatment time, i.e., exposing patients to radiation for a shorter time, and with lower radiation leakage from the MLCs (Broderick et al, 2009). In fact, near-optimal plans that can be delivered efficiently and accurately are often preferable to complex optimal solutions in practice (Carlsson and Forsgren, 2014). Earlier studies on deliverable multi-objective optimisation in radiotherapy limit the number of segments requiring that all computed plans use either the same segments (Salari and Unkelbach, 2013), that the number of segments in each plan is limited (Craft and Richter, 2013), or that the plans use both a subset of segments from a common pool and some individual ones (Fredriksson and Bokrantz, 2013). We use column generation within the RNBI method to control the number of segments that are generated for each plan, making no restrictions on the set from which these segments are drawn. Because the method provides quality guarantees for the computed plans, the planner is then able to decide how many segments to allow solely based on plan quality without setting any a-priori limits.

In this paper, we apply column generation within the RNBI framework for multi-objective radiotherapy treatment design. This approach produces a representative set of plans that are deliverable and are close to efficient plans that do not consider the deliverability of intensity patterns. The approach therefore combines the advantage of considering multiple objectives in the FMO problem with the advantage of producing deliverable intensity maps without much deterioration of treatment quality, which column generation delivers.

4.1 Formulation

To formulate a mathematical model of the FMO problem, the patient volume is discretised into m small volume elements called voxels. The radiation fields for all beams are discretised into n small rectangular fields called bixels, as explained above. The dose for each voxel is calculated by

$$d = Ax, \quad (11)$$

where $d \in \mathbb{R}^m$ is the dose vector in which d_i represents the dose delivered to voxel i . Vector $x \in \mathbb{R}^n$ is the radiation intensity vector in which x_j describes the radiation intensity for bixel j . $A \in \mathbb{R}^{m \times n}$ is the dose deposition matrix. Element a_{ij} of A represents the dose deposited to voxel i from bixel j under unit intensity. A is specific

to the radiation source (modality, energy) used for the treatment and the patient volume. Different dose calculation algorithms can be used to calculate A (Reynaert et al, 2007; Jeleń et al, 2005; Keall and Hoban, 1996). For convenience, A can be partitioned and re-ordered into submatrices 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 PTV T with m_T voxels, for critical organs C with m_C voxels and for normal tissue N with m_N voxels, respectively.

In the treatment planning process, the oncologist determines a prescription dose to be delivered to the PTV, respectively not to be exceeded for organs at risk and normal tissue. The planner attempts to achieve the prescription dose by setting the appropriate optimisation parameters. The formulation used in this study is based on the model of Holder (2003), as shown in (12), but with slight variation. The parameters include the dose lower bound $LB_T \in \mathbb{R}^{m_T}$ for the tumour and upper bounds $UB_T \in \mathbb{R}^{m_T}$, $UB_C \in \mathbb{R}^{m_C}$ and $UB_N \in \mathbb{R}^{m_N}$ for the tumour, critical organs and normal tissue, respectively. Variables $\alpha \in \mathbb{R}^{m_T}$, $\beta \in \mathbb{R}^{m_C}$ and $\gamma \in \mathbb{R}^{m_N}$ are voxel-wise one-sided dose deviations from tumour lower bound, critical organ upper bound and normal tissue upper bound, respectively.

$$\begin{aligned} \min \left(\frac{1}{m_T} e^T \alpha, \frac{1}{m_C} e^T \beta, \frac{1}{m_N} e^T \gamma \right) \\ \text{s.t. } LB_T - \alpha \leq A_T x \leq UB_T, \\ A_C x \leq UB_C + \beta, \\ A_N x \leq UB_N + \gamma, \\ \alpha \leq UB_\alpha, \\ \beta \leq UB_\beta, \\ \gamma \leq UB_\gamma, \\ 0 \leq x, \alpha, \beta, \gamma. \end{aligned} \quad (12)$$

Different from Holder's model, we introduce upper bounds $UB_\alpha \in \mathbb{R}^{m_T}$, $UB_\beta \in \mathbb{R}^{m_C}$ and $UB_\gamma \in \mathbb{R}^{m_N}$ for α , β and γ , respectively. These upper bounds can easily be set so that the Y is bounded and thus allows us to compute y^{AI} . Note that the unit of the objective functions is Gray (unit for radiation dose), since we are trying to minimise dose deviations. The RNBI subproblem of (12) is simply (7) with constraints (7c) and (7d) replaced by the constraints of (12) and with the objective functions of (12) incorporated in the form of constraint (7b).

An optimal fluence map obtained by solving the RNBI subproblem of (12) with variables representing bixel intensities may not be practically deliverable. Alternatively, deliverable plans can be generated from a reformulation in which the bixel intensity variables x are replaced by segment intensity variables \bar{x} and the dose deposition matrix A based on bixel columns is replaced by the dose deposition matrix \bar{A} using segment columns. Column s of \bar{A} , denoted as $\bar{a}^s \in \mathbb{R}^m$, represents the dose deposited to the m voxels by segment s at unit intensity. Column \bar{a}^s is derived by

$$\bar{a}^s = A u^s, \quad (13)$$

where $u^s \in \{0, 1\}^n$ is a vector defining segment s with $u_j^s = 1$ if bixel j is open and $u_j^s = 0$ if bixel j is closed in segment s . Note that the two formulations have the same

optimal values since feasible solutions in terms of variables x can be represented by variables \bar{x} and vice versa, through the relationship

$$x = U\bar{x}, \quad (14)$$

with U being a matrix containing all feasible segment columns u .

Due to the large number of feasible segments, \bar{A} has a much larger number of columns than A , which makes the reformulation hard to solve. Therefore it is beneficial to use column generation to solve the reformulation in which we only consider a subset of segments in the RMP. By solving the RMP, we obtain a vector of dual values π^* , which is passed to the subproblem to find a nonbasic variable (representing the radiation intensity for a segment) with the most negative reduced cost. Here the subproblem is

$$\min \{-\pi^{*T} Au : u \in U\}, \quad (15)$$

where, as before, $u \in \{0, 1\}^n$ is a vector defining a segment and U , in slight abuse of notation introduced above, is the set of all feasible segment columns satisfying the MLC constraints. Note that the objective function coefficients of the segment intensity variables, with model (12) reformulated to the form of RNBISub, are zero, thus they are not included in subproblem (15). Let u^* be an optimal solution of (15). Given u^* , we can derive the dose deposition column $a^* = Au^*$ to be added to the RMP-RNBISub reformulation of (12).

Since we consider only the leaf collision constraint and the constraint that the opening for each row of collimator leaves must be contiguous, all MLC rows are independent of one another. Therefore, (15) can be further decomposed by MLC row. For a given row, the objective of the decomposed problem is to find the leaf positions that result in the lowest reduced cost for the MLC row. To decompose (15) by MLC row, one needs to change the index of each bixel to its corresponding beam index, MLC row position and MLC column position. Let $\tau = -(\pi^*)^T A$ and denote by $\tau_{b,r,c}$ the objective function coefficient of (15) corresponding to the bixel at beam b , MLC row position r and MLC column position c . We assume that each MLC row consists of t bixels, with MLC column positions indexed incrementally from left to right. Denote by t_1 the column index of the right-most bixel covered by the left leaf and by t_2 the column index of the left-most bixel covered by the right leaf. Then the decomposed subproblem (15) for beam b and MLC row r is

$$\begin{aligned} \min \quad & \sum_{c=t_1+1}^{t_2-1} \tau_{b,r,c} \\ \text{s.t.} \quad & t_1 \leq t_2 - 1, \\ & t_1 \in \{0, \dots, t\}, \\ & t_2 \in \{1, \dots, t+1\}. \end{aligned} \quad (16)$$

This decomposed problem is then solved by the algorithm described in Section 3.1.1 of Romeijn et al (2005). It is important to note, that the non-dominated set of the original formulation of (12) and its column generation reformulation are identical.

4.2 The test case

We apply both the original RNBI method and the column generation RNBI method to a prostate radiotherapy treatment design problem. The RNBI method solves the RNBISub reformulation of (12) with bixel intensity variables x , producing a set of (not necessarily deliverable) intensity patterns that define a representative set of non-dominated points for MOLP (12). The column generation RNBI method solves the RMP-RNBISub reformulation of (12) with a subset of segment intensity variables \bar{x} . We consider three objective functions: one for the PTV (objective 1), one for the rectum (objective 2) and one for the bladder (objective 3). Other clinically relevant structures such as the prostate, the right and left femoral head and normal tissues, are involved in the formulation as constraints, e.g. voxels of the prostate are given a lower bound and an upper bound on the delivered dose and voxels of femoral heads and normal tissues are given structure specific upper bounds. By only involving three objective functions, we are able to illustrate the results graphically.

The dose deposition matrix A consists of 593 columns (corresponding to bixels) for 11 equi-spaced coplanar beam angles and 20000 rows (corresponding to voxels). Both methods use the set of reference points generated by the standard RNBI method. Therefore, we are able to identify feasible reference points, and we apply the column generation RNBI method only to feasible reference points. The column generation process terminates when any one of the following termination conditions is satisfied:

- no variable with a negative reduced cost can be found
- the number of segments assigned with a positive intensity in a solution exceeds 100
- the number of column generation iterations (or equivalently the number of segments) exceeds 150 .

We implement the Phase-1 approach, the Big- M approach and Farkas pricing to handle infeasibility of RMP-RNBISub. The initialisation phase stops when the feasibility of RMP-RNBISub without artificial variables is guaranteed or when the termination condition is reached. The column generation model starts with only one coefficient column representing fully closed MLC segments (i.e. $u = 0$) in all beam directions. If a segment is assigned a positive intensity, we refer to this segment as a “positive segment”. Positive segments are those segments that will be delivered in a treatment plan. This is in contrast to the zero-intensity segments, which will not be part of the treatment plan. We limit the number of positive segments in each plan so that the plans can be delivered in a reasonable treatment time. Radiotherapy treatment plan quality depends on the number of segments involved in a treatment, therefore only solutions of similar number of positive segments should be compared. A solution is separately recorded when the number of positive segments in the solutions first exceeds 40, 50, 60, 70, 80, 90 and 100. These solutions are grouped into representative sets according to the number of positive segments.

4.3 Results

For convenience, representative points generated using RNBI and column generation RNBI will be referred to as RNBI points and the CG-RNBI points, respectively. The representative sets of the CG-RNBI points, grouped according to the number of positive segments, will be denoted as *CG-number* with *number* being the corresponding number of positive segments.

Using the standard RNBI method with increment $\eta = 0.08$ (see Section 2.2.1) or a distance of 3.2153 Gray between closest reference points, we identify that 17 of 91 reference points are feasible. Figure 5 illustrates the reference points and the RNBI points. We then initialise the column generation RNBI subproblems with the Phase-1 approach, the *Big-M* approach or Farkas pricing, followed by RMP-RNBISub when feasibility of RMP-RNBISub is guaranteed.




Fig. 5 Illustration of the reference points and standard RNBI intersection points. The blue lines indicate the bounding for reference points that lead to infeasible RNBISub.

The colour of the intersection points indicates the value of bladder deviation.

Figure 6 shows RMP-RNBISub objective function values versus column generation iteration for the first 4 of the 17 reference points after the initialisation stage. In each column generation iteration, one newly generated column, which represents a segment, is added to RMP-RNBISub. The red dashed line, blue solid line and green dash-dot line indicate the objective function values of the corresponding RMP-RNBISub problem initialised with the *Big-M* approach, the Phase-1 approach and Farkas pricing, respectively. The starting point of these lines indicates the number of iterations used to reach RMP-RNBISub feasibility using the three approaches. The dark line parallel to the horizontal axis represents the objective function value obtained using the standard RNBI method for the same reference point. The results show that, for all 17 cases, an initialisation using the *Big-M* approach reaches feasibility with the same or a smaller number of iterations compared to an initialisation with the




Fig. 6 RMP-RNBISub objective function values (vertical axis) versus column generation iteration (horizontal axis) for the first four reference points.

other two approaches. In addition, during the early stages of the column generation process, initialisation with the Big- M approach generally produces a lower objective function value compared to initialisations with the other two approaches. The results indicate that the Big- M approach is superior to the other two approaches in terms of identifying columns that contribute to the objective function value. However, the difference in objective function values among different initialisation strategy diminishes as more columns are added to the model.

The next results are based on the solutions obtained from runs that employ an initialisation with the Big- M approach. Figure 7 shows the RNBI points (solid circle) and the points of CG-40 (asterisk) and CG-100 (empty circle). The figure illustrates how intersection points gradually move toward the boundary of the feasible set in objective space during the column generation process. Notice that in each of the column generation representative sets, there can be points dominated by other points. Our results show that the number of dominated points in a representative set ranges from one point in CG-60 to 5 points in CG-90.

The objective function values and the average computation time of the RNBI points, CG-40 and CG-100 are shown in Table 1. The objective function values of the points in CG-40 and CG-100 are on average 0.3647 and 0.1264 Gray higher than the objective function values of the RNBI points. The average computation time used to obtain the RNBI points, CG-40 and CG-100 are approximately 17, 43 and 553 seconds, respectively. We observe that, as the number of generated columns increases, the computation effort for solving RMP-RNBISub increases as well.

We apply the segmentation algorithm by Engel (2005) to the optimal intensity patterns generated by the standard RNBI method. The results show that the intensity patterns (after discretisation by rounding to integers) can be reproduced with an average of 91.1 segments. In addition, as shown in Table 2, the total monitor units for delivering the segmented patterns (an average of 408.8 MUs) is much higher than that

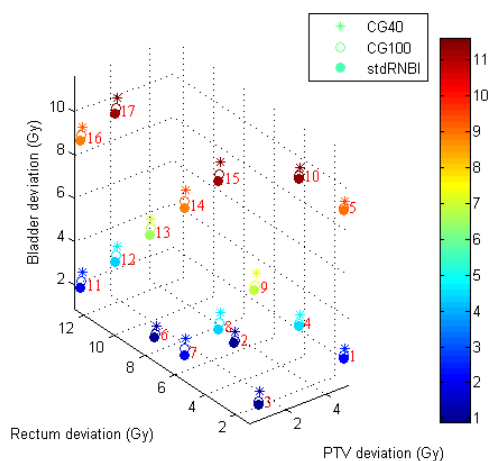


Fig. 7 The RNBI points (solid circle) and the points in CG-40 (asterisk) and CG-100 (empty circle). The colour indicates the value of bladder deviation.

of the segment patterns generated by column generation (an average of 76.7 MUs and 123.8 MUs for CG-40 and CG-100, respectively).

Table 1 Objective values and average computation time (rounded to seconds) of RNBIsub and CG-RNBIsub with 40 positive segments (CG-40) and 100 positive segments (CG-100).

Reference point	Standard	CG-40	CG-100
1	1.7450	1.9880	1.8140
2	2.6806	2.9661	2.7525
3	0.4216	0.7765	0.5019
4	1.7451	2.0407	1.8215
5	4.0186	4.2616	4.0875
6	2.6941	2.9738	2.7695
7	1.5498	1.9892	1.7157
8	1.5417	1.9638	1.7104
9	1.8497	2.2986	1.9810
10	4.0186	4.2959	4.0916
11	3.8230	4.2212	3.9890
12	3.8108	4.1933	3.9770
13	3.8108	4.2176	3.9943
14	3.8112	4.2856	3.9865
15	3.8153	4.3023	3.9846
16	6.0843	6.4450	6.2286
17	6.0843	6.4846	6.2481
Average time	17	43	553

We use (8) to measure the uniformity of the representative sets. The results show that the uniformity levels for all representative non-dominated sets are the same up to 4 decimal places, with a value of 3.2153 Gray, which is the same as the distance be-

Table 2 Total monitor units for delivering the intensity patterns generated by RNBI and CG-RNBI with 40 positive segments (CG-40) and 100 positive segments (CG-100).

Reference point	Standard	CG-40	CG-100
1	494	97.8	155.4
2	474	82.2	146.6
3	444	79.1	138.0
4	493	97.2	152.3
5	494	97.8	155.4
6	469	72.9	136.8
7	341	66.4	119.7
8	354	70.0	116.3
9	446	79.7	135.6
10	493	87.1	144.8
11	343	65.8	103.6
12	350	67.2	94.6
13	350	60.8	106.5
14	351	68.9	95.9
15	354	77.9	108.0
16	350	65.7	98.8
17	350	66.7	96.3
Average MUs	408.8	76.7	123.8

Table 3 Minimum and maximum value for each objective based on the reference point bounding solved to optimality and solved by column generation (CG).

	Optimal Minimum	Optimal Maximum	CG Minimum	CG Maximum
Objective 1	-7.5352	3.9529	-7.2641	3.5035
Objective 2	-4.4038	9.159	-4.0101	8.9119
Objective 3	-4.5379	9.1226	-4.2902	8.8587

tween any two closest reference points. However, the two closest intersection points that define the uniformity level are different for the different representative sets.

Next, we apply the reference point bounding method described in Section 3.2 to the column generation RNBI method. Column generation is typically used for problems that cannot be practically solved by standard linear optimisation algorithms due to a large number of variables. When applying column generation RNBI to these problems, one also needs to solve the reference point bounding problem using column generation. Therefore, to assess how reference point bounding can be affected by column generation, reference point bounding is firstly solved to optimality (with the bixel intensity based formulation) and then solved using column generation (with the segment intensity based formulation). The column generation process terminates when the number of positive segments exceeds 40. The objective function values for reference point bounding are shown in Table 3. We can see that the minima and maxima produced using column generation are very close to the corresponding minima and maxima solved to optimality, with a maximum absolute difference of 0.4494. In fact, using either set of minima and maxima, we are able to eliminate 67 out of 91 (73.63%) reference points. Out of the remaining 24 reference points, only 7 lead to RNBIsub infeasibility.

Table 4 Number of iterations required for Farkas pricing to identify RNBISub infeasibility.

Iterations used to identify infeasibility	2	3	4	5	10	>150
Number of RNBISub instances	48	10	1	6	1	8

We also test the performance of Farkas pricing in concluding the infeasibility of RNBISub instances. Note that we have 91 reference points in total, with 74 reference points leading to RNBISub infeasibility. Table 4 shows that Farkas pricing is capable of concluding the infeasibility of 66 out of 74 RNBISub instances using 10 or fewer iterations. The average computation time for these 66 instances is 0.4 seconds. However, Farkas pricing is incapable of concluding infeasibility within 150 iterations for the remaining 8 instances. The computation time for each of these 8 instances ranges from 1179 to 7084 seconds, with an average of 4396 seconds and a standard deviation of 1676 seconds. For comparison, we apply column generation with the big-M initialisation to 10 reference points leading to RNBISub infeasibility. With a termination condition of 150 column generation iterations, the average computation time for solving each of the 10 reference points is 1213 seconds, with a standard deviation of 71 seconds. The results suggest that Farkas pricing can potentially be quite time consuming. Thus, if Farkas pricing cannot identify the infeasibility of a RNBISub instance in a small number of iterations, it would be beneficial to change the initialisation method to another approach.

5 Discussion and conclusion

In this paper we propose the use of column generation in the revised normal boundary intersection method to compute a finite representative subset of the non-dominated set of a multi-objective linear programme. We introduced a reference point bounding procedure to eliminate the investigation of infeasible reference points. In terms of the quality of a discrete representation computed with the column generation RNBI method we showed that the uniformity level is at least dq , the distance between closest reference points, and therefore the same as that of standard RNBI representative set. The coverage error is bounded by the the distance of CG-RNBI points to the non-dominated set plus $\sqrt{p}dq/2$. This feature allows one to choose a value of dq to produce a representative set that suits decision making in the application considered.

To illustrate the method and demonstrate the advantages of using column generation to solve the RNBI subproblems, we apply the method to an MOLP formulation of a radiotherapy treatment design problem, which can be solved by both the standard RNBI and the column generation RNBI method. In agreement with the remark in Lübbbecke (2010) that column generation is in general not a competitive technique in solving linear programmes, we observe that computation times using column generation are longer. However, column generation allows us to use variables representing segment intensities, as opposed to bixel intensities which are used in the conventional formulation. As a consequence, the number of variables involved in the model is much greater than for a model formulated with bixel intensities, since the number

of possible segment shapes, formed by all possible combinations of opening bixels, is much greater than the number of bixels. On the other hand, the column generation formulation avoids the segmentation step which deteriorates treatment quality. Our results show that plans generated by column generation are near-optimal and can be delivered with dramatically lower monitor units than the corresponding optimal ones followed by segmentation. This reduced delivery complexity is desirable in practice due to shorter treatment time and lower radiation leakage from the MLCs (hence better delivery accuracy) (Broderick et al, 2009; Carlsson and Forsgren, 2014).

Fredriksson and Bokrantz (2013) introduce a concept of non-dominance called the “ n -aperture Pareto set”, which is a set of efficient plans given that each plan is formed by only n segments. However, to our knowledge, there is no practical method available to generate the n -aperture Pareto set. The concept of the n -aperture Pareto set can be generalised to the n -column non-dominated set for problems solved by column generation. Further research is required to extend the column generation RNBI method to ensure n -column non-dominance. Another topic for future research is the extension of the column generation RNBI method to nonlinear multi-objective optimisation problems. This will, e.g., allow us to consider other formulations of the radiotherapy treatment design problem.

Data statement Due to the sensitive nature of the data, we are not in a position to share data as we do not have permission to do so.

Conflict of interest The authors declare that they have no conflict of interest.

References

- Andersen E (2001) Certificates of primal or dual infeasibility in linear programming. *Comput Optim Appl* 20(2):171–183
- Baatar D, Hamacher H, Ehrgott M, Woeginger G (2005) Decomposition of integer matrices and multileaf collimator sequencing. *Discret Appl Math* 152(1-3):6–34
- Benson HP, Sayin S (1997) Towards finding global representations of the efficient set in multiple objective mathematical programming. *Nav Res Log* 44(1):47–67
- Bokrantz R, Forsgren A (2013) An algorithm for approximating convex pareto surfaces based on dual techniques. *INFORMS J Comput* 25(2):377–393
- Bokrantz R, Miettinen K (2015) Projections onto the Pareto surface in multicriteria radiation therapy optimization. *Med Phys* 42(10):5862–5870
- Breedveld S, Storchi PR, Heijmen BJ (2009) The equivalence of multi-criteria methods for radiotherapy plan optimization. *Phys Med Biol* 54(23):7199–7209
- Broderick M, Leech M, Coffey M (2009) Direct aperture optimization as a means of reducing the complexity of intensity modulated radiation therapy plans. *Radiat Oncol* 4(8):1–7
- Carlsson F, Forsgren A (2014) On column generation approaches for approximate solutions of quadratic programs in intensity-modulated radiation therapy. *Ann Oper Res* 223(1):471–481

- Chvátal V (1983) *Linear Programming*. W.H. Freeman and Company, New York
- Craft D, Richter C (2013) Deliverable navigation for multicriteria step and shoot IMRT treatment planning. *Phys Med Biol* 58(1):87–103
- Craft D, Halabi T, Bortfeld T (2005) Exploration of tradeoffs in intensity-modulated radiotherapy. *Phys Med Biol* 50(24):5857–5868
- Craft DL, Halabi TF, Shih HA, Bortfeld TR (2006) Approximating convex Pareto surfaces in multiobjective radiotherapy planning. *Med Phys* 33(9):3399–3407
- Das I, Dennis JE (1998) Normal-boundary intersection: A new method for generating the Pareto surface in nonlinear multicriteria optimization problems. *SIAM J Optimiz* 8(3):631–657
- Ehrgott M (2005) *Multicriteria Optimization*, 2nd edn. Springer-Verlag, Berlin Heidelberg
- Ehrgott M, Güler Ç, Hamacher HW, Shao L (2008a) Mathematical optimization in intensity modulated radiation therapy. *4OR-Q J Oper Res* 6(3):199–262
- Ehrgott M, Hamacher H, Nußbaum M (2008b) Decomposition of matrices and static multileaf collimators: A survey. In: Alves C, Pardalos P, Vicente L (eds) *Optimization in Medicine*, Springer Verlag, Berlin, Springer Optimization and its Applications, vol 12, pp 27–48
- Ehrgott M, Holder A, Reese J (2008c) Beam selection in radiotherapy design. *Linear Algebra Appl* 428(5-6):1272–1312
- Engel K (2005) A new algorithm for optimal multileaf collimator field segmentation. *Discrete Appl Math* 152(1):35–51
- Falkinger M, Schell S, Müller J, Wilkens JJ (2012) Prioritized optimization in intensity modulated proton therapy. *Z Med Phys* 22(1):21–28
- Faulkenberg SL, Wiecek MM (2010) On the quality of discrete representations in multiple objective programming. *Optim Eng* 11(3):423–440
- Fredriksson A, Bokrantz R (2013) Deliverable navigation for multicriteria IMRT treatment planning by combining shared and individual apertures. *Phys Med Biol* 58(21):7683–7697
- Hamacher HW, Küfer KH (2002) Inverse radiation therapy planning - a multiple objective optimization approach. *Discrete Appl Math* 118(1-2):145–161
- Hoffmann AL, Siem AYD, Den Hertog D, Kaanders JHAM, Huizenga H (2006) Derivative-free generation and interpolation of convex Pareto optimal IMRT plans. *Phys Med Biol* 51(24):6349–6369
- Holder A (2003) Designing radiotherapy plans with elastic constraints and interior point methods. *Health Care Management Science* 6(1):5–16
- Jee KW, McShan DL, Fraass BA (2007) Lexicographic ordering: Intuitive multicriteria optimization for IMRT. *Phys Med Biol* 52(7):1845–1861
- Jeleń U, Söhn M, Alber M (2005) A finite size pencil beam for IMRT dose optimization. *Phys Med Biol* 50(8):1747–1766
- Keall PJ, Hoban PW (1996) Superposition dose calculation incorporating Monte Carlo generated electron track kernels. *Med Phys* 23(4):479–485
- Küfer KH, Scherrer A, Monz M, Alonso F, Trinkaus H, Bortfeld T, Thieke C (2003) Intensity-modulated radiotherapy - a large scale multi-criteria programming problem. *OR Spectrum* 25(2):223–249

- Lin KM, Ehrgott M (2016) Multiobjective navigation of external radiotherapy plans based on clinical criteria. Tech. rep., Lancaster University
- Lübbecke ME (2010) Column generation. In: Cochran JJ, Cox LA, Keskinocak P, Kharoufeh JP, Smith JC (eds) Wiley Encyclopedia of Operations Research and Management Science, John Wiley & Sons, Inc., pp 1–14
- Messac A, Ismail-Yahaya A, Mattson CA (2003) The normalized normal constraint method for generating the Pareto frontier. *Struct Multidiscip O* 25(2):86–98
- Monz M, Küfer K, Bortfeld T, Thieke C (2008) Pareto navigation—algorithmic foundation of interactive multi-criteria IMRT planning. *Phys Med Biol* 53(4):985–998
- Moradi S, Raith A, Ehrgott M (2015) A bi-objective column generation algorithm for the multi-commodity minimum cost flow problem. *Eur J Oper Res* 244(2):369–378
- Preciado-Walters F, Rardin R, Langer M, Thai V (2004) A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer. *Math Program* 101(2):319–338
- Rennen G, van Dam ER, Hertog Dd (2011) Enhancement of sandwich algorithms for approximating higher-dimensional convex Pareto sets. *INFORMS J Comput* 23(4):493–517
- Reynaert N, van der Marck SC, Schaart DR, Van der Zee W, Van Vliet-Vroegindewij C, Tomsej M, Jansen J, Heijmen B, Coghe M, De Wagter C (2007) Monte Carlo treatment planning for photon and electron beams. *Radiat Phys Chem* 76(4):643–686
- Rocha H, Dias JM, Ferreira BC, Lopes MC (2012) Discretization of optimal beamlet intensities in IMRT: A binary integer programming approach. *Math Comput Model* 55(7-8):1969–1980
- Romeijn HE, Ahuja RK, Dempsey JF, Kumar A (2005) A column generation approach to radiation therapy treatment planning using aperture modulation. *SIAM J Optimiz* 15(3):838–862
- Salari E, Unkelbach J (2013) A column-generation-based method for multi-criteria direct aperture optimization. *Phys Med Biol* 58(3):621–639
- Sayin S (2000) Measuring the quality of discrete representations of efficient sets in multiple objective mathematical programming. *Math Program* 87(3):543–560
- Schlegel W, Mahr A (2002) 3D-Conformal Radiation Therapy: A Multimedia Introduction to Methods and Techniques. Springer Verlag, Berlin
- Shao L, Ehrgott M (2007) Finding representative nondominated points in multiobjective linear programming. In: 1st IEEE Symp. on Comput. Intell. in Multicriteria Decision Making, MCDM 2007, pp 245–252
- Shao L, Ehrgott M (2008) Approximately solving multiobjective linear programmes in objective space and an application in radiotherapy treatment planning. *Math Method Oper Res* 68(2):257–276
- Shao L, Ehrgott M (2016) Discrete representation of non-dominated sets in multi-objective linear programming. *Eur J Oper Res* 255(3):687–698
- Webb S (2001) Intensity-modulated radiation therapy. Series in Med Phys, Institute of Physics Publishing, Bristol and Philadelphia
- Wilkens JJ, Alaly JR, Zakarian K, Thorstad WL, Deasy JO (2007) IMRT treatment planning based on prioritizing prescription goals. *Phys Med Biol* 52(6):1675–1692