

A Matheuristic Approach to Solve the Multi-objective Beam Angle Optimisation Problem in Intensity Modulated Radiation Therapy

Guillermo Cabrera G.^{1,2,*}, Matthias Ehrgott³, Andrew Mason¹,
and Andrea Raith¹

¹Department of Engineering Science, University of Auckland (New Zealand)

²Escuela de Ingeniería Informática, Pontificia Universidad Católica de Valparaíso (Chile)

³Management School, Lancaster University (United Kingdom)

Received: date / Accepted: date

Abstract

Selecting a suitable set of beam angles is an important but difficult task in intensity modulated radiation therapy (IMRT) for cancer treatment. From a single objective point of view this problem, known as beam angle optimisation (BAO) problem, is solved by finding a beam angle configuration (BAC) that leads to the best dose distribution, according to some objective function. Because there exists a trade-off between the main goals in IMRT (to irradiate the tumour according to some prescription and to avoid surrounding healthy tissue) it makes sense to solve this problem from a multi-objective (MO) point of view. When doing so, a solution of the BAO problem is no longer a single BAC but instead a set of BACs which lead to a set of dose distributions that, depending on both dose prescription and physician preferences, can be selected as the preferred treatment.

We solve this MO problem using a two-phase strategy. During the first phase, a deterministic local search algorithm is used to select a set of locally optimal BACs, according to a single objective function. During this search, an optimal dose distribution for each BAC, with respect to the single objective function, is calculated using an exact non-linear programming algorithm. During the second phase a set of non-dominated

*Please address correspondence to guillermo.cabrera@ucv.cl

points is generated for each promising locally optimal BAC and a dominance analysis among them is performed. The output of the procedure is a set of (approximately) efficient BACs that lead to good dose distributions. To demonstrate the viability of the method, the two-phase strategy is applied to a prostate case.

Keywords: Multi-objective Beam Angle Optimisation, Deterministic Local Search, Mathematical Programming, Intensity Modulated Radiation Therapy

1 Introduction

Intensity modulated radiation therapy (IMRT) is a common technique applied in external radiation therapy. Its goal is to eradicate tumour cells by delivering ionising radiation from an external source to the tumour without compromising surrounding normal tissue and organs at risk (OARs). Unfortunately, because of the physics of radiation delivery, there is a trade-off between control of the tumour and sparing surrounding normal tissue. A distinctive feature of IMRT is that radiation can be modulated using a device called a *multi-leaf collimator* (Romeijn and Dempsey, 2008). Modulation is achieved by moving the metal leaves of the multi-leaf collimator into the radiation fields, thereby partially blocking the radiation. Several of these partially blocked radiation fields are superimposed to generate a *fluence map* which leads to a dose distribution that delivers more radiation to the tumour while better sparing surrounding OARs than unmodulated conformal radiotherapy. Modulation is particularly important for tumours with non-convex shapes in difficult anatomical situations (Ehrgott et al, 2009; Romeijn and Dempsey, 2008).

Due to the complexity of designing a treatment plan, the IMRT planning problem is usually divided into three sequential sub-problems, namely, beam angle optimisation (BAO), fluence map optimisation (FMO) and multi-leaf collimator sequencing. A solution of the BAO problem determines the number and directions of radiation beam angles, i.e. a beam angle configuration (BAC). Then, the optimal fluence of radiation for each beam angle needs to be computed (FMO problem) to give a high quality dose distribution in the organs. Finally, a sequencing problem needs to be solved to control the movement of the multi-leaf collimator leaves during delivery of the optimised fluence (Ehrgott et al, 2009). This sequential approach implies that the quality of the solution obtained at each step depends to a possibly large extent on the solution obtained in the previous steps. Therefore, beam angle selection will have a considerable impact on the fluence map obtained when solving the FMO problem using a specific BAC.

In this paper, we focus on the BAO problem and the associated FMO problems, but ignore the collimator sequencing problem, i.e. we consider the problem of selecting the appropriate beam angles that lead to a clinically acceptable dose distribution (Pugachev et al, 2000). In radiation therapy practice, treatment planners usually define the set of beam angles manually based on their ex-

perience and intuition as well as considering some geometrical features of the problem. However, as has been pointed out previously in the literature, manual selection may lead to sub-optimal treatment plans (Ehrgott and Johnston, 2003; Pugachev and Xing, 2002; Pugachev et al, 2001; Rowbottom et al, 1998). Currently, the planning process proceeds as follows: The planner selects a BAC based on his/her experience and generates the fluence map (i.e. the planner uses specialised software to solve the FMO problem) for that BAC. If the resulting fluence map is clinically unacceptable, the treatment planner often defines a new BAC and repeats the fluence map optimisation for the new BAC. This process is repeated until a clinically acceptable fluence map is obtained (Ehrgott et al, 2008). It is not uncommon that this process can take several hours. Thus, solving the BAO problem using automated beam angle selection techniques has the potential to generate good quality treatment plans independently of planner experience and to alleviate the tedious repetitive process described above.

Mathematically, the BAO problem can be formulated as follows. Let K be the set of all possible beam angles around the patient. In this work we consider $K = \{k\pi/180 : k = 0, 1, 2, \dots, 359\}$. Coarse angular resolutions can also be used within our framework to reduce the number of angles in K . Let $\mathcal{A} \in \mathcal{P}^N(K)$ be a BAC where $\mathcal{P}^N(K)$ is the set of all N -element subsets of K , with $N > 0$ the a priori determined number of angles. We denote the i -th angle of \mathcal{A} by \mathcal{A}_i for $i = 1, \dots, N$. The BAO problem seeks to find a BAC \mathcal{A} which minimises some objective function $h : \mathcal{P}^N(K) \rightarrow \mathbb{R}$,

$$\min_{\mathcal{A} \in \mathcal{P}^N(K)} h(\mathcal{A}). \quad (1)$$

The BAO problem is usually stated as a mixed integer (or binary) optimisation problem (Bangert et al, 2012; Ehrgott and Johnston, 2003; Preciado-Walters et al, 2004, 2006), as the set of all possible angles K is discrete and only a subset of them is considered in a BAC \mathcal{A} . Furthermore, for most choices of h , BAO is a non-convex problem with possibly many local optima (Ehrgott and Johnston, 2003; Ehrgott et al, 2008; Lim and Cao, 2012; Pugachev et al, 2001). The feasible set for the BAO problem has $\binom{|K|}{N}$ elements, and since clinically relevant values of N range from 5 to 15 beam angles, there are between 4.9×10^{10} and 1.3×10^{26} subsets of K (Ehrgott et al, 2008). Moreover, when BAO and FMO are posed together (i.e. when evaluating $h(\mathcal{A})$ requires solving the FMO problem) the set of feasible solutions is highly enlarged (Bortfeld and Schlegel, 1993; Pugachev and Xing, 2001). Thus, it is not possible to explore the entire set of feasible solutions in reasonable computational time. Hence, we need to find some efficient strategy to identify those BACs that minimise h .

The general framework applied so far in the literature considers hybrid strategies combining mathematical programming and meta-heuristics to solve the BAO problem. While meta-heuristics are used because of their ability to avoid getting trapped in local optima, mathematical programming is used to identify the optimal solution of the corresponding FMO problem for a specific set of angles. For instance, Li et al (2004) present a genetic algorithm to solve the BAO problem. In their paper, the objective function of the corresponding FMO

problem is a commonly used quadratic penalty function which penalises those sub-volumes in the tumour (OARs) below (above) lower (upper) bounds. The genetic algorithm selects a subset of angles and a conjugate gradient method is used to solve the FMO problem, which is modelled as a quadratic optimisation problem. According to the authors, their genetic algorithm obtains very good solutions in clinical cases, reaching high radiation uniformity in the tumour and low radiation in most of the OARs, outperforming manual beam angle selection consistently. Instead of the genetic algorithm used in Li et al (2004), Li et al (2005a) present a particle swarm optimisation algorithm to solve the BAO problem. The corresponding FMO problem in each BAC considered by the particle swarm optimisation algorithm is solved, again, by a conjugate gradient algorithm. The particle swarm optimisation algorithm is applied to both clinical and non-clinical cases. According to the authors, the particle swarm optimisation approach seems to be at least as good as the genetic algorithm strategy presented in Li et al (2004). Dias et al (2014) present a genetic algorithm combined with a neural network to solve the single objective BAO problem. The authors make use of a surrogate function to speed up the algorithm. To do that, the genetic algorithm uses the neural network as a surrogate model to calculate an approximation of the judgement function value of most of the visited BACs. The authors highlight that their approach is able to find improved BACs when compared to the ones used in clinical practice. Another heuristic used to solve the BAO problem is an ant colony system. While in Li et al (2005b) a pure ant colony system is combined with a conjugate gradient algorithm, the same authors present a hybrid strategy combining a genetic algorithm and an ant colony system to solve the BAO problem in Li and Yao (2006). The authors propose this hybrid approach in order to obtain a balance between exploration and exploitation attributes of the hybrid algorithm. While the genetic algorithm is used to explore the set of feasible BACs to initialise the ant colony system (exploration), the ant colony system is used to fine tune the (locally) optimal BAC (exploitation). According to the authors, the combined genetic algorithm and ant colony system algorithm is a more efficient approach than each algorithm separately. Another meta-heuristic that has been widely implemented to solve the BAO problem is simulated annealing (Bortfeld and Schlegel, 1993; Djajaputra et al, 2003; Pugachev and Xing, 2002; Stein et al, 1997). Recently, Bertsimas et al (2013) presented a hybrid implementation of simulated annealing and a gradient descent method. While the gradient descent method is used to quickly find a local optimal solution, simulated annealing is used to get the algorithm out of locally optimal solutions and to search in different parts of the solution set. One distinctive feature of this approach is the use of gradient information to refine the BAC. Because they consider a linear objective function, they can obtain the gradient information using linear programming duality theory (Craft, 2007; Bertsimas et al, 2013).

Local search strategies have also been applied to the BAO problem. Das et al (2003) present a simple local search strategy based on a “beam picker” that selects one beam angle at a time from the current BAC and replaces it by a beam angle from a set of beam angle candidates. Aleman et al (2008) imple-

ment two different approaches to solve the BAO problem: simulated annealing and add/drop algorithms. They also propose different neighbourhood definitions for the local search and all of them are tested in both algorithms. The objective function for their corresponding FMO problem is a convex penalty function which is solved to optimality by means of an interior point method. Results show that the neighbourhood choice can be important in order to improve the performance of a particular algorithm implementation. More recently, Lim et al (2014) present a hybrid framework combining local search and other optimisation methods such as simulated annealing, genetic algorithms, nested partitions, to solve the single objective BAO problem. Their framework consists of two sequential phases. During the first phase, either a heuristic or an exact method is used to quickly find a good feasible BAC. During the second phase the BAC obtained during the previous phase is used to warm-start a local search algorithm. According to the authors, their framework overcomes the weakness observed when using the optimization methods separately and consistently finds good solutions.

All these approaches have some characteristics in common. They all use heuristics to search for new BACs and mathematical programming to find an optimal solution to the corresponding FMO problem. Solving the associated FMO problem to optimality is an important feature of all these approaches, as it allows a fair comparison of different subsets of beam angles. Another shared characteristic is that they evaluate the quality of a BAC by considering the dose distribution obtained by the associated FMO problem. Furthermore, they all model the problem as a single-objective optimisation problem and both the BAO and the corresponding FMO problem are treated as separate problems.

Unlike the approaches mentioned above, in this paper we propose a heuristic method to solve the BAO problem from a multi-objective (MO) point of view. While most of the efforts in multi-objective IMRT optimisation have been focused on the MO-FMO problem (Breedveld et al, 2007; Craft et al, 2007; Hamacher and Küfer, 2002; Kalantzis and Apte, 2014; Küfer et al, 2003; Lahanas et al, 2003) only little attention has been paid to the MO-BAO problem. The MO-BAO was first addressed by Schreibmann et al (2004), who propose an *a posteriori* method that consists of a MO genetic algorithm, namely NSGAIIC (Deb et al, 2002), and an exact solver based on the Broyden-Fletcher-Goldberg-Shanno algorithm (L-BFGS). While the NSGAIIC algorithm is used to generate new BACs, the L-BFGS solver is used to calculate one fluence map, optimizing a weighted sum of objectives, for each individual of the population. Because genetic algorithms work with a population of solutions, the evaluation of each individual (BAC) in the population by solving an FMO problem cannot be allowed to consume much computational time. For this reason, Schreibmann et al (2004) limit the number of iterations the L-BFGS solver is allowed to perform, leading to sub-optimal fluence maps. Furthermore, only one fluence map is produced for each BAC. Another genetic algorithm is proposed by Fiege et al (2011). Their algorithm, called *Ferret*, simultaneously optimises fluence maps and beam directions. This is a very distinctive feature as most approaches proposed so far consider an outer loop to find good BACs and an inner loop which

optimises the corresponding fluence map. Fiege et al (2011) point out that due to the enlarged feasible set, solving both problems simultaneously is harder than solving them consecutively. Because of that, and similar to Schreibmann et al (2004), their algorithm only approximates optimal fluence maps and simplified objective functions are considered in order to speed up the algorithm. More recently, Breedveld et al (2012) propose an *a priori* method to solve the MO-BAO problem. Their method, called *iCycle*, starts with an empty BAC, i.e., no beams are selected. Then, the best possible beam angle is selected from a predefined set of candidate beams at each iteration and added to the BAC that is being generated. The method stops when no further (relevant) improvement in the obtained BAC is made. The method can also stop if the number of beams in the current BAC has become impractically large. The *iCycle* method determines the best beam angle to be added at each iteration by solving a single objective optimization problem where the objective function corresponds to a weighted sum of the objectives from the original MO problem, as explained in Breedveld et al (2009). As the method obtains one BAC at each iteration the treatment planner can select the BAC with the best trade-off between the number of beams used and the treatment plan quality. A similar approach is presented in Azizi-Sultan (2010). He presents an algorithm that consists of two steps that are repeated until the BAC obtained at the end of the second step has the desired number of beams. During the first step an FMO problem is solved and the dose deviation from the prescribed dose is calculated for the current N -BAC (a BAC considering N beam angles). During the second step, the contribution of every beam angle not in the current N -BAC is calculated, using a predefined score function. The beam angle with the best contribution is then added to the new $(N + 1)$ -BAC. The algorithm ends when the desired number of beam angles in the current BAC is achieved. Finally, a MO-FMO problem is solved using the obtained BAC.

Unlike previously proposed strategies, in this paper we use a two-phase strategy to solve the MO-BAO problem. In the first phase, we use local search to select a set of locally optimal BACs (and their associated optimal fluence maps) according to a pre-defined single objective function. Each of these locally optimal BACs is represented by a *sample point* in the objective space of the MO-FMO problem. Each sample point is the image of the fluence map associated with its corresponding locally optimal BAC. Each coordinate of the sample point corresponds to one of the objective function values in the MO-FMO problem, where each objective function is associated with either the tumour or a specific organ at risk considered in the MO-FMO problem. One key feature of this local search is that an optimal fluence map of each BAC leads to a non-dominated point of the associated MO-FMO problem. This is important as we can ensure we are comparing solutions that are not only optimal in the single objective sense but are also efficient from a MO point of view, i.e. we cannot further improve any objective without impairing the value of at least one other objective. Finally, after the set of locally optimal BACs has been generated, a dominance analysis is performed to find those BACs whose sample points are not dominated by any other sample point. We expect that these BACs will be used in the set of

efficient treatment plans. We say that a BAC that occurs in the set of efficient treatment plans is efficient. A formal definition for efficient BACs is given in Section 3. As we cannot evaluate all possible BACs we cannot ensure that a BAC is truly efficient. Thus BACs found by the two-phase approach presented here are called *approximately* efficient BACs. During the second phase we generate a large set of non-dominated points for all the obtained (approximately) efficient locally optimal BACs by means of the well-known ε -constraint method (Chankong and Haimes, 1983; Haimes et al, 1971). Some such points will be dominated by points generated by other BACs. Those dominated points are deleted leaving a final set of non-dominated points for the treatment planner to choose from. Hereafter, we will refer to the objective function(s) of the BAO problem as a *judgement function*.

To formally describe the BAO problem, Ehrgott et al (2008, p. 1276) define a judgement function as a function that “describes how well a patient can be treated with a set of N angles”. Then, they explain that the goal of any judgement function is “to capture the nature of how the treatment planner decides between good and bad treatments, and the value of such a function is the [value of the optimal] solution to an optimisation problem that decides a fluency [map]” (Ehrgott et al, 2008, p. 1276). Therefore, solving the BAO problem requires solving the FMO problem many times, once for every BAC that is considered by the BAO problem. Thus, we can say that judgement function h of the general BAO problem in (1) can be expressed as

$$h(\mathcal{A}) = \min \{z(x) : x \in X(\mathcal{A})\}. \quad (2)$$

Equation (2) shows the general form of a judgement function, where z maps a solution $x \in X(\mathcal{A})$ of the corresponding FMO problem into $\mathbb{R}_+^* := \{v \in \mathbb{R} : v \geq 0\} \cup \{\infty\}$ by finding the best fluence map that can be delivered for given BAC \mathcal{A} , according to function z . Fluence map $x \in \mathbb{R}^n$ is a non-negative vector with each component x_i representing the length of time that a patient is exposed to sub-beam (or *beamlet*) i and where n is the total number of beamlets summed over all 360 possible beam angles. The set $X(\mathcal{A})$ is the set of all feasible solutions of the FMO problem when BAC \mathcal{A} is considered. Only beamlets x_i that belong to a beam angle in \mathcal{A} are allowed to be greater than zero. See Ehrgott et al (2008) for a comprehensive explanation of the judgement function concept.

This paper is organised as follows. Section 2 presents the IMRT problem focusing on the BAO problem. In Section 2 we also present the generalised equivalent uniform dose (*gEUD*) concept which is a key element of our model. In Section 3 the MO-BAO problem is presented. After introducing a mathematical formulation of the MO-BAO problem in Section 3.1, single objective judgement functions considered in this work are presented in Section 3.2. Then, in Section 4, the two-phase approach is outlined. The main differences compared to other approaches for solving the BAO problem are presented in this section. In Section 5 we introduce a test instance and the obtained results are discussed. Finally, in Section 6 some conclusions are drawn and future work is outlined.

2 An Overview of IMRT Optimisation Problems

In IMRT, radiation can be delivered from any beam angle $k \in K$. As we mentioned before, vector $x \in \mathbb{R}^n$ denotes the fluence map, where n corresponds to the total number of beamlets summed over all possible 360 beam angles and its elements $x_i \geq 0$ correspond to the fluence at beamlet i . To mathematically model the IMRT problem, each organ is discretised into small sub-volumes called *voxels*. A crucial component of all mathematical optimisation models of the FMO problem is the radiation dose deposited into each voxel j of the tumour and OARs by fluence map x . This dose d is calculated using expression (3) (Ehrgott et al, 2009)

$$d_j^r = \sum_{i=1}^n A_{ji}^r x_i \text{ for all } j = 1, 2, \dots, m^r, \quad (3)$$

where $r \in R = \{O_1, \dots, O_Q, T\}$ is an element of the index set of regions (we use the term region to denote either the tumour, any organ at risk, or normal tissue), with the tumour indexed by $r = T$ and the organs at risk and normal tissue indexed by $r = O_q$ with $q = 1, \dots, Q$. m^r is the total number of voxels in region r , j corresponds to a specific voxel in region r , $d^r \in \mathbb{R}^{m^r}$ is a dose vector and its elements d_j^r give the total dose delivered to voxel j in region r by the fluence map $x \in X(\mathcal{A})$. Here, dose deposition matrix $A^r \in \mathbb{R}^{m^r \times n}$ is a given matrix where $A_{ji}^r \geq 0$ defines the rate at which radiation dose along beamlet i is deposited into voxel j in region r .

Based on the dose distribution in formula (3), a large number of formulations for the FMO problem for a specific BAC have been proposed in the literature during the last decade (see Ehrgott et al (2009) for a survey). On the one hand there are the so-called physical models, also known as *dose-volume* models, which are focused on the percentage of the total volume of either the tumour or the OARs that is irradiated at a certain level. Dose-volume constraints of the form *no more than x% of the volume of the tumour may receive a dose of less than y* or, for the OAR case, *no more than x% of the volume of the OAR may receive a dose of more than y* are usually considered in these models. The objective function in dose-volume models is usually the maximisation of the dose delivered to the tumour and/or the minimisation of the dose delivered to the OARs subject to both bound constraints and dose-volume constraints. Most of these physical models give rise to linear, mixed-integer, or quadratic optimisation models (Ehrgott et al, 2009). This allows researchers and treatment planners to find clinically acceptable treatment plans using well known exact optimisation techniques.

On the other hand, biological models, also known as *dose-response* models, relate the delivered dose to the biological response of the irradiated structures. In general, the objective function of a dose-response model is to maximise the *tumour control probability* (i.e. the probability of completely eradicating all cancerous cells from the local tumour site (Webb, 2010)) while minimising *normal tissue complication probabilities* of OARs (i.e. the probability of

inducing some particular complication in a healthy organ (Webb, 2010)). Some researchers have pointed out advantages of dose-response models over physical ones (Dirscherl et al, 2011; Nahum and Uzan, 2012; Thomas et al, 2005; Wu et al, 2002). Although conceptually more appealing (the patient is, after all, a biological system and not a dosimeter (Amols and Ling, 2002)), they are not yet widely used because of a lack of consolidated clinical data and also because the corresponding FMO problems are difficult to solve mathematically.

A well-known biological model is the generalised equivalent uniform dose (*gEUD*). Originally proposed in Niemierko (1997) the *gEUD* can be defined as the biologically equivalent dose that, if delivered uniformly, would lead to the same response as the actual non-uniform dose distribution (Niemierko, 1997). One advantage of *gEUD*-based models over other dose-response models is that they do not require many parameters, making the optimisation process less dependent on parameters that could be difficult to determine. Moreover, *gEUD* penalises less (more) irradiated voxels in tumour (OAR) regions which leads to a more homogeneous dose distribution in the tumour and to avoidance of overdosed voxels in OARs, two very desirable properties in IMRT. Therefore, *gEUD* allows us to evaluate and compare dose distributions even if they are non-uniform. Several research articles have been devoted to the study of *gEUD*-based IMRT planning. Most of them highlight the ability of *gEUD*-based models to obtain better OAR sparing while keeping the same or even better tumour coverage (Thomas et al, 2005; Wu et al, 2000, 2003).

The mathematical expression for *gEUD* is

$$gEUD^r(x) = \left(\frac{1}{m^r} \sum_{j=1}^{m^r} (d_j^r)^{a^r} \right)^{1/a^r}, \quad (4)$$

where a^r is a region-dependent parameter and d_j^r (which depends on x) is as given in Equation (3). For the tumour, we put $a^r < 0$, whereas for OARs we choose $a^r > 1$. As $|a^r|$ increases, the *gEUD* becomes more sensitive to less (more) irradiated voxels in the tumour (OARs). Therefore, for those OARs that allow certain levels of radiation without a functional compromise (also called parallel organs), parameter a^r is set close to 1, whereas for serial organs (those that must be irradiated as little as possible) values for parameter a^r are chosen to be large.

The objective functions of both the BAO and FMO problems considered in this paper are *gEUD*-based functions. Mathematical formulations that consider such functions are presented in the next section. We consider convex *gEUD*-based functions which allows us to find optimal solutions using standard mathematical programming solvers. However other formulations could be used within the two-phase framework that is presented in this work. In case the chosen objective functions are non-convex, one could seek a suitable decomposition of the original non-convex objective function into a convex objective function and an increasing function (Romeijn et al, 2004). Romeijn et al (2004) present a list of different objectives (both convex and non-convex) and their corresponding

decompositions into a convex objective and an increasing function. They also demonstrate that equivalent sets of non-dominated points can be found using other MO-FMO formulations, i.e. considering other objective functions. Hence, although our work focuses on gEUD, the solutions we produce are efficient under any of these objectives as well.

3 The Multi-Objective Beam Angle Optimisation Problem

In this paper we solve an optimisation problem which we denote the multi-objective beam angle optimisation problem (MO-BAO). Throughout the paper we will use the following notation for the comparison of vectors. Let $y^1, y^2 \in \mathbb{R}^p$. We write $y^1 \leq y^2$ if $y_k^1 \leq y_k^2$ for all $k = 1, \dots, p$; $y^1 < y^2$ if $y^1 \leq y^2$ but $y^1 \neq y^2$; and $y^1 < y^2$ if $y_k^1 < y_k^2$ for all $k = 1, \dots, p$. Thus, we say that a solution $x \in X$ is *efficient* if there is no $x' \in X$ such that $z(x') \leq z(x)$ (assuming that we want to minimise all the objective functions). Equivalently, the image $z(x)$ of the efficient solution x is called *non-dominated*.

3.1 Multi-objective beam angle optimisation

Before we introduce the MO-BAO problem we need to define the associated MO-FMO problem. For a fixed BAC $\mathcal{A} \in \mathcal{P}^N(K)$, the general MO-FMO problem can be formulated as

$$f(\mathcal{A}) = \min_{x \in X(\mathcal{A})} z(x), \quad (5)$$

where $z(x) \in \mathbb{R}^{|R|}$ is a vector of $|R|$ objective functions z_r , $r = 1, \dots, |R|$ and $|R|$ is the total number of regions considered in the problem. Unlike the single objective formulation (2), which requires the determination of a single optimal solution, the solution to this multi-objective problem is a set $X_E^{\mathcal{A}}$ containing efficient fluence maps of MO-FMO problem (5). We define $Y_N^{\mathcal{A}} = f(\mathcal{A})$ as the set of associated non-dominated points given by $Y_N^{\mathcal{A}} = \{z(x) : x \in X_E^{\mathcal{A}}\}$.

The MO-BAO problem we are investigating in this paper is

$$\min_{\mathcal{A} \in \mathcal{P}^N(K)} \min_{x \in X(\mathcal{A})} z(x), \quad (6)$$

the solution of which is a set containing all efficient BACs which use exactly N angles. Additionally, (6) asks for the generation of a set X_E containing the efficient fluence maps which belong to those efficient BACs and that lead to the set $Y_N = \{z(x) : x \in X_E\}$, the associated set of all non-dominated points. We also write problem (6) as

$$\min_{\mathcal{A} \in \mathcal{P}^N(K)} f(\mathcal{A}), \quad (7)$$

to emphasise the practical need to solve the MO-FMO problem (5) for different beam angle configurations $\mathcal{A} \in \mathcal{P}^N(K)$.

Unlike in the single objective BAO problem, solving the MO-BAO problem means finding the set of efficient BACs in $\mathcal{P}^N(K)$. As we mentioned before, a BAC \mathcal{A} is efficient if $X(\mathcal{A}) \cap X_E \neq \{\}$, or equivalently if there is a fluence map $x \in X(\mathcal{A})$ such that there is no BAC \mathcal{B} and fluence map $x' \in X(\mathcal{B})$ with $z(x') \leq z(x)$.

Unfortunately we cannot evaluate all possible BACs in reasonable time. Therefore, we will only find a set $\hat{\mathcal{A}}$ that approximates the actual set of efficient BACs. Similarly, \hat{X}_E will denote the approximation to the set of all efficient fluence maps X_E . Images of solutions $x \in \hat{X}_E$ are denoted by $y \in \hat{Y}_N$.

Although our two-phase algorithm can be adapted for other models of the FMO problem, in particular constrained optimization problems, as we explained in Section 2, *gEUD*-based models are considered in this work. A special case of the MO-FMO problem in (5) that is based on the *gEUD* is

$$\begin{aligned} \max \quad & gEUD^T(x) \\ \min \quad & gEUD^{O_q}(x) \quad \text{for } q = 1, \dots, Q, \\ \text{s.t.} \quad & x \in X(\mathcal{A}), \end{aligned} \tag{8}$$

where $X(\mathcal{A})$ is, again, the set of feasible fluence maps which is defined by $x \geq 0$ and beamlets $x_i = 0$ for all beamlets not belonging to beam angles in \mathcal{A} . The MO-BAO problem considered in this paper is the minimisation (in the sense explained above) of (8) over all $\mathcal{A} \in \mathcal{P}^N(K)$

The *gEUD* function has several favourable optimisation properties such as convexity and positive homogeneity. As Cabrera G. et al (2014) demonstrate, positively homogeneous multi-objective optimisation problems, such as the one from Equation (8), with p objective functions can be solved as multi-objective optimisation problems with $p - 1$ objectives. This can be done by transforming one of the objectives into a constraint. Following the procedure presented in Cabrera G. et al (2014) we can generate an infinite number of efficient solutions located on a finite number of rays, each of which corresponds to an efficient solution of the reduced $p - 1$ objective problem. Thus, the MO-BAO problem considered in this paper can be re-stated as follows:

$$\min_{\mathcal{A} \in \mathcal{P}^N(K)} f_{O_q}(\mathcal{A}) = \min_{\mathcal{A} \in \mathcal{P}^N(K)} \min_{x \in X(\mathcal{A}): gEUD^T(x) \geq t} gEUD^{O_q}(x) \quad \text{for } q = 1, \dots, Q, \tag{9}$$

where t is equal to the prescribed *gEUD* of the tumour.

The corresponding MO-FMO problem for a fixed BAC $\mathcal{A} \in \mathcal{P}^N(K)$ based on the *gEUD* is

$$\begin{aligned} \min \quad & gEUD^{O_q}(x) \quad \text{for } q = 1, \dots, Q, \\ \text{s.t.} \quad & gEUD^T(x) \geq t, \\ & x \in X(\mathcal{A}). \end{aligned} \tag{10}$$

It is important to note that the problem in (10) is convex, and so we can find optimal solutions using non-linear solvers such as IPOPT (Wächter and Biegler,

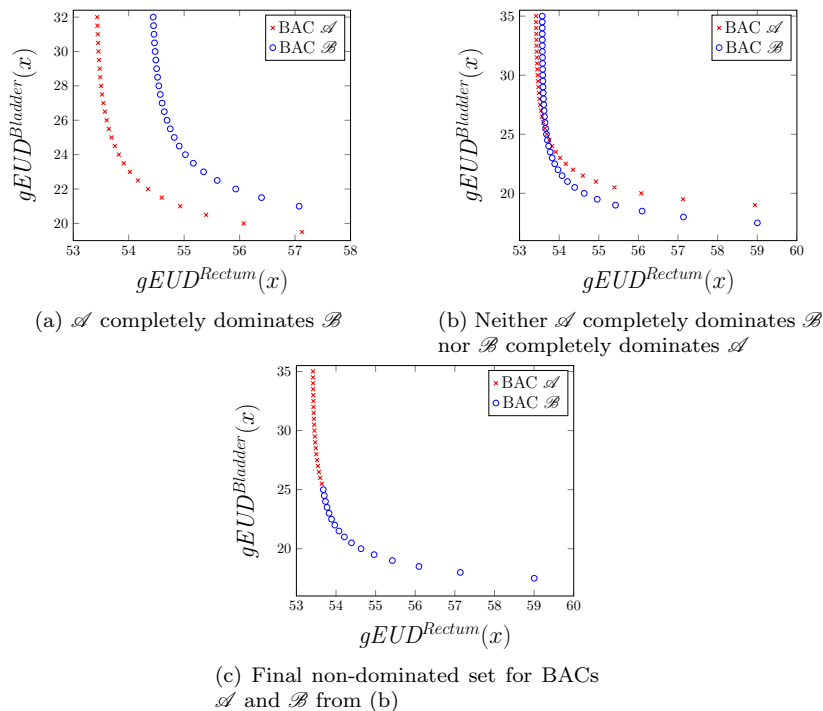


Figure 1: Dominance relations between sets of non-dominated points. $gEUD$ of the tumour is fixed

2006). Scalarisation methods such as the well-known ϵ -constraint method mentioned before also give sub-problems that are convex, and so we can use IPOPT to obtain a set of truly efficient solutions for this problem. Having this guarantee of optimality for this sub-problem is likely to result in better quality solutions being produced for the full MO-BAO problem.

In the framework of judgement functions, we are optimising the multi-objective judgement function $f_{O_q}(\mathcal{A})$ with $q = 1, \dots, Q$ defined in (9) over all N -element BACs $\mathcal{A} \in \mathcal{P}^N(K)$. One difficulty in solving the BAO problem is to determine whether a BAC is better than another one. When this problem is tackled from a single objective point of view, a decision about which BAC is better is made using the single objective judgement function value. Thus, a BAC will be better than another one if its single objective optimal fluence map leads to a better judgement function value. However, when this problem is tackled from a MO point of view, comparison of different BACs is not a trivial task. In the MO case, we need to decide between two sets $Y_N^{\mathcal{A}}$ and $Y_N^{\mathcal{B}}$ rather than single optimal values. Figures 1(a) and 1(b) show two sets $Y_N^{\mathcal{A}}$ and $Y_N^{\mathcal{B}}$ for the MO-BAO problem (9) with two OARs. Figure 1(a) shows a situation where it is easy to decide between \mathcal{A} (crosses) and \mathcal{B} (circles). There is no reason to select a vector $y' \in Y_N^{\mathcal{B}}$ as there is always a vector $y \in Y_N^{\mathcal{A}}$ such that

$y \leq y'$. Therefore, we can say that BAC \mathcal{A} completely dominates BAC \mathcal{B} and, consequently, it leads to better fluence maps.

Unfortunately this situation is not as common as the one presented in Figure 1(b). In this case there is at least one $y \in Y_N^{\mathcal{A}}$ that dominates at least one $y' \in Y_N^{\mathcal{B}}$ and vice versa. Thus, we need to keep all vectors $y \in Y_N^{\mathcal{A}} \cup Y_N^{\mathcal{B}}$ that are non-dominated as in Figure 1(c).

As mentioned in Section 2, generating a set containing many non-dominated points $Y_N^{\mathcal{A}}$ for a large number of BACs $\mathcal{A} \in \mathcal{P}^N(K)$ is not possible because of the time that it requires. Therefore, we propose to consider only one non-dominated point $s \in Y_N^{\mathcal{A}}$ to represent the BAC \mathcal{A} and make the assumption that the quality of BAC \mathcal{A} corresponds to the quality of s . We call s a *sample point*. The sample point corresponds to the optimal solution of a single objective FMO problem. This means that $s = z(x)$ with z being the objective function of the MO-FMO problem in (10) whereas x is the optimal solution of a single objective FMO problem using BAC \mathcal{A} .

Comparison between two BACs is done by using a single objective judgement function h as in (2). Thus, a BAC \mathcal{A} will be considered better than BAC \mathcal{B} iff $h(\mathcal{A}) < h(\mathcal{B})$. We need to state at this point that the single objective judgement function h in (2) is only used to determine a subset of “potentially” good quality BACs. However, the final approximation to the set of efficient BACs $\hat{\mathcal{A}}$ and its corresponding set of (approximately) efficient fluence maps \hat{X}_E are calculated using the MO formulation based on the $gEUD$ function presented in Equations (9) and (10). We would expect that our heuristic approach is more likely to find useful sample points if h has a strictly monotone objective function, i.e. if $z(x)$ dominates $z(x')$ then $h(\mathcal{A}) < h(\mathcal{A}')$ with $x \in X(\mathcal{A})$ and $x' \in X(\mathcal{A}')$, so its optimal solution is also an efficient solution of the MO-FMO problem in (10) (Miettinen, 1999).

It is clear that using different single objective judgement functions h may lead to different sample points for the same BAC. The next sub-section presents two $gEUD$ -based judgement functions that are considered in this paper for the generation of locally optimal BACs.

3.2 Judgement Functions

In this paper we consider two single objective judgement functions. The first one is related to an unconstrained FMO model presented originally by Wu et al (2002) and reformulated by Olafsson et al (2005). This model, also known as the *logistic model*, has been widely studied in the literature (Cabrera G. et al, 2012; Choi and Deasy, 2002; Olafsson et al, 2005). According to the majority of these studies, solutions obtained using this model irradiate the tumour in a similar way to dose-volume models but with better OAR sparing.

Equations (11), (12), and (13) show this model.

$$\min_{x \in X(\mathcal{A})} z_1(x) = -\ln L(x; T, v_T, eud_0^T) - \sum_{q=1}^Q \ln U(x; O_q, v_{O_q}, eud_0^{O_q}), \quad (11)$$

where

$$L(x; T, \nu_T, eud_0^T) = \left(1 + \left(\frac{eud_0^T}{gEUD^T(x)} \right)^{\nu_T} \right)^{-1}, \text{ and} \quad (12)$$

$$U(x; O_q, \nu_{O_q}, eud_0^{O_q}) = \left(1 + \left(\frac{gEUD^{O_q}(x)}{eud_0^{O_q}} \right)^{\nu_{O_q}} \right)^{-1}. \quad (13)$$

Parameters eud_0^T and $eud_0^{O_q}$ correspond to the prescribed $gEUD$ values for tumour and OARs, respectively, and $\nu_T, \nu_{O_q} > 0$ are user-defined parameters that indicate the importance of the tumour and the q -th OAR, respectively. While $-\ln(L)$ is a convex function of $gEUD^T(x)$, function $-\ln(U)$ is monotonically increasing in $gEUD^{O_q}(x)$, but not convex. However, we observe that the second derivative of $-\ln(U)$ is strictly positive for $0 < gEUD^{O_q}(x) < U_+''$ where

$$U_+'' = \left((\nu_{O_q} - 1)^{1/\nu_{O_q}} \right) eud_0^{O_q}.$$

In practice, we would not be interested in values of $gEUD^{O_q}(x)$ that exceeded this upper bound, and so $-\ln(U)$ and hence (11) is convex for all x values of practical interest. We see this in our computational results (see Section 5) where all the $gEUD$ values obtained by our optimisation runs (see Table 2 and Figure 4a) fall within the convex regions of the U functions we used (see Table 1). For a comprehensive analysis of the logistic function readers are referred to Olafsson et al (2005).

Because the MO-FMO problem (10), and consequently the MO-BAO problem (9), have a constraint on the $gEUD$ value of the tumour, we need to make some changes to the logistic model presented in Equations (11), (12), and (13), in order to include this constraint. We must set the $gEUD$ of the tumour to be greater than or equal to its prescribed dose eud_0^T , thus we do not need to consider L in (11) and (12) but, instead we need to include a constraint for the $gEUD^T$. Due to the conflict between $gEUD^T(x)$ and $gEUD^{O_q}(x)$, this constraint will always be active at an optimal solution. Thus, the logistic FMO problem given some BAC \mathcal{A} is now

$$\begin{aligned} h_1(\mathcal{A}) = \min z_1(x) &= - \sum_{q=1}^Q \ln U(x; O_q, \nu_{O_q}, eud_0^{O_q}) \\ \text{s.t. } gEUD^T(x) &\geq eud_0^T \\ x &\in X(\mathcal{A}), \end{aligned} \quad (14)$$

where h_1 is the judgement function associated with FMO problem with objective function z_1 . Since $gEUD^T(x)$ is a convex function of x and $-\ln(U)$ is convex in $gEUD^T(x)$ for the region of interest, the problem in (14) is also convex, and so solutions found by the IPOPT solver will be optimal for problems of practical interest.

The second judgement function used in this work is associated with an optimisation model which minimises the $gEUD$ value of one of the OARs. The idea

is to identify the OAR O_r which the planner would like to spare the most among all OARs considered in the treatment planning process. We call this judgement function *lexicographic*. Equation (15) shows this FMO model:

$$\begin{aligned} h_2(\mathcal{A}) = \min z_2(x) &= gEUD^{O_r}(x) \\ \text{s.t. } gEUD^T(x) &\geq eud_0^T \\ x &\in X(\mathcal{A}), \end{aligned} \tag{15}$$

where h_2 is the judgement function associated with FMO problem with objective function z_2 .

Using mathematical programming techniques we can obtain optimal fluence maps of a BAC for both single objective FMO problems (14) and (15). As we mentioned before, one key feature of these problems is that their optimal solutions are also efficient solutions of the MO-FMO problem in (10), as they both have strongly decreasing objective functions (Miettinen, 1999). Thus, based on their single objective judgement function value, we can hopefully identify good BACs. The fact that optimal solutions of (14) and (15) are also efficient solutions to (10) gives us some confidence that we can compare different BACs from a MO point of view. Sample points are obtained evaluating optimal solutions of (14) and (15) with the MO function z in (10).

In this paper we propose a two-phase approach which allows us, firstly, to find a set of locally optimal BACs based on the obtained sample points and, secondly, to generate a set of approximately efficient BACs of the MO-BAO problem in (9) and its corresponding set of approximately efficient fluence maps. This two-phase approach is presented in the next section.

4 The Two-phase Approach

One challenge that algorithms solving the BAO problem must face is the large number of possible combinations of beam angles to explore. Several strategies have been proposed in order to improve the computational efficiency of the algorithms proposed so far in the literature. Most of them try to reduce the number of alternatives by means of rankings or score functions. For instance, Pugachev and Xing (2002) rank beam angles prior to the optimisation process based on beam's-eye-view dosimetrics (Pugachev and Xing, 2001), i.e. the portion of the target that can be irradiated by each beam angle. Then, based on both geometric and dosimetric information that is available before the optimisation process starts, the beam's-eye-view can guide the search. Although the authors do not ban any beam angle a priori, they avoid the FMO step for those configurations that have a low beam's-eye-view score.

In this paper we do not consider any reduction of the set of possible BACs. Instead, we propose a two-phase strategy that is based on the concept of sample points. Before presenting the two-phase approach we need to introduce some specific notation we will use here. Let \mathcal{A}^* be a set of locally optimal BACs with respect to some single objective judgement function. Equivalently, let X^*

be a set containing an optimal fluence map for each BAC in \mathcal{A}^* . Let $\mathcal{S} = \{z(x) : x \in X^*\}$, with z being the objective function of the MO-FMO problem in (10), be the set of all generated sample points. It is important to note here that we record only one optimal fluence map for each $\mathcal{A} \in \mathcal{A}^*$. Consequently, there is only one sample point in \mathcal{S} for each BAC in \mathcal{A}^* . In case two BACs have the same sample point (i.e. the same objective function values) both points will be considered. Let $\mathcal{S}_N \subseteq \mathcal{S}$ be the set of non-dominated sample points with respect to points in \mathcal{S} under the objectives of the MO-FMO problem in Equation (10). Thus, all elements of \mathcal{S}_N are pairwise non-dominated.

Our two-phase approach is presented in the next sections.

4.1 First Phase: Deterministic Local Search

The main goal of the first phase of our strategy is to generate a set of promising BACs, \mathcal{A}^* , using one of the single objective judgement functions presented before (h_1 or h_2). In order to generate such a set we perform a deterministic local search algorithm. As with any other local search approach, our algorithm needs a neighbourhood \mathcal{N} to be defined by means of a neighbourhood movement. Our local search defines its neighbourhood by a one-degree-move in one of the beam angles. Mathematically, neighbourhood $\mathcal{N}(\mathcal{A})$, i.e. the neighbourhood of BAC \mathcal{A} , is as follows:

$$\mathcal{N}(\mathcal{A}) = \{\mathcal{B} \in \mathcal{P}^N(K) : \mathcal{A}_j = \mathcal{B}_j \pm \pi/180 \text{ for some } j \in \{1, \dots, N\} \text{ and } \mathcal{A}_i = \mathcal{B}_i \text{ for all } i = 1, \dots, N, \text{ where } i \neq j\}.$$

Although in this paper we only consider coplanar angles, our approach can be extended to non-coplanar angles by changing the neighbourhood definition \mathcal{N} to, for example, one of the neighbourhood definitions used in Lim and Cao (2012) or Mišić et al (2010).

Although we start with a randomly generated initial BAC, the local search implemented here is deterministic in the sense that, given an initial BAC, it always converges to the same local optima (see Algorithm 1). The initial BAC meets some geometric constraints such as having a minimum distance between any two beam angles and not having any two beam angles with a difference of 180 degrees. Elements of the neighbourhood $\mathcal{N}(\mathcal{A})$ do not need to meet these constraints. We could also start the local search using BACs commonly used in clinical practice. If that is the case, the two-phase approach will lead to a result that is at least as good as the one provided by well-known BACs. In other words, treatment plans obtained by those BACs commonly used in clinical practice can be considered as upper bounds of our two-phase approach.

After generating the initial BAC, the algorithm proceeds with the generation of the entire neighbourhood of the current BAC and the calculation of corresponding single objective judgement function values ($h()$ in Algorithm 1). This value is calculated by solving the associated FMO problem in either (14) or (15). For instance, if we consider five beam angles then the neighbourhood size will be equal to ten. Once the entire neighbourhood is generated we select the

BAC with the best judgement function value (\mathcal{A}' in Algorithm 1). If this value is better than the judgement function value of the current BAC then we set the best neighbour as the current BAC. The neighbourhood of the current BAC is then generated and we repeat this process until the single objective judgement function value of the best neighbour is no better than that of the current BAC. In this case, we have found a locally optimal BAC with respect to the judgement function that was considered. This BAC is added to the set of locally optimal BACs \mathcal{A}^* . We repeat this process several times starting from different randomly generated initial BACs. Each fluence map $x \in X^*$ is represented as a sample point $s \in \mathcal{S}$ in the objective space of the MO-FMO problem. Once we have generated the entire set of locally optimal BACs (and its corresponding sample points) we perform a dominance analysis over the set of sample points \mathcal{S} so only sample points that are not dominated by any other sample point are considered in the next phase (*isDominated()* function in Algorithm 1). Algorithm 1 shows the pseudocode of the deterministic local search.

Algorithm 1: Phase 1: Deterministic Beam Angle Local Search

Input: N (Number of angles in a BAC), ℓ (Number of local optima to be generated)

Output: \mathcal{A}^* (Set of locally optimal approximately efficient BACs)

```

begin
   $\mathcal{A}^* = \{\}$ ; // initially empty
  repeat
     $\mathcal{A} = \text{initialRandomSolution}(\mathcal{P}^N(K))$ ;
    localOptimum = false;
    repeat
       $\mathcal{N} = \text{generateNeighbourhood}(\mathcal{A})$ ;
       $\mathcal{A}' = \underset{\mathcal{B} \in \mathcal{N}}{\text{argmin}} h(\mathcal{B})$ ;
      if  $h(\mathcal{A}') < h(\mathcal{A})$  then
        |  $\mathcal{A} = \mathcal{A}'$ 
      else
        | localOptimum = true;
    until localOptimum;
    Add  $\mathcal{A}$  to  $\mathcal{A}^*$ ;
  until  $\ell$  local optima have been generated;
  foreach  $\mathcal{A} \in \mathcal{A}^*$  do
    | if isDominated( $\mathcal{A}, \mathcal{A}^*$ ) then
    | | Remove  $\mathcal{A}$  from  $\mathcal{A}^*$ ;
  return ( $\mathcal{A}^*$ );

```

We need to point out that the BACs in \mathcal{A}^* obtained during Phase 1 do not necessarily correspond to the ones with the best single objective judgement function values. This represents one important difference from single-objective

approaches that select only the BAC with the best judgement function value.

4.2 Second Phase: Exact Optimisation of the MO-FMO Problem

During the second phase of our approach we solve the associated MO-FMO problem for the set of locally optimal BACs found in Phase 1. Figure 2 shows an example of a set of sample points \mathcal{S} generated during Phase 1. This set includes $\ell = 100$ sample points. Each point corresponds to an optimal solution of a single objective FMO problem and is associated with a locally optimal BAC obtained in the previous phase. The area dominated by each non-dominated sample point $s \in \mathcal{S}_N$ is delimited by dashed lines. BACs associated with the sample points are locally optimal with respect to the single objective judgement function considered in Phase 1 of the procedure.

The goal of the second phase is to find a set of BACs $\hat{\mathcal{A}}$ that approximates the actual set of efficient BACs. As Algorithm 2 shows, we need, firstly, to generate a set $X_E^{\mathcal{A}}$ of efficient fluence maps of the MO-FMO problem in (10) for each BAC $\mathcal{A} \in \mathcal{A}^*$. To do that we can use scalarisation methods such as the ε -constraint method (Haimes et al, 1971) or the adaptive ε -constraint method (Eichfelder, 2008) (*SolveMO-FMO()* method in Algorithm 2). In this paper we use a procedure similar to the one presented in Cabrera G. et al (2014) to generate the set $X_E^{\mathcal{A}}$, which is based on the ε -constraint method and is specific to unconstrained MO-FMO problems with positively homogeneous objectives. This procedure begins by computing lexicographic solutions for the MO-FMO problem. After that, a set of non-dominated points is generated using




Figure 2: An example of a set of sample points. Each point is associated with a locally optimal BAC found by the deterministic local search. Axes correspond to the objective functions of the MO-FMO problem in (10). Dashed lines indicate the area dominated by each of the non-dominated points.

Algorithm 2: Phase 2: Efficient Set Generation.

Input: \mathcal{A}^* (Set of locally optimal approximately efficient BACs)

Output: $\hat{\mathcal{A}} \subseteq \mathcal{A}^*$ (Refined set of approximately efficient BACs)

begin

$\hat{\mathcal{A}} = \{\};$ // initially empty set of BACs

$\hat{X}_E = \{\};$ // initially empty set of fluence maps

foreach $\mathcal{A} \in \mathcal{A}^*$ **do**

$X_E^{\mathcal{A}} = \text{SolveMO-FMO}(\mathcal{A});$

$\hat{X}_E = \text{merge}(\hat{X}_E, X_E^{\mathcal{A}});$

$\text{Update}\hat{\mathcal{A}}(\hat{X}_E);$

return $\hat{\mathcal{A}};$

the ε -constraint method given a predefined set of bounds for each objective to be used as constraints during the optimisation process. Again, since the MO-FMO problem in (10) is a convex MO problem, treatment plans found by the ε -constraint method are truly efficient treatment plans for the corresponding BAC. Depending on the number of objectives we consider, the ε -constraint method might need to be performed several times minimising different objective functions. As a result of applying this procedure on a given BAC \mathcal{A} we have a set of efficient treatment plans $X_E^{\mathcal{A}}$. For each BAC \mathcal{A} , we merge both sets of fluence maps $X_E^{\mathcal{A}}$ and \hat{X}_E (*merge()* method in Algorithm 2) by eliminating those fluence maps that are dominated by another one in the union of the sets to be merged. After we merge these two sets, only potentially efficient fluence maps remain in the resulting set \hat{X}_E . The set \hat{X}_E is initialised empty. Since fluence maps in \hat{X}_E might change at each iteration, set $\hat{\mathcal{A}}$ is updated so only BACs that have at least one fluence map $x \in \hat{X}_E$ will be in $\hat{\mathcal{A}}$ (*Update $\hat{\mathcal{A}}$ ()* in Algorithm 2). Set $\hat{\mathcal{A}}$ approximates the actual set of efficient BACs of the MO-BAO problem in Equation (9).

For the example of Figure 2, set \mathcal{A}^* generated during the first phase will only consider the four non-dominated sample points that have dashed lines indicating their dominated area.

5 Computational Experiments

A prostate case has been considered in this study. This case is included in the CERR package (Deasy et al, 2003). Figure 3 shows this case. Boundaries of the target volume (tumour + margin), rectum and bladder (OARs) are highlighted as the regions of interest in this study. Table 1 shows the desired $gEUD$ values (eud_0^r column) for both tumour and OARs and the parameters for model z_1 . It also shows, for each OAR r , an upper bound U_+'' on $gEUD^r(x)$ below which (11) is guaranteed to be convex. These parameters are similar to the ones




Figure 3: Prostate case from CERR. Two OARs (bladder and rectum) are considered

used in Olafsson et al (2005). Different values might be needed in clinical practice. In this prostate case, the target volume has more than 7,000 voxels,

Table 1: Parameters for FMO models

r	a^r	ν^r	eud_0^r	U_+''
Tumour	-10	12	75	-
Rectum	8	8	50	63.77
Bladder	2	5	50	65.98

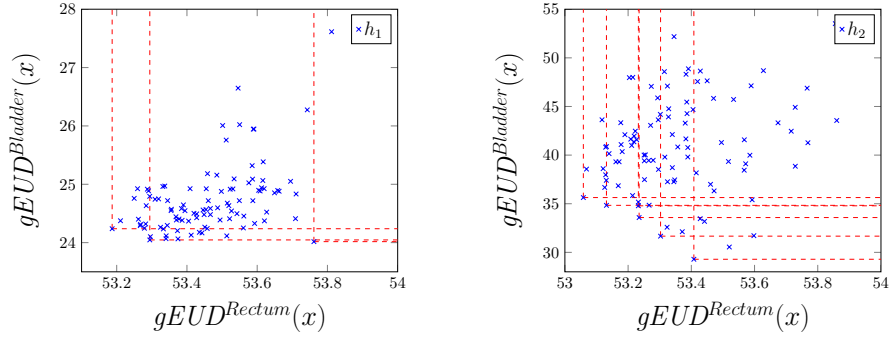
the rectum about 5,500 and the bladder around 9,500. Although the same multi-leaf collimator is used for all beam angles the number of beamlets considered in the optimisation process depends on each beam angle. This is because those beamlets that do not irradiate at least one voxel of the target are not considered as decision variables in the optimisation process. Thus, the number of decision variables (beamlets) depends on the BAC and ranges between 160 and 220. The number of beam angles N considered in a BAC is equal to 5. The dose deposition matrix A is given. We consider 360 beam angles, all of which are on the same plane. As we mentioned before, we use IPOPT as a solver for all non-linear optimisation problems.

During the first phase of our approach a set of locally optimal BACs and its corresponding set of fluence maps X^* is generated. The associated set of sample points is $\mathcal{S} = \{z(x) : x \in X^*\}$. As stated before, one can expect that different judgement functions lead to different sets \mathcal{S} . Figure 4 shows this situation. The axes correspond to the objectives of the MO-FMO problem (10). Because of

the positive homogeneity of the $gEUD$, as was explained in Section 3, one of the objectives (maximisation of the $gEUD$ of the tumour) was converted into an inequality constraint which we know to be binding and, consequently, it is not included as an axis.

Figure 4(a) shows the set \mathcal{S} when h_1 is used as judgement function. In this case there are only three non-dominated sample points in \mathcal{S}_N . Therefore only three BACs are passed on to the second phase of our approach. Table 2 lists these three BACs and the corresponding $gEUD$ values for the rectum and the bladder. Figure 4b shows the set \mathcal{S} when the $gEUD$ value of the rectum h_2 is used as a judgement function. In this case there are six non-dominated sample points in \mathcal{S}_N . Therefore, only six BACs are considered. These BACs are also listed in Table 2. In both cases, non-dominated sample points in \mathcal{S}_N represent only a small fraction ($\leq 6\%$) of the entire set \mathcal{S} .

Table 3 shows the set of (approximately) efficient BACs obtained by the



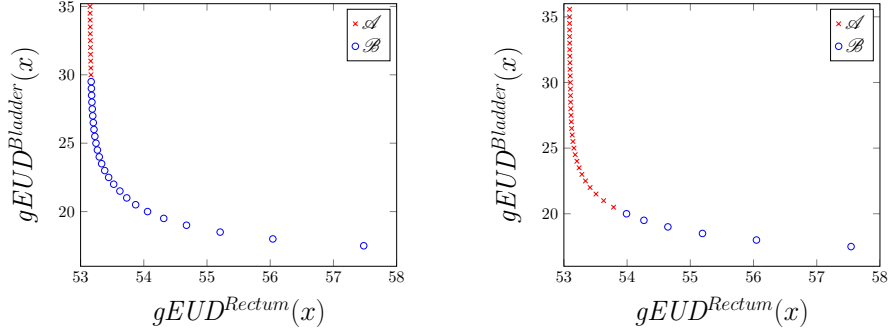
(a) Sample points found using judgement function h_1

(b) Sample points found using judgement function h_2

Figure 4: Sample points corresponding to the locally optimal BACs found by the local search algorithm for the prostate case using five beam angles

Table 2: BACs corresponding to the non-dominated sample points for judgement function h_1 and h_2 and the $gEUD$ values of the OARs (i.e. the sample points).

	\mathcal{A}_1	\mathcal{A}_2	\mathcal{A}_3	\mathcal{A}_4	\mathcal{A}_5	$gEUD^{Rectum}(x)$	$gEUD^{Bladder}(x)$
h_1	64°	85°	91°	201°	278°	53.76	24.01
	100°	151°	207°	269°	311°	53.29	24.04
	31°	91°	109°	136°	222°	53.18	24.23
h_2	104°	134°	136°	268°	298°	53.40	29.28
	78°	105°	168°	218°	274°	53.30	31.66
	91°	140°	161°	219°	343°	53.23	33.57
	49°	102°	126°	170°	221°	53.23	34.78
	37°	100°	131°	163°	269°	53.13	34.83
	49°	91°	126°	179°	255°	53.05	35.63



(a) Approximately non-dominated points obtained using h_1 . \mathcal{A} (crosses) corresponds to the first BAC in Table 3 ($100^\circ - 151^\circ - 207^\circ - 269^\circ - 311^\circ$) and \mathcal{B} (circles) corresponds to the second one

(b) Approximately non-dominated points obtained using h_2 . \mathcal{A} (crosses) corresponds to the first BAC in Table 3 ($37^\circ - 100^\circ - 131^\circ - 163^\circ - 269^\circ$) and \mathcal{B} (circles) corresponds to the second one

Figure 5: Final sets of (approximately) non-dominated points obtained using the two-phase approach and the resulting approximately efficient BACs

two-phase approach using judgement functions h_1 and h_2 introduced before. Table 3 also shows the number of non-dominated points (column NDP) for each (approximately) efficient BAC.

Results in Table 3 show also that (approximately) efficient BACs contribute five or more non-dominated points to the final set \hat{Y}_N . This means that those BACs with a significant number of non-dominated points $y \in \hat{Y}_N$ cover an important part of the approximation in the objective space of the FMO problem.

We can see this situation in Figure 5. We need to point out that we only found a finite number of non-dominated points (efficient solutions) of the MO-FMO problem (10) for each locally optimal BAC found during Phase 1. We want to recall the fact that the set of efficient BACs we found at the end of our algorithm is just an approximation of the truly efficient solution set of the MO-BAO problem, as it is not possible to evaluate the entire set of feasible BACs.

As we can see in Figure 6, neither does the set of non-dominated points generated using judgement function h_1 completely dominate the one generated by

Table 3: Final set of (approximately) efficient BACs for h_1 and h_2 and the number of non-dominated points (NDP) belonging to the corresponding BAC

Function	\mathcal{A}_1	\mathcal{A}_2	\mathcal{A}_3	\mathcal{A}_4	\mathcal{A}_5	NDP
h_1	100°	151°	207°	269°	311°	9
	31°	91°	109°	136°	222°	27
h_2	37°	100°	131°	163°	269°	26
	49°	91°	126°	179°	255°	11




Figure 6: Comparison between sets \hat{Y}_N obtained using h_1 and h_2

using judgement function h_2 nor does the set of non-dominated points generated using h_2 completely dominates the one generated using h_1 . Thus we cannot say one judgement function (either h_1 or h_2) is better than the other.

5.1 ε -dominance

It is clear that including more BACs during Phase 2 will increase our chances of producing a better approximation to the set of efficient BACs. However, solving the MO-FMO problem for the entire set of locally optimal BACs obtained during Phase 1 is quite time consuming. Moreover, our results above suggest that most of these BACs do not lead to any improvement in the final approximation obtained at the end of Phase 2. Although the dominance criterion considered in this study is very effective at finding those BACs that are part of the final approximation, it could happen that the sample point of a BAC \mathcal{A} that should be part of the final approximation is dominated by the sample point of some other BAC \mathcal{B} and consequently the BAC \mathcal{A} is not passed on to Phase 2. Hence we may not generate the best possible approximation given the set of locally optimal BACs generated during Phase 1. For this reason we need to find a way to include BACs that lead to improvements in the final approximation without making the number of BACs passed on to Phase 2 too large. We seek to achieve this using the concept of ε -dominance.

In the ε -dominance concept (Loridan, 1984; White, 1986) the main idea is to make the Pareto dominance criterion presented in Section 3 more (or less) restrictive. For this study, we want to modify the Pareto dominance test so that more sample points $s \in S$ are passed on to Phase 2. White (1986) and Laumanns et al (2002) define ε -dominance as follows: Let $f, g \in \mathbb{R}^p$ and $\varepsilon > 0$. Then, f is said to ε -dominate g , if

$$(1 - \varepsilon)f_r \leq g_r \text{ for all } r = 1, \dots, |R|. \quad (16)$$

We denote the set of sample points that are not ε -dominated by any other sample point as S_N^ε . Unfortunately using a single ε value is not possible in the context of radiotherapy due to the differences in magnitude among the objectives. Because of that, we consider a vector $\varepsilon \in \mathbb{R}^{|R|}$ and we slightly modify the definition in Equation 16 as follows: f is said to ε -dominate g if

$$(1 - \varepsilon_r)f_r \leq g_r \text{ for all } r = 1, \dots, |R|. \quad (17)$$

Table 4 shows the ε vectors used in this study to show how ε -dominance can be used to improve the final approximation obtained by the two-phase approach. In order to measure the improvement obtained when using each ε vector, we consider the well-known hypervolume metric proposed in Zitzler (1999). The hypervolume can be defined as the p -dimensional area below some reference point that is dominated by a set of points in objective space (Zitzler, 1999). The hypervolume of a set of non-dominated points is expressed as a percentage of the maximum hypervolume that can be covered. A set with a larger hypervolume is preferred as it is likely to present a better set of trade-offs than one with a smaller hypervolume. It is clear that if we set all ε_r , $r = 1, \dots, |R|$, large enough we would obtain the same set of BACs we would get if all sample points were passed to Phase 2, whereas for $\varepsilon_r = 0$ this will be the set of nondominated sample points as suggested in Algorithm 1.

As mentioned before Table 4 shows the hypervolume value obtained for five different ε vectors. For each ε vector, ε -values for rectum and bladder are shown (columns $\varepsilon_{gEUDRectum}$ and $\varepsilon_{gEUDBladder}$ respectively). $|\mathcal{A}^*|$ is the number of BACs in the set of locally optimal BACs \mathcal{A}^* found at the end of the Phase 1. $|\hat{\mathcal{A}}|$ is the number of BACs in the set of (approximately) efficient BACs $\hat{\mathcal{A}}$ of the MO-BAO problem.

As Table 4 shows, when h_2 is used as the judgement function, the final approximation does not improve by increasing values of ε . For this case, Pareto dominance is able to find those BACs that lead to the best possible approximation given the set of locally optimal BACs found in Phase 1, which is the same as shown in Figure 5b. A different situation occurs when we consider h_1 as the judgement function. As Tables 2 and 3 show, when h_1 is used as the judgement function three locally optimal BACs are passed on to Phase 2. Two of these three BACs are part of the final approximation obtained by the two-phase approach. If we now consider an ε large enough to include the all 100 locally optimal BACs we have that the number of BACs included in the final approximation is equal to five. Figure 7b shows the difference between the hypervolume obtained by the two-phase approach when using Pareto dominance and the one obtained when using the ε -dominance concept. Table 4 shows how the hypervolume increases as ε values increase and, consequently, the final approximation improves. In spite of that, the final difference between the best possible hypervolume value and the one obtained with $\varepsilon = 0$ is very small (less than 1% of the maximum hypervolume) which demonstrates the efficiency of our

Table 4: Different ε vectors and their corresponding hypervolumes for both judgement functions h_1 and h_2 . For each ε vector the number of BACs in the set of locally optimal BACs $|\mathcal{A}^*|$ and the number of (approximately) efficient BACs $|\hat{\mathcal{A}}|$ are shown.

Logistic Function (h_1)				
$\varepsilon_{gEUDRectum}$	$\varepsilon_{gEUDBladder}$	$ \hat{\mathcal{A}} $	$ \mathcal{A}^* $	Hypervolume
0	0	2	3	87.87%
0.00025	0.005	3	14	88.49%
0.00050	0.010	4	24	88.59%
0.00130	0.028	5	66	88.62%
≥ 0.015	≥ 0.16	5	100	88.62%
Lexicographic Function (h_2)				
$\varepsilon_{gEUDRectum}$	$\varepsilon_{gEUDBladder}$	$ \hat{\mathcal{A}} $	$ \mathcal{A}^* $	Hypervolume
0	0	2	5	89.55%
≥ 0.017	≥ 0.34	2	100	89.55%

approach in finding good quality BACs based only on non-dominated sample points obtained during the Phase 1.

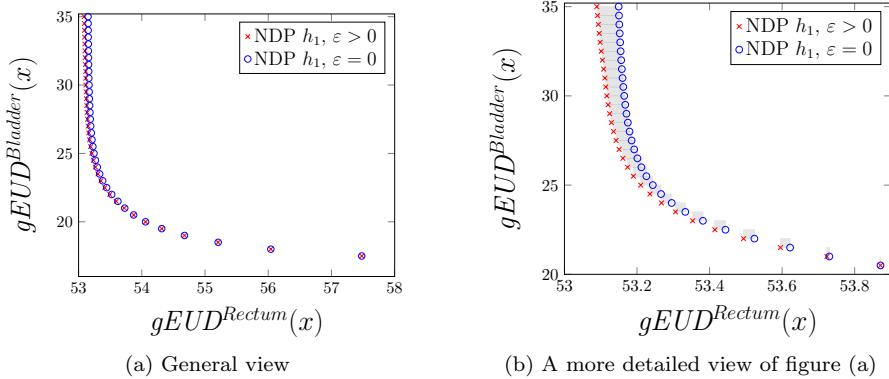


Figure 7: Comparison between final approximations using $\varepsilon_{gEUDRectum} = 0$, $\varepsilon_{gEUDBladder} = 0$ ($\varepsilon = 0$, circles) and $\varepsilon_{gEUDRectum} \geq 0.015$, $\varepsilon_{gEUDBladder} \geq 0.16$ ($\varepsilon > 0$, crosses) based on \mathcal{S} obtained using judgement function h_1 . Figure 7b is a more detailed view of Figure 7a that also shows the difference in the hypervolumes for these two cases.

6 Conclusions and Future Work

In this paper the MO-BAO problem is solved using a two-phase approach. During the first phase a set of locally optimal BACs is determined using a single

objective judgement function. Each locally optimal BAC is represented by a sample point in the objective space of the MO-FMO problem. The first phase ends by determining the BACs with pairwise non-dominated sample points among the set of locally optimal BACs. Then, during the second phase we generate a large set of non-dominated points of the MO-FMO problem for all the BACs found in the first phase. As a result, we obtain a set of (approximately) non-dominated points belonging to one or more BACs. These BACs are our approximation to the set of truly efficient BACs for the MO-BAO problem. This is substantially different from previous approaches where only one BAC is presented as the “best one” based only on its judgement function value. Unlike those single objective approaches, in this work once a set of locally optimal BACs is obtained by means of a single objective local search, a subset of them is selected based on a dominance criterion. Selected BACs are then further processed to solve their associated MO-FMO problem to determine a final approximation to the set of efficient BACs. An ε -dominance strategy was also considered to relax the dominance-based selection criterion allowing the algorithm to pass a larger number of promising locally optimal BACs onto the second phase, which leads to a better quality final set of approximately efficient BACs.

In our example, we have shown that the non-dominated points we generate are associated not with just a single BAC, but instead with a set of BACs. Thus, single-objective optimisation is not enough to find the “right BAC”. By considering multiple objectives during its search, our algorithm can better describe the trade-off between objectives and can therefore provide more (and better quality) alternatives for treatment planners to choose from.

6.1 Future Work

In future work, we intend to use MO judgement functions instead of single objective ones to find a set of good quality BACs. Using a MO judgement function would allow us to identify a set of potentially efficient BACs directly in a heuristic algorithm. In that case, MO strategies such as Pareto local search (Paquete et al, 2004) could be used instead of single objective approaches. Moreover, results show that different judgement functions lead to different sets of non-dominated points Y_N . Thus, other judgement functions should be evaluated using the two-phase approach. Also, other MO-BAO models such as dosimetric ones could be tested within our framework.

Positive initial results reported in this paper should motivate more experiments using clinical data. Moreover, testing our approach for different types of cancer would also be quite useful in order to validate the applicability of our approach, as location and shape of both tumours and OARs can affect the result obtained using our approach. Finally, other heuristic methods such as tabu search, evolutionary algorithms, and their MO variants could also be considered instead of the deterministic local search.

References

- Aleman D, Kumar A, Ahuja R, Romeijn H, Dempsey J (2008) Neighborhood search approaches to beam orientation optimization in intensity modulated radiation therapy treatment planning. *Journal of Global Optimization* 42(4):587–607
- Amols H, Ling C (2002) EUD but not QED. *International Journal of Radiation Oncology Biology Physics* 52(1):1–2
- Azizi-Sultan AS (2010) Automatic selection of beam orientations in intensity-modulated radiation therapy. *Electronic Notes in Discrete Mathematics* 36:127 – 134
- Bangert M, Ziegenhein P, Oelfke U (2012) Characterizing the combinatorial beam angle selection problem. *Physics in Medicine and Biology* 57(20):6707–6723
- Bertsimas D, Cacchiani V, Craft D, Nohadani O (2013) A hybrid approach to beam angle optimization in intensity-modulated radiation therapy. *Computers & Operations Research* 40(9):2187–2197
- Bortfeld T, Schlegel W (1993) Optimization of beam orientations in radiation therapy: some theoretical considerations. *Physics in Medicine and Biology* 38(2):291–304
- Breedveld S, Storchi P, Keijzer M, Heemink A, Heijmen B (2007) A novel approach to multi-criteria inverse planning for IMRT. *Physics in Medicine and Biology* 52:6339–6353
- Breedveld S, Storchi PRM, Heijmen BJM (2009) The equivalence of multi-criteria methods for radiotherapy plan optimization. *Physics in Medicine and Biology* 54(23):7199
- Breedveld S, Storchi PRM, Voet PWJ, Heijmen BJM (2012) icycle: Integrated, multicriterial beam angle, and profile optimization for generation of coplanar and noncoplanar IMRT plans. *Medical Physics* 39(2)
- Cabrera G G, Chica M, Ehrgott M, Mason A (2012) Mathematical programming and metaheuristic approaches applied to biological-based fluence map optimization in radiotherapy. In: Haywood J (ed) *Proceedings of the ORSNZ 46th Annual Conference Proceedings*, Wellington, New Zealand, pp 47–56
- Cabrera G G, Ehrgott M, Mason A, Philpott A (2014) Multi-objective optimization of positively homogeneous functions and an application in radiation therapy. *Operations Research Letters* 42:268–272
- Chankong V, Haimes Y (1983) *Multiobjective Decision Making Theory and Methodology*. Elsevier Science, New York

- Choi B, Deasy J (2002) The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning. *Physics in Medicine and Biology* 47(20):3579–3589
- Craft D (2007) Local beam angle optimization with linear programming and gradient search. *Physics in Medicine and Biology* 52:127–135
- Craft D, Halabi T, Shih HA, Bortfeld T (2007) An approach for practical multi-objective IMRT treatment planning. *International Journal of Radiation Oncology Biology Physics* 69(5):1600 – 1607
- Das S, Cullip T, Tracton G, Chang S, Marks L, Anscher M, Rosenman J (2003) Beam orientation selection for intensity-modulated radiation therapy based on target equivalent uniform dose maximization. *International Journal of Radiation Oncology Biology Physics* 55(1):215–224
- Deasy J, Blanco A, Clark V (2003) CERR: A computational environment for radiotherapy research. *Medical Physics* 30(5):979–985
- Deb K, Pratap A, Agarwal S, Meyarivan T (2002) A fast and elitist multiobjective genetic algorithm: Nsga-ii. *IEEE Transactions on Evolutionary Computation* 6(2):182–197
- Dias J, Rocha H, Ferreira B, Lopes M (2014) A genetic algorithm with neural network fitness function evaluation for IMRT beam angle optimization. *Central European Journal of Operations Research* 22(3):431–455
- Dirscherl T, Alvarez-Moret J, Bogner L (2011) Advantage of biological over physical optimization in prostate cancer? *Zeitschrift für Medizinische Physik* 21(3):228–235
- Djajaputra D, Wu Q, Wu Y, Mohan R (2003) Algorithm and performance of a clinical IMRT beam-angle optimization system. *Physics in Medicine and Biology* 48(19):3191
- Ehrgott M, Johnston R (2003) Optimisation of beam directions in intensity modulated radiation therapy planning. *OR Spectrum* 25:251–264
- Ehrgott M, Holder A, Reese J (2008) Beam selection in radiotherapy design. *Linear Algebra and its Applications* 428(5):1272–1312
- Ehrgott M, Güler C, Hamacher HW, Shao L (2009) Mathematical optimization in intensity modulated radiation therapy. *Annals of Operations Research* 175(1):309–365
- Eichfelder G (2008) Adaptive scalarization methods in multiobjective optimization. Springer-Verlag, Berlin
- Fiege J, McCurdy B, Potrebko P, Champion H, Cull A (2011) Pareto: A novel evolutionary optimization approach to multiobjective IMRT planning. *Medical Physics* 38(9)

- Haimes YY, Lasdon L, Wismer D (1971) On a bicriterion formulation of the problems of integrated system identification and system optimization. *IEEE Transactions on Systems, Man and Cybernetics* 1(3):296–297
- Hamacher H, Küfer KH (2002) Inverse radiation therapy planning a multiple objective optimization approach. *Discrete Applied Mathematics* 118(1):145–161
- Kalantzis G, Apte A (2014) A novel reduced-order prioritized optimization method for radiation therapy treatment planning. *IEEE Transactions on Biomedical Engineering* 61(4):1062–1070
- Küfer KH, Scherrer A, Monz M, Alonso F, Trinkaus H, Bortfeld T, Thiele C (2003) Intensity-modulated radiotherapy a large scale multi-criteria programming problem. *OR Spectrum* 25(2):223–249
- Lahanas M, Schreibmann E, Milickovic N, Baltas D (2003) Intensity modulated beam radiation therapy dose optimization with multiobjective evolutionary algorithms. In: Fonseca C, Fleming P, Zitzler E, Thiele L, Deb K (eds) *Evolutionary Multi-Criterion Optimization*, Lecture Notes in Computer Science, vol 2632, Springer Berlin Heidelberg, pp 648–661
- Laumanns M, Thiele L, Deb K, Zitzler E (2002) Combining convergence and diversity in evolutionary multiobjective optimization. *Evolutionary Computation* 10(3):263–282
- Li Y, Yao D (2006) Accelerating the radiotherapy planning with a hybrid method of genetic algorithm and ant colony system. In: Jiao L, Wang L, Gao X, Liu J, Wu F (eds) *Advances in Natural Computation*, Lecture Notes in Computer Science, vol 4222, Springer Berlin Heidelberg, pp 340–349
- Li Y, Yao J, Yao D (2004) Automatic beam angle selection in IMRT planning using genetic algorithm. *Physics in Medicine and Biology* 49(10):1915
- Li Y, Yao D, Chen W (2005a) A particle swarm optimization algorithm for beam angle selection in intensity-modulated radiotherapy planning. *Physics in Medicine and Biology* 50(15):3491–3514
- Li Y, Yao D, Chen W, Zheng J, Yao J (2005b) Ant colony system for the beam angle optimization problem in radiotherapy planning: A preliminary study. In: Corne D (ed) *Proceedings of the IEEE Congress on Evolutionary Computation (CEC)*, IEEE, pp 1532–1538
- Lim G, Cao W (2012) A two-phase method for selecting IMRT treatment beam angles: Branch-and-prune and local neighborhood search. *European Journal of Operational Research* 217(3):609–618
- Lim G, Kardar L, Cao W (2014) A hybrid framework for optimizing beam angles in radiation therapy planning. *Annals of Operation Research* 217:357–383

- Loridan P (1984) ε -solutions in vector minimization problems. *Journal of Optimization Theory and Applications* 43(2):265–276
- Miettinen K (1999) *Nonlinear Multiobjective Optimization*, International Series in Operations Research and Management Science, vol 12. Kluwer Academic Publishers, Dordrecht
- Mišić V, Aleman D, Sharpe M (2010) Neighborhood search approaches to non-coplanar beam orientation optimization for total marrow irradiation using IMRT. *European Journal of Operational Research* 205(3):522–527
- Nahum A, Uzan J (2012) (Radio)biological optimization of external-beam radiotherapy. *Computational and Mathematical Methods in Medicine* 1:13
- Niemierko A (1997) Reporting and analyzing dose distributions: A concept of equivalent uniform dose. *Medical Physics* 24(1):103–110
- Olafsson A, Jeraaj R, Wright S (2005) Optimization of intensity-modulated radiation therapy with biological objectives. *Physics in Medicine and Biology* 50(22):5357–5379
- Paquete L, Chiarandini M, Stützle T (2004) Pareto local optimum sets in the biobjective traveling salesman problem: An experimental study. In: Gandibleux X, Sevaux M, Srensen K, Tkindt V (eds) *Metaheuristics for Multiobjective Optimisation*, Lecture Notes in Economics and Mathematical Systems, vol 535, Springer Berlin Heidelberg, pp 177–199
- Preciado-Walters F, Rardin RL, Langer M, Thai V (2004) A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer. *Mathematical Programming* 101(2):319–338
- Preciado-Walters F, Langer M, Rardin R, Thai V (2006) Column generation for IMRT cancer therapy optimization with implementable segments. *Annals of Operations Research* 148(1):65–79
- Pugachev A, Xing L (2001) Pseudo beams-eye-view as applied to beam orientation selection in intensity-modulated radiation therapy. *International Journal of Radiation Oncology Biology Physics* 51(5):1361–1370
- Pugachev A, Xing L (2002) Incorporating prior knowledge into beam orientation optimization in IMRT. *International Journal of Radiation Oncology Biology Physics* 54(5):1565–1574
- Pugachev A, Boyer A, Xing L (2000) Beam orientation optimization in intensity-modulated radiation treatment planning. *Medical Physics* 27(6):1238–1245
- Pugachev A, Li JG, Boyer A, Hancock S, Le Q, Donaldson S, Xing L (2001) Role of beam orientation optimization in intensity-modulated radiation therapy. *International Journal of Radiation Oncology Biology Physics* 50(2):551–560

- Romeijn H, Dempsey J (2008) Intensity modulated radiation therapy treatment plan optimization. *TOP* 16(2):215–243
- Romeijn H, Dempsey J, Li J (2004) A unifying framework for multi-criteria fluence map optimization models. *Physics in Medicine and Biology* 49(10):1991–2013
- Rowbottom C, Webb S, Oldham M (1998) Beam orientation optimization in intensity-modulated radiation treatment planning. *Medical Physics* 25(7):1171–1179
- Schreibmann E, Lahanas M, Xing L, Baltas D (2004) Multiobjective evolutionary optimization of the number of beams, their orientations and weights for intensity-modulated radiation therapy. *Physics in Medicine and Biology* 49(5):747
- Stein J, Mohan R, Wang X, Bortfeld T, Wu Q, Preiser K, Ling C, Schlegel W (1997) Number and orientation of beams in intensity-modulated radiation treatments. *Medical Physics* 24:149–160
- Thomas E, Chapet O, Kessler M, Lawrence T, Ten Haken R (2005) Benefit of using biologic parameters (EUD and NTCP) in IMRT optimization for treatment of intrahepatic tumors. *International Journal of Radiation Oncology Biology Physics* 62(2):571–578
- Wächter A, Biegler L (2006) On the implementation of a primal-dual interior point filter line search algorithm for large-scale nonlinear programming. *Mathematical Programming* 106(1):25–57
- Webb S (2010) *The Physics of Three Dimensional Radiation Therapy: Conformal Radiotherapy, Radiosurgery and Treatment Planning*. Series in Medical Physics and Biomedical Engineering, Taylor & Francis
- White D (1986) Epsilon efficiency. *Journal of Optimization Theory and Applications* 49(2):319–337
- Wu Q, Mohan R, Niemierko A (2000) IMRT optimization based on the generalized equivalent uniform dose (EUD). In: Enderle J (ed) *Engineering in Medicine and Biology Society, 2000. Proceedings of the 22nd Annual International Conference of the IEEE, IEEE*, vol 1, pp 710–713
- Wu Q, Mohan R, Niemierko A, Schmidt-Ullrich R (2002) Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose. *International Journal of Radiation Oncology Biology Physics* 52(1):224–235
- Wu Q, Djajaputra D, Wu Y, Zhou J, Liu H, Mohan R (2003) Intensity-modulated radiotherapy optimization with gEUD-guided dose-volume objectives. *Physics in Medicine and Biology* 48(3):279–291
- Zitzler E (1999) *Evolutionary Algorithms for Multiobjective Optimization: Methods and Applications*. PhD thesis, ETH Zurich, Switzerland