A two-phase method for selecting IMRT treatment beam angles: Branch-and-Prune and local neighborhood search

Gino J. Lim[a,*], Wenhua Cao[1]

[a]*Department of Industrial Engineering, University of Houston, 4800 Calhoun Road, Houston, TX 77204,*

Abstract

This paper presents a new two-phase solution approach to the beam angle and fluence map optimization problem in *Intensity Modulated Radiation Therapy* (IMRT) planning. We introduce *Branch-and-Prune* (B&P) to generate a robust feasible solution in the first phase. A local neighborhood search algorithm is developed to find a local optimal solution from the Phase I starting point in the second phase. The goal of the first phase is to generate a clinically acceptable feasible solution in a fast manner based on a Branch-and-Bound tree. In this approach, a substantially reduced search tree is iteratively constructed. In each iteration, a merit score based branching rule is used to select a pool of promising child nodes. Then pruning rules are applied to select one child node as the branching node for the next iteration. The algorithm terminates when we obtain a desired number of angles in the current node. Although Phase I generates quality feasible solutions, it does not guarantee optimality. Therefore, the second phase is designed to converge Phase I starting solutions to local optimality. Our methods are tested on two sets of real patient data. Results show that not only can B&P alone generate clinically acceptable solutions, but the two-phase method consistently generates local optimal solutions, some of which are shown to be globally optimal.

Keywords: IMRT, Beam Angle Optimization, Fluence Map Optimization, Radiation Treatment, Local Neighborhood Search

*Corresponding author

Email addresses: ginolim@uh.edu (Gino J. Lim), wcao@mail.uh.edu (Wenhua Cao)

1. Introduction

Intensity Modulated Radiation Therapy (IMRT) is a high-precision radiotherapy that has been extensively used to treat a range of cancer types since the 1990s [1]. As an advanced form of the state-of-the-art three-dimensional conformal radiation treatment (3DCRT), IMRT is able to deliver radiation that precisely conforms to the three-dimensional shape of the tumor while sparing the surrounding normal tissues. While a patient is lying on a treatment couch, a linear accelerator (LINAC) generates photon beams to irradiate the tumor target. The movable arm, called the *gantry*, of the LINAC can rotate 360° in a plane perpendicular to the couch. It stops at one angle to deliver multiple shapes of radiation beams (*beam apertures*) and moves on to the next angle. This paper considers *coplanar* beams where the only movement in the treatment session is made by rotating the gantry. Other forms of movement of the LINAC and treatment couch are possible (see [2, 3, 4] for *non-coplanar* beam treatment planning). Note that the *multi-leaf collimator* (MLC) is capable of changing the beam shapes by moving its "leaves" back and forth to block a portion of each beam's radiation dose. Therefore, various two-dimensional beam shapes can be constructed to perform intensity modulation. The MLC functions by dividing its *aperture* into rectangular grids and thus partitioning each beam into *sub-beams*, also called *beamlets* or *pencil beams*.

Before an IMRT plan is designed, physicians create a precise geometry of the tumor. Several scanning procedures are used, such as computed tomography (CT), magnetic resonance imaging (MRI), or positron emission tomography (PET). Those diagnostic images help physicians determine the three-dimensional shape of the tumor and surrounding normal tissues. Conventionally, there are three types of volumes considered: *planning target volume* (PTV), *organs-at-risk* (OARs) or *critical structures*, and normal tissues. The PTV includes the suspected cancerous area; the OARs are organs located close to the PTV and all remaining portions are normal tissues. The physician creates a *prescription* to treat the tumor while limiting radiation on the OARs and the normal tissues. These are conflicting goals, and one may expect a tradeoff between these objectives. Thereby, upper and lower bounds of radiation doses are employed to guarantee a successful treatment plan according to the radiation treatment planner's experience. Optimiza-

tion methods are often used to find a clinically acceptable solution that is a compromise between the two objectives.

Typically, the IMRT treatment planning optimization system contains three sequential problems: beam angle optimization (BAO), fluence map optimization (FMO), and beam segmentation (BS). The BAO problem is to find a subset $\mathcal{A}^*$ of a given set $\mathcal{A}$ of candidate beam angles ($\mathcal{A}^* \subset \mathcal{A}$) to deliver a satisfactory treatment. FMO solves for an optimal beam intensity map, also called *fluence map*, for each treatment angle. The optimal FMO solution illustrates the treatment plan quality based on the angle selection and the other treatment planning parameters. The BS problem is to find an optimal schedule for MLC to fulfill the delivery of fluence maps. The theme of this paper is to solve the BAO problem. Since the optimal beam angle set is inherently related to the optimal fluence map for the selected angle set, our proposed approaches focus on obtaining solutions for both BAO and FMO, simultaneously.

Various optimization models have been developed for optimizing beam angles and fluence intensity maps in the literature [5, 6, 7]. In particular, a comprehensive survey on optimization models in IMRT has been provided by Erghott et al. [8]. Finding an optimal set of beam angles is a combinatorial optimization problem [9]. However, obtaining global optimality using an exact solution algorithm such as Branch and Bound is practically unreasonable. Therefore, previous research in the literature has concentrated on clever heuristics to find good beam angles efficiently. Pugachev and Xing [10] applied the approach of beam's-eye-view (BEV) to score candidate beam angles with their dose characteristics on PTV and OARs to select quality beam angles. D'Souza et al. [11] utilized the mean OAR dose (MOD) concept to introduce an integer programming (IP) model to solve for optimal beam angles.

Methods introduced in recent studies for the BAO problem often use objective values of exact FMO models in search of optimal beam angles [6, 7, 12, 13, 14, 15]. FMO models in the literature are either linear or nonlinear. The objective functions generally minimize the deviations from the prescribed dose on the PTV and the dose delivered to surrounding normal tissues. In our proposed method, an LP FMO model with a dose-and-organ-based objective function is adopted to evaluate beam angle selections in such a way that the function penalizes the portion of organs deviating from its desired radiation dose [6].

In that BAO is fundamentally a non-convex problem and multiple local

optima may exist, the majority of published studies on the BAO problem have focused on global optimization algorithms including simulated annealing [2, 16, 17, 18], genetic algorithms [14, 19, 20, 21], exhaustive search [3, 22, 23], evolutionary algorithms [15], particle swarm algorithms [24] and response surface method [25]. Although those global algorithms can avoid local optima theoretically, globally optimal or even "good" solutions can not be guaranteed unless a significant number of iterations are allowed. Local search algorithms have also been developed in [12, 26]. Aleman et al. [12] discussed an angle proximity based discrete neighborhood search algorithm and Craft [26] discussed an approximation based continuous gradient search algorithm. However, local algorithms may be restricted by their starting points, and hence, they may require many starting points to guarantee good solutions. Therefore, in our approach, we first develop a robust feasible solution generation approach whose solution quality is near local optimal. We then develop a local neighborhood search algorithm to guarantee local optimality.

Lim *et al.* [6] discussed an *LP based iterative beam angle elimination* (IBAELP) algorithm to obtain quality beam angles with optimized fluence maps. Based on the FMO solution, IBAELP iteratively eliminates one beam angle whose merit score is the minimum in each iteration until a reduced beam angle set is obtained. They use a MIP model with the reduced beam angle set to obtain the final solution with the desired number of beam angles. However, IBAELP's major drawback is that it tends to lose accuracy as the number of candidate beam angles increases. Selecting six angles from 12 candidate angles may give a better result than selecting six out of 36 (see Table 3 as an example). One intuitive way to remedy the situation is to examine the outcome of eliminating a beam angle one step ahead before making the angle elimination decision. Therefore, we introduce a new solution approach that outperforms the IBAELP algorithm in terms of strengthening robustness and solution quality. Our solution approach consists of two sequential phases: *Branch-and-Prune* and *Local Neighborhood Search*. A new robust feasible solution generation method *Branch-and-Prune* is used in the first phase. A local search algorithm is then applied in the second phase to further improve the feasible solution to local optimality. In our proposed neighborhood search algorithm, we use a reduced neighborhood definition similar to [12]. Furthermore, the approach in [12] stops each iteration once the first improving angle set in the neighborhood is found, whereas the algorithm presented here performs an exhaustive search in the neighborhood and then selects the best performing solution.

The rest of this paper is organized as follows. Section 2 describes two optimization models that are adopted from [6]. One is an LP model for optimizing a fluence map (FMO), and the other is a MIP model for finding the optimal beam angle set and the fluence maps simultaneously (BAO/FMO). Sections 3 and 4 present our proposed two-phase algorithms. First, *Branch-and-Prune* is described in Section 3. To further improve the feasible solution, a local neighborhood search method is developed to find local optimal solutions in Section 4. We reveal the advantages of our approach through numerical experiments on two cancer cases with real patient data in Section 5. We conclude the paper in Section 6. Although the B&P method is based on a linear FMO model in this paper, with a minor change, our algorithm can be easily applied to others.

2. Optimization Models

The primary decision variables of the optimization models are the beamlet weights $\omega_{a,l,p}$, which are the intensities of radiation. Specifically $\omega_{a,l,p}$ is delivered by the beam angle a and through the $(l,p)^{th}$ beamlet. Thus, the total dose $D_{x,y,z}$ deposited in voxel (x,y,z) is

$$D_{x,y,z} = \sum_{a,l,p} (\omega_{a,l,p} \cdot d_{x,y,z,a,l,p}), \tag{1}$$

where $d_{x,y,z,a,l,p}$ denotes the dose contribution to voxel (x,y,z) from beamlet (a,l,p) at unit weight. Table 1 provides a complete list of model parameters used in this paper. We define D_T as the collection of doses $D_{x,y,z}$ for $(x,y,z) \in T$, D_S for $(x,y,z) \in S$, D_N for $(x,y,z) \in N$.

Note that a voxel is a discrete cube of treatment volume, and our approach for constructing voxels is described in [27]. Due to the *streaking* issue in IMRT treatment planning, we impose a dose upper bound on normal tissues that are not close to the tumor region. Let $\bar{N}$ be defined as the modified N that excludes voxels that are close to the PTV, and let $D_{\bar{N}}$ be the collection of doses $D_{x,y,z}$ for $(x,y,z) \in \bar{N}$. We consider normal voxels that are three voxel layers beyond the PTV belonging to $\bar{N}$ in this paper. Based on these definitions, we discuss two optimization models for radiation therapy planning in the following sub-sections.

2.1. Linear Programming Model

Our proposed heuristic (see Section 3) chooses to solve an LP FMO model with different beam angle sets, $\mathcal{A}'$ ($\mathcal{A}' \subseteq \mathcal{A}$), iteratively to find a promising

Notation	Definition
$\mathcal{A}$	A set of candidate beam angles
T	A set of voxels in planning target volume (PTV)
S	A set of voxels in organ-at-risk (OAR)
N	A set of voxels in normal tissues
$\bar{N}$	A subset of N that excludes voxels close to PTV
θ_L	Cold spot control parameter on PTV
θ_U	Hot spot control parameter on PTV
ϕ	Hot spot control parameter on OAR
η	Maximum number of treatment beam angles
L_T	Lower reference bound on PTV
U_T	Upper reference bound on PTV
$U_{\bar{N}}$	Upper reference bound on normal structure
λ_t^+	Penalty coefficient for hot spots on PTV
λ_t^-	Penalty coefficient for cold spots on PTV
λ_s	Penalty coefficient for hot spots on OAR
λ_n	Penalty coefficient for normal structure

Table 1: Optimization model parameters

beam angle set. The LP allows us to evaluate the treatment plan quality given by any subset of beam angles, $\mathcal{A}'$. Using the calculated upper bounds ($M_{a,l,p}$) for beamlet weights given in [6], our LP formulation for solving FMO is

$$\begin{aligned} \min_{\omega} \quad & f(D) \\ & L_T \leq D_T \leq U_T \\ & D_{\bar{N}} \leq U_{\bar{N}} \\ & 0 \leq \omega_{a,l,p} \leq M_{a,l,p} \quad \forall a \in \mathcal{A}', \;\; l = 1,2,...,m, \;\; p = 1,2,...,n. \end{aligned} \tag{2}$$

Note that $\omega_{a,l,p} = 0$ if $a \notin \mathcal{A}'$. The objective function of the LP model estimates the quality of a treatment plan. Three basic criteria for evaluating a treatment plan include *conformity*, *uniformity* and *homogeneity* [27]. We use a penalty based objective function [6]:

$$\begin{aligned} f(D) \;=\; & \lambda_t^+ \parallel (D_T - \theta_U \cdot e_T)_+ \parallel_\infty \; + \lambda_t^- \parallel (\theta_L \cdot e_T - D_T)_+ \parallel_\infty \\ & + \frac{\lambda_s \parallel (D_S - \phi \cdot e_S)_+ \parallel_1}{\mid S \mid} \; + \; \frac{\lambda_n \parallel D_N \parallel_1}{\mid N \mid}, \end{aligned} \tag{3}$$

where $(\cdot)_+$ represents $max\{\cdot, 0\}$ and e_T and e_S represent the vectors of ones. For this objective function, maximum violations on the PTV are penalized while average violations are penalized on the OARs and other normal tissues. In the constraints, a lower bound is imposed on T, and upper bounds are imposed on T and $\bar{N}$. Although assigning upper bounds on OARs could be useful in certain cancer cases, the constraints on S are omitted here because they did not seem to contribute to solution quality in our tests. Model (2) can be easily expanded to a mixed integer programming model for selecting the optimal angle set.

2.2. Mixed Integer Programming Model

The MIP model is designed to select η beam angles out of a given set of angles $\mathcal{A}$, $|\mathcal{A}| >> \eta$, and to provide corresponding optimal beamlet weights $\omega_{a,l,p}$. We add binary variables ψ_a to LP (2), where ψ_a denotes whether the angle a is selected or not. The corresponding MIP formulation is

$$\begin{array}{lrcll} \min\limits_{\omega,\psi} & f(D) & & & \\ & L_T & \leq & D_T \leq U_T & \\ & D_{\bar{N}} & \leq & U_{\bar{N}} & \\ & \sum\limits_{a\in\mathcal{A}} \psi_a & \leq & \eta & \\ & 0 & \leq & \omega_{a,l,p} \leq M_{a,l,p} \cdot \psi_a & \forall a \in \mathcal{A},\ \ l = 1,2,...,m,\ \ p = 1,2,...,n \\ & \psi_a & \in & \{0,1\} & \forall a \in \mathcal{A}. \end{array} \tag{4}$$

Note that solving the MIP model (4) can be computationally challenging if there are a large number of candidate beam angles, $\mathcal{A}$. This motivates us to develop a heuristic algorithm for selecting treatment angles. In the next section, we introduce a *Branch and Prune* algorithm that utilizes both a score function and a linear programming objective value for branching and pruning the solution search tree.

3. Branch and Prune for Selecting Treatment Angles

3.1. Branch and Prune (B&P)

B&P is an iterative IP Branch-and-Bound tree node elimination algorithm that is developed for finding a good feasible solution for an MIP model (4). We describe a general B&P procedure using Figure 1. The branching rule

is based on a score function while the pruning rule is based on the objective value of the LP relaxation.

Suppose we are given a finite angle set $\mathcal{A} = \{a_1, a_2, \ldots, a_{12}\}$ and the value of $\eta = 6$. Our goal is to select a subset $\bar{\mathcal{A}} \subseteq \mathcal{F}$ and $|\bar{\mathcal{A}}| = 6$, where $\mathcal{F}$ is a set of feasible subsets of $\mathcal{A}$, that minimizes the objective function. The B&P tree is constructed starting from the root node at iteration 1. The root node is associated with a linear programming (LP) relaxation of the MIP model with $\mathcal{A}^{(1)} = \{a_1, a_2, \ldots, a_{12}\}$.

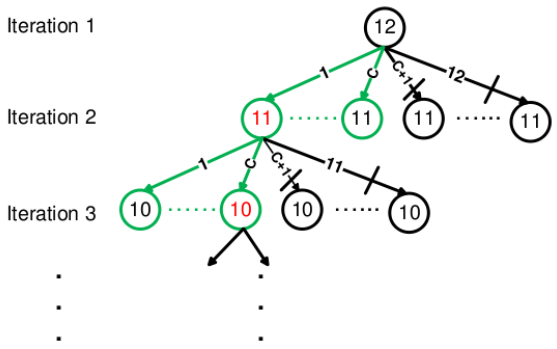

Figure 1: An overview of Branch and Prune

3.2. Branching Rule: Merit Score Function

In Figure 1, the B&P process starts from the root node. The number 12 at the root node represents that the LP relaxation problem with all 12 elements (or angles) is solved in iteration 1. In iteration 2, if we remove one element from the root node at a time, there are 12 possible branches in such a way that each child node consists of the remaining 11 elements. The number associated with the arc indexes the element that will be removed from $\mathcal{A}^{(1)}$ to form a new child node with $\mathcal{A}^{(2)}$. Therefore, child node j is associated with a set $\mathcal{A}^{(2)} = \{\mathcal{A}^{(1)} \backslash \{a_j\}\}$, $j = 1, 2, \ldots, 12$. (In a general form, $\mathcal{A}^{(i+1)} = \{\mathcal{A}^{(i)} \backslash \{a_j\}\}$, where i represents iteration number and j represents beam angle index.) In iteration 3, eleven branches are constructed for the selected node in iteration 2. Each branch has a node with 10 elements. This branching continues until the leaf node contains a reduced number of elements. For example, eight elements are contained in each node in the later iteration in the B&P process.

Ideally, we would test all nodes that are considered at the current iteration. But there will be a substantial burden on computation because each

node corresponds to an LP relaxation problem of the MIP model (4). Thus, we select a subset of nodes for branching and the rest of nodes are pruned without further examination. The selection of nodes to examine is made by a *merit score function.* In the example tree shown in Figure 1, once the LP relaxation problem on the root node with 12 angles is solved, merit scores are calculated for the 12 angles. The merit scores are sorted in ascending order, i.e., $\mathcal{S}_1 \leq \mathcal{S}_2 \leq \cdots \leq \mathcal{S}_{12}$. Then the first child node with 11 angles is formed by eliminating the lowest scored angle from the root node; the second child node is formed by eliminating the second lowest scored angle; and so forth. If we are given the value of $c = 3$, only the first three child nodes will remain in the tree for a further examination and the remaining nine nodes are pruned. We now describe our merit score function used for our application.

Beam Angle Score Function for Branching. A beam angle score function is introduced for evaluating a particular angle's merit to its resulting treatment. Beyond solving (2) with a beam angle set $\mathcal{A}'$, $\mathcal{A}' \subseteq \mathcal{A}$, the total amount of dose deposited on the PTV, the OARs, and the normal structure by beam angle a is

$$D(\Omega)_a = \sum_{(x,y,z)\in\Omega} \sum_{l,p} (\omega_{a,l,p} \cdot d_{x,y,z,a,l,p}) \qquad a \in \mathcal{A}',\ \Omega \in \{T, S, N\}. \quad (5)$$

We obtain three sub-scores: a contribution score s_a^T for PTV, an average non-negative dose score s_a^S for OAR and s_a^N for the normal structure. These values are based on a given beam angle and are defined as

$$s_a^T = \frac{D(T)_a}{uD(T)_a}, \qquad s_a^S = \frac{D(S)_a}{C_a^S}, \qquad s_a^N = \frac{D(N)_a}{C_a^N}, \quad a \in \mathcal{A}'. \quad (6)$$

where C_a^S and C_a^N are the numbers of voxels on the OAR and the normal structure, respectively, that receive a positive dose by the beam angle a; and $uD(T)_a$ is the total PTV dose from angle a by a uniform weight of 1. Then the merit score of beam angle a ($\mathcal{S}_a$ for $a \in \mathcal{A}'$) is defined as a linear combination of three normalized sub-scores in (6):

$$\mathcal{S}_a = \left(\frac{s_a^T}{\sum_{a\in\mathcal{A}'} s_a^T}\right) - \kappa_S \left(\frac{s_a^S}{\sum_{a\in\mathcal{A}'} s_a^S}\right) - \kappa_N \left(\frac{s_a^N}{\sum_{a\in\mathcal{A}'} s_a^N}\right), \quad 0 \leq \kappa_S, \kappa_N \leq 1. \quad (7)$$

This value measures the quality of each beam angle that will be potentially used. Beam angles with a better quality are identified if they have

higher score function values. The weighting coefficients, $\kappa_{_S}$ and $\kappa_{_N}$, are used as avoidance factors for OARs and normal tissues. For example, if the treatment planner emphasizes a low dose on an OAR, the corresponding $\kappa_{_S}$ should be set to a higher value. Obtaining good values for $\kappa_{_S}$ and $\kappa_{_N}$ requires a trial-and-error process. Although it is not a trivial task to determine appropriate weighting coefficients, the score function (7) provides an important feature to investigate the relationship of various treatment plans in terms of beam angle sets (See *Lemma 3.1*).

Lemma 3.1. *The average beam angle score is determined only by the size of a beam angle set, not by the angles that are contained in* $\mathcal{A}' = \{a_1, a_2, ..., a_\eta\}$. *That is, for any* $\mathcal{A}'$, *we have*

$$\bar{\mathcal{S}}_a = \frac{1 - \kappa_{_S} - \kappa_{_N}}{|\mathcal{A}'|}, \qquad a \in \mathcal{A}', \;\; \mathcal{A}' \neq \emptyset. \tag{8}$$

Proof. Let $\mathcal{A}' \subseteq \mathcal{A}$ be nonempty. Then,
$\sum\limits_{a\in\mathcal{A}'} \mathcal{S}_a = \Big(\sum\limits_{a\in\mathcal{A}'} s_a^T \Big/ \sum\limits_{a\in\mathcal{A}'} s_a^T\Big) - \kappa_{_S}\Big(\sum\limits_{a\in\mathcal{A}'} s_a^S \Big/ \sum\limits_{a\in\mathcal{A}'} s_a^S\Big) - \kappa_{_N}\Big(\sum\limits_{a\in\mathcal{A}'} s_a^N \Big/ \sum\limits_{a\in\mathcal{A}'} s_a^N\Big) =$
$1 - \kappa_{_S} - \kappa_{_N}$,
and the average of beam angle scores is $\bar{\mathcal{S}}_a = \frac{\sum\limits_{a\in\mathcal{A}'} \mathcal{S}_a}{|\mathcal{A}'|} = \frac{1-\kappa_{_S}-\kappa_{_N}}{|\mathcal{A}'|}$. □

We make the following remark based on Lemma 3.1.

Remark 3.1. *The average beam angle score* $\bar{\mathcal{S}}_a$ *decreases when the size of angle set* $|\mathcal{A}'|$ *increases. It is reasonable that if more treatment angles are used, then treatment quality improves.*

The scoring measure of an individual beam angle is crucial to B&P. The scores are calculated based on the LP (FMO) relaxation. The more accurately the score reflects the corresponding angle's contribution to a treatment plan, the more likely we may obtain a better final solution from the beam angle elimination process. The mean and standard deviation of angle scores in each iteration of the BP-C with $c = 1$ procedure (described in Section 3.4) are shown in Figure 2. This example is based on the prostate cancer case with 36 candidate beam angles. Firstly, the average angle score increases monotonically as the size of the angle set decreases (*Lemma 3.1*). Secondly, the standard deviations of angle scores fluctuates toward the end of the iterations, which indicates that the variation of scores (or significance) of individual beam angles turns notable if the size of the angle set becomes

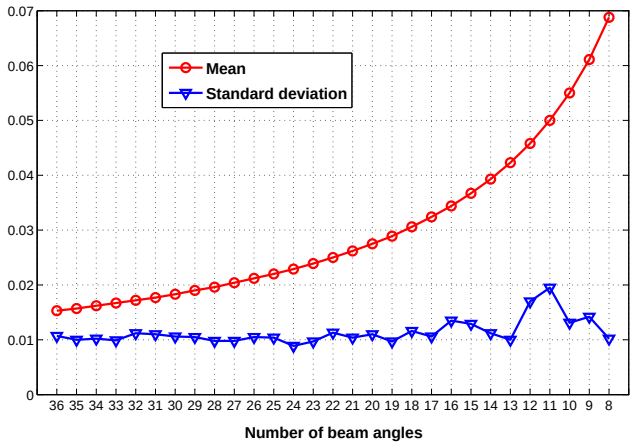

Figure 2: An example of beam angle score statistics during the beam angle elimination process

small. This example represents a general pattern of angle scores based on our observation of several cases. It is insufficient to compare one set of beam angles from another solely based on either beam angle scores or their average value, since one beam angle may yield various levels of contribution in different $\mathcal{A}'$s. Therefore, just because one beam angle scored less in one set of beam angles, it does not necessarily mean that the angle will have a lower score in another set. This leads us to use the objective function value of the optimization model to compare the quality of different angle sets and select the next branching node whose objective value is lowest among those in $\mathcal{A}'$.

3.3. Pruning Rule: LP Objective Value

B&P is designed to select only one node from the current iteration with c nodes for branching. This selection is made by comparing objective values of the corresponding LP relaxations. A greedy approach is used. Since our goal is to minimize the objective value, the node with the lowest objective value is selected as the branching node. Suppose that the objective value of the first node is smaller than the rest. Then, the first node is selected as the branching node and the remaining two nodes are pruned.

A question here is "how many nodes, $c \in \{1, 2, \ldots, |\mathcal{A}'|\}$, should be examined at each iteration?" If c is too large, it increases the computation time, while a small c is prone to prematurely removing optimal angles before reaching the final solution. However, if more nodes are examined in each iteration,

there is a better opportunity to obtain good interim solutions. Therefore, two branching rules are explored for selecting c: *BP-C* (see Section 3.4) and *BP-D* (see Section 3.5). BP-C examines a fixed number of nodes in each iteration, whereas, BP-D examines a dynamic (variable) number of nodes in each iteration, i.e., $c^{(i)}$ can be different from one iteration to the next. Overall, both BP-C and BP-D share the same algorithmic structure with an exception of eliminating candidates, as explained in Algorithm 1.

3.4. B&P with a Constant Branching Scheme (BP-C)

The basic idea of BP-C is to branch a constant number of child nodes from a parent node in each iteration. The child nodes are compared based on their objective values and the child node with the minimum objective value is kept and the rest ones are pruned.

Algorithm 1 Branch-and-Prune

begin

1. $i := 1$; $\mathcal{A}^{(i)} := \mathcal{A}$;
2. ***while*** $(|\mathcal{A}^{(i)}| > \eta + \alpha)$
   - (a) Solve LP (2) with $\mathcal{A}^{(i)}$;
   - (b) Compute $\mathcal{S}_a^{(i)}$, $\forall a \in \mathcal{A}^{(i)}$;
   - (c) Select c^i angles, $\{\bar{a}_1, ..., \bar{a}_{c^i}\} \subset \mathcal{A}^{(i)}$, that correspond to the c^i lowest scores;
   - (d) Branch c^i child nodes with angle sets $B^j = \{A^i \backslash \bar{a}_j\}$, $j = 1, \ldots, c^i$;
   - (e) ***for*** $j = 1 : c^i$
     Solve LP (2) with B^j and store objective value z_j;
     end
   - (f) $j^* \leftarrow arg\min_j z_j$;
   - (g) $\mathcal{A}^{(i+1)} \leftarrow B^{j^*}$; $i \leftarrow i + 1$;

   end
3. Solve MIP (4) with $\mathcal{A}^{(i)}$ to obtain optimal solution $\bar{\mathcal{A}}$ for B&P.

end

Algorithm 1 illustrates the general structure of B&P. After solving the LP relaxation model in (2), the beam angle scores are calculated for all angles in $\mathcal{A}^{(i)}$. Since BP-C examines a fixed number of child nodes in each iteration, c^i is the constant c. In Step 2(c), c angles, $\{\bar{a}_1, \ldots, \bar{a}_c\} \subset \mathcal{A}^{(i)}$ that correspond to the c lowest scores are selected for further examination. Each node in the

search tree represents a beam angle set. Hence, c child nodes are formed from the parent node by removing each of the c angles from the current beam angle set $\mathcal{A}^{(i)}$ one at a time. In Step 2(e), c iterations of solving (2) are performed and each of the c nodes is associated with its objective value z_j. We then select one promising node whose objective value is minimum. The remaining nodes are pruned. A tie is broken arbitrarily. Step 2 of the algorithm stops when $|\mathcal{A}^{(i)}| \leq \eta + \alpha$, $\alpha \in \mathbb{Z}_+$. At the end of Step 2, there are only $\eta + \alpha$ candidate beam angles left and the MIP model is then solved to select η out of $\eta+\alpha$ beam angles as the final solution, and the algorithm terminates. Note that α defines the number of elimination iterations ($|\mathcal{A}|-\eta-\alpha$) to perform in Step 2 and the size of the input beam angle set ($\eta+\alpha$) in Step 3 (solving the final MIP). A relatively small value of α is recommended because the MIP model may be difficult time to solve if α is large. On the other hand, a small α, say $\alpha < 2$, may diminish the potential of the MIP model to obtaining an improved solution.

3.5. B&P with a Dynamic Branching Scheme (BP-D)

BP-C may require a trial-and-error approach to obtain a good value of c. Also, it is reasonable to expect that BP-C may perform unnecessary branching if the value of c is large. For example, let us assume that c is set to 4. Suppose that there is one beam angle in $\mathcal{A}^{(i)}$ whose merit score at Step 2 is far smaller than the others. It is reasonable to remove this angle from the current beam angle set to form only one child node in this case. However, in BP-C, we still choose four angles based on the scores in Step 2(c) and construct four child nodes. We then perform four LP solves to make the pruning decisions. Therefore, one may use a variable number of elimination candidates (BP-D) as an alternative. The value of c^i in BP-D can change from one iteration to the next. In BP-D, the merit scores are analyzed to determine the value of c^i for the i^{th} iteration (Step 2(c)). We construct a lower bound (a threshold value) for beam angle scores at each iteration, i.e., $LB^{(i)} = u^{(i)} - \tau \cdot s^{(i)}$, where τ ($\tau > 0$) is an interval adjustment parameter and $u^{(i)}$ and $s^{(i)}$ are the mean and standard deviation of all $\mathcal{S}_a^{(i)}$, respectively. Note that a larger value of τ relates to a smaller c^i, and vice versa. Any beam angle with a score below $LB^{(i)}$ is selected as one of the c^i tentative angle elimination candidates. If there are no angles whose scores are less than their lower bounds $LB^{(i)}$, a constant number of c angles with lowest scores are selected ($c = 2$ is used in our implementation).

4. Local Neighborhood Search Algorithm

We obtain feasible solutions to the beam angle optimization problem by the iterative beam angle elimination heuristics discussed in Section 3. Both IBAELP and B&P generates good feasible solutions. However, such methods do not guarantee optimality. Therefore, we use a local neighborhood search (LNS) algorithm to converge a feasible solution (found in $B\&P$) to its local optimum. For illustration, let $\bar{\mathcal{A}}$ be a starting solution, i.e. a beam angle set, $\bar{\mathcal{A}} \subset \mathcal{A}$. We search for a beam angle set $\bar{\mathcal{A}}' \in \mathcal{N}(\bar{\mathcal{A}}) \subseteq \mathcal{A}$, such that $z(\bar{\mathcal{A}}') < z(\bar{\mathcal{A}})$, where $\mathcal{N}(\bar{\mathcal{A}})$ is a neighborhood of $\bar{\mathcal{A}}$ and z is the objective value of the optimization model. If $\bar{\mathcal{A}}'$ does not exist, $\bar{\mathcal{A}}$ is a local optimum; otherwise, set $\bar{\mathcal{A}} = \bar{\mathcal{A}}'$, and repeat the search process.

Neighborhood Definition. The decision variable to the beam angle selection problem is a set of η angles: $\bar{\mathcal{A}}$ ($\eta << |\mathcal{A}|$). The neighborhood of $\bar{\mathcal{A}}$ is obtained by altering one or more angles of the set $\bar{\mathcal{A}}$. Allowing more angles to be exchanged at one time can introduce additional complexity to our problem due to the increased neighborhood size. Therefore, we use a one-angle-exchange algorithm [28] to construct a neighborhood of a beam angle set, i.e., one angle of the current set $\bar{\mathcal{A}}$ is swapped with another angle that is not in the set, $\bar{\mathcal{A}}^C (:= \mathcal{A} \backslash \bar{\mathcal{A}})$, one at a time. The neighborhood definition is

$$\mathcal{N}_0(\bar{\mathcal{A}}) = \{\bar{\mathcal{A}}' : \bar{\mathcal{A}}' = (\bar{\mathcal{A}} \cup \{a_j\}) \backslash \{a_i\}, \ \ for\ a_j \in \bar{\mathcal{A}}^C,\ a_i \in \bar{\mathcal{A}}\}. \tag{9}$$

The number of elements in a neighborhood can be large for a real clinical problem. For example, a problem instance with $\eta = 6$ and $|\mathcal{A}| = 72$ will have $|\bar{\mathcal{A}}| * |\bar{\mathcal{A}}^C| = 6 * 66 = 396$ elements in the neighborhood. In this case, the resulting local search procedure is non-trivial. A reasonably sized neighborhood allows the local search to converge quickly, but if the neighborhood is unnecessarily small, it may easily degrade the search quality. In order to address the tradeoff between the solution time and the quality of the local neighborhood search, we further restrict our search. Instead of replacing an angle in $\bar{\mathcal{A}}$ by any angle in $\bar{\mathcal{A}}^C$, we restrict $\bar{\mathcal{A}}^C$ to a selected subset. Based on our observations, the geometrically adjacent angles and the opposite angle of a given beam angle contribute similarly to the treatment plan, and those angles consist of a neighborhood of a single beam angle. Thus, we introduce a neighborhood $\delta(a_i)$ of a given beam angle a_i, where $\delta(a_i) \subset \bar{\mathcal{A}}^C$, to replace $\bar{\mathcal{A}}^C$ in (9):

$$\mathcal{N}(\bar{\mathcal{A}}) = \{\bar{\mathcal{A}}' : \bar{\mathcal{A}}' = (\bar{\mathcal{A}} \cup \{a_j\}) \backslash \{a_i\}, \ \ for\ a_j \in \delta(a_i),\ a_i \in \bar{\mathcal{A}}\}, \tag{10}$$

where

$$\delta(a_i) = \{a_i' : a_i' \in \{a_i \pm (\rho - u)\Theta\} \cup \{a_i + 180^\circ\}\} \; mod \; 360^\circ, \qquad u = 1, 2, ..., \rho - 1, \tag{11}$$

in which the positive integer ρ ($1 \leq \rho \leq |\mathcal{A}| - 1$) is a user-defined parameter. The set $\{a_i \pm (\rho - u)\Theta\}$ includes adjacent angles of a_i and $a_i + 180^\circ$ represents the opposite angle of a_i.

Note that it is common to use equally spaced coplanar beam angles in clinical practice. For example, a 360° circumference is equally divided by an incremental unit (Θ), i.e., if $\Theta = 5^\circ$.

The size of the neighborhood set $\mathcal{N}(\bar{\mathcal{A}})$ is determined by ρ. If the available beam angles and the number of decision beam angles (η) are both large, small ρ is efficient. Otherwise, large ρ is appropriate to construct a sufficiently large neighborhood set to ensure better local solutions. The restricted neighborhood definition (10) is used in our local search algorithm.

Algorithm 2 Local Neighborhood Search

begin

1. $i := 0$; $\bar{\mathcal{A}}_i^* := \bar{\mathcal{A}}$; $\bar{z}_i^* := f_{\bar{\mathcal{A}}}(D)$;
2. ***do***
   - (a) Generate $\mathcal{N}(\bar{\mathcal{A}}_i)$;
   - (b) ***for*** $j = 1 : |\mathcal{N}(\bar{\mathcal{A}}_i)|$
     Solve LP (2) with A_j^N and store the objective value f_j;
     end
   - (c) $\bar{\mathcal{A}}_{i+1}^* \leftarrow argmin_j\{f_j\}$; $\bar{z}_{i+1}^* \leftarrow min_j\{f_j\}$;
   - (d) $i \leftarrow i + 1$;

   while $(\bar{z}_i^* < \bar{z}_{i-1}^*)$;
3. Stop. $(\bar{z}_{i-1}^*, \bar{\mathcal{A}}_{i-1}^*)$ is the optimal solution.

end

Methodology (Local Neighborhood Search Algorithm). As described in Algorithm 2, our Local Neighborhood Search (LNS) algorithm starts by enumerating the neighbors of a given beam angle vector and iteratively searches for improving solutions within the restricted neighborhood. First, a starting feasible solution ($\bar{\mathcal{A}}_0$) is obtained by the heuristic method discussed in Section 3. The input data to the algorithm includes the feasible solution $\bar{\mathcal{A}}_0$, the corresponding objective value $\bar{z}_0 = f_{\bar{\mathcal{A}}_0}(D)$, and the value of $\eta = |\bar{\mathcal{A}}_0|$. At each

iteration $i = \{1, 2, \ldots\}$, a neighborhood of the current solution $\bar{\mathcal{A}}_i$ is constructed: $\mathcal{N}(\bar{\mathcal{A}}_i)$. Suppose the cardinality of $\mathcal{N}(\bar{\mathcal{A}}_i)$ is $\bar{m}$, i.e., $\bar{m} = |\mathcal{N}(\bar{\mathcal{A}}_i)|$. Let A_j^N be the j^{th} neighborhood set of $\mathcal{N}(\bar{\mathcal{A}}_i)$, i.e., $\mathcal{N}(\bar{\mathcal{A}}_i) = \{A_j^N\}_{j=1,\ldots,\bar{m}}$. In Step 2($b$), the LP model in (2) with the j^{th} angle set A_j^N is solved. Then the results are saved; $[f_j,\ A_j^N]$. We repeat this process until all neighborhood angle sets are evaluated. In Step 2(c), the solution with the minimum objective value is selected as the next best feasible solution. If the current best solution $\bar{\mathcal{A}}_i^*$ yields a better objective value ($\bar{z}_i^* < \bar{z}_{i-1}^*$), then indexes are updated accordingly and the search algorithm continues. Otherwise, the current solution ($\bar{z}_i^*$, $\bar{\mathcal{A}}_i^*$) is optimal and the algorithm terminates.

5. Numerical Experiments and Result

Our experiments are implemented on two clinical data sets: a prostate cancer case and a pancreas cancer case. The prostate cancer case contains 5,246 PTV voxels, 1,936 OAR voxels, 477,193 normal tissue voxels. There are 1,244 voxels in the pancreas (PTV), 50,391 voxels in the liver (OAR), 9,116 voxels in the left kidney (OAR), 5,920 voxels in the right kidney (OAR), 489 voxels in the spinal cord (OAR), and 530,863 voxels in normal tissues for the pancreas case.

Table 2: Dose volume requirements for OARs

Case	OAR	Dose volume requirement
Prostate	Rectum	40 Gy < 60%
		60 Gy < 40%
		70 Gy < 25%
		76 Gy < 15%
		78-80 Gy < 5%
Pancreas	Liver	Mean dose < 20 Gy
	Left kidney	20 Gy < 60%
	Right kidney	20 Gy < 60%
	Spinal cord	45 Gy < 100%

The prescription dose for the PTV is 76 Gy for the prostate case and 54 Gy for the pancreas cases. Dose volume requirements for OARs in both cases are listed in Table 2. Scenarios with 12 and 36 input beam angles are primarily considered for both the prostate cancer ($\eta_{prostate} = 6$) and the pancreas data ($\eta_{pancreas} = 4$). Parameter α is set to 2 for both problems. Other model

parameters and the data pre-processing are based on our previous study [6]. All computations were performed on a Linux computer with a Pentium 4 3.6 GHz Dual Xeon processor. Both the LP model and the MIP model are solved using CPLEX 11.0. We used the interior point method for solving the LP model and the branch-and-bound method for solving the MIP model.

5.1. B&P vs. Random Starting Point Strategies on Local Neighborhood Search

There are different ways for finding a starting point to feed our local neighborhood search algorithm. We compare three starting point generation strategies: Random (R), Equi-spaced Random (RE), and B&P schemes. Random (R) uses a random starting solution method while Equi-spaced Random (RE) selects a desired number of angles that are approximately equally spread out over the discrete search space. These two strategies are compared against B&P in the prostate case with twelve candidate beam angles of which we select six. The results are shown in Figures 3 and 4.

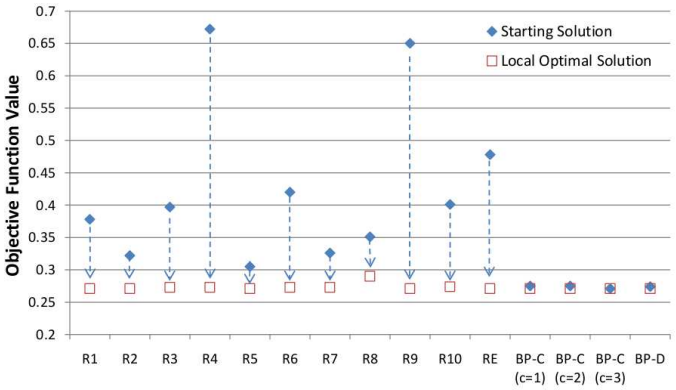

Figure 3: Solution quality comparison: the first ten solutions are based on R (random starting points) followed by one for RE and the last four solutions are based on different B&P schemes.

Figure 3 compares objective values from the LNS model that is solved with different starting point strategies, i.e., R, RE, BP-C and BP-D. Except for the case R8 with the local optimum objective value of 0.290, local optimum values in all other cases are within a small range between 0.271 to 0.274, where the global optimum is known as 0.271. In fact, case R8 brings up an interesting point about the starting point for LNS. LNS may converge to a poor local optimal solution if the quality of the starting solution is not

good. The local optimal solution corresponding to R8 was even worse than all starting solutions generated by B&P schemes. This is one of the strengths of our two-phase method, i.e., it is able to produce good local optimum solutions since the starting solutions obtained from B&P are consistently within a small gap from the global optimal solution in all of our experiments.

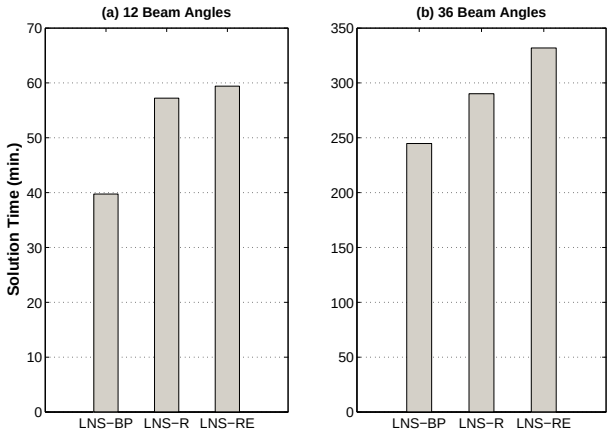

Figure 4: Solution time comparison between 12 and 36 initial candidate beam angle configurations. Selecting 6 treatment angles for the prostate cancer case

Figure 4 shows the solution time comparison of the LNS algorithm for the three different starting point generation methods. LNS-BP is LNS initialized with a B&P method. Similarly, LNS-R uses R as the starting point generation method and and LNS-RE uses RE. For the 12 initial beam angle configuration, solution times are 39.75, 57.22, and 59.41 minutes for LNS-BP, LNS-R, and LNS-RE, respectively. Note that the total solution time of B&P and LNS is recorded for the LNS-BP case. LNS-BP had a 45% speed gain against the other two. When we repeat the experiment with 36 candidate beam angles to select 6 treatment angles, LNS-BP made an 18.5% and a 36.5% speed gain over LNS-R and LNS-RE, respectively.

5.2. B&P vs. IBAELP

BP-C with $c = 1$ performs the same procedure as the IBAELP algorithm proposed by Lim et al. [6]. The main difference between the two algorithms is that BP-C is used for generating a starting feasible solution while the solution of IBAELP is considered final. Figure 5 shows Phase I and Phase

II solutions of different B&P schemes on the prostate cancer case. In Phase I, BP-C ($c = 2$ and 3) and BP-D found lower objective values than the ones from IBAELP (i.e., BP-C with $c = 1$) in all scenarios. It is also true for the pancreas case that Phase I B&P scheme outperformed IBAELP. Especially, for the pancreas case with 12 candidate beams, BP-C found global optimal solutions in all scenarios (see Table 4).

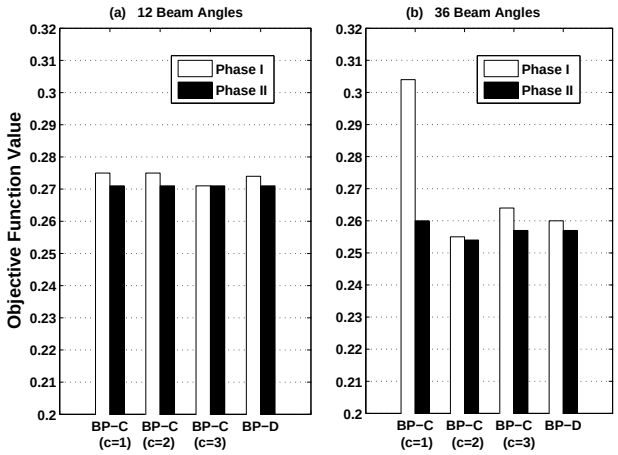

Figure 5: Objective function value comparison between Phase I and Phase II on the prostate cancer case

Another advantage of B&P is its ability to find a good feasible solution. To show this, we compare objective function values of B&P and IBAELP solutions based on different settings of score function parameters ($\kappa_{_S}$ and $\kappa_{_N}$) on the prostate cancer case in Figure 6. Although $\kappa_{_S}$ and $\kappa_{_N}$ theoretically can take any positive value between 0 and 1, it is more appropriate to restrict their sum less than 1, i.e., $\kappa_{_S} + \kappa_{_N} < 1$, so that the average angle scores remain positive as discussed in Lemma 3.1. In this test, we first fix $\kappa_{_N} = 0.1$ and vary $\kappa_{_S} \in \{0.1, 0.15, 0.2, ..., 0.5\}$. This is because the priority on the OARs is normally at least the same as the priority on normal tissues. Figure 6 shows that IBAELP is sensitive to score function parameters while B&P is relatively stable. Note that B&P solutions are average values of solutions from different B&P schemes, i.e., BP-C ($c = 2$ and 3) and BP-D. Those average values are used to represent individual objective values of B&P schemes because the variances of all value groups are small (at most 10^{-5}). Another observation is that B&P outperforms IBAELP with respect to objective value in most of

these settings. However, weighting factors such as κ_S and κ_N are subjective, and one may need a trial-and-error approach to find optimal settings that lead to good solutions.

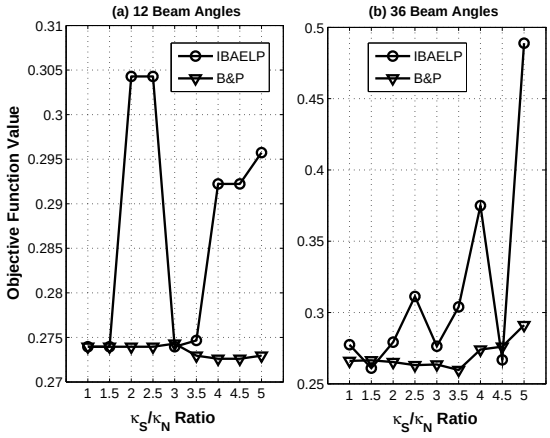

Figure 6: Impact of score function parameters on solution quality on the prostate cancer case

5.3. IBAELP vs. Two-phase

A comparison between IBAELP over two-phase approach with respect to objective values in the prostate cancer case is shown in Figure 7, where 12, 24 and 36 candidate beam angle configurations are used. Since IBAELP has only one scheme, one objective value for each configuration is displayed in the box plot (0.275 for 12 beam angles, 0.287 for 24 beam angles, and 0.304 for 36 beam angles). On the other hand, a group of objective values are displayed in the box plot for the two-phase method, which are based on different B&P schemes, i.e., BP-C (c = 1, 2 and 3) and BP-D. The 12 candidate beam configuration produced the same objective value of 0.271, and the 24 candidate set gave 0.265. A value was found between 0.254 and 0.260 for the 36 angle case. Note that the objective value decreases as we consider more candidate beam angles when the B&P scheme is used, but the result is opposite for IBAELP.

In our experiments, the solution quality of the two-phase method monotonically improved as the number of candidate angles increased. The solution time increased accordingly. However, IBAELP's objective value degraded as more candidates were allowed. This result supports that our method is efficient since a well designed algorithm is expected to find a better treatment

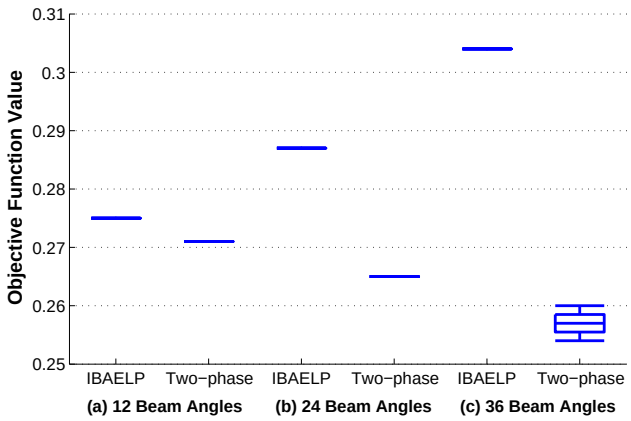

Figure 7: A box plot for comparing objective values from IBAELP and two-phase method (prostate cancer case)

plan if there are more angles to choose from. Furthermore, the two-phase method consistently found local optimal solutions that are within a small fraction of the global optimal solution.

5.4. Effects of Increasing Candidate Beam Angles

Although B&P typically obtains better feasible solutions than IBAELP, the local neighborhood search further improves the solution. Numerical results of our two-phase approach are tabulated in Table 3 for the prostate cancer case and Table 4 for the pancreas case.

$\eta = 6$	$\|\mathcal{A}\|$=12		$\|\mathcal{A}\|$=36	
Model	Obj. (% opt. gap)	Time (% gain over MIP)	Obj.	Time
MIP	0.271	313.2	-	-
IBAELP	0.275	24.6	0.304	68.0
BP-C($c = 1$)+LNS	0.271 (0.0%)	41.7 (87%)	0.260	222.0
BP-C($c = 2$)+LNS	0.271 (0.0%)	37.6 (88%)	0.254	159.4
BP-C($c = 3$)+LNS	0.271 (0.0%)	33.7 (89%)	0.257	239.3
BP-D+LNS	0.271 (0.0%)	39.2 (88%)	0.257	174.8

Table 3: Solutions obtained from the 12 and 36 candidate beam angle configurations (prostate cancer case). CPU time is measure in minutes

For the prostate case with 12 candidate beam angles (Figure 5 and Table 3), most B&P schemes (Phase I) found solutions close to the global optimal

solution (0.271) with less than 15% of the computation time required by the MIP solver. A speed gain of up to 89% was observed compared to the MIP solution approach (Branch-and-Bound).

In the case of 36 candidate beam angles, CPLEX failed to find an optimal solution for the MIP model within seven days and the process was manually terminated due to an unexpected machine shutdown. The last recorded dual gap was 43%. The Phase II running times varied from 43 minutes to 154 minutes. It is not surprising that there is an increase in computation time for Phase II if there are more candidate angles in the model. It is due to the fact that there are more iterations to run and there are larger LP models to solve in each iteration. In comparison, up to 7.7% improvement was observed in the objective value by moving from 12 candidate angles to 36.

$\eta = 4$	$\|\mathcal{A}\|$=12		$\|\mathcal{A}\|$=36	
Model	Obj.	Time (% gain over MIP)	Obj.	Time
MIP	0.110	60.9	-	-
IBAELP	0.110	12.8 (79%)	0.084	299.5
BP-C($c = 1$)+LNS	0.110	17.8 (71%)	0.072	311.6
BP-C($c = 2$)+LNS	0.110	24.3 (60%)	0.072	559.9
BP-C($c = 3$)+LNS	0.110	31.2 (49%)	0.072	817.0
BP-D+LNS	0.110	25.4 (58.3%)	0.072	478.5

Table 4: Solutions obtained based on 12 and 36 candidate beam angle configurations (Pancreas) CPU time is measure in minutes

For the pancreas with a 12 angle configuration (Table 4), all Phase I heuristic approaches found the global optimal solution while utilizing less than 45% of the computation time required by the MIP solver. Up to a 71% speed gain was observed. Since Phase I found the optimal solution, Phase II was terminated after the first iteration. When there are 36 candidate angles, the objective value was improved by 34.5% compared to that of the 12 angle setup. Furthermore, the LNS algorithm improved the starting solution up to 16.7%. This shows a substantial increase in the computation time due to the large organ volumes.

Table 5 shows the maximum and mean doses on OARs if the two-phase method is applied to both the prostate and pancreas cases in 12 and 36 beam configurations. The results support that 1) our treatment plans are clinically acceptable, i.e., the dose volume requirements for OARs listed in

Case	Prostate		Pancreas							
	Max Rect	Mean Rect	Max Liver	Mean Liver	Max Spin	Mean Spin	Max L Kidn	Mean L Kidn	Max R Kidn	Mean R Kidn
12	85.8	24.5	31.6	4.8	11.5	4.6	10.8	1.8	21.2	2.8
36	85.0	20.0	29.0	3.8	11.3	4.0	10.8	0.1	16.8	2.5

Table 5: Maximum and mean OAR dose (Gy) of solutions with 12 and 36 candidate beam angles

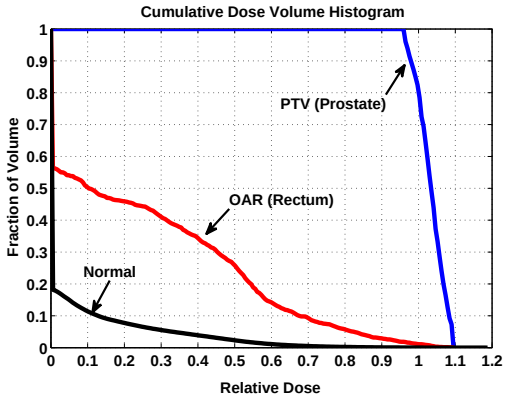

(a) Dose volume histogram

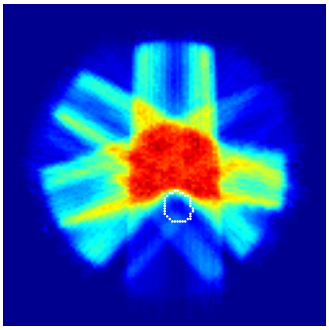

(b) Dose distribution plot

Figure 8: Dose volume histogram and axial dose distribution plot for the prostate case with $|\mathcal{A}| = 36 and \eta = 6$

Table 2 are satisfied; and 2) plans obtained from the 36 candidate beam angles consistently outperform plans obtained from 12 candidate beam angles. Dose volume histograms (DVH) and dose distribution plots in a axial slice for plans obtained from 36 candidate beam angles are shown in Figures 8 and 9. The DVHs demonstrate that our solutions have achieved a relatively uniform dose distribution on the PTV while avoiding an excessive dose on the OARs and the normal tissue. The solution quality was confirmed by inspecting the dose distribution plots. A high radiation dose conforms around the tumor while the organs close to PTV receive limited radiation doses.

6. Conclusion

We introduced a two-phase solution approach to beam angle and fluence map optimization for IMRT treatment planning. First, a B&P algorithm

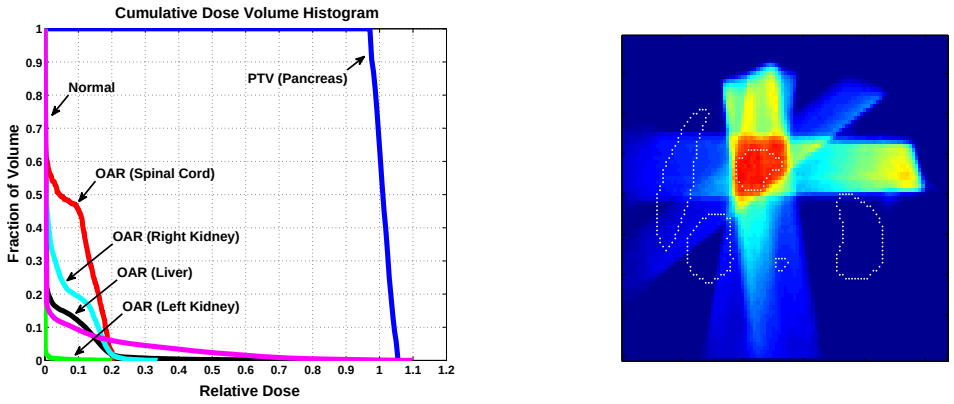

(a) Dose volume histogram (b) Dose distribution plot

Figure 9: Dose volume histogram and axial dose distribution plot for Pancreas with $|\mathcal{A}| = 36, \eta = 4$

was introduced to generate quality feasible solutions in Phase I. This is an extension of the IBAELP method in [6]. B&P carefully selects a candidate beam angle for elimination. One-step-ahead consequence estimation due to an elimination of an angle is made. Then one angle with the smallest impact due to elimination is selected for removal in each iteration. We discussed branching and pruning rules for B&P in order to keep the search tree small. The branching decision is made based on a merit score and the pruning rule is based on the LP relaxation of the FMO problem. Two different branching schemes were also discussed to select the number of branches to consider in each iteration. Since B&P does not guarantee optimality, a local neighborhood search algorithm was developed to improve a B&P solution. This Phase II is an exhaustive search within a neighborhood set to select the best performing angle set. Numerical results on two clinical data sets show that B&P alone can generate clinically acceptable solutions. We have demonstrated that the two-phase method found global optimal solutions for both twelve candidate angle configurations. We also made a numerical comparison among different starting point generation methods for the second phase. B&P consistently outperformed random starting solution strategies in terms of overall CPU time as well as final solution quality.

References

[1] S. Webb, Intensity-modulated radiation therapy, Institute of physics, Series in medical physics, IOP Publishing Ltd, 2001.

[2] T. Bortfeld, W. Schlegel, Optimization of beam orientations in radiation-therapy - some theoretical considerations, Physics in Medicine and Biology 38 (2) (1993) 291–304.

[3] G. Meedt, M. Alber, F. Nüsslin, Non-coplanar beam direction optimization for intensity-modulated radiotherapy, Physics in Medicine and Biology 48 (18) (2003) 2999–3019.

[4] V. V. Misic, D. M. Aleman, M. B. Sharpe, Non-coplanar extensions to neighborhood search methods for beam orientation optimization for total marrow irradiation using imrt, European Journal of Operational Research 205 (3) (2010) 522–527.

[5] E. K. Lee, T. Fox, I. Crocker, Optimization of radiosurgery treatment planning via mixed integer programming, Medical Physics 27 (5) (2000) 995–1004.

[6] G. J. Lim, J. Choi, R. Mohan, Iterative solution methods for beam angle and fluence map optimization in Intensity Modulated Radiation Therapy planning, OR Spectrum 30 (2) (2008) 289–309.

[7] H. H. Zhang, L. Shi, R. R. Meyer, D. Nazareth, W. D. D'Souza, Solving beam-angle selection and dose optimization simultaneously via high-throughput computing, INFORMS Journal on Computing 21 (3) (2009) 427–444.

[8] M. Ehrgott, Ç. Güler, H. W. Hamacher, L. Shao, Mathematical optimization in intensity modulated radiation therapy, Annals of Operations Research 175 (1) (2010) 309–365.

[9] M. Ehrgott, A. Holder, J. Reese, Beam selection in radiotherapy design, Linear Algebra and Its Applications 428 (2008) 1272–1312.

[10] A. Pugachev, L. Xing, Pseudo beam's-eye-view as applied to beam orientation selection in intensity-modulated radiation therapy, International Journal of Radiation Oncology: Biology, Physics 51 (5) (2001) 1361–1370.

[11] W. D. D'Souza, R. R. Meyer, L. Shi, Selection of beam orientations in intensity modulated radiation therapy using single-beam indices and integer programming, Physics in Medicine and Biology 49 (15) (2004) 3465–3481.

[12] D. M. Aleman, A. Kumar, R. K. Ahuja, H. E. Romeijn, J. F. Dempsey, Neighborhood search approaches to beam orientation optimization in intensity modulated radiation therapy treatment planning, Journal of Global Optimization 42 (4) (2008) 587–607.

[13] S. K. Das, L. B. Marks, Selection of coplanar or noncoplanar beams using three-dimensional optimization based on maximum beam separation and minimized nontarget irradiation, International Journal of Radiation Oncology: Biology, Physics 38 (3) (1997) 643–655.

[14] O. C. Haas, K. J. Burnham, J. A. Mills, Optimization of beam orientation in radiotherapy using planar geometry, Physics in Medicine and Biology 43 (8) (1998) 2179–2193.

[15] E. Schreibmann, M. Lahanas, L. Xing, D. Baltas, Multiobjective evolutionary optimization of the number of beams, their orientations and weights for intensity-modulated radiation therapy, Physics in Medicine and Biology 49 (5) (2004) 747–770.

[16] D. Djajaputra, Q. Wu, Y. Wu, R. Mohan, Algorithm and performance of a clinical IMRT beam angle optimization system, Physics in Medicine and Biology 48 (19) (2003) 3191–3212.

[17] A. Pugachev, J. G. Li, A. L. Boyer, S. L. Hancock, Q. T. Le, D. S. S., L. Xing, Role of beam orientation optimization in intensity-modulated radiation therapy, International Journal of Radiation Oncology: Biology, Physics 50 (2) (2001) 551–560.

[18] J. Stein, R. Mohan, X. H. Wang, T. Bortfeld, Q. Wu, K. Preiser, C. C. Ling, W. Schlegel, Number and orientation of beams in intensity-modulated radiation treatments, Medical Physics 24 (2) (1997) 149–160.

[19] Y. Li, J. Yao, D. Yao, Automatic beam angle selection in IMRT planning genetic algorithm, Physics in Medicine and Biology 49 (10) (2004) 1915–1932.

[20] D. Nazareth, S. Brunner, M. D. Jones, H. K. Malhotra, M. Bakhtiari, Optimization of beam angles for intensity modulated radiation therapy treatment planning using genetic algorithm on a distributed computing platform, Journal of Medical Physics 34 (3) (2009) 129–132.

[21] Q. Wu, J. Wang, Y. Chen, J. M. Galvin, Beam orientation optimization for imrt by a hybrid method of the genetic algorithm and the simulated dynamics, Medical Physics 30 (9) (2003) 2360–2367.

[22] H. H. Liu, M. Jauregui, X. D. Zhang, X. C. Wang, L. Dong, R. Mohan, Beam angle optimization and reduction for intensity-modulated radiation therapy of non-small-cell lung cancers, International Journal of Radiation Oncology Biology Physics 65 (2) (2006) 561–572.

[23] X. Wang, X. Zhang, L. Dong, H. Liu, Q. Wu, R. Mohan, Development of methods for beam angle optimization for IMRT using an accelerated exhaustive search strategy, International Journal of Radiation Oncology: Biology, Physics 60 (4) (2004) 1325–1337.

[24] Y. Li, J. Yao, D. Yao, W. Chen, A particle swarm optimization algorithm for beam angle selection in intensity-modulated radiotherapy planning, Physics in Medicine and Biology 50 (15).

[25] D. M. Aleman, H. E. Romeijn, J. F. Dempsey, A response surface approach to beam orientation optimization in intensity modulated radiation therapy treatment planning, INFORMS Journal on Computing: Computational Biology and Medical Applications 21 (1) (2009) 62–76.

[26] D. Craft, Local beam angle optimization with linear programming and gradient search, Physics in Medicine and Biology 52 (7) (2007) N127–135.

[27] G. J. Lim, M. C. Ferris, S. J. Wright, D. M. Shepard, M. A. Earl, An optimization framework for conformal radiation treatment planning, INFORMS Journal on Computing 19 (3) (2007) 366–380.

[28] L. A. Wolsey, Integer Programming, John Wiley & Sons, Inc., 1998.