J Glob Optim (2008) 42:587–607
DOI 10.1007/s10898-008-9286-x

Neighborhood search approaches to beam orientation optimization in intensity modulated radiation therapy treatment planning

Dionne M. Aleman · Arvind Kumar · Ravindra K. Ahuja · H. Edwin Romeijn · James F. Dempsey

Received: 7 June 2007 / Accepted: 18 February 2008 / Published online: 18 March 2008

Abstract The intensity modulated radiation therapy (IMRT) treatment planning problem consists of several subproblems which are typically solved sequentially. We seek to combine two of the subproblems: the beam orientation optimization (BOO) problem and the fluence map optimization (FMO) problem. The BOO problem is the problem of selecting the beam orientations to deliver radiation to the patient. The FMO problem is the problem of determining the amount of radiation intensity, or fluence, of each beamlet in each beam. The solution to the FMO problem measures the quality of a beam set, but the majority of previous BOO studies rely on heuristics and approximations to gauge the quality of the beam set. In contrast with these studies, we use an exact measure of the treatment plan quality attainable using a given beam set, which ensures convergence to a global optimum in the case of our simulated annealing algorithm and a local optimum in the case of our local search algorithm. We have also developed a new neighborhood structure that allows for faster convergence using our simulated annealing and local search algorithms, thus reducing the amount of time required to obtain a good solution. Finally, we show empirically that we can generate clinically

D. M. Aleman (✉) · R. K. Ahuja · H. E. Romeijn
Department of Industrial and Systems Engineering, University of Florida, 303 Weil Hall, P.O. Box 116595,Gainesville, FL 32611-6595, USA
e-mail: daleman1@ufl.edu

R. K. Ahuja
e-mail: ahuja@ise.ufl.edu

H. E. Romeijn
e-mail: romeijn@ise.ufl.edu

A. Kumar
Innovative Scheduling, Gainesville Technology Enterprise Center (GTEC), 2153 Hawthorne Road, Suite 128, Gainesville, FL 32641, USA
e-mail: arvind@innovativescheduling.com

J. F. Dempsey
Department of Radiation Oncology, University of Florida, P.O. Box 100385, Gainesville, FL 32610-0385, USA
e-mail: dempsey@ufl.edu

acceptable treatment plans that require fewer beams than in current practice. This may reduce the length of treatment time, which is an important clinical consideration in IMRT.

Keywords Intensity modulated radiation therapy · IMRT · Beam orientation optimization · BOO · Neighborhood search · Add/Drop · Stimulated annealing

1 Introduction

In the United States alone, approximately 1.4 million people are newly diagnosed with cancer each year (American Cancer Society, 2006), of which more than half will be treated with some form of radiation therapy. During this therapy, beams of radiation are passed through the patient in an attempt to erradicate the cancerous cells. Although the radiation beams will in fact kill the cancerous cells, they will also kill normal cells. To avoid potentially serious side effects which may significantly decrease the patient's quality of life, the treatment plan must be carefully designed so that while the target cells receive the prescribed radiation dose, nearby organs and tissues (called critical structures) are spared. Typically, patients are irradiated from several beams spaced around the patient so that the targets lie in the intersection of the beams, and thus receive the highest dose, whereas the critical structures receive significantly less radiation and thus may be spared.

Currently, a technique called *intensity modulated radiation therapy* (IMRT) is considered to be the most effective radiation therapy for many forms of cancer. The problem of designing an IMRT treatment plan for an individual patient is a very large-scale mathematical programming problem that is not yet solved satisfactorily. Current treatment planning systems decompose the planning problem into several stages, and the corresponding subproblems are solved sequentially. This paper addresses the integration of the beam orientation optimization (BOO) and fluence map optimization (FMO) subproblems based on a convex voxel-based penalty function-based formulation of the latter and associated efficient algorithm for solving it.

IMRT is so-named because each beam of radiation can be discretized in hundreds of smaller *beamlets*, the radiation intensities (or, more accurately, *fluences*) of which can be modulated independently of the other beamlets. These fluence values are known as a fluence map, and optimization of these fluences given a fixed set of beams is known as fluence map optimization (FMO). In order to model the IMRT problem as a mathematical programming problem, a means of quantifying the quality of a treatment plan, the ability of the plan to deliver the prescribed radiation dose to the specified target structures while ensuring that important function organs receive an acceptably low amount of radiation, is required. The optimal solution of the FMO problem performed with a given vector of beam orientations measures the quality of the treatment plan using those beams. Thus, our approach to the problem of selecting the best directions from which to deliver radiation to the patient (the BOO problem) is based on the treatment plan quality indicated by the optimal solution value to the corresponding FMO problem.

The FMO problem has been extensively studied in the literature. FMO models previously studied include linear programming (LP)-based multi-criteria optimization [22] and mixed-integer linear programming (MILP) [5,27–30,50]. Both "physical" [9] and "biological" [1,23,24,35,37,38,54,55] bases for the development of the objective function and constraints have been used. Homogeneity of the target dose distribution [22], *full-volume* constraints requiring that the dose in every voxel of a structure be within pre-determined upper and lower bounds [5,22,29,30,43,46] and *partial volume* constraints requiring that dose in only a subset of voxels be within pre-determined upper and/or lower bounds [29,30,43,45,50]

have all been used as criteria for developing constraints in FMO models. Lim et al. [33] use structure-based penalty functions.

In addition to varying constraints, there are many competing methods of formulating the FMO objective function to reflect the quality of the treatment plan. Romeijn et al. [44] showed that most of the treatment plan evaluation criteria proposed in the medical physics literature are equivalent to convex penalty function criteria when viewed as a multicriteria optimization problem, in the sense that for each set of treatment plan evaluation criteria from a very large class there exists a class of convex penalty functions that produces an identical Pareto efficient frontier.

Because of this equivalence, we use a convex penalty function-based approach to evaluate the FMO problem, and the FMO model is used to investigate the BOO problem. Although this approach could be used in a multicriteria setting, Romeijn et al. [43,46] suggest that it is possible to quantify a trade-off between the different evaluation criteria that produces high-quality treatment plans for a population of patients, eliminating the need to solve the FMO problem as a multicriteria optimization problem for each individual patient.

The beam orientation optimization problem has also been extensively studied in the literature. The approaches in previous works largely consist of applying various global optimization algorithms including evolutionary algorithms [2,3,49], genetic algorithms [15,21,31], particle swarm algorithms [32], quasi-Newton methods [13], artificial neural network algorithms [48] and simulated annealing [10,14,34,42,47,53]. Heuristic measures including beam's-eye-view (BEV—a "bird's-eye view" of the patient's tissues and structures as seen from the beam) [11,12,19,34,40–42], path of least resistance [20] and entropy [52]. Meedt et al. [36] use a fast exhaustive search to obtain a beam vector solution.

We test the simulated annealing algorithm on the BOO problem, as well as existing and new variants of a greedy neighborhood search heuristic called the *Add/Drop algorithm* (see Kumar [26]) to obtain a good solution to the BOO problem. In each step of the Add/Drop algorithm, a beam in the current beam set is replaced by a neighboring beam that yields an improving solution. As with the simulated annealing implementation, we also apply our new neighborhood to the Add/Drop algorithm and compare its performance to a commonly used neighborhood structure.

With the exception of Das and Marks [13], Haas et al. [21] and Schreibmann [49], the previous studies do not select beam vectors using the FMO problem as a model for determining the beam vector's quality; instead, the beam vectors are chosen based on scoring methods (e.g., BEV, path of least resistance), on approximations to the FMO or on locally optimal or sub-optimal FMO solutions. Without basing BOO on the optimal FMO solutions, the resulting beam vector solutions have no guarantee of optimality, or even of local optimality. Das and Marks [13] strive to maximize angle separation, which, although intuitively indicative of a solution that may spare critical organs, may bias the objective function toward solutions that are less capable of treating the targets and sparing critical structures than other solutions containing beams that are close together. Haas et al. [21] also use intuitive geometric-based measures to reduce the objective function weighting of beams in certain orientations to the targets, which may prevent their algorithm from finding an optimal solution. Schreibmann [49] use a multiobjective evolutionary algorithm to optimize both the number of beams and their orientations.

In contrast to the previous studies, both our simulated annealing and local search approaches to the BOO problem measure the quality of a beam vector, given the number of beams to be used, according to the exact FMO solution, thus ensuring convergence to a local optimum in the case of the Add/Drop algorithm and a global optimum in the case of the simulated annealing algorithm run for a sufficient period time. Additionally, we have

developed a new neighborhood structure which allows for faster convergence of the FMO solution.

The paper is organized as follows. Section 2 describes the BOO model, Sect. 3 describes neighborhood search methods for BOO (the Add/Drop and simulated annealing algorithms) and Sect. 3.3 describes the new neighborhood structure applied to the neighborhood search algorithms. Section 4 contains the results of the neighborhood search methods and Sect. 5 contains the conclusions on the effectiveness of our simulated annealing and local search approaches, as well as the effectiveness of the flip neighborhood in improving the convergence rate of the two algorithms.

2 Beam orientation optimization

Beam orientation optimization (BOO) is the problem of selecting from all of the available beam orientations the best set to use in delivering a treatment plan. Coplanar beams are those beams obtained from the rotation of only the gantry of the linear accelerator, the machine which delivers radiation beams to the patient. Figure 1 shows a linear accelerator with arrows indicating the movements available to its components; the gantry rotation is highlighted. If all other components of the linear accelerator are fixed, the rotation of the gantry sweeps out a set of coplanar beams. In practice, usually only coplanar beams are used to deliver radiation.

This study focuses on obtaining high-quality coplanar beam vectors for head-and-neck cancer site treatment plans. A typical head-and-neck treatment plan consists of radiation delivered from 5 to 9 nominally-spaced coplanar orientations around the patient.

2.1 BOO model

To determine the quality of a given beam set, a quantitative measure for assessing the quality of the beam set must be used. Let $F(\boldsymbol{\theta})$ be a black-box function that quantifies the quality of the treatment plan if radiation is delivered from beam vector $\boldsymbol{\theta} = (\theta_1, \ldots, \theta_k)$, where k is the user-specified number of orientations that may be used. We assume that F is formulated in such a way that the minimum function value corresponds to the optimal plan. F must thus be able to appropriately make the trade-offs between the contradictory goals of delivering a prescribed radiation dose of targets and an acceptably low radiation dose to critical structures so that they may be spared. This is especially problematic for certain cancers, such as tumors in the head-and-neck area, which are often located very close to, for instance, the spinal cord, brain stem and salivary glands.

To determine the quality of the beam vector, the decision vector $\boldsymbol{\theta}$ is used as input into the black-box function $F(\boldsymbol{\theta})$. For a treatment plan with k beams, the BOO problem is formulated as

$$\begin{aligned} &\min F(\boldsymbol{\theta}) \\ &\text{subject to } \boldsymbol{\theta} \in \mathcal{B}^k, \end{aligned}$$

where $\mathcal{B}$ is the set of candidate beams for a single beam. The solution vector $\boldsymbol{\theta}$ lies in the solution space $\mathcal{B}^k = \mathcal{B} \times \cdots \times \mathcal{B}$. The set of candidate beams $\mathcal{B}$ can be selected according to any user-specified criteria. Although the linear accelerator is able to deliver a continuous set of beam orientations from nearly any orientation in 3-dimensional space by rotating/translating the various components indicated in Fig. 1, it is common to only consider a discretized set of beam orientations due to limitations in machine tolerances. Because the many degrees of freedom in a linear accelerator result in a very large set of orientations, it is also common to

Fig. 1 The available movements of a linear accelerator; the gantry rotation is highlighted

further restrict the set of candidate beams by considering only a subset of the discretized candidate beams. The most common restriction is to allow only *coplanar* beams, beams that are obtained from rotating only the gantry component of the linear accelerator (whose rotation sweeps out a set of coplanar orientations). Thus, for f candidate beam orientations, let $\mathcal{B}$ be a finite set defined by $\mathcal{B} = \{360(k/f) : k = 0, \ldots, f-1\}$. $\mathcal{B}$ is therefore a set of equi-spaced beams containing orientation 0°. For example, if $f = 72$, then $\mathcal{B} = \{0, 5, \ldots, 355\}$.

This formulation of the BOO problem is fundamentally nonlinear because the physics of dose deposition change with each beam orientation; that is, the effect of a beam on each patient can be drastically different than the effect of a neighboring beam. In most of the previous BOO studies, the black-box function F is an approximation to the FMO solution, whereas in our BOO model, the black-box function $F(\boldsymbol{\theta})$ is the optimal value function of the convex FMO problem given the beam orientations $\boldsymbol{\theta}$ described in Sect. 2.2, which allows for an exact measure of the quality of each beam vector. Because of the nonlinearity in this formulation, a simulated annealing algorithm, an algorithm capable of escaping local minima, is proposed and applied to solve the problem. Additionally, a deterministic local search method called the Add/Drop algorithm, which obtains a local minimum solution, is also enhanced and applied.

2.2 FMO model

Given a fixed set of beams, different fluence maps (radiation intensities of beamlets) can yield treatment plans with very different qualities. Thus, the quality of an optimized fluence map must be considered when selecting a set of beam orientations to use in a treatment plan. In this study, the convex penalty function-based approach to the FMO model described by Romeijn et al. [43] is used. This formulation yields a quadratic programming problem with linear constraints.

The set critical structures is denoted by S and the set of target structures is denoted by T. Each structure $s \in S \cup T$ is discretized into a finite number v_s of volume cubes, known as *voxels*. Approximately 350,000 voxels are typically required to accurately represent the targets and surrounding structures of a head-and-neck cancer site.

In FMO, the patient is irradiated with k predetermined beam orientations ($\boldsymbol{\theta} \in \mathcal{B}^k$), where k is a user-specified value. In IMRT, each beam can be discretized into a rectangular grid of typically 100–400 smaller beamlets. The set of all beamlet intensities, or *bixels*, in beam θ_h is denoted by B_{θ_h}, $h = 1, \ldots, k$. Fluence map optimization treats each bixel $x_i \in B_{\theta_h}$, $h = 1, \ldots, k$, as a decision variable used to optimize the objective function chosen to make the trade-offs in FMO that have been previously discussed.

The goal of IMRT treatment planning is to limit the dose received by each structure so that the critical structures may be spared, but this is not always possible to do while still delivering enough radiation to the targets to ensure that the cancer is eradicated. For example, in head-and-neck cases, it is very common for one or more saliva glands to be partially encased by a target volume. As organs in such a position may not be able to be spared, hard constraints cannot be used to control the dose received in the critical structures as a feasible solution may not exist. The patient must be treated, so the FMO model is formulated so that there is always a feasible solution. To ensure that feasible solutions exist, we instead impose penalties on each voxel for the amount of radiation dose received using a given set of beamlet intensities. These voxel-based penalty functions ensure that the solution obtained provides a quality treatment.

We introduce a threshold dose value $T_s \geq 0$ for $s \in S \cup T$ that indicates the point at which dose in a structure is penalized for being too high or too low. Define $\underline{w}_s$ and $\underline{p}_s$ to be weighting factors for underdosing, and $\overline{w}_s$ and $\overline{p}_s$ to be weighting factors for overdosing in structure $s \in S \cup T$. Let F_{js} denote a convex penalty function for voxel j in structure s of the following form:

$$F_{js}(z_{js}) = \frac{1}{v_s}\left(\underline{w}_s\left[(T_s - z_{js})_+\right]^{\underline{p}_s} + \overline{w}_s\left[(z_{js} - T_s)_+\right]^{\overline{p}_s}\right),$$

where z_{js} represents the cumulative dose received by voxel j in structure s from the treatment plan. The function is normalized over the number of voxels in the structure using the coefficient $1/v_s$. This normalization allows each structure to be represented equally, regardless of how large or small it may be. By setting $\underline{w}_s, \overline{w}_s \geq 0$ and $\overline{p}_s, \underline{p}_s \geq 1$, convexity is ensured because $(T_s - z_{js})_+$ and $(z_{js} - T_s)_+$ are always nonnegative. Since it is desirable to deliver as little dose as possible to the critical structures, there is no penalty for underdosing, so $\underline{w}_s = 0$ for $s \in S$.

The cumulative dose for a voxel can be expressed as the sum of the dose received in the voxel by each beamlet in each beam. Each beamlet deposits some amount of energy into every voxel along its path. Therefore, the amount of energy deposited into voxel j by beamlet i may not be the full fluence of beamlet i, x_i. For each bixel and voxel combination, there is a dose deposition coefficient, D_{ijs}, that indicates the amount of dose received by voxel j in structure s by beamlet i at unit intensity. The total dose received by voxel j in structure s is thus given by

$$z_{js} = \sum_{h=1}^{k} \sum_{i \in B_{\theta_h}} D_{ijs} x_i \quad j = 1, \ldots, v_s, \quad s = 1, \ldots, S.$$

A basic formulation of the FMO problem is then:

$$\begin{aligned} \text{minimize} \quad & \sum_{s=1}^{S} \sum_{j=1}^{v_s} F_{js}(z_{js}) \\ \text{subject to} \quad & z_{js} = \sum_{h=1}^{k} \sum_{i \in B_{\theta_h}} D_{ijs} x_i \quad j = 1, \ldots, v_s, \quad s = 1, \ldots, S \\ & x_i \geq 0 \qquad\qquad\qquad\qquad i \in B_{\theta_h}, h = 1, \ldots, k. \end{aligned}$$

The FMO model is solved using a primal-dual interior point algorithm which returns a near-optimal solution [4,18,39]. A detailed explanation of the algorithm is beyond the scope of this paper.

3 Neighborhood search approaches

Neighborhood search approaches are common methods of obtaining solutions to global optimization problems. For a vector of decision variables, a neighbor is obtained by perturbing one or more of the decision variables. A neighborhood for a particular vector of decision variables is the set of all its neighbors for a given method of perturbating the decision variable vector. A solution is considered to be locally optimal if there are no improving solutions in its neighborhood.

Both deterministic and stochastic neighborhood search algorithms have been applied to a wide variety of optimization problems. A deterministic neighborhood search algorithm is one in which the entire neighborhood, or a pre-defined subset of the neighborhood, is enumerated in each iteration to find an improving solution. Stochastic versions of neighborhood search approaches, for example, simulated annealing, randomly select neighboring solutions in an attempt to find an improving solution in each iteration.

For the BOO problem, we consider two neighborhood search methods. The first is a deterministic neighborhood search algorithm that finds a locally optimal solution, and the second is the simulated annealing algorithm, which, although based on neighborhood searches, provably converges to the globally optimal solution for certain neighborhood structures. Depending on computing power, it may be impractical to run the simulated annealing until global convergence in practice; however, the algorithm can find good solutions quickly.

3.1 A deterministic neighborhood search method for BOO

Deterministic neighborhood search methods are optimization algorithms that start from a given solution and then iteratively select the best point in the current neighborhood as the next iterate. The best point in the neighborhood can be found by complete enumeration if the neighborhood is small, or by optimization if the neighborhood is large or if objective function evaluations are expensive. Due to the complexity of the BOO problem, even when only a subset of available orientations is considered, we will focus on smaller neighborhoods and use enumeration. However, the neighborhood could alternatively be searched heuristically, for example by searching the neighborhood until the first improving solution is found, rather than the best improving solution. If no improved solution can be found, the current solution is a local optimum.

In our implementation of the Add/Drop algorithm, a small neighborhood is desired for enumeration purposes. In each iteration, a neighborhood for just a single beam is considered.

Say a beam set consisting of k beams is desired. Letting the neighborhood of a single beam θ_h in $\boldsymbol{\theta}$ be denoted as $\mathcal{N}_h(\boldsymbol{\theta})$, the Add/Drop algorithm is as follows:

- *Initialization:*
  - Choose an initial starting solution $\boldsymbol{\theta}^{(0)}$.
  - Set $\boldsymbol{\theta}^* = \boldsymbol{\theta}^{(0)}$ and $i = 0$.
- *Iteration:*
  1. Select $h \in \{1, \ldots, k\}$, then generate $\bar{\boldsymbol{\theta}} \in \mathcal{N}_h(\boldsymbol{\theta}^{(i)})$.
  2. If $F(\bar{\boldsymbol{\theta}}) < F(\boldsymbol{\theta}^*)$, set $\boldsymbol{\theta}^* = \boldsymbol{\theta}^{(i+1)} = \bar{\boldsymbol{\theta}}$ and set $i \leftarrow i + 1$.
  3. If all points in $\cup_{h=1}^{k} \mathcal{N}_h(\boldsymbol{\theta}^{(i)})$ have been sampled without improvement, stop with $\boldsymbol{\theta}^*$ as a local minimum. Otherwise, repeat Step 1.

3.1.1 Neighborhood definition

In each step of the Add/Drop algorithm, a beam in the current solution is replaced with an improving beam in its neighborhood. Rather than define a neighbor as related to an entire beam vector, the neighborhoods of individual beams are considered. The neighborhood of a single beam θ_h in $\boldsymbol{\theta}$ is defined as $\mathcal{N}_h(\boldsymbol{\theta}) = \{(\theta_1, \ldots, \theta_{h-1}, \theta \bmod 360, \theta_{h+1}, \ldots, \theta_k) \in \mathcal{B}^k : \theta_h - \delta \leq \theta \leq \theta_h + \delta\}$. In other words, the neighborhood of a beam is all beams within $\pm\delta$ degrees taking into account the cyclic nature of the angles. The cyclicality of the angles refers to the fact that all angles can be represented by degrees in [0,360]. For example, $400^\circ = 40^\circ$ and $-100^\circ = 260^\circ$. The expression $\theta \bmod 360$ captures this cyclicality.

3.1.2 Neighbor selection

The process of selecting a neighboring point in each iteration consists of two steps: selecting the index h to change and then selecting an improving angle in $\mathcal{N}_h(\boldsymbol{\theta})$ to replace θ_h. If h is selected as $i \bmod k + 1$, the algorithm will cycle through each index sequentially, similar to a Gibbs Sampler (see, for example, Geman and Geman [17] and Gelfand and Smith [16]). The Gibbs Sampler also uses a similar two-step approach to generating a new point by sequentially generating a new value for each variable in turn. If h is selected randomly in each iteration, the resulting algorithm is similar to a Hit-and-Run method (see, for example, Smith [51] and Bélisle [7]), in which a variable to be changed is selected randomly, and then a new value for that variable is also selected randomly within a neighborhood.

Once h is selected, the new value for θ_h can be generated by enumeration or by a heuristic method. The Add/Drop algorithm compares the quality of the new solution to the current solution, and then only accepts improving solutions. This greedy approach results in a locally optimal solution.

3.1.3 Implementation

The index of beam angle to be changed in each iteration, h in Step 1 of the algorithm in Sect. 3.1, is chosen as $h = i \bmod k + 1$ to cycle through each index in a sequential manner. In the Add/Drop implementation, once h is determined, $\bar{\boldsymbol{\theta}}$ in iteration i is chosen as $\bar{\boldsymbol{\theta}} = \arg\min_{\boldsymbol{\theta} \in \mathcal{N}_h(\boldsymbol{\theta}^{(i)})} \{F(\boldsymbol{\theta})\}$. By replacing each beam by the most improving neighbor, the Add/Drop algorithm is a greedy heuristic which terminates when there is no improving neighbor for any beam.

Additionally, a multi-start aspect is added by repeating the algorithm with multiple initial starting points. Because the only constraint in our FMO model is that beamlet fluences are

nonnegative, any beam solution will provide a feasible FMO solution, although clearly some beam solutions will provide better optimal FMO solution values than others. For example, one strategy to select starting points would be to select a random starting point according to some distribution. Another strategy would be to select an equi-spaced solution and rotate it a fixed number of times to obtain new starting points until the initial equi-spaced solution is repeated. Equi-spaced beam solutions are common in clinical practice for an odd number of beams. The reason that such a method is not generally used in practice for even-numbered beams is that the resulting beam set would contain parallel-opposed beams (beams that lie on the same line), which are not used by convention as it is believed that the effect of a parallel-opposed beam is very similar to simply doubling the radiation delivered from a beam. If an equi-spaced solution is not possible given a beam set of k beams and the discretization level of the candidate beam set $\mathcal{B}$, then the solution can be rounded so that $\boldsymbol{\theta}_h^{(0)} \in \mathcal{B}$, $h = 1, \ldots, k$.

3.2 Simulated annealing

The simulated annealing algorithm used is similar to the classical simulated annealing approach proposed in Kirkpatrick et al. [25]. The simulated annealing algorithm is based on the Metropolis algorithm, wherein a neighboring solution to the current iterate is generated, and if it is an improving point, it becomes the current iterate. Otherwise, it becomes the current iterate with probability $\exp\{\Delta F/T\}$, where ΔF is the difference in FMO value between the current iterate and the newly generated point and T is the *temperature*, a measure of the randomness of the algorithm. If $T = 0$, then only improving points are selected. If T is very large, then any move is accepted, which is essentially a random search.

The simulated annealing algorithm starts with an initial temperature T_0 and performs a number of iterations of the Metropolis algorithm using $T = T_0$. Then, the temperature is decreased according to some *cooling schedule* such that $\{T_i\} \to 0$.

Obvious parallels can be drawn between the simulated annealing algorithm and the Add/Drop neighborhood search method described in Sect. 3.1. While the Add/Drop algorithm deterministically searches the neighborhood for improving solutions, the simulated annealing algorithm randomly selects neighboring solutions. Rather than being limited by the ability to only move to improving solutions, the simulated annealing algorithm may still move to a non-improving solution with a certain probability, thus allowing for the escape from local minima. The Add/Drop algorithm, on the other hand, is a greedy algorithm that is specifically designed to find local minima.

The simulated annealing algorithm is essentially a randomization of the Add/Drop algorithm. In addition to the added randomness, the possibility of changing more than one beam in each iteration is allowed by selecting a set of indices $H \subseteq \{1, \ldots, k\}$ to change, rather than just selecting a single index h. The simulated annealing algorithm is as follows:

- *Initialization:*
  - Choose an initial beam set $\boldsymbol{\theta}^{(0)}$ and calculate its FMO objective function value F_0.
  - Set $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}^{(0)}$, $\hat{F} = F_0$, $i = 0$.
- *Iteration:*
  1. Select $H \subseteq \{1, \ldots, k\}$, generate $\boldsymbol{\theta} \in \cup_{h \in H} \mathcal{N}_h(\boldsymbol{\theta}^{(i)})$, and calculate its FMO objective function value F.
  2. If $F < \hat{F}$, set $\hat{F} = F$, $F_{i+1} = F$, $\boldsymbol{\theta}^{(i+1)} = \boldsymbol{\theta}$ and $\hat{\boldsymbol{\theta}} = \boldsymbol{\theta}$. Otherwise, set $F_{i+1} = F$ and $\boldsymbol{\theta}^{(i+1)} = \boldsymbol{\theta}$ with probability $\exp\{(F_i - F)/T_i\}$.
  3. Set $i \leftarrow i + 1$ and repeat Step 1.

The simulated annealing algorithm has been previously applied to the BOO problem. Bortfeld and Schlegel [10] use the "fast" simulated annealing algorithm described by Szu and Hartley [8] which employs a Cauchy distribution in generating neighboring points. Stein et al. [53], Rowbottom [47] and Djajaputra et al. [14] also use a Cauchy distribution in generating neighboring solutions. Lu et al. [34] randomly select new points satisfying BEV and conventional wisdom criteria and Pugachev and Xing [42] randomly generate new points and then vary them according to an exponential distribution. All accept improving solutions, and with the exception of Rowbottom et al. [47] who only accept improving solutions (essentially $T_i = 0$ for all i), all accept non-improving solutions with a Boltzmann probability. Additionally, none of the previous BOO studies employing simulated annealing use the exact FMO as a measure of the quality of a beam set.

3.2.1 Neighborhood definition

Two neighborhood structures are explored. The first neighborhood is similar to that described in Sect. 3.1.1 in that a neighborhood $\mathcal{N}_h(\boldsymbol{\theta})$ is considered for only a single beam index $h \in \{1, \ldots, k\}$, just as in the Add/Drop method.

As an extension to changing a single angle in each iteration, we also consider a neighborhood that involves changing all beams in each iteration, corresponding to $H = \{1, \ldots, k\}$ in Step 1 of the simulated annealing algorithm in Sect. 3.2. This neighborhood is defined as $\mathcal{N}(\boldsymbol{\theta}) = \cup_{h=1}^{k} \mathcal{N}_h(\boldsymbol{\theta})$. Again, the neighborhoods for the individual beams are defined as in the first method, with bounds of $\pm\,\delta$ degrees.

3.2.2 Neighbor selection

The method of selecting a neighbor depends on the neighborhood structure as described in Sect. 3.2.1. In the first method where only one beam is changed at a time, a neighbor is selected using the randomized approach described in Sect. 3.1.2. Once h is selected, the probability of selecting a particular solution in $\mathcal{N}_h(\boldsymbol{\theta})$ where the new θ is d degrees from θ_h is $P\{D = d\}$, where D is the realization of a random variable of some probability distribution defined on the interval $[-\delta, -\delta + 1, \ldots, \delta]$.

For the neighborhood $\mathcal{N}(\boldsymbol{\theta})$ where all beams are changed in an iteration, the new value for each beam $h \in \{1, \ldots, k\}$ is generated from $\mathcal{N}_h(\boldsymbol{\theta})$ in the same manner described above.

3.2.3 Implementation

In addition to basing our algorithms on the exact FMO solution rather than on heuristics or scoring measures, our simulated annealing approach differs from the previous studies in the distribution used to generate neighbors, the definition of the neighborhood, the cooling schedule and the number of iterations/restarts used. Not only do we use a new neighborhood structure, but also a geometric probability distribution rather than a uniform or Cauchy distribution on the neighborhood. The geometric distribution is similar in shape to the Cauchy distribution in that they both can have fat tails depending on the choice of probability parameters. The fat tails of these distributions allow for points far away from the current solution to be selected as successive iterates, which potentially increases the likelihood of finding a globally optimal solution. The geometric distribution has the added attractiveness of producing discrete solutions, which is desirable for the BOO problem in which discrete solutions are preferred.

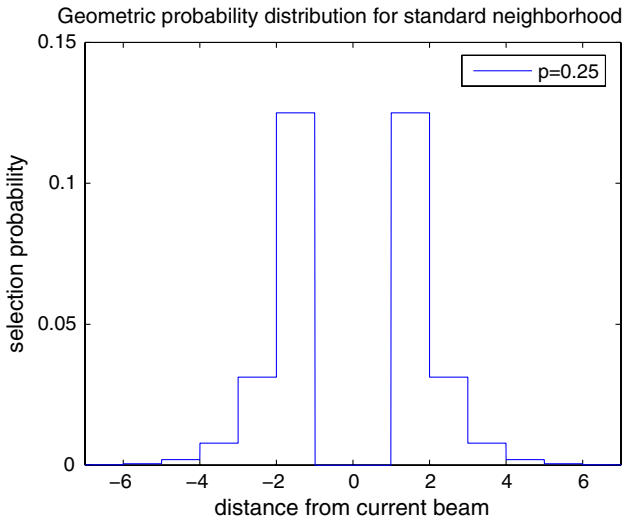

Fig. 2 Selection probability in $\mathcal{N}_h(\boldsymbol{\theta})$

By using the cooling schedule $T_{i+1} = \alpha T_i$ with $\alpha < 1$, the sequence of temperatures $\{T_i\}$ converges to zero as the number of iterations increases. In our approach, the neighborhood of a beam for both the $\mathcal{N}_h(\boldsymbol{\theta})$ and $\mathcal{N}(\boldsymbol{\theta})$ neighborhoods is defined using $\delta = 180$, that is, $\mathcal{N}_h(\boldsymbol{\theta}) = \mathcal{B}$. By defining the neighborhood of each beam to be the entire single-beam solution space, the simulated annealing algorithm converges to the global optimum when using the neighborhood $\mathcal{N}(\boldsymbol{\theta})$ defined in Sect. 3.2.1. Though $\mathcal{N}_h(\boldsymbol{\theta})$ is large, each beam in $\mathcal{N}_h(\boldsymbol{\theta})$ is assigned a probability so that only the beams closest to θ_h have a significant probability of being selected. Figure 2 shows the probability of replacing θ_h with beams at varying distances using probability $p = 0.25$ for the geometric distribution. Note that the current beam cannot be selected as a replacement.

Additionally, as with the Add/Drop method, a multi-start aspect is added to the simulated annealing algorithm by repeating the algorithm using several different starting points.

3.2.4 Convergence

The type of neighborhood structure can affect the global convergence property of the simulated annealing algorithm. Unlike many previously proposed simulated annealing algorithms, our algorithm converges to the globally optimal solution to the BOO problem under mild conditions. The following theorem summarizes these conditions and shows that our simulated annealing algorithm retains the global convergence property for some of the tested neighborhood structures.

Theorem 3.1 *Suppose that*

- $H = \{1, \ldots, k\}$
- $\lim_{i \to \infty} T_i = 0$
- $\delta = 180$
- *There is a positive probability of generating any solution in the neighborhood.*

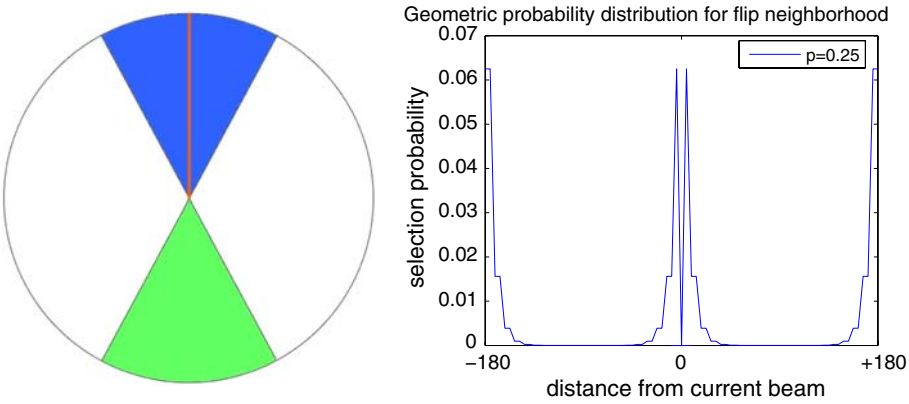

Fig. 3 The flip neighborhood. (**a**) Graphical depiction of $\mathcal{N}_h(\boldsymbol{\theta})$ (top shaded area) and $\mathcal{N}_h^F(\boldsymbol{\theta})$ (top and bottom shaded areas) for $\theta_h = 0$. (**b**) Selection probability in $\mathcal{N}_h^F(\boldsymbol{\theta})$

Then our simulated annealing algorithm converges to the global optimum solution in the sense that

$$\lim_{i\to\infty} F_i = F^* \qquad \textit{in probability}$$

where F^* is the global optimum value of the BOO problem.

Proof This follows from Theorem 1 in Bélisle et al. [6]. □

3.3 A new neighborhood structure

For the BOO problem, the neighborhood structure that is typically used for a vector of beam orientations is simply the collection of beam vectors obtained from changing one or more of the beams to a neighboring beam, where each beam has its own neighborhood $\mathcal{N}_h(\boldsymbol{\theta})$.

In addition to $\mathcal{N}_h(\boldsymbol{\theta})$, we consider a new neighborhood which we call a "flip" neighborhood. The flip neighborhood of a beam index h consists of $\mathcal{N}_h(\boldsymbol{\theta})$ plus a neighborhood around the parallel opposed beam of h (the beam 180° away, $h' = (\theta_h + 180) \bmod 360$). The flip neighborhood can be defined as $\mathcal{N}_h^F(\boldsymbol{\theta}) = \{(\theta_1, \ldots, \theta_{h-1}, \theta \bmod 360, \theta_{h+1}, \ldots, \theta_k) \in \mathcal{B}^k : \theta \in [\theta_h - \delta, \theta_h + \delta] \cup [\theta_h + 180 - \delta^F, \theta + 180 + \delta^F]\}$. Note that the values δ and δ^F may be different. Figure 3a depicts a flip neighborhood for a beam located at 0° degrees, the center of the top shaded wedge representing $\mathcal{N}_h(\boldsymbol{\theta})$, where $\theta_h = 0$.

The motivation for the flip neighborhoods arises from the observation that many of the 3-beam simulated annealing plans generated using the regular neighborhood contained two beams very close to two beams in the optimal solution (obtained by explicit enumeration), while the third beam was very close to the parallel opposed beam of the third beam in the optimal solution. Given this observation, it is intuitive that the inclusion of the neighborhood around the parallel beam should provide improved solutions.

The neighborhoods $\mathcal{N}_h(\boldsymbol{\theta})$ and $\mathcal{N}_h^F(\boldsymbol{\theta})$ with varying δ^F values are applied to both the Add/Drop and the simulated annealing frameworks. For the geometric probability distribution used in the simulated annealing method, Fig. 3b shows the probability of selecting beams at different distances using a flip neighborhood with probability $p = 0.25$. Note that the current beam cannot be selected as its own neighbor.

4 Results

The simulated annealing method was tested on six head-and-neck cases using a Windows XP computer with a 2.13 GHz Pentium M processor and 2 GB of RAM. On average, $\approx$ 340 FMOs were calculated in the 30-min run time allowed for the simulated annealing and Add/Drop algorithms. It should be noted that while 30 min is enough time for the simulated annealing algorithm to obtain a good solution, it is not enough time to obtain the globally optimal solution. Because coplanar beams are generally used exclusively in head-and-neck cancer treatments, only coplanar beam solutions are investigated here. Beams were selected on a 5-degree grid, yielding 72 candidate coplanar beams.

The simulated annealing and Add/Drop algorithms were used to obtain 4-beam coplanar plans using regular and flip neighborhoods. Plans consisting of four beams were obtained because previous tests on 3-beam plans indicate that it is unlikely that only three coplanar beams will deliver a quality treatment plan. We will show that quality treatment plans can be obtained using as few as four coplanar beams, which can lead to shorter treatment times.

In order to compare the quality of the treatment across different plans, the plans are compared to equi-spaced treatment plans of four, five and seven beams whose fluences are optimized using the same FMO model described in Sect. 2.2. The beams in the equi-spaced plans are obtained by dividing 360° by the number of beams. For 5-beam and 7-beam equi-spaced plans, this type of treatment plan is typical for head-and-neck cases in clinical practice in the Davis Cancer Center at Shands Hospital at the University of Florida. The 5- and 7-beam equi-spaced plans are chosen for comparison as they are common treatment plans for the head-and-neck site. Additionally, this comparison will show that the algorithms are successful enough to compete with equi-spaced plans with nearly double the number of beams. The 4-beam equi-spaced plan is shown to illustrate that in order to use as few as four beams, BOO must be applied.

Figure 4 demonstrates the improved convergence times possible using the flip neighborhood.

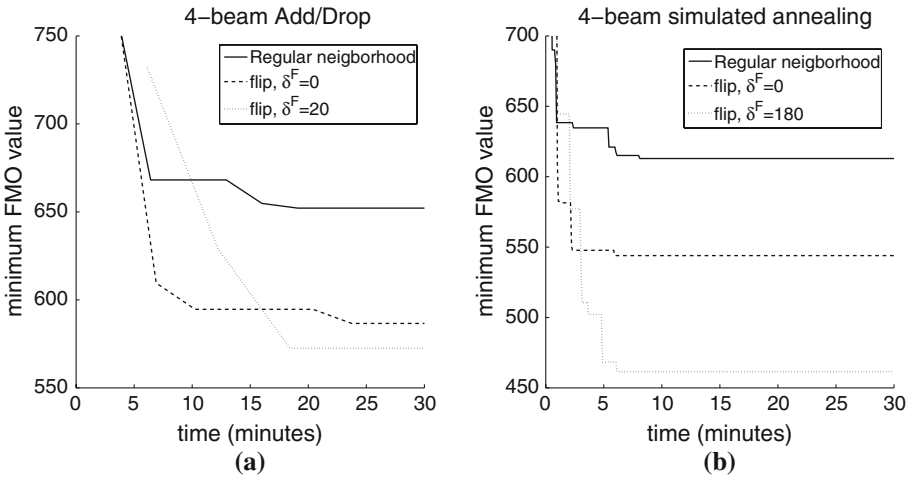

Fig. 4 Add/Drop (**a**) and simulated annealing (**b**) comparison of FMO convergence

4.1 Add/Drop algorithm results

The Add/Drop algorithm was allowed to run for 30 min to generate a 4-beam plan. The $\mathcal{N}_h(\boldsymbol{\theta})$ neighborhood with $\delta = 20$ and the $\mathcal{N}_h^F(\boldsymbol{\theta})$ with $\delta^F = 0$ and $\delta^F = 20$ neighborhoods are tested for the Add/Drop algorithm. The value $\delta = 20$ is chosen to approximate the neighborhood size that is expected from the simulated annealing implementation using a large flip neighborhood, where $\delta^F = 180$. More details on the simulated annealing implementations are provided in Sect. 4.2.

Using $\mathcal{N}_h(\boldsymbol{\theta})$, the 4-beam Add/Drop solution is nearly identical to the 7-beam equi-spaced plan, while the flip neighborhoods allow the Add/Drop algorithm to find 4-beam solutions that exceed the quality of the 7-beam plans. Figure 4a illustrates that the flip neighborhoods provide faster FMO convergence than that of $\mathcal{N}_h(\boldsymbol{\theta})$, while Figs. 7 and 9 demonstrates the quality of the solutions.

4.2 Simulated annealing results

Several parameter sets were tested for the simulated annealing algorithm. For simplicity, each of the parameter sets and methods of generating a neighboring solution are numbered according to Table 1. The value m refers to the number of Metropolis iterations performed before modifying the temperature, while n is the number of times the temperature is modified. The value N is the number of beams that are changed to neighboring beams in each iteration.

For the cooling schedule, we update the temperature according to an exponential cooling schedule, $T_{i+1} = \alpha T_i$, where $\alpha < 1$. Due to the random nature of the algorithm, the algorithm is restarted five times, each time with a different initial starting point. The first initial starting point is an equi-spaced solution, and each subsequent starting point is the previous initial solution rotated by d degrees, where candidate angles are considered on a d-degree grid, that is, every dth angle is considered. The number of simulated annealing and Metropolis iterations are chosen such that the total number of iterations is 500. To ensure clinical practicality, the algorithm was allowed to run for a maximum of 30 min or 500 iterations, whichever came first.

Table 1 Definitions of implementations

Number	n	m	N	α	T_0
1	100	1	1	0.9	0
2	10	10	1	0.9	0
3	100	1	1	0.99	0
4	10	10	1	0.99	0
5	100	1	1	0.9	75
6	10	10	1	0.9	75
7	100	1	1	0.99	75
8	10	10	1	0.99	75
9	100	1	all	0.9	0
10	10	10	all	0.9	0
11	100	1	all	0.99	0
12	10	10	all	0.99	0
13	100	1	all	0.9	75
14	10	10	all	0.9	75
15	100	1	all	0.99	75
16	10	10	all	0.99	75

The initial temperature values tested are $T_0 = 0$ and $T_0 = 75$. $T_0 = 0$ results in the acceptance of only improving solutions, while the initial temperature value 75 was selected as the value that would approximately yield a 50% probability of selecting a non-improving solution for the initial iterations of the algorithm.

For both the $\mathcal{N}_h(\boldsymbol{\theta})$ and $\mathcal{N}_h^F(\boldsymbol{\theta})$ neighborhoods, $\delta = \delta^F = 180$ is used so that the entire solution space is considered as a neighborhood. As shown in Fig. 2, the probability of selecting a beam 20° away using the $\mathcal{N}_h(\boldsymbol{\theta})$ neighborhood with geometric distribution with $p = 0.25$ is only 0.39% on a 5° grid. We consider this sufficiently small to not consider neighborhoods larger than $\delta = 20$ for $\mathcal{N}_h(\boldsymbol{\theta})$ and $\delta^F = 20$ for $\mathcal{N}_h^F(\boldsymbol{\theta})$ in the Add/Drop algorithm. Additionally, just as in the Add/Drop implementation, the neighborhood $\mathcal{N}_h^F(\boldsymbol{\theta})$ with $\delta^F = 0$ is also considered.

Figure 4b shows that the flip neighborhoods converge in FMO value significantly faster than does the $\mathcal{N}_h(\boldsymbol{\theta})$ neighborhood, while Figures 8 and 10 shows that the flip neighborhoods provide comparable solution quality to both the non-flip simulated annealing and 7-beam equi-spaced solutions.

4.3 Clinical implementation

Because there is no fundamental way of quantifying a treatment plan, a tool commonly used by physicians to judge the quality of a treatment plan is the *dose-volume histogram* (*DVH*). A DVH is a graphical measure of the cumulative dose received by a given structure. It specifies the percentage of each structure's volume that receives at least a certain amount of dose, thus providing an intuitive means of assessing the quality of a treatment plan. Regardless of the methods used to formulate and solve IMRT optimization problems, DVHs must be examined as a fail-safe check that the solutions are clinically acceptable. For a given beam solution, there may be many clinically satisfactory treatment plans, however, some treatment plans will be better than others. For example, a physician may decide that a treatment plan which spares three out of four saliva glands is acceptable, but a better treatment plan using a different FMO solution may spare all four saliva glands. For this reason, we use the optimal FMO solution to measure the BOO solutions.

The plans tested each contain two target structures. The *gross tumor volume* (*GTV*) is the tumor mass observed from imaging scans. The *clinical tumor volume* (*CTV*) is the GTV plus some margin specified by the physician. The CTV is used by physicians in case there are elements of the tumor mass that cannot be seen from the imaging scans, and the dose prescribed for the CTV is less than the dose prescribed by the GTV.

For target structures, we require that at least 95% of the target receives the full prescription dose. Typically, a good plan ensures that the target receives a homogeneous dose at the prescription level indicated by the physician. A *hotspot* occurs if a significant percentage of the target receives more than the prescription dose. Similarly, a *coldspot* occurs where less than a certain percent of the target receives the prescription dose. Because the GTV is contained inside the CTV, the CTV will necessarily have a sizable, but less important, hotspot. Thus, we are only concerned with GTV overdose.

The sparing criteria for each of the common critical structures in head-and-neck cases are listed in Table 2. The saliva glands, the left and right submandibular and parotid glands, are of particular importance in head-and-neck treatment planning. Because of their close proximity to the usual tumor locations, the four saliva glands can be difficult to spare. Studies show that a patient can lead a relatively normal life with three of the four glands spared. The loss of other organs, especially the spinal cord, will also greatly affect the patient's quality of life, but head-and-neck tumors are usually situated in such a way that other organs can be

Table 2 Sparing criteria varies for each critical structure

Structure	Sparing criteria
Brain stem	100% $\leq$ 55 Gy
Eyes	50% $\leq$ 30 Gy
Mandible	100% $\leq$ 70 Gy
Optic chiasm	100% $\leq$ 55 Gy
Optic nerves	100% $\leq$ 50 Gy
Parotid glands	100% $\leq$ 30 Gy
Skin	100% $\leq$ 60 Gy
Spinal cord	100% $\leq$ 45 Gy
Submandibular glands	100% $\leq$ 30 Gy

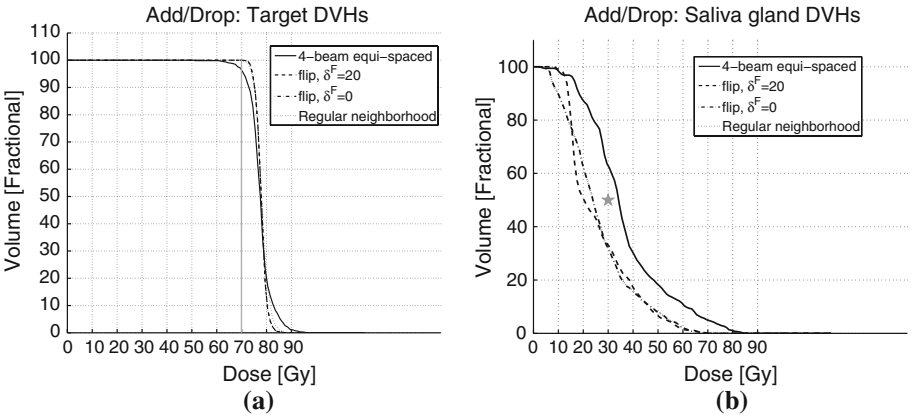

Fig. 5 (**a**) The Add/Drop plans obtain superior target coverage to the 4-beam equi-spaced plan and still achieve a smaller hotspot. (**b**) The Add/Drop plans spare all saliva glands, while the 4-beam equi-spaced plan spares only one

easily spared in the FMO optimzation. For this reason, the DVH results provided include only target structures and saliva glands for clarity, although the model does in fact account for and spare the organs listed in Table 2.

DVHs for a representative case comparing the 4-beam equi-spaced plan with the Add/Drop and simulated annealing plans are shown in Figs. 5 and 6. DVHs for the same case comparing the 5-beam equi-spaced plan are shown in Figs. 7 and 8. Likewise, DVHs comparing the 7-beam equi-spaced plan with the Add/Drop and simulated annealing plans are shown in Figs. 9 and 10.

The simulated annealing plans are obtained using a regular neighborhood and flip neighborhoods with $\delta^F = 0$ and $\delta^F = 180$. The Add/Drop plans are obtained using a regular neighborhood and flip neighborhoods with $\delta^F = 0$ and $\delta^F = 20$. The sparing criteria used for the saliva glands, no more than 50% of the gland receiving 30 Gy, is marked by the star in the saliva gland DVHs. Note that from a clinical standpoint, the amount by which an organ is spared is unimportant; it is only important that the organ is spared. The prescription dose for the GTV is 73.8 Gy, which is marked by the vertical line in the target DVHs. As previously stated, for target structures, we require that at least 95% of the target receives the full prescription dose.

Figures 5 and 6 illustrate that a 4-beam equi-spaced treatment plan yields a far from satisfactory treatment. The targets are underdosed and the saliva gland sparing is poor. The 4-beam

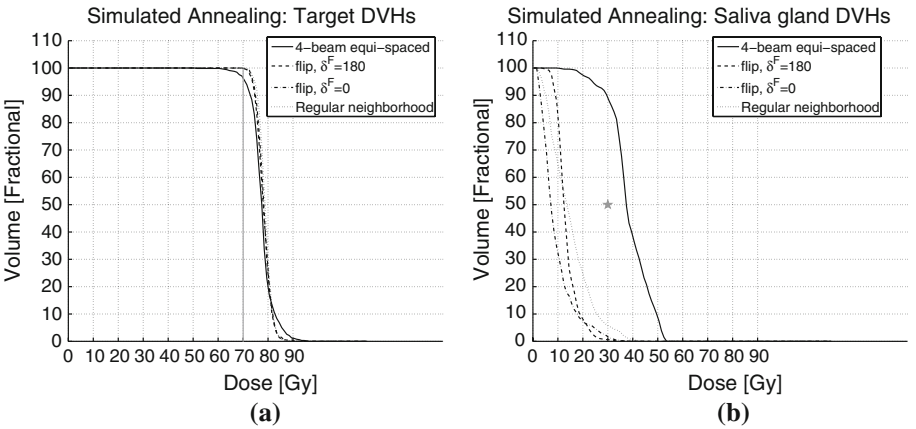

Fig. 6 (**a**) The simulated annealing plans obtain superior target coverage to the 4-beam equi-spaced plan and still achieve a smaller hotspot. (**b**) The simulated annealing plans spare all saliva glands, while the 4-beam equi-spaced plan saves only one

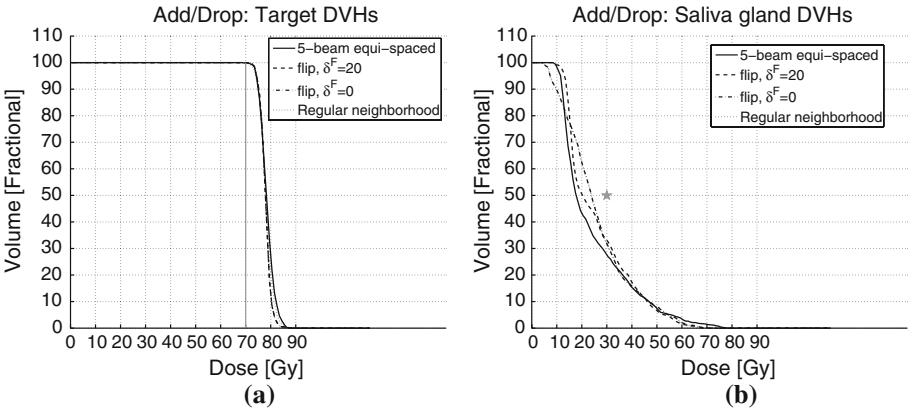

Fig. 7 (**a**) The Add/Drop plans achieve nearly identical target coverage when compared to the 5-beam equi-spaced plan. (**b**) The saliva gland sparing in the Add/Drop plans and the 5-beam equi-spaced plan is also clinically equivalent

Add/Drop and simulated annealing plans, however, achieve acceptable target coverage with a smaller hotspot than the 4-beam equi-spaced plan. The Add/Drop and simulated annealing plans also spare more saliva glands than the 4-beam equi-spaced plan.

In comparison to the 5-beam equi-spaced plan, the 4-beam Add/Drop and simulated annealing plans obtained remarkably similar target coverage. Figure 7 shows that the Add/Drop plans were able to achieve a smaller hotspot, whereas Fig. 8 shows that the simulated annealing plans are indistinguishable from the 5-beam equi-spaced plans in terms of target coverage. For both Add/Drop and simulated annealing, the 4-beam plans show clinically equivalent saliva gland sparing, although the simulated annealing plans deliver less dose to the saliva glands.

Figure 9 shows that the 4-beam Add/Drop plans obtain very similar treatments when compared to the 7-beam equi-spaced DVHs. Each of the Add/Drop plans are comparable to the 7-beam equi-spaced plan in terms of both saliva gland sparing and target coverage.

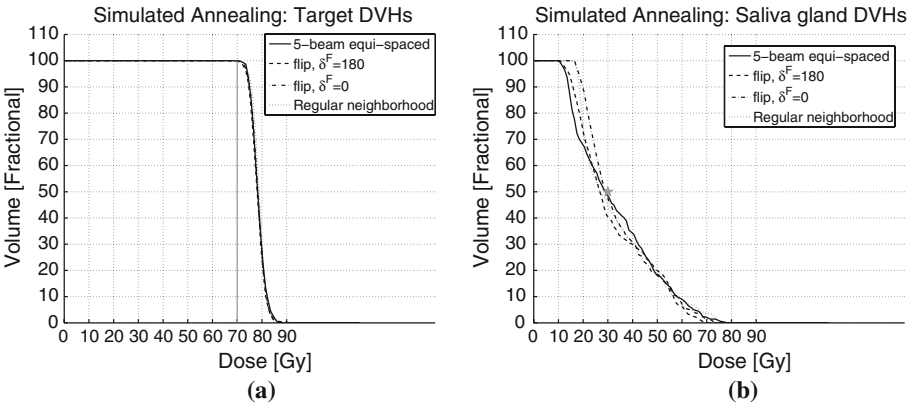

Fig. 8 (**a**) The simulated annealing plans achieve virtually indistinguishable target coverage with the 5-beam equi-spaced plan. (**b**) The saliva gland sparing in the simulated annealing plans and the 5-beam equi-spaced plan is also clinically equivalent

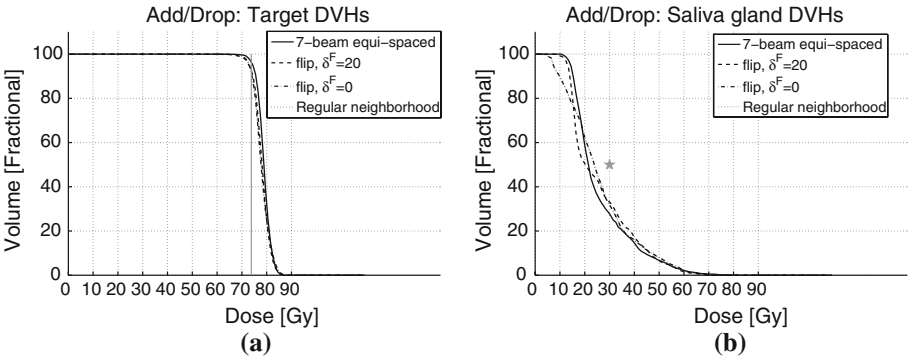

Fig. 9 (**a**) The Add/Drop plans achieve nearly identical target coverage when compared to the 7-beam equi-spaced plan. (**b**) The saliva gland sparing in the Add/Drop plans and the 7-beam equi-spaced plan is also clinically equivalent

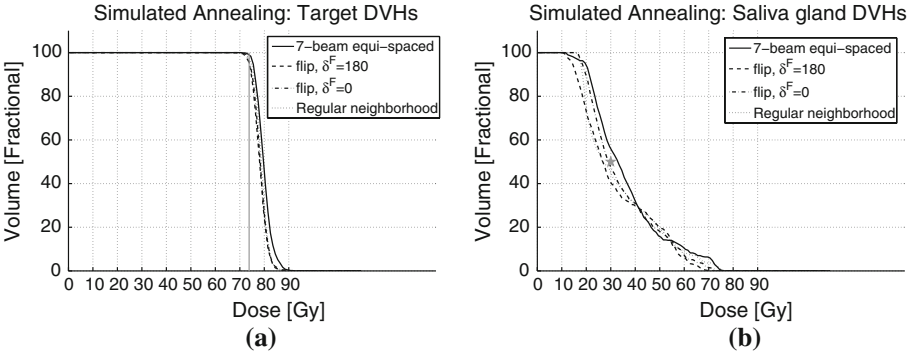

Fig. 10 (**a**) Unlike the 7-beam equi-spaced plan, the 4-beam simulated annealing plans do not overdose the target. (**b**) The simulated annealing plans are also capable of sparing more saliva glands than the 7-beam equi-spaced plan

Figure 10 reveals that the 7-beam equi-spaced plan actually overdoses the target and has a larger hotspot than the 4-beam simulated annealing plans. The 7-beam equi-spaced plan only spares three of the four saliva glands, whereas the 4-beam simulated annealing plans spare three or more saliva glands. The simulated annealing plans obtained using the flip neighborhoods spare all four saliva glands, while the plan obtained by the $\mathcal{N}_h(\boldsymbol{\theta})$ neighborhood only spares three saliva glands, indicating that the flip neighborhoods do in fact find superior solutions in terms of clinical quality.

5 Conclusions

We have shown that for head-and-neck cases, quality plans with fewer beams than a standard treatment plan can be obtained if BOO is applied. The simulated annealing and Add/Drop algorithms both regularly obtained quality treatment plans with as few as four beams in only 30 min. The use of the flip neighborhood improves the rate of FMO convergence in both algorithms, and even has the ability to improve organ sparing as shown in the simulated annealing results. The simulated annealing and Add/Drop algorithms performed comparably to each other, with neither algorithm indicating a significant benefit over the other.

It is possible to incorporate flip neighborhoods into other BOO algorithms that rely on neighborhood searches to yield improved treatment plans in clinically acceptable time frames.

Acknowledgements The work of D. M. Aleman was supported by The Alliance for Graduate Education and the Professoriate and a Graduate research Fellowship of the National Science Foundation. The work of C. J. F. Dempsey was supported by the National Science Foundation under grant no. DMI-0457394.

References

1. Alber, M., Nusslin, F.: An objective function for radiation treatment optimization based on local biological measures. Phys. Med. Biol. **44**, 479–493 (1999)
2. Aleman, D.M., Romeijn, H.E., Dempsey, J.F.: Beam orientation optimization methods in intensity modulated radiation therapy treatment planning. IIE Conference Proceedings (2006)
3. Aleman, D.M., Romeijn, H.E., Dempsey, J.F.: A response surface approach to beam orientation optimization in intensity modulated radiation therapy treatment planning. INFORMS J. Comput: Coumputat. Biol. Med. Appl. (2008, to appear)
4. Aleman, D.M., Glaser, D., Romeijn, H.E., Dempsey, J.F.: A primal-dual interior point algorithm for fluence map optimization in intensity modulated radiation therapy treatment planning. Work in progress (2008)
5. Bednarz, G., Michalski, D., Houser, C., Huq, M.S., Xiao, Y., Anne, P.R., Galvin, J.M.: The use of mixed-integer programming for inverse treatment planning with pre-defined field segments. Phys. Med. Biol. **47**, 2235–2245 (2002)
6. Bélisle, C.J.P.: Convergence theorems for a class of simulated annealing algorithms on $\mathbb{R}^d$. J. Appl. Prob. **29**, 885–895 (1992)
7. Bélisle, C.J.P., Romeijn, H.E., Smith, R.L.: Hit-and-run algorithms for generating multivariate distributions. Math. Oper. Res. **18**, 255–266 (1993)
8. Bomze, I.: Fast simulated annealing. Phys. Lett. **122A**, 157–162 (1987)
9. Bortfeld, T.: Optimized planning using physical objectives and constraints. Semin. Radiat. Oncol. **9**, 20–34 (1999)
10. Bortfeld, T., Schlegel, W.: Optimization of beam orientations in radiation therapy: some theoretical considerations. Phys. Med. Biol. **38**, 291–304 (1993)

11. Chen, G.T., Spelbring, D.R., Pelizzari, C.A., Balter, J.M., Myrianthopoulos, L.C., Vijayakumar, S., Halpern, H.: The use of beams eye view volumetrics in the selection of non-coplanar radiation portals. Int. J. Radiat. Oncol. Biol. Phys. **23**, 153–163 (1992)
12. Cho, B.C.J., Roa, H.W., Robinson, D., Murray, B.: The development of target-eye-view maps for selection of coplanar or noncoplanar beams in conformal radiotherapy treatment planning. Med. Phys. **26**, 2367–2372 (1999)
13. Das, S.K., Marks, L.B.: Selection of coplanar or noncoplanar beams using three-dimensional optimization based on maximum beam separation and minimized nontarget irradiation. Int. J. Radiat. Oncol. Biol. Phys. **38**, 643–655 (1997)
14. Djajaputra, D., Wu, Q., Wu, Y., Mohan, R.: Algorithm and performance of a clinical IMRT beam-angle optimization system. Phys. Med. Biol. **48**, 3191–3212 (2003)
15. Ezzell, G.A.: Genetic and geometric optimization of three-dimensional radiation therapy treatment planning. Med. Phys. **23**, 293–305 (1996)
16. Gelfand, A.E., Smith, A.F.M.: Sampling based approaches to calculating marginal densities. J. Am. Stat. Assoc. **85**, 398–409 (1990)
17. Geman, S., Geman, D.: Stochastic relaxation, gibbs distributions, and the bayesian restoration of images. IEEE Trans. Pattern Analy. Mach. Intell. **6**, 721–741 (1984)
18. Glaser, D.: A primal-dual interior-point method for convex Fluence Map Optimization problems. Master's thesis, Royal Institute of Technology (2005)
19. Goitein, M., Abrams, M., Rowell, D., Pollari, H., Wiles, J.: Multi-dimensional treatment planning: Ii. beams eye-view, back projection, and projection through CT sections. Int. J. Radiat. Oncol. Biol. Phys. **9**, 789–797 (1983)
20. Gokhale, P., Hussein, E.M., Kulkarni, N.: The use of beams eye view volumetrics in the selection of non-coplanar radiation portals. Med. Phys. **23**, 153–163 (1994)
21. Haas, O.C., Burnham, K.J., Mills, J.: Optimization of beam orientation in radiotherapy using planar geometry. Phys. Med. Biol. **43**, 2179–2193 (1998)
22. Hamacher, H.W., Küfer, K.-H.: Inverse radiation therapy planning a multiple objective optimization approach. Discr. Appl. Math. **118**, 145–161 (2002)
23. Jones, L.C., Hoban, P.W.: Treatment plan comparison using equivalent uniform biologically effective dose (eubed). Phys. Med. Biol. **45**, 159–170 (2000)
24. Kallman, P., Lind, B.K., Brahme, A.: An algorithm for maximizing the probability of complication–free tumor–control in radiation-therapy. Phys. Med. Biol. **37**, 871–890 (1992)
25. Kirkpatrick, S., Gelatt, C.D.: Optimization by simulated annealing. Science **220**, 671–680 (1983)
26. Kumar, A.: *Novel methods for intensity-modulated radiation therapy treatment planning*. PhD thesis, University of Florida (2005)
27. Langer, M., Brown, R., Urie, M., Leong, J., Stracher, M., Shapiro, J.: Large-scale optimization of beam weights under dose-volume restrictions. Int. J. Radiat. Oncol. Biol. Phys. **18**, 887–893 (1990)
28. Langer, M., Morrill, S., Brown, R., Lee, O., Lane, R.: A comparison of mixed integer programming and fast simulated annealing for optimizing beam weights in radiation therapy. Med. Phys. **23**, 957–964 (1996)
29. Lee, E.K., Fox, T., Crocker, I.: Simultaneous beam geometry and intensity map optimization in intensity-modulated radiation therapy treatment planning. Ann. Oper. Res. **119**, 165–181 (2003)
30. Lee, E.K., Fox, T., Crocker, I.: Integer programming applied to intensity-modulated radiation therapy treatment planning. Int. J. Radiat. Oncol. Biol. Phys. **64**, 301–320 (2006)
31. Li, Y., Yao, J., Yao, D.: Automatic beam angle selection in IMRT planning using genetic algorithm. Phys. Med. Biol. **49**, 1915–1932 (2004)
32. Li, Y., Yao, J., Yao, D., Chen, W.: A particle swarm optimization algorithm for beam angle selection in intensity-modulated radiotherapy planning. Phys. Med. Biol. **50**, 3491–3514 (2005)
33. Lim, G.J., Choi, J., Mohan, R.: Iterative solution methods for beam angle and fluence map optimization in intensity modulated radiation therapy planning. OR Spectrum **30**, 289–309 (2008)
34. Lu, H.M., Kooy, H.M., Leber, Z.H., Ledoux, R.J.: Optimized beam planning for linear accelerator-based stereotactic radiosurgery. Int. J. Radiat. Oncol. Biol. Phys. **39**, 1183–1189 (1997)
35. Mavroidis, P., Lind, B.K., Brahme, A.: Biologically effective uniform dose for specification, report and comparison of dose response relations and treatment plans. Phys. Med. Biol. **46**, 2607–2630 (2001)
36. Meedt, G., Alber, M., Nüsslin, F.: Non-coplanar beam direction optimization for intensity-modulated radiotherapy. Phys. Med. Biol. **48**, 2999–3019 (2003)
37. Niemierko, A.: Reporting and analyzing dose distributions: a concept of equivalent uniform dose. Med. Phys. **24**, 103–110 (1997)
38. Niemierko, A., Urie, M., Goitein, M.: Optimization of 3d radiation-therapy with both physical and biological end-points and constraints. Int. J. Radiat. Oncol. Biol. Phys. **23**, 99–108 (1992)

39. Nocedal, J., Wright, S.: Numerical Optimization. Springer-Verlag, Inc., New York (1999)
40. Pugachev, A., Xing, L.: Computer-assisted selection of coplanar beam orientations in intensity-modulated radiation therapy. Phys. Med. Biol. **46**, 2467–2476 (2001)
41. Pugachev, A., Xing, L.: Pseudo beam's-eye-view as applied to beam orientation selection in intensity-modulated radiation therapy. Int. J. Radiat. Oncol. Biol. Phys. **51**, 1361–1370 (2001)
42. Pugachev, A., Xing, L.: Incorporating prior knowledge into beam orientation optimization in IMRT. Int. J. Radiat. Oncol. Biol. Phys. **54**, 1565–1574 (2002)
43. Romeijn, H.E., Ahuja, R.K., Dempsey, J.F., Kumar, A., Li, J.G.: A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning. Phys. Med. Biol. **38**, 3521–3542 (2003)
44. Romeijn, H.E., Dempsey, J.F., Li, J.G.: A unifying framework for multi-criteria fluence map optimization models. Phys. Med. Biol. **49**, 1991–2013 (2004)
45. Romeijn, H.E., Ahuja, R.K., Dempsey, J.F., Kumar, A., Li, J.G.: A column generation approach to radiation therapy treatment planning using aperature modulation. SIAM J. Optim. **15**, 838–862 (2005)
46. Romeijn, H.E., Ahuja, R.K., Dempsey, J.F., Kumar, A.: A new linear programming approach to radiation therapy treatment planning problems. Oper. Res. **54**, 201–216 (2006)
47. Rowbottom, C.G., Oldham, M., Webb, S.: Constrained customization of non-coplanar beam orientations in radiotherapy of brain tumours. Phys. Med. Biol. **44**, 383–399 (1999)
48. Rowbottom, C.G., Webb, S., Oldham, M.: Beam-orientation customization using an artificial neural network. Phys. Med. Biol. **44**, 2251–2262 (1999)
49. Schreibmann, E., Lahanas, M., Xing, L., Baltas, D.: Multiobjective evolutionary optimization of the number of beams, their orientations and weights for intensity-modulated radiation therapy. Phys. Med. Biol. **49**, 747–770 (2004)
50. Shepard, D.M., Ferris, M.C., Olivera, G.H., Mackie, T.R.: Optimizing the delivery of radiation therapy to cancer patients. SIAM Rev. **41**, 721–744 (1999)
51. Smith, R.L.: A monte carlo procedure for the random generation of feasible solutions to mathematical programming problems. Bulletin of the TIMS/ORSA Joint National Meeting, p. 101 (1980)
52. Söderstrom, S., Brahme, A.: Selection of suitable beam orientations in radiation therapy using entropy and fourier transform measures. Phys. Med. Biol. **37**, 911–924 (1992)
53. Stein, J., Mohan, R., Wang, X.H., Bortfeld, T., Wu, Q., Preiser, K., Ling, C.C., Schlegel, W.: Number and orientations of beams in intensity-modulated radiation treatments. Med. Phys. **24**, 149–160 (1997)
54. Wu, Q.W., Mohan, R., Niemierko, A., Schmidt-Ullrich, R.: Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose. Int. J. Radiat. Oncol. Biol. Phys. **52**, 224–235 (2002)
55. Wu, Q.W., Djajaputra, D., Wu, Y., Zhou, J.N., Liu, H.H., Mohan, R.: Intensity-modulated radiotherapy optimization with geud-guided dose-volume objectives. Phys. Med. Biol. **48**, 279–291 (2003)