



Continuous Optimization

Neighborhood search approaches to non-coplanar beam orientation optimization for total marrow irradiation using IMRT

V.V. Mišić^a, D.M. Aleman^{a,*}, M.B. Sharpe^b^a Department of Mechanical and Industrial Engineering, University of Toronto, 5 King's College Road, Toronto, Canada, ON M5S 3G8^b Princess Margaret Hospital, Department of Radiation Oncology, University of Toronto, 610 University Avenue, Toronto, Canada, ON M5G 2M9

ARTICLE INFO

Article history:

Received 6 October 2009

Accepted 12 February 2010

Available online 17 February 2010

Keywords:

Total marrow irradiation

Total body irradiation

Intensity modulated radiation therapy

Radiotherapy optimization

Local search

Beam orientation optimization

ABSTRACT

We consider the beam orientation optimization (BOO) problem for total marrow irradiation (TMI) treatment planning using intensity modulated radiation therapy (IMRT). Currently, IMRT is not widely used in TMI treatment delivery; furthermore, the effect of using non-coplanar beam orientations is not known. We propose and implement several variations of a single neighborhood search algorithm that solves the BOO problem effectively when gantry angles and couch translations are considered. Our work shows that the BOO problem for TMI cases can be solved in a clinically acceptable amount of time and leads to treatment plans that are more effective than the conventional approach to TMI.

© 2010 Elsevier B.V. All rights reserved.

1. Introduction

Total marrow irradiation (TMI) is a technique that delivers radiation to the entire body of a patient. The goal of TMI is to destroy the patient's bone marrow in preparation for a bone marrow transplant (BMT). Currently, TMI is most commonly performed by first irradiating one side of the patient's body using two beams, then rotating the patient and irradiating the other side using the same beams. Although it is possible to protect some of the patient's organs using organ shields, this therapy generally does not allow for effective sparing of most of the patient's organs, as most of the patient's body receives the TMI dose intended for the bone marrow. As a result, the higher dose levels needed to ensure complete marrow elimination cause greater toxicity effects in healthy organs [1]. Additionally, the patient is placed far from isocenter in order to cover as much of the patient's body as possible with radiation: as a result, there is a high degree of uncertainty in how much dose is delivered to the patient's body. In contrast, in typical site-specific treatments, the patient is positioned at isocenter to ensure accuracy of the delivery.

In order to make TMI more accurate, we propose using IMRT to deliver the treatments. There has been little research into using IMRT with TMI. In [2,3], the authors consider the delivery of TMI using IM-TMI (intensity modulated total marrow irradiation) and

show that using standard commercial planning systems, large reductions in dose to organs such as the liver, kidneys and heart can potentially be achieved. In [4,5], the authors consider TMI using helical tomotherapy, and similarly show that the dose delivered to critical organs can be significantly reduced from conventional TMI levels. Unlike these previous studies, we consider the problem of TMI treatment planning within a mathematical framework that has been successfully applied to head-and-neck cancer cases (for example [6,7]). We also consider far more critical structures than [2]. To our knowledge, there are no other studies that approach TMI treatment planning from a mathematical perspective or that have considered as many critical structures as we do.

To bring the patient up to isocenter to allow for IMRT, we use non-coplanar beams in our beam orientation optimization (BOO) problem, thereby reducing the uncertainty in delivered dose. We integrate the BOO and fluence map optimization (FMO) problems by assigning the objective function of BOO to be the optimal solution to the FMO problem. The convexity of our FMO model facilitates this integration. While our FMO model is solved with a simple projected gradient algorithm, we modify the Add/Drop BOO algorithm presented in [6] to address the need for non-coplanar beams in BOO. In particular, we develop a new neighborhood structure specifically to address the nature of non-coplanar beam selection. This neighborhood has not been proposed in the BOO literature, non-coplanar or otherwise. Using this approach, we are able to develop non-coplanar beam solutions for TMI that deliver high quality treatments in a clinically feasible period of time.

* Corresponding author.

E-mail addresses: velibor.misic@utoronto.ca (V.V. Mišić), aleman@mie.utoronto.ca (D.M. Aleman), michael.sharpe@rmp.uhn.on.ca (M.B. Sharpe).

The paper is organized as follows. The remainder of this section introduces our BOO and FMO formulations. Section 2 discusses the Add/Drop algorithm and variations. Section 3 presents the results of applying the Add/Drop algorithm to our BOO and FMO models, and Section 4 discusses conclusions and future directions.

1.1. BOO model

The BOO problem for IMRT treatment planning has been well studied (see [8–10] for examples of BOO formulations) although little research has been done in the consideration of non-coplanar beams in BOO, especially in the context of TMI+IMRT.

Our formulation of the BOO problem is as follows. Let θ be a vector containing the orientations of k beams defined as $\theta = (\theta_1, \theta_2, \dots, \theta_k)^\top$. Each orientation θ is itself a vector of two components—the gantry angle (G) and the couch z -translation (z)—and is defined as

$$\theta = \begin{bmatrix} \theta_G \\ \theta_z \end{bmatrix}.$$

The gantry angle and the couch z -translation are restricted to their respective sets of permissible values, \mathcal{S}_G and \mathcal{S}_z , respectively:

$$\theta_G \in \mathcal{S}_G = \{L_G, L_G + \eta_G, L_G + 2\eta_G, \dots, U_G\},$$

$$\theta_z \in \mathcal{S}_z = \{L_z, L_z + \eta_z, L_z + 2\eta_z, \dots, U_z\},$$

where L and U represent the lower and upper limits of the components' values, respectively, and η is the discretization increment for each component. In our study, the gantry angle was allowed to range from 0 to 350° in 10° increments. Similarly, the couch z -translation was allowed to range from –160 cm (placing the isocenter below the pelvis) to –60 cm (placing the isocenter at the top of the head), in 10 cm increments. Although it is possible to extend this range and cover the patient's legs as well, it was deemed unnecessary as the legs do not contain any critical structures and can be treated separately from the rest of the body.

The set of all possible orientations for a single beam is represented by $\mathcal{B} = \{\theta : \theta_G \in \mathcal{S}_G, \theta_z \in \mathcal{S}_z\}$; the set of all possible sets of orientations for k beams is given by $\mathcal{B}^k = \mathcal{B} \times \dots \times \mathcal{B}$ (k times). The BOO model is then

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

where the function $\mathcal{F}(\theta)$ is the optimal FMO value that results from using the beams specified by θ in the FMO problem. The function $\mathcal{F}(\cdot)$ is formulated in such a way that lower values of F correspond to better treatments.

1.2. FMO model

The FMO problem optimizes the fluences of the beamlets of a given set of beams. In previous studies, the FMO problem has been modelled in many ways (see [9,11] for examples of FMO models). We use a convex FMO model which is amenable to convex optimization solution techniques, such as the projected gradient and interior point methods. It has been proved that formulating the problem as a convex optimization problem with voxel-based penalty functions leads to a model that is essentially the same as a multi-criteria optimization model with convex treatment plan criteria [12]. Thus, we use the model presented in [12] and in our previous work [6,7,13].

The set of critical structures is denoted by S ; the set of target structures is denoted by T . We represent the fluence of beamlet i by x_i , and we denote the set of beamlets for a particular beam orientation θ by B_θ . The dose deposited in a voxel j in structure s is denoted z_{js} and is defined in the optimization model. D_{ijs} in the model

is the dose deposition coefficient for beamlet i and voxel j in structure s ; ν_s is the number of voxels in structure s .

The penalty function F_{js} penalizes the amount of overdose or underdose to each voxel, with T_s being the ideal dose for structure s :

$$F_{js}(z_{js}) = \frac{1}{\nu_s} \left(\underline{w}_s [(T_s - z_{js})_+]^{p_s} + \bar{w}_s [(z_{js} - T_s)_+]^{\bar{p}_s} \right).$$

The function $(\cdot)_+$ represents $\max(\cdot, 0)$. The quantities \bar{w}_s and \underline{w}_s are the coefficients for overdosing and underdosing respectively, while \bar{p}_s and p_s are the powers for overdosing and underdosing respectively. To ensure convexity, we set $\bar{w}_s, \underline{w}_s \geq 0$ and $\bar{p}_s, p_s \geq 1$. The penalty function is normalized by $1/\nu_s$.

The problem for a given set of beams θ is then:

$$\begin{aligned} &\text{minimize} && \sum_{s \in S \cup T} \sum_{j=1}^{\nu_s} F_{js}(z_{js}), \\ &\text{subject to} && z_{js} = \sum_{\theta \in \Theta} \sum_{i \in B_\theta} D_{ijs} x_i \quad j = 1, \dots, \nu_s, \quad s \in S \cup T, \\ &&& x_i \geq 0 \quad i \in B_\theta, \quad \theta \in \Theta. \end{aligned}$$

To solve our FMO problem we use a standard projected gradient algorithm. Although the accuracy of the solution cannot be guaranteed, empirical testing indicates that projected gradient methods consistently return quality treatment plans.

2. Materials and methods

The Add/Drop algorithm used for BOO is a deterministic local search procedure. A single iteration of the Add/Drop algorithm is performed by enumerating all of the solutions (and their corresponding FMO values) in some neighborhood of the current solution, and moving to the most improving solution in the neighborhood if there are any improving solutions. (In the case that two or more solutions are tied for the most improving solution, the choice is made arbitrarily, although we have never observed such a situation occurring in practice.) The neighborhood that is searched changes from iteration to iteration, and the algorithm stops if all of the neighborhoods of the current solution have been searched without leading to an improvement. Once the algorithm reaches a solution, it is repeated until a specific number of executions has been performed.

Search techniques have been applied to coplanar BOO in the past (see for example [14,15]) as well as to non-coplanar BOO (see for example [16,17]). In most studies which apply search techniques to coplanar BOO, the search method that is proposed can be applied with some modification to the non-coplanar TMI case; however, adapting the method of objective function evaluation to a larger patient anatomy poses a much greater challenge. For example, in [14], a local search method is proposed where the quality of a set of beams is obtained by solving a linear program with two constraints per voxel. Applying this method to our problem (i.e., using the same number of beams and the same patient) would result in very large linear programs (approximately 1.3 million constraints and 90,000 variables) that would require much more memory to solve than the convex formulation we use. The approach in [15], which involves solving a linear programming formulation of the FMO problem iteratively with fewer and fewer beams, is affected by the same problem.

Different types of neighborhood structures have been proposed in prior BOO work. In [17], non-coplanar beams are obtained by changing both the gantry and couch angles, and the neighborhood of beams considered for addition to the solution consists of all possible combinations of gantry and couch angles. In [16], non-coplanar beams are obtained by changing both the gantry and couch angles: the neighborhood of a beam in the solution consists of all

beams within a certain angular distance of the current beam, and at least some angular distance away from the other beams in the solution. A neighborhood of a single beam can therefore contain candidate beams which differ in couch angle as well as beams which differ in gantry angle. In [14], which studies coplanar BOO, the beams are numbered from 1 to R (which is the number of beams) and the neighborhood of a beam θ_i consists of beams from θ_i to θ_{i+1} (i.e., the beam can only be rotated in one direction).

In contrast to these studies, the neighborhood structure that we propose for the non-coplanar BOO problem only allows for a single component of a single beam in the solution to be altered at a time in more than one direction. Such a neighborhood is desirable because it strikes a balance between how effectively the search space is explored and the number of FMO evaluations that are performed in each iteration which, due to the large number of voxels, are much more computationally intensive than they are for site-specific treatment planning. (In contrast, neighborhoods where more than one component of a beam and/or more than one beam can change would lead to a higher number of solutions and thus a higher number of FMO evaluations in each iteration.) Furthermore, the new beam that replaces the old beam in the solution is not restricted in how close it can be to the other beams in the solution; this allows the algorithm to access more of the solution space earlier on and increases the likelihood of finding a good solution early in the search. This type of neighborhood structure, to the best of our knowledge, has not been previously studied for non-coplanar BOO.

2.1. Neighborhood definition

The Add/Drop algorithm in each iteration searches a neighborhood of the current solution which is defined relative to a single beam and a single component/degree of freedom. An arbitrary degree of freedom is represented by d ; the set of all possible components is denoted by D . In our study, $D = \{G, z\}$, where G represents the gantry angle and z represents couch z -translation.

For clearer exposition, we define the neighborhood in three steps. We first define the neighborhood for a single component of a single beam; in the case of an angular component (i.e., the gantry rotation), we have

$$N_d(\theta_d) = \{\theta'_d \bmod 360 \in S_d : \theta_d - \delta_d \leq \theta'_d \leq \theta_d + \delta_d\},$$

and in the case of a translational component (i.e., couch translation), we have the expression:

$$N_d(\theta_d) = \{\theta'_d \in S_d : \theta_d - \delta_d \leq \theta'_d \leq \theta_d + \delta_d\}.$$

The expression for the angular component incorporates the cyclic nature of angular movements – for example, 370° is equivalent to 10° . The δ variables specify the size of the neighborhood around the current value of the component.

Next, we define the neighborhood of a single beam, relative to a single component:

$$\bar{N}_d(\theta) = \{\theta' \in \mathcal{B} : \theta'_d \in N_d(\theta_d) \wedge \theta'_a = \theta_a, \quad \forall d \in D, \bar{d} \neq d\}.$$

In other words, only one component of the beam is allowed to move. The other component is fixed.

We now provide the final definition of a neighborhood for a solution vector of beams relative to a single component d and a single beam b :

$$\mathcal{N}_{bd}(\Theta) = \{(\theta_1, \theta_2, \dots, \theta'_b, \dots, \theta_k) \in \mathcal{B}^k : \theta'_b \in \bar{N}_d(\theta_b)\}.$$

In other words, only one component of the b th beam is allowed to move at a time.

2.2. Basic Add/Drop algorithm definition

The basic Add/Drop algorithm is defined as Algorithm 1. Several implementations of the algorithm are tested for BOO in TMI. Each method is a variation of the basic version, created by modifying the selection of beam-component pairs in Step 4 or the generation of starting points in Step 1.

Algorithm 1. Basic Add/Drop

```

1: Generate initial starting point  $\Theta_0$ .
2: Set  $\Theta^* := \Theta_0$  and  $i := 0$ .
3: while Stopping criterion is not met do
4:   Select  $d \in D$  and  $b \in \{1, \dots, |\Theta|\}$ .
5:   Set  $\bar{\Theta} \in \arg \min \mathcal{F}(\mathcal{N}_{bd}(\Theta_i))$ .
6:   if  $\mathcal{F}(\bar{\Theta}) < \mathcal{F}(\Theta^*)$  then
7:     Set  $\Theta^* := \bar{\Theta}$  and  $\Theta_{i+1} := \bar{\Theta}$ .
8:     Set  $i := i + 1$ .
9:   end if
10:  if All points in  $\bigcup_{b=1}^k \bigcup_{d \in D} \mathcal{N}_{bd}(\Theta_i)$  have been sampled
    without improvement then
11:     $\Theta^*$  is a local minimum; go to Step 16.
12:  else
13:    Go to Step 4.
14:  end if
15: end while
16: return  $\Theta^*$ .

```

2.3. Variations on the Add/Drop algorithm

Along with the basic Add/Drop algorithm, our study includes several variations on the basic version. Variations include how the beam-component pair considered in a particular iteration is chosen and how starting points are generated after each execution.

2.3.1. Generation of starting points

We consider two methods of starting point generation. The first method involves randomly generating a starting point in the solution space for which the objective function \mathcal{F} has not yet been evaluated. The second method involves randomly generating a single starting point for the first execution and rotating the gantry angle of each beam by a certain amount for each subsequent execution.

2.3.2. Simple cycling

The Simple Cycling Add/Drop (SCAD) algorithm selects neighborhoods and components in a sequential fashion. It first goes through all of the components for a particular beam before selecting the next beam and going through all of the components for that beam.

2.3.3. Probabilistic selection of beam-component pairs

In this variation, the beam-component pair is selected probabilistically based on previous improvements. The algorithm keeps track of the most recent improvement in the FMO value that is realized from the neighborhood of each beam-component pair, and uses this set of improvements to construct probabilities which are then used to randomly select the next beam and component. Beam-component pairs that have recently led to higher improvements have a higher probability of being selected than pairs that led to lower improvements. Due to this difference, the algorithm has the potential to converge more quickly to a local minimum by having a tendency to explore promising neighborhoods first, rather than being constrained to go through the neighborhoods sequentially.

Selection based on historical improvements. The algorithm uses a joint probability mass function (PMF) that specifies the probability of selecting a particular beam-component pair, which is defined in terms of

$$\bar{p}(b, d) = \begin{cases} 0 & \text{if } (b, d) \in C, \\ \frac{1}{k|D|-|C|} + \frac{\alpha}{k|D|-|C|} \left(\frac{\Delta_{bdr} - \bar{\Delta}_{mC}}{\bar{\Delta}_{mC}} \right) & \text{if } (b, d) \notin C \text{ and } \bar{\Delta}_{mC} \neq 0, \\ \frac{1}{k|D|-|C|} & \text{otherwise,} \end{cases}$$

where the random variables **B** and **C** are the beam and the component, respectively; Δ_{bdr} is the average of the r most recent improvements associated with beam b and component d ; α is a weighting parameter in $[0, 1]$ that controls the emphasis of recent improvements; the set C is the set of all pairs that are not to be sampled; and $\bar{\Delta}_{mC}$ is the average of the m most recent improvements of pairs not in C . The formula assigns each (b, d) pair that can be sampled a uniform probability $(1/(k|D| - |C|))$ and then adds or subtracts an additional amount that depends on the improvement of each pair relative to the overall average improvement. The values generated in this manner are non-negative; the reason for this is that Δ_{bdr} is always greater than or equal to 0, so the quotient $(\Delta_{bdr} - \bar{\Delta}_{mC})/\bar{\Delta}_{mC}$ is at least -1 . With some algebraic manipulation it can then be shown that

$$\frac{1}{k|D| - |C|} + \frac{\alpha}{k|D| - |C|} \left(\frac{\Delta_{bdr} - \bar{\Delta}_{mC}}{\bar{\Delta}_{mC}} \right) \geq \frac{1}{k|D| - |C|} - \frac{\alpha}{k|D| - |C|}.$$

The lowest value that the right hand side can attain is 0, obtained by setting $\alpha = 1$; we therefore must have that $\bar{p}(b, d) \geq 0$.

The reason for incorporating the set C of pairs of that are not to be sampled into the above expression for $\bar{p}(b, d)$ is two-fold. If the algorithm selects a non-improving beam-component pair then intuitively, that pair should not be selected again until the iterate changes. Similarly, if the algorithm finds an improving pair, it may be desirable to exclusively explore other pairs in the subsequent iterates. This is based on the intuition that once a beam component is improved, it may be unlikely to find significant further improvement in that beam component until the other beams and components are in new positions.

The selection probabilities $p(b, d)$ are re-normalized so that the total of all probabilities sums to 1:

$$p(b, d) = \frac{\bar{p}(b, d)}{\sum_{b'=1}^k \sum_{d' \in D} \bar{p}(b', d')}.$$

3. Results

The algorithms were tested on a single patient case. Only one patient was tested due to the time involved in contouring such a large area. The Add/Drop algorithms were executed on a 64-bit, 32-node CentOS cluster, with each node having 8 GB of memory and eight AMD Opteron 2354 processors. Each algorithm was executed with 30 beams. Preliminary testing indicated that solutions with fewer beams were unable to control the degree of overdosage in the target structure due to the very large physical size of the target.

3.1. Computational results

The average time required to evaluate a single FMO value is approximately 17.1 min for a 30-beam plan. Each execution of the Add/Drop algorithm was started with a randomly generated starting set of beams. Each execution was also allowed to run for a maximum of 12 h, as this amount of time is the most an optimization algorithm would be allowed to run in a clinical setting at the

Princess Margaret Hospital where this research is performed. For both the SCAD and the probabilistic Add/Drop, δ_G and δ_z were set to 20. The probabilistic Add/Drop variant was tested with r and m both set to 5 and α values of 0, 0.25, 0.5, 0.75 and 1.

The plots of FMO value versus time and iteration shown in Figs. 1 and 2 respectively suggest that the SCAD is slower to converge to its final FMO value than the probabilistic Add/Drop. This may be because the SCAD in each execution samples approximately the same number of beams, while the probabilistic Add/Drop is not constrained in this way and has the potential to change more beams in the solution over the course of one execution. Also, SCAD checks both the gantry angle and the couch-z neighborhood for the same beam, which may be unnecessary.

3.2. Treatment plan quality

Due to the relative absence of treatment planning in TMI, there is a lack of clear criteria on what level of underdosage and overdosage is acceptable in the target and in the critical structures. To ad-

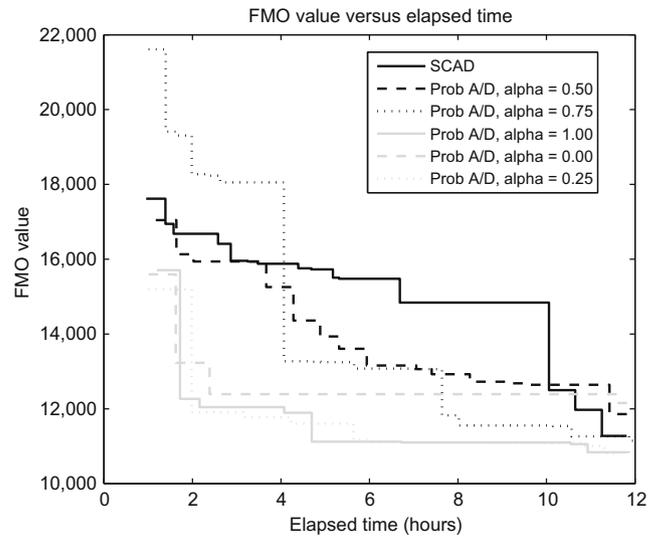


Fig. 1. Improvement in FMO value by time for each Add/Drop implementation.

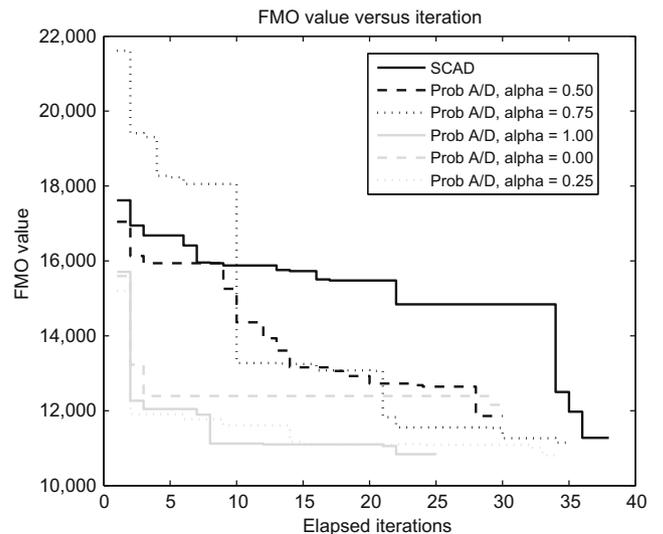


Fig. 2. Improvement in FMO value by iteration number for each Add/Drop implementation.

dress the quality of the treatment, we developed treatment plan criteria with expert medical physicists and physicians from Princess Margaret Hospital. For the target structure, at least 95% of the structure must receive at least 12 Gy, while a strict 25 Gy maximum must be observed. Any cells receiving more than 25 Gy undergo fibrosis, making it impossible for the new bone marrow to take hold. At most 20% of the bone marrow can receive more than 20 Gy.

Because TMI has yet to be performed with the accuracy of IMRT, studies have not yet been done to indicate the amount of radiation that each critical structure can tolerate under TMI conditions. Based on consultations with our collaborators, we classify an organ as spared if the majority of each organ should receive less than 8 Gy.

Dose-volume histograms (DVHs) are used to assess treatment quality. The DVHs obtained from the SCAD and the probabilistic Add/Drop variants in general met all of the goals, and do not appear to exhibit significant differences in organ sparing and marrow elimination. One representative set of DVHs, obtained from an execution of the probabilistic Add/Drop with $\alpha = 0.75$, is shown in Fig. 3 (the DVH curves for the organs are displayed on two graphs for clarity of presentation). From this DVH we can see that our methods are able to attain acceptable dose levels within the bone marrow. We can also see that for most organs, the majority of the

volume receives less than 8 Gy, suggesting that our methods lead to an improvement over conventional TMI (where most of the organs would receive around 12 Gy).

With regard to prior IM-TMI work, the solution in Fig. 3 compares favorably to the solution given in [2]: in particular, the median doses for the lungs, liver, kidneys and heart are all lower than 5 Gy and are lower than the corresponding reported values in [2]. However, the solution given in [2] is able to control target and organ overdose to a greater extent than the solution we provide in Fig. 3. A reason for this is the larger set of organs that is considered in this study; the more organs that are accounted for, the harder it is to simultaneously spare all of the organs while achieving a desired level of dose in the target.

One specific organ where the solution in [2] and our solution in Fig. 3 greatly differ in sparing is the spinal cord. The spinal cord is problematic because it is encased in a structure which contains bone marrow (the spinal column) and is thus very difficult to spare. The spinal cord is not considered explicitly in [2], and no numerical results are provided about the amount of dose delivered to the spinal cord. The only indication of the dose in the spinal cord comes from the colour washes, which indicate that the spine is irradiated to a dose between 12 Gy and 13 Gy—essentially the entire target prescription dose. This amount is significantly greater

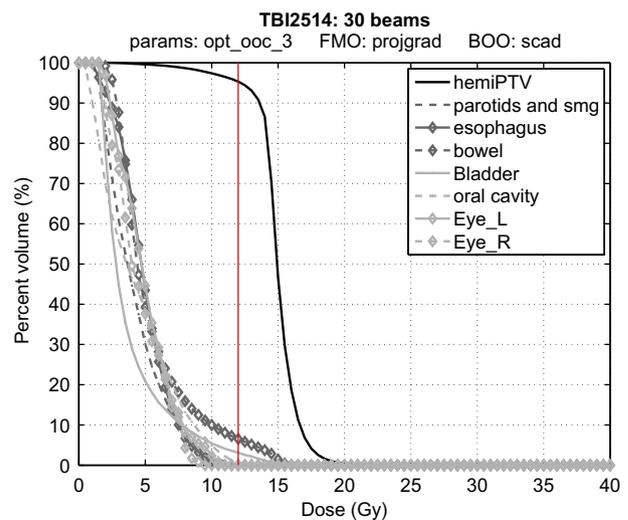
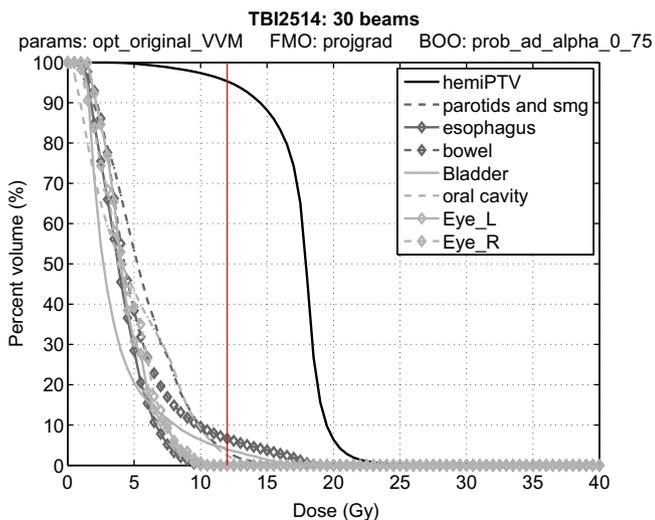
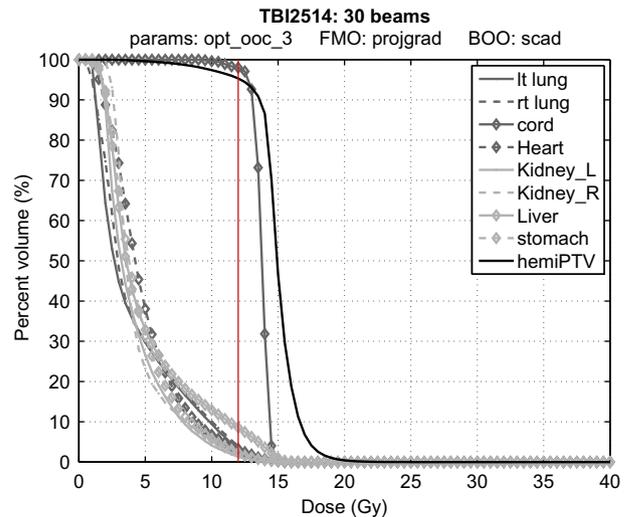
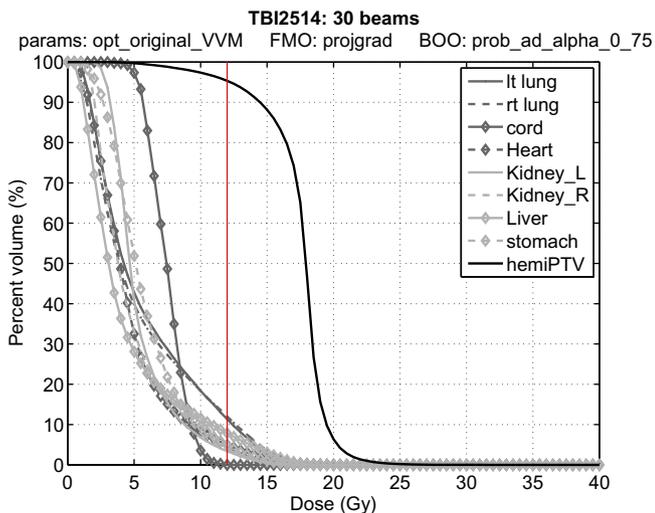


Fig. 3. DVHs for a representative 30-beam treatment from probabilistic Add/Drop. The two graphs correspond to the same solution, but the critical structures are split between the two graphs for clarity.

Fig. 4. DVHs for a representative 30-beam treatment from SCAD, with reduced spinal cord importance to provide a more direct comparison to [2]. Again, the two graphs correspond to the same solution.

than the amount delivered to the spinal cord by our solution in Fig. 3.

With this in mind, we tested our algorithms with a different set of FMO parameters which reduced the importance of the spinal cord. A representative set of DVHs, obtained from an execution of SCAD with these modified parameters, is shown in Fig. 4. (As before, the curves are shown on two graphs for clarity of presentation.) This solution shows reduced overdosage in the bone marrow and most other organs, including the lungs, and with the exception of the spinal cord, meets the treatment plan criteria we have set. More significantly, this solution demonstrates the versatility of our approach. Our approach can exploit the reduced importance of the spinal cord to deliver a solution that exhibits better overall organ sparing than the solution in [2], while still taking into account many more structures than [2].

4. Discussion

From our tests the SCAD and the probabilistic Add/Drop algorithms are both capable of obtaining solutions that meet the recommended treatment plan criteria in 12 h from a single execution and do not differ significantly in solution quality. More testing on the current patient data and other patient data is required to reduce overdosage to the bone marrow and critical structures.

In addition to the probabilistic Add/Drop, there are other variants of the Add/Drop algorithm that will be implemented. One variant we are studying is a “dynamic δ ” Add/Drop, where the neighborhood size is modified after each iteration in response to how much the iterate shifted in that iteration. We are also studying the Add/Drop algorithm in a multiple execution context where points sampled in one execution of the Add/Drop algorithm are used to generate the starting point for the next execution of the Add/Drop. In particular, the concept of expected improvement from response surface methods will be examined as a means to generate new starting points.

To bring TMI+IMRT closer to a clinical viability, future research will use this work as a starting point to explore optimal intensity modulated arc therapy (IMAT) treatment plans. This will involve designing an algorithm to “connect” the beams identified by our algorithms, and to fill in the individual bixel time profiles along the connecting arcs. Other future work will factor patient motion into the FMO formulation and to use a robust optimization approach in order to make the treatment plans more suitable for implementation with conventional IMRT (for examples of robust optimization applied to IMRT treatment planning see [18–20]).

Acknowledgements

This work was supported by an Undergraduate Research Award of the Natural Sciences and Engineering Research Council of Canada (NSERC).

References

- [1] R.A. Clift, C.D. Buckner, F.R. Appelbaum, E. Bryant, S.I. Bearman, F.B. Petersen, L.D. Fisher, C. Anasetti, P. Beatty, W.I. Bensinger, K. Doney, R.S. Hill, G.B. McDonald, P. Martin, J. Meyers, J. Sanders, J. Singer, P. Stewart, K.M. Sullivan, R. Witherspoon, R. Storb, J.A. Hansen, E.D. Thomas, Allogeneic marrow transplantation in patients with chronic myeloid leukemia in the chronic phase: A randomized trial of two irradiation regimens, *Blood* 77 (1991) 1660–1665.
- [2] B. Aydogan, A.J. Mundt, J.C. Roeske, Linac-based total marrow irradiation (IM-TMI), *Technology in Cancer Research and Treatment* 5 (5) (2007) 513–519.
- [3] B. Aydogan, J.C. Roeske, Feasibility study of linac-modulated total marrow irradiation (IM-TMI), *Medical Physics* 34 (6) (2007) 2455.
- [4] T.E. Schultheiss, J. Wong, A. Liu, G. Olivera, G. Somlo, Image-guided total marrow and total lymphatic irradiation using helical tomotherapy, *International Journal of Radiation Oncology, Biology, Physics* 67 (2007) 1259–1267.
- [5] J.Y.C. Wong, A. Liu, T. Schultheiss, L. Popplewell, A. Stein, J. Rosenthal, M. Essensten, S. Forman, G. Somlo, Targeted total marrow irradiation using three-dimensional image-guided tomographic intensity-modulated radiation therapy: An alternative to standard total body irradiation, *Biology of Blood and Marrow Transplantation* 12 (3) (2006) 306–315.
- [6] 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.
- [7] 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.
- [8] D.M. Aleman, H.E. Romeijn, J.F. Dempsey, Beam orientation optimization in intensity modulated radiation therapy treatment planning, in: G.J. Lim, E.K. Lee (Eds.), *Optimization in Medicine and Biology*, Auerbach Publications, 2008, pp. 223–251.
- [9] M. Ehr Gott, Ç. Güler, H.W. Hamacher, L. Shao, Mathematical optimization in intensity modulated radiation therapy, *4OR* 6 (3) (2008) 199–262.
- [10] M. Ehr Gott, A. Holder, J. Reese, Beam selection in radiotherapy design, *Linear Algebra and its Applications* 428 (2007) 1272–1312.
- [11] H.E. Romeijn, J.F. Dempsey, Intensity modulated radiation therapy treatment plan optimization, *TOP* 16 (2008) 215–243.
- [12] H.E. Romeijn, J.F. Dempsey, J.G. Li, A unifying framework for multi-criteria fluence map optimization models, *Physics in Medicine and Biology* 49 (2004) 1991–2003.
- [13] H.E. Romeijn, R.K. Ahuja, J.F. Dempsey, A. Kumar, A new linear programming approach to radiation therapy treatment planning problems, *Operations Research* 54 (2) (2006) 201–216.
- [14] M. Ehr Gott, R. Johnston, Optimization of beam directions in intensity modulated radiation therapy, *OR Spectrum* 25 (2003) 251–264.
- [15] 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 (2008) 289–309.
- [16] S. Das, T. Cullip, G. Tracton, S. Chang, Lawrence Marks, M. Anscher, J. Rosenman, Beam orientation selection for intensity-modulated radiation therapy based on target equivalent uniform dose maximization, *International Journal of Radiation Oncology, Biology, Physics* 55 (1) (2003) 215–224.
- [17] G. Meedt, M. Alber, F. Nüsslin, Non-coplanar beam direction optimization for intensity modulated radiotherapy, *Physics in Medicine and Biology* 48 (2003) 2999–3019.
- [18] A. Ólafsson, S.J. Wright, Efficient schemes for robust IMRT treatment planning design, *Physics in Medicine and Biology* 51 (2006) 5621–5642.
- [19] M. Chu, Y. Zinchenko, S.G. Henderson, M.B. Sharpe, Robust optimization for intensity modulated radiation therapy treatment planning under uncertainty, *Physics in Medicine and Biology* 50 (2005) 5463–5477.
- [20] T.C.Y. Chan, T. Bortfeld, J.N. Tsitsiklis, A robust approach to IMRT optimization, *Physics in Medicine and Biology* 51 (2006) 2567–2583.