

Comparing Beam Selection Strategies in Radiotherapy Treatment Design: The Influence of Dose Point Resolution

Ryan Acosta^a Matthias Ehrgott^{b,*} Allen Holder^c Daniel Nevin^d
Josh Reese^e Bill Salter^f

November 3, 2005

Abstract

The design of a radiotherapy treatment includes the selection of beam angles (geometry problem), the computation of a fluence pattern for each selected beam angle (intensity problem), and finding a sequence of configurations of a multileaf collimator to deliver the treatment (realization problem). While many mathematical optimization models and algorithms have been proposed for the intensity problem and (to a lesser extent) the realization problem, this is not the case for the geometry problem. In clinical practice, beam directions are manually selected by a clinician and are typically based on the clinician's experience. Solving the beam selection problem optimally is beyond capability of current optimization algorithms and software. However, heuristic methods have been proposed. In this paper we compare various heuristic approaches on a clinical case. In particular, we study the influence of dose point resolution on the performance of these heuristics. We also compare the solutions obtained by the heuristics with those achieved by a clinician using a commercial planning system.

Keywords: Optimization, Radiotherapy, Heuristics, Set Covering, Vector Quantization, Medical Physics.

^a Institute for Computational and Mathematical Engineering, Stanford University, Stanford, California, USA. raacosta@stanford.edu

^b Department of Engineering Science, The University of Auckland, Auckland, New Zealand. m.ehrgott@auckland.ac.nz.

^c Department of Mathematics, Trinity University, and Department of Radiation Oncology, University of Texas Health Science Center, San Antonio, Texas, USA. aholder@trinity.edu.

^d Department of Computer Science, Texas A&M University, College Station, Texas, USA. dnevin@tamu.edu.

^e Department of Mathematics, Trinity University, San Antonio, Texas, USA. jreese@trinity.edu.

^f Department of Radiology, University of Texas Health Science Center, Department of Radiation Oncology, University of Utah Huntsman Cancer Institute. bsalter@ctrc.net.

* Corresponding author.

1 Introduction

Radiotherapy is the treatment of cancerous and displasiac tissues with ionizing radiation that can damage the DNA of cells. While non-cancerous cells are able to repair slightly damaged DNA, the heightened state of reproduction that non-cancerous cells are in means that small amounts of DNA damage render them incapable of reproducing. The goal of radiotherapy is to exploit this therapeutic advantage by focussing the radiation so that enough dose is delivered to the targeted region to kill the cancerous cells while surrounding anatomical structures are maintained at minimal damage levels so that they are spared.

In the past, it was reasonable for a clinician to design radiotherapy treatments manually due to the limited capabilities of radiotherapy equipment. However, with the advent of intensity modulated radiotherapy (IMRT), the number of possible treatment options and the number of parameters have become so immense that they exceed the capabilities of even the most experienced treatment planner. Therefore, optimization methods and computer assisted planning tools have become a necessity. IMRT treatments use multileaf collimators to shape the beam and control, or modulate, the dose that is delivered along a fixed direction of focus. IMRT allows beams to be decomposed into a (large) number of sub-beams, for which the intensity can be chosen individually. In addition, movement of the treatment couch and gantry allows radiation to be focused from almost any point on a (virtual) sphere around the target volume. For background on radiotherapy and IMRT we refer to Schlegel and Mahr (2002) and Webb (2001).

Designing an optimal treatment means deciding on a huge number of parameters. The design process is therefore usually divided into three phases, namely 1) the selection of directions from which to focus radiation on the patient, 2) the selection of fluence patterns (amount of radiation delivered) for the directions selected in phase one, and 3) the selection of a mechanical delivery sequence that efficiently administers the treatment. Today there are many optimization methods for the intensity problem, suggested models include linear (e.g. Romeijn *et al.* (2003); Rosen *et al.* (1991)), integer (e.g. Lee *et al.* (2003); Preciado-Walters *et al.* (2004)), and nonlinear (e.g. Löf (2000); Spirou and Chui (1998)) models as well as models of multiobjective optimization (e.g. Hamacher and Küfer (2002); Holder (2001); Romeijn *et al.* (2004)).

Similarly, algorithms have been proposed to find good multileaf collimator sequences to reduce treatment times and minimize between-leaf leakage and background dose (Bortfeld *et al.*, 1994; Siochi, 1999; Xia and Verhey, 1999). Such algorithms are in use in existing radiotherapy equipment. Moreover, researchers have studied the mathematical structure of these problems to improve algorithm design or to establish the optimality of an algorithm (Ahuja and Hamacher, 2004; Baatar *et al.*, 2004; Kamath *et al.*, 2003).

In this paper we consider the geometry problem. The literature on this topic reveals a different picture than that of the intensity and realization problems. While a number of methods were proposed, there was a lack of understanding of the underlying mathematics. Ehrgott *et al.* (2005) propose a mathematical framework that unifies the approaches found in the literature. The focus of this paper is on how different approximations of the anatomical dose affect beam selection.

The beam selection problem is important for several reasons. First, changing beam directions during treatment is time consuming, and the number of directions is limited to reduce the overall treatment time. Moreover, short treatments are desirable because lengthy procedures increase the likelihood of a patient altering his or her position on the couch, which can lead to inaccurate and potentially dangerous treatments. Additionally, since most clinics treat patients steadily throughout the day patients are usually treated in daily sessions of 15 – 30 minutes to make sure that demand is satisfied. Lastly, and perhaps most importantly, beam directions must be judiciously selected so as to minimize the radiation exposure to life-critical tissues and organs, while maximizing the dose to the targeted tumor.

Selecting the beam directions is currently done manually, and it typically requires several trial-and-error iterations between selecting beam directions and calculating fluence patterns until a satisfactory treatment is designed. Hence, the process is time intensive and subject to the experience of the clinician. Finding a suitable collection of directions often takes several hours. The goal of using an optimization method to identify quality directions is to remove the dependency

on a clinician’s experience and to alleviate the tedious repetitive process of selecting angles.

To evaluate the dose distribution in the patient, it is necessary to calculate how radiation is deposited into the patient. There are numerous dose models in the literature, with the gold standard being a Monte Carlo technique that simulates each particle’s path through the anatomy. We use an accurate 3D dose model developed in Nizin (1998) and Nizin and Mooij (1998). Positions within the anatomy where dose is calculated are called *dose-points*. Each patient image represents a slice of the anatomy of varying thickness, and hence, each dose-point represents a 3D hyper-rectangle whose dimensions are decided by the spacing of the dose-points. The goal of this paper is to evaluate the influence of dose point spacing on automated beam selection.

2 The Beam Selection Problem

First we note that throughout this paper the terms beam, direction, and angle are used interchangeably. The beam selection problem is to find N positions for the patient and gantry from which the treatment will be delivered. The gantry of a linear accelerator can rotate around the patient in a great circle and the couch can rotate in the plane that keeps it flat. There are physical restrictions on the directions that can be used because some couch and gantry positions result in collisions.

In this paper we consider co-planar treatments. That is, beam angles are chosen on a great circle around the CT-slice of the body that contains the center of the tumour. We let $\mathcal{A} = \{a_j : j \in J\}$ be a candidate collection of angles from which we will select N to treat the patient, where we typically consider $\mathcal{A} = \{i\pi/36 : i = 0, 1, 2, \dots, 71\}$. To evaluate a collection of angles, a judgment function is needed that describes how well a patient can be treated with N angles (Ehrgott *et al.*, 2005).

We denote the power set of \mathcal{A} by $\mathcal{P}(\mathcal{A})$ and the nonnegative extended reals by \mathbb{R}_+^* . A *judgment function* is a function $f : \mathcal{P}(\mathcal{A}) \rightarrow \mathbb{R}_+^*$ with the property that $\mathcal{A}' \supseteq \mathcal{A}''$ implies $f(\mathcal{A}') \leq f(\mathcal{A}'')$. The value of $f(\mathcal{A}')$ is the optimal value of an optimization problem that decides a fluence pattern for angles \mathcal{A}' , i.e. for any $\mathcal{A}' \in \mathcal{P}(\mathcal{A})$,

$$f(\mathcal{A}') = \min\{z(x) : x \in X(\mathcal{A}')\}, \tag{1}$$

where z maps a fluence pattern $x \in X(\mathcal{A}')$ into \mathbb{R}_+^* .

We assume that if a feasible treatment cannot be achieved with a given set of angles \mathcal{A}' ($X(\mathcal{A}') = \emptyset$) then $f(\mathcal{A}') = \infty$. We further assume that x is a vector in $\mathbb{R}^{|\mathcal{A}'| \times I}$, where I is the number of sub-beams of a beam, and make the tacit assumptions that $x_{(a,i)} = 0$ for all sub-beams i of any angle $a \in \mathcal{A} \setminus \mathcal{A}'$. The non-decreasing behaviour of f with respect to set inclusion is then modeled via the set of feasible fluence patterns $X(\mathcal{A})$ by assuming that $X(\mathcal{A}'') \subseteq X(\mathcal{A}')$ whenever $\mathcal{A}'' \subseteq \mathcal{A}'$. We say that the fluence pattern x is optimal for \mathcal{A}' if $f(\mathcal{A}') = z(x)$ and $x \in X(\mathcal{A}')$.

A judgment function is defined by the data that forms the optimization problem in (1). This data includes a dose operator D , a prescription P , and an objective function z . We let $d_{(k,a,i)}$ be the rate at which radiation along sub-beam i in angle a is deposited into dose-point k , and we assume that $d_{(k,a,i)}$ is nonnegative for each (k, a, i) . These rates are patient-specific constants and the operator that maps a fluence pattern into anatomical dose (measured in Grays, Gy) is linear. We let D be the matrix whose elements are $d_{(k,a,i)}$, where the rows are indexed by k and the columns by (a, i) . The linear operator $x \mapsto Dx$ maps the fluence pattern x to the dose that is deposited into the patient. To avoid unnecessary notation we use \sum_i to indicate that we are summing over the sub-beams in an angle. So, $\sum_i x_{(a,i)}$ is the total exposure (or fluence) for angle a , and $\sum_i d_{(k,a,i)}$ is the aggregated rate at which dose is deposited into dose-point k from angle a .

As pointed out above, there is a large amount of literature on modeling and calculating f , i.e. solving the intensity problem. In fact, all commercial planning systems use an optimization routine to decide a fluence pattern, but the model and calculation method differ from system to system (Winz, 2004). All these methods share the property that the quality of a treatment cannot deteriorate if more angles are used. The result that a judgment function is non-decreasing with

respect to the number of angles follows from the definition of a judgment function and the above assumptions, see Ehrgott *et al.* (2005).

There are a variety of forms that a prescription can have, each dependent on what the optimization problem is attempting to accomplish. Since the purpose of this paper is to compare the effect of dose point resolution on various approaches to the beam selection problem, we focus on one particular judgment function. Let us partition the set of dose-points into those that are targeted, those that are within a critical structure, and those that represent normal tissue. We denote the set of targeted dose points by T , the collection of dose-points in the critical regions by C , and the remaining dose-points by N . We further let D_T , D_C , and D_N be the submatrices of D such that $D_T x$, $D_C x$, and $D_N x$ map the fluence pattern x into the targeted region, the critical structures, and the normal tissue, respectively. The prescription consists of TLB and TUB , which are vectors of lower and upper bounds on the targeted dose points, CUB , which is a vector of upper bounds on the critical structures, and NUB , which is a vector of upper bounds on the normal tissue. The judgment function is defined by the following linear program (Holder, 2003).

$$\left. \begin{aligned} f(\mathcal{A}') &= \min \omega \alpha + \beta + \gamma \\ TLB - e\alpha &\leq D_T x \\ D_T x &\leq TUB \\ D_C x &\leq CUB + e\beta \\ D_N x &\leq NUB + e\gamma \\ TLB &\geq e\alpha \\ -CUB &\leq e\beta \\ x, \gamma &\geq 0 \\ \sum_i x_{(a,i)} &= 0 \text{ for all } a \in \mathcal{A} \setminus \mathcal{A}'. \end{aligned} \right\} \quad (2)$$

Here e is the vector of ones of appropriate dimension. The scalars α, β , and γ measure the worst deviation from TLB, CUB , and NUB for any single dose point in the target, the critical structures, and the normal tissue, respectively.

For a fixed judgment function such as (2), the N -beam selection problem is

$$\begin{aligned} &\min \{f(\mathcal{A}') - f(\mathcal{A}) : \mathcal{A}' \in \mathcal{P}(\mathcal{A}), |\mathcal{A}'| = N\} \\ &= \min \{f(\mathcal{A}') : \mathcal{A}' \in \mathcal{P}(\mathcal{A}), |\mathcal{A}'| = N\} - f(\mathcal{A}). \end{aligned} \quad (3)$$

This minimization problem can be stated as an extension of the optimization problem that defines f using binary variables. Let

$$y_a = \begin{cases} 1 & \text{angle } a \text{ is selected} \\ 0 & \text{otherwise.} \end{cases}$$

Then the beam selection problem becomes

$$\left. \begin{aligned} \min z(x) \\ \sum_{a \in \mathcal{A}} y_a &= N \\ \sum_i x_{(i,a)} &\leq M y_a \text{ for all } a \in \mathcal{A} \\ x &\in X(\mathcal{A}), \end{aligned} \right\} \quad (4)$$

where M is a sufficiently large constant.

While (4) is a general model that combines the optimal selection of beams with the optimization of their fluence patterns, such problems are currently intractable because they are beyond modern solution capabilities. Note that there are between 1.4×10^7 and 5.4×10^{11} subsets of $\{i\pi/36 : i = 0, 1, 2, \dots, 71\}$ for clinically relevant values of N ranging from 5 to 10. In any study where the solution of these MIPs is attempted (Ehrgott and Johnston, 2003; Lee *et al.*, 2003; Lim *et al.*, 2002; Preciado-Walters *et al.*, 2004; Wang *et al.*, 2003) the set $|\mathcal{A}|$ is severely restricted so that the number of binary variables is manageable. This fact has led researchers to investigate heuristics.

In the following section we present the heuristics that we include in our computational results in the framework of beam selectors introduced in Ehrgott *et al.* (2005). The function $g : \mathcal{W} \rightarrow \mathcal{V}$

is a *beam selector* if \mathcal{W} and \mathcal{V} are subsets of $\mathcal{P}(\mathcal{A})$ and $g(W) \subseteq W$ for all $W \in \mathcal{W}$. A beam selector $g : \mathcal{W} \rightarrow \mathcal{V}$ maps every collection of angles in \mathcal{W} to a subcollection of *selected angles*. An *N-beam selector* is a beam selector with $|\cup_{W \in \mathcal{W}} g(W)| = N$. A beam selector is *informed* if it is defined in terms of the value of a judgment function and it is *weakly informed* if it is defined in terms of the data (D, P, z) . A beam selector is otherwise *uninformed*. If g is defined in terms of a random variable, then g is *stochastic*.

An important observation is that for any collection of angles $\mathcal{A}' \subset \mathcal{A}$ there is not necessarily a unique optimal fluence pattern, which means that informed beam selectors are solver dependent. An example in Section 5 of Ehr Gott *et al.* (2005) shows how radically different optimal fluence patterns obtained by different solvers for the same judgment function can be.

There are several heuristic beam selection techniques in the literature. Each heuristic approach to the problem can be interpreted as choosing a best beam selector of a specified type as described in Ehr Gott *et al.* (2005). Additional references on methods not used in this study and methods for which the original papers do not provide sufficient detail to reproduce their results can be found in Ehr Gott *et al.* (2005).

3 The Beam Selection Methods

We first present the set covering approach found in Ehr Gott and Johnston (2003). An angle a *covers* the dose-point k if $\sum_i d_{(k,a,i)} \geq \varepsilon$, and for each $k \in T$, let $\mathcal{A}_k^\varepsilon = \{a \in \mathcal{A} : a \text{ covers dose-point } k\}$. A (set-covering) SC- N -beam selector is an N -beam selector having the form

$$g_{sc} : \{\mathcal{A}_k^\varepsilon : k \in T\} \rightarrow \bigcup_{k \in T} (\mathcal{P}(\mathcal{A}_k^\varepsilon) \setminus \{\emptyset\}).$$

Two observations are important:

1. We have $\mathcal{A}_k^\varepsilon = \mathcal{A}$ for all $k \in T$ if and only if $0 \leq \varepsilon \leq \varepsilon^* := \min\{\sum_i d_{(k,a,i)} : k \in T, a \in \mathcal{A}\}$. The most common scenario is that each targeted dose-point is covered by every angle.
2. Since g_{sc} cannot map to \emptyset , the mapping has to select at least one angle to cover each targeted dose-point.

It was shown in Ehr Gott *et al.* (2005) that for $0 \leq \varepsilon \leq \varepsilon^*$, the set covering approach to beam selection is equivalent to the beam selection problem (3). This equivalence means that we cannot solve the set-covering beam selection problem efficiently. However, heuristically it is possible to restrict the optimisation to subsets of SC- N -beam selectors. This was done in Ehr Gott and Johnston (2003). The second observation allows the formulation of a traditional set covering problem to identify a single g_{sc} . For each targeted dose-point k , let $q_{(k,a,i)}$ be 1 if sub-beam i in angle a covers dose-point k and 0 otherwise. For each angle a , define

$$c_a = \begin{cases} \sum_{k \in C} \sum_i \frac{q_{(k,a,i)}}{CUB_k} & \text{if } C \neq \emptyset \\ 0 & \text{if } C = \emptyset \end{cases} \quad (5)$$

and

$$\hat{c}_a = \begin{cases} \sum_{k \in C} \sum_i \frac{q_{(k,a,i)} \cdot d_{(k,a,i)}}{CUB_k} & \text{if } C \neq \emptyset \\ 0 & \text{if } C = \emptyset \end{cases}, \quad (6)$$

where CUB is part of the prescription in (2). The costs c_a and \hat{c}_a are large if sub-beams of a intersect a critical structure that has a small upper bound. Cost coefficients \hat{c}_a are additionally scaled by the rate at which dose is deposited into dose-point k from sub-beam (a, i) .

The associated set covering problems are

$$\min \left\{ \sum_a c_a y_a : \sum_a q_{(k,a)} y_a \geq 1, k \in T, \sum_a y_a = N, y_a \in \{0, 1\} \right\} \quad (7)$$

and

$$\min \left\{ \sum_a \hat{c}_a y_a : \sum_a q_{(k,a)} y_a \geq 1, k \in T, \sum_a y_a = N, y_a \in \{0, 1\} \right\}. \quad (8)$$

The angles for which $y_a^* = 1$ in an optimal solution y^* of (7) or (8) are selected and define a particular SC- N -beam selector. Note that such N -beam selectors are weakly informed, but not informed, as they use the data but do not evaluate f .

These particular set covering problems are generally easy to solve. In fact, in the common situation of $\mathcal{A}_k^\varepsilon = \mathcal{A}$ for $k \in T$, (7) and (8) reduce to selecting N angles in order of increasing c_a or \hat{c}_a , respectively. This leads us to scoring techniques for the beam selection problem.

We can interpret c_a or \hat{c}_a as a score of angle a . A (scoring) S- N -beam selector is an N -beam selector $g_s : \{\mathcal{A}\} \rightarrow \mathcal{P}(\mathcal{A})$. It is not surprising that the scoring approach is equivalent to the beam selection problem. The difficulty here lies in defining scores that accurately predict angles that are used in an optimal treatment.

The first scoring approach we consider is found in Pugachev and Xing (2001), where each angle is assigned the score

$$c_a = \frac{1}{|T|} \sum_{k \in T} \sum_i \left(\frac{d_{(k,a,i)} \cdot \hat{x}_{(a,i)}}{TG} \right)^2, \quad (9)$$

where

$$\hat{x}_{(a,i)} = \min\{\min\{CUB_k/d_{(k,a,i)} : k \in C\}, \min\{NUB_k/d_{(k,a,i)} : k \in N\}\}$$

and TG is a goal dose to the target and $TLB \leq TG \leq TUB$. An angle's score increases as the sub-beams that comprise the angle are capable of delivering more radiation to the target without violating the restrictions placed on the non-targeted region(s). Here, high scores are desirable. The scoring technique uses the bounds on the non-targeted tissues to form constraints, and the score represents how well the target can be treated while satisfying these constraints. This is the reverse of the perspective in (7) and (8). Nevertheless, mathematically, every scoring technique is a set covering problem (Ehrgott *et al.*, 2005).

Another scoring method is found in Söderström and Brahme (1992). Letting x^* be an optimal fluence pattern for \mathcal{A} , the authors of Söderström and Brahme (1992) define the entropy of an angle by $e_a := -\sum_i x_{(a,i)}^* \ln x_{(a,i)}^*$ and the score of a is

$$c_a = 1 - \frac{e_a - \min\{e_a : a \in \mathcal{A}\}}{\max\{e_a : a \in \mathcal{A}\}}. \quad (10)$$

In this approach, an angle's score is high if the optimal fluence pattern of an angle's sub-beams is uniformly high. So, an angle with a single high-fluence sub-beam would likely have a lower score than an angle with a more uniform fluence pattern. Unlike the scoring procedure in Pugachev and Xing (2001), this technique is informed since it requires an evaluation of f .

The last of the techniques we consider is based on the image compression technique called vector quantization (Holder and Salter, 2004) (see Gersho and Gray (1991) for further information on vector quantization). \mathcal{A}' is a contiguous subset of \mathcal{A} if \mathcal{A}' is an ordered subset of the form $\{a_j, a_{j+1}, \dots, a_{j+r}\}$. A contiguous partition of \mathcal{A} is a collection of contiguous subsets of \mathcal{A} that partition \mathcal{A} , and we let $\mathcal{W}_{vq}(N)$ be the collection of N element contiguous partitions of \mathcal{A} . A VQ- N -beam selector is a function of the form

$$g_{vq} : \{W_j : j = 1, 2, \dots, N\} \rightarrow \{\{a_j\} : a_j \in W_j\},$$

where $\{W_j : j = 1, 2, \dots, N\} \in \mathcal{W}_{vq}(N)$.

The image of W_j is a singleton $\{a_j\}$, and we usually write a_j instead of $\{a_j\}$. The VQ- N -beam selector relies on the probability that an angle is used in an optimal treatment. Letting $\alpha(a)$ be this probability, we have that the *distortion* of a quantizer is

$$\sum_{j=1}^N \sum_{a \in W_j} \alpha(a) \cdot \|a - g_{vq}(W_j)\|_2.$$

Once the probability distribution α is known, a VQ- N beam-selector is calculated to minimize distortion. In the special case of a continuous \mathcal{A} , Gersho and Gray (1991) show that the selected angles are the centers-of-mass of the contiguous sets. We mimic this behavior in the discrete setting by defining

$$g_{vq}(W_j) = \frac{\sum_{a \in W_j} a \cdot \alpha(a)}{\sum_{a \in W_j} \alpha(a)}. \quad (11)$$

This center-of-mass calculation is not exact for discrete sets since the center-of-mass may not be an element of the contiguous set. Therefore angles not in \mathcal{A} are mapped to their nearest neighbor, with ties being mapped to the larger element of \mathcal{A} .

Vector quantization heuristics select a contiguous partition from which a single VQ- N -beam selector is created according to condition (11). The process in Holder and Salter (2004) starts by selecting the zero angle as the beginning of the first contiguous set. The endpoints of the contiguous sets are found by forming the cumulative density and evenly dividing its range into N intervals. To improve this, we could use the same rule and rotate the starting angle through the 72 candidates. We could then evaluate f over these sets of beams and take the smallest value.

The success of the vector quantization approach directly relies on the ability of the probability distribution to accurately gauge the likelihood of an angle being used in an optimal N -beam treatment. An immediate idea is to make a weakly informed probability distribution by normalizing the scoring techniques in (5), (6) and (9). Additionally, the scores in (10) are normalized to create an informed model of α . We test these methods in Section 4. An alternative informed probability density is suggested in Holder and Salter (2004), where the authors assume that an optimal fluence pattern x^* for $f(\mathcal{A})$ contains information about which angles should and should not be used. Let

$$\alpha(a) = \frac{\sum_i x_{(a,i)}^*}{\sum_{a \in \mathcal{A}} \sum_i x_{(a,i)}^*}.$$

Since optimal fluence patterns are not unique, these probabilities are solver-dependent. In Ehr Gott *et al.* (2005) an algorithm is given to remove this solver dependency. The algorithm transforms an optimal fluence x^* into a balanced probability density α , i.e. one that is as uniform as possible, by solving the problem

$$\text{lexmin} \left(z(x), \text{sort} \left(\sum_i x_{(a,i)} : a \in \mathcal{A} \right) \right), \quad (12)$$

where sort is a function $\mathbb{R}^{\mathcal{A}} \mapsto \mathbb{R}^{\mathcal{A}}$ that reorders the components of the vector $(\sum_i x_{(a,i)})_{a \in \mathcal{A}}$ in a nonincreasing order. The algorithm that produces the balanced solution iteratively reduces the maximum exposure time of the angles that are not fixed, which intuitively means that we are re-distributing fluence over the remaining angles. As the maximum exposure time decreases, the exposure times for some angles needs to increase to guarantee an optimal treatment. The algorithm terminates as soon as the variables that are fixed by this “equalizing” process attain one of the bounds that describe an optimal treatment. At the algorithm’s conclusion, we have that $\alpha(a) = 0$ if and only if the exposure time of angle a is forced to zero when the other angles are set at their ‘smallest’ exposure time (smallest relative to the iterative process of reducing the maximum exposure time).

4 Numerical Comparisons

In this section we numerically compare how the resolution of the dose points affects set cover (SC), scoring (S), and vector quantization (VQ) 9-beam selectors. The Radiotherapy optimAl Design software (\mathcal{RAD}) at www.trinity.edu/aholder/research/oncology/rad.html was altered to accommodate the different beam selectors. This system is written in Matlab[©] and links

to the CPLEX[®] solvers (CPLEX v. 6.6. was used). The code, except for commercial packages, and all figures used in this paper (and more) are available at lagrange.math.trinity.edu/tumath/research/reports/misc/report97/.

The clinical example is an acoustic neuroma in which the target is encroaching the brain stem and is asked to receive between 48.08 and 59.36 Gy. The brain stem is restricted to no more than 50 Gy and the eye sockets to less than 5 Gy. Each image represents a 1.5 mm swath of the patient, and the 7 images in Figure 1 were used, creating a 10.5 mm thickness. The full clinical set contained 110 images, but we were unable to handle the full complement because of inherent memory limitations in Matlab. Angles are selected from $\{i\pi/36 : i = 1, 2, \dots, 71\}$. These




Figure 1: The target is immediately to the left of the brainstem. The critical structures are the brain stem and the two eye sockets.

candidate angles were assigned 12 different values as follows. An optimal treatment (according to judgment function (2)) for the full set of candidate angles was found with CPLEX’s primal, dual, and interior-point methods and a balanced solution according to (12) was also calculated. The angle values were either the average sub-beam exposure or the maximal sub-beam exposure. So, “BalancedAvg” indicates that the angle values were created from the balanced solution of a 72-angle optimal treatment, where the angle values were the average sub-beam exposure. Similar nomenclature is used for “DualMax”, “PrimalAvg”, and so on. This yields eight values. The scaled and unscaled set cover values in (5) and (6) were also used and are denoted by “SC1” and “SC2”. The informed entropy measure in (10) is denoted by “Entropy”, and the scoring technique in (9) is denoted by “S”. We used $TG = 0.5(TLB + TUB)$ in (9). So, in total we tested 12 different angle values for each of the beam selectors.

The dose points were placed on 3 mm and 5 mm grids throughout the 3D patient space, and each dose point was classified by the type of tissue it represented. Since the images were spaced at 1.5mm, we point out that dose points were not necessarily located on the images. The classification of whether or not a dose point was targeted, critical, or normal was accomplished by relating the dose point to the nearest image. In a clinical setting, the anatomical dose is approximated by a 2 mm or less spacing, so the experiments approach clinical practice. However, as with the number of images, Matlab’s memory limitation did not allow us to further increase the resolution.

Treatments are judged by viewing the level curves of the radiation per slice, called *isodose curves*, and by their cumulative dose-volume histogram (DVH). A dose-volume histogram is a plot of percent dose (relative to TLB) versus the percent volume. The isodose curves and DVHs for the balanced 72-angle treatment are shown for the 3 mm and 5 mm resolutions in Figures 2 through 5. An ideal DVH would have the target at 100% for the entire volume and then drop immediately to zero, indicating that the target is treated exactly as specified with no under or over dosing. The curves for the critical structures would instead drop immediately to zero, meaning that they

receive no radiation. The DVHs in Figures 3 and 5 are fairly good and follow this trend. The curves from upper-right to lower left are for the target, the brain stem, normal tissue, and the eye sockets. The eye socket curves drop immediately to zero as intended and appear on the axes. The 3 mm brain stem curve indicates that this structure is receiving more radiation than with the 5 mm resolution. While the fluence maps generated for these two treatments are different, the largest part of this discrepancy is likely due to the 3 mm spacing more accurately representing the dose variation.

Figures 6 and 7 are from Nomos’ tomotherapy system, which also uses 72 equally spaced angles (the curve for the normal tissue is not displayed). Two observations are important. First, the similarity between the DVHs of our computed solutions and Nomos’ DVHs suggests that our dose model and judgment function are accurate. Second, if our resolutions were decreased to 2 or 1.5 mm, it is likely that we would observe a brain stem curve akin to that in Nomos’ DVH. We point out the judgment function and solution procedure are different for Nomos’ system (and are proprietary).

A natural question is whether or not the dose point resolution affects the angle values. We expected some differences, but we generally thought that the values would remain intact when altering the resolution. We were surprised to find that some of the differences were rather dramatic. The 3 mm and 5 mm “average” values are shown in Table 1. The selected angles and solution

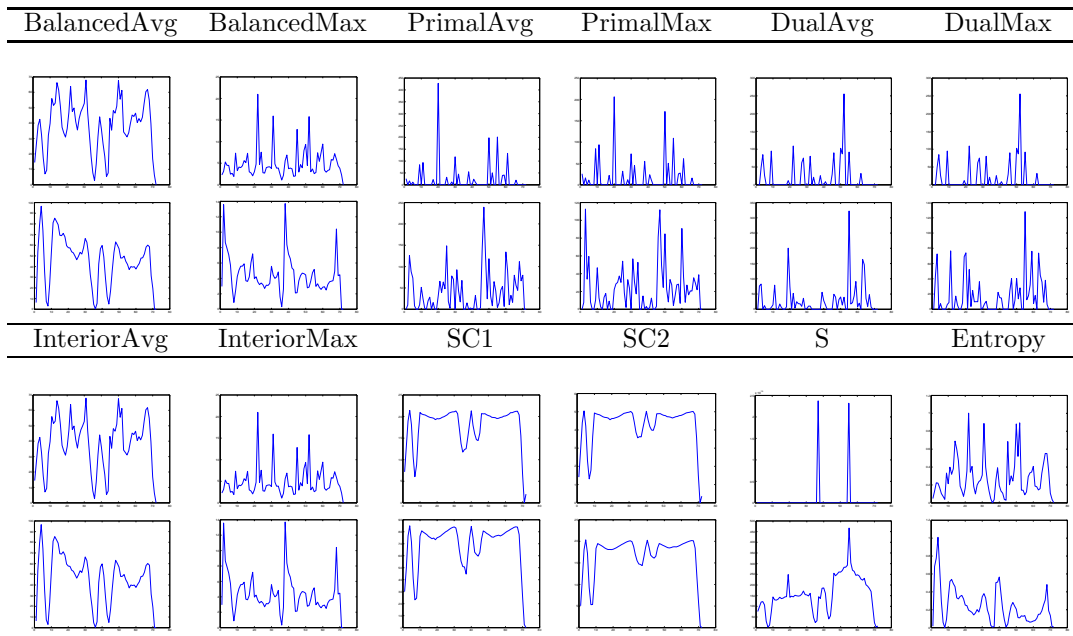


Table 1: The angle values. The top rows are with 5 mm resolution and the bottom rows are with 3 mm resolution.

times are shown in Tables 2 and 3. The angles vary significantly from beam selector to beam selector and for the same beam selector with different resolutions.

Measuring the quality of the selected angles is not obvious. One measure is of course the value of the judgement function. This information is shown in Table 4. The judgment values indicate that the 5 mm spacing is too coarse for the fluence model to adequately address the trade-offs between treating the tumor and not treating the brain stem. The 5 mm spacing so crudely approximates the anatomical structures that it was always possible to design a 9-beam treatment that treated the patient as well as a 72-beam treatment. The problem is that the boundaries between structures, which is where over and under irradiating typically occurs, are not well defined, and hence, the regions that are of most importance are largely ignored. These boundaries are better defined by the 3 mm grid, and a degradation in the judgment value is




Figure 2: The isodose contours for the balanced 72-angle treatment with 5 mm spacing.




Figure 3: The DVH for the balanced 72-angle treatment with 5 mm spacing.




Figure 4: The isodose contours for the balanced 72-angle treatment with 3 mm spacing.




Figure 5: The DVH for the balanced 72-angle treatment with 3 mm spacing.




Figure 6: The isodose contours for a clinical tomotherapy treatment.




Figure 7: The DVH for a clinical tomotherapy treatment.

Selector	Angle Value	Selected Angles									Time
Set Cover	BalancedAvg	15	20	25	55	60	65	70	85	240	113.51
	BalancedMax	10	15	20	95	190	195	200	275	340	126.23
	PrimalAvg	15	125	155	230	235	240	250	300	340	43.96
	PrimalMax	15	25	125	155	170	230	235	250	300	45.52
	DualAvg	10	15	55	95	100	275	295	315	320	34.02
	DualMax	15	55	95	100	110	275	295	315	320	68.80
	InteriorAvg	15	20	25	55	60	65	70	85	240	115.75
	InteriorMax	10	15	20	95	190	195	200	275	340	128.66
	SetCover 1	20	145	150	155	200	320	325	330	335	90.91
	SetCover 2	20	140	145	150	155	200	325	330	335	134.43
	Scoring	245	255	260	265	270	275	280	285	290	108.19
Entropy	10	15	20	25	55	60	195	200	240	144.43	
Scoring	BalancedAvg	15	20	25	55	60	65	70	75	85	104.93
	BalancedMax	10	15	20	25	95	190	195	200	340	108.29
	PrimalAvg	15	125	155	230	235	240	250	300	340	48.59
	PrimalMax	15	25	125	155	170	230	235	250	300	46.22
	DualAvg	10	15	55	95	100	275	295	315	320	36.24
	DualMax	15	55	95	100	110	275	295	315	320	66.56
	InteriorAvg	15	20	25	55	60	65	70	75	85	105.91
	InteriorMax	10	15	20	25	95	190	195	200	340	107.92
	SetCover1	20	145	150	155	200	320	325	330	335	83.87
	SetCover2	20	140	145	150	155	200	325	330	335	104.36
	Scoring	245	255	260	265	270	275	280	285	290	122.59
Entropy	10	15	20	25	55	60	190	195	200	235.84	
VQ	BalancedAvg	30	60	90	120	155	205	255	295	340	197.62
	BalancedMax	20	50	85	130	175	205	245	295	345	71.93
	PrimalAvg	35	90	135	190	235	250	280	320	350	55.27
	PrimalMax	20	70	125	160	205	245	275	305	340	121.91
	DualAvg	35	80	115	180	255	280	290	310	345	115.53
	DualMax	35	80	105	155	225	265	290	310	340	126.94
	InteriorAvg	30	60	90	120	155	205	255	295	340	198.43
	InteriorMax	20	50	85	130	175	205	245	295	345	71.98
	SetCover1	40	75	115	150	190	225	265	300	340	52.56
	SetCover2	40	75	110	145	185	230	265	300	340	187.10
	Scoring	50	95	135	185	230	260	285	305	340	134.33
Entropy	15	40	65	90	130	175	220	275	340	56.14	

Table 2: The angles selected by the different beam selectors with 3 mm resolution. The times are in seconds and include the time needed to select angles and design a treatment with these angles.

observed.

Judgment values do not tell the entire story and are only one of many ways to evaluate treatments. The mean judgment values of the different techniques all approach the goal value of -5.0000 , and claiming that one technique is better than another based on these values is tenuous. However, there are some outliers, and most significantly the scoring values did poorly with a judgment value of 3.0515 in the scoring and set cover beam selectors. The resulting 3 mm isodose curves and DVH for the scoring 9-beam selector are seen in Figures 8 and 9. These treatments are clearly inappropriate, especially when compared to Figures 4 and 5.

Besides the judgment value, another measure is to see how well the selected angles represent the intentions of the angle values. If we think of the angle values as probability densities, then the expected value of the 9 selected angles represents the likelihood of the angle collection being optimal. These expected values are found in Table 5. The trend to observe is that the set cover and scoring techniques select angles with higher expected values than the vector quantization technique, meaning that the angles selected more accurately represent the intent of the angle values. This is not surprising, as the set cover and scoring methods can be interpreted as attempting to maximize their expected value. However, if the angle assignments do not accurately gauge the intrinsic value of an angle, such accuracy is miss-leading. As an example, both the set cover and scoring methods have an expected value of 1 with respect to the scoring angle values in the 5 mm case. In this case, the only angles with nonzero values are 185 and 275, and the perfect expected value

Selector	Angle Value	Selected Angles									Time
Set Cover	BalancedAvg	55	70	75	110	155	250	260	330	335	4.32
	BalancedMax	110	120	155	225	245	250	260	295	300	4.46
	PrimalAvg	45	55	100	150	190	250	260	275	305	4.81
	PrimalMax	45	55	100	150	190	250	260	275	305	4.89
	DualAvg	20	45	110	160	230	250	255	260	275	4.96
	DualMax	20	45	110	160	230	250	255	260	275	5.04
	InteriorAvg	55	70	75	110	155	250	260	330	335	4.67
	InteriorMax	110	120	155	225	245	250	260	295	300	4.90
	SetCover 1	20	145	150	155	200	320	325	330	335	5.43
	SetCover 2	20	140	145	150	155	200	325	330	335	5.79
	Scoring	95	185	230	260	265	270	275	280	320	5.10
	Entropy	70	75	110	155	225	250	260	335	340	5.32
Scoring	BalancedAvg	55	70	75	110	155	250	260	330	335	2.12
	BalancedMax	110	120	155	225	245	250	260	295	300	2.34
	PrimalAvg	45	55	100	150	190	250	260	275	305	2.68
	PrimalMax	45	55	100	150	190	250	260	275	305	2.72
	DualAvg	20	45	110	160	230	250	255	260	275	2.88
	DualMax	20	45	110	160	230	250	255	260	275	2.94
	InteriorAvg	55	70	75	110	155	250	260	330	335	2.48
	InteriorMax	110	120	155	225	245	250	260	295	300	2.78
	SetCover1	20	145	150	155	200	320	325	330	335	3.31
	SetCover2	20	140	145	150	155	200	325	330	335	3.53
	Scoring	95	185	230	260	265	270	275	280	320	3.01
	Entropy	70	75	110	155	225	250	260	335	340	3.24
VQ	BalancedAvg	40	75	105	140	185	230	270	305	345	3.77
	BalancedMax	40	80	115	145	190	235	270	300	340	3.41
	PrimalAvg	40	85	105	130	175	225	265	290	330	3.32
	PrimalMax	30	80	105	130	175	225	260	270	320	4.11
	DualAvg	20	75	130	160	205	245	260	265	315	3.99
	DualMax	20	75	130	160	205	245	260	265	315	4.11
	InteriorAvg	40	75	105	140	185	230	270	305	345	4.40
	InteriorMax	40	80	115	145	190	235	270	300	340	4.03
	SetCover1	40	75	110	145	185	225	265	300	340	4.70
	SetCover2	40	75	110	145	185	230	265	300	340	4.88
	Scoring	185	190	195	200	240	280	285	290	330	5.58
	Entropy	45	75	105	140	195	240	270	305	345	4.75

Table 3: The angles selected by the different beam selectors with 5 mm resolution. The times are in seconds and include the time needed to select angles and design a treatment with these angles.

only indicates that these two angles are selected. A scoring technique that only scores 2 of the 72 possible angles is not meaningful, and in fact, the other 7 angles could be selected at random.

The expected values in Table 5 highlight how the angle assignments differ in philosophy. The weakly informed angle values attempt to measure each angle’s individual worth in an optimal treatment, regardless of which other angles are selected. The informed values allow the individual angles to compete through the optimization process for high values, and hence, these values are tempered with the knowledge that other angles will be used. The trend in Table 5 is that informed expected values are lower than weakly informed values, although this is not a perfect correlation.

From the previous discussions, it is clear that beam selectors depend on the dose point resolution, but none of this discussion attempts to quantify the difference. We conclude with such an attempt. For each of the selected sets of angles we calculated (in degrees) the difference between consecutive angles. These distances provide a measure of how the angles are spread around the great circle without a concern about specific angles. These values were compared in the 3 mm and 5 mm cases. For example, the 9 angles selected by the VQ selector with the BalancedAvg angle values were $\{30, 60, 90, 120, 155, 205, 255, 295, 340\}$ and $\{40, 75, 105, 140, 185, 230, 270, 305, 345\}$ for the 3 mm and 5 mm cases, respectively. The associated relative spacings are $\{30, 30, 30, 35, 50, 50, 40, 45, 50\}$ and $\{35, 30, 35, 45, 45, 40, 35, 40, 55\}$. This information allows us to ask whether or not one set of angles can be rotated to obtain the other. We begin by taking the absolute value of the corresponding relative spacings, so for this example the differences are

	SC		S		VQ		3 mm Mean
	3 mm	5 mm	3 mm	5 mm	3 mm	5 mm	
BalancedAvg	-5.0000	-5.0000	-5.0000	-5.0000	-4.9194	-5.0000	-4.9731
BalancedMax	-4.8977	-5.0000	-4.8714	-5.0000	-5.0000	-5.0000	-4.9230
PrimalAvg	-5.0000	-5.0000	-5.0000	-5.0000	-5.0000	-5.0000	-5.0000
PrimalMax	-5.0000	-5.0000	-5.0000	-5.0000	-5.0000	-5.0000	-5.0000
DualAvg	-5.0000	-5.0000	-5.0000	-5.0000	-3.5214	-5.0000	-4.5071
DualMax	-5.0000	-5.0000	-5.0000	-5.0000	-4.8909	-5.0000	-4.9636
InteriorAvg	-5.0000	-5.0000	-5.0000	-5.0000	-4.9194	-5.0000	-4.9731
InteriorMax	-4.8977	-5.0000	-4.8714	-5.0000	-5.0000	-5.0000	-4.9230
SC1	-4.9841	-5.0000	-4.9841	-5.0000	-5.0000	-5.0000	-4.9894
SC2	-4.9820	-5.0000	-4.9820	-5.0000	-4.9984	-5.0000	-4.9875
S	3.0515	-5.0000	3.0515	-5.0000	-4.9967	-5.0000	0.3688
Entropy	-5.0000	-5.0000	-5.0000	-5.0000	-5.0000	-5.0000	-5.0000
Mean	-4.3092	-5.0000	-4.3048	-5.0000	-4.8538	-5.0000	

Table 4: The judgment values of the selected angles.

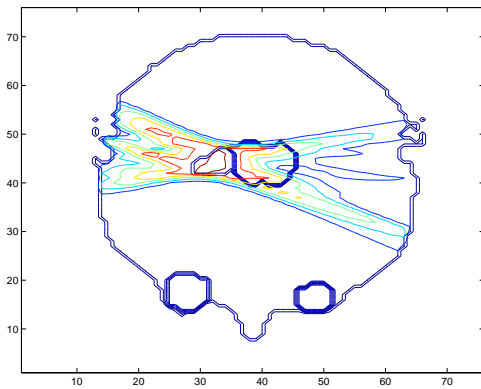


Figure 8: The 3 mm isodose contours for the balanced treatment when 9 angles were selected with a scoring method and scoring angle values.




Figure 9: The 3 mm DVH for the balanced treatment when 9 angles were selected with a scoring method and scoring angle values.

3 mm Relative Spacing	30	30	30	35	50	50	40	45	50
5 mm Relative Spacing	35	30	35	45	45	40	35	40	55
Difference	5	0	5	10	5	10	5	5	5

Depending on how the angles from the 3 mm and 5 mm cases interlace, we rotate (or shift) the first set to either the left or the right and repeat the calculation. In our example, the first angle in the 3 mm selection is 30, which is positioned between angles 40 and 345 in the 5 mm case. So we shift the 3 mm relative spacings to the left to obtain the following differences (notice that the first 30 of the 3 mm above is now compared to the last 55 of the 5 mm case).

3 mm Relative Spacing	30	30	35	50	50	40	45	50	30
5 mm Relative Spacing	35	30	35	45	45	40	35	40	55
Difference	5	0	0	5	5	0	10	10	25

The smallest aggregate difference, which is 50 in the first comparisons versus 60 in the second, is used in our calculations. We do not include all possible shifts of the first set because some spacial

	SC		S		VQ	
	3 mm	5 mm	3 mm	5 mm	3 mm	5 mm
BalancedAvg	0.2157	0.2059	0.2176	0.2059	0.1506	0.1189
BalancedMax	0.2613	0.3045	0.2673	0.3045	0.1234	0.1344
PrimalAvg	0.4191	0.8189	0.4191	0.8189	0.1600	0.0487
PrimalMax	0.4194	0.7699	0.4194	0.7699	0.1362	0.0429
DualAvg	0.6144	0.7443	0.6144	0.7443	0.0394	0.3207
DualMax	0.5264	0.7443	0.5264	0.7443	0.0359	0.3207
InteriorAvg	0.2157	0.2059	0.2176	0.2059	0.1506	0.1189
InteriorMax	0.2613	0.3045	0.2673	0.3045	0.1234	0.1344
SC1	0.1492	0.1461	0.1492	0.1461	0.1251	0.1248
SC2	0.1523	0.1491	0.1523	0.1491	0.1234	0.1273
S	0.2352	1.0000	0.2352	1.0000	0.1673	0.5058
Entropy	0.3176	0.3320	0.3303	0.332	0.1399	0.1402

Table 5: The expected values of the selected angles.

positioning should be respected, and our calculation honors this by comparing spacing between neighboring angles.

Table 6 contains the means and standard deviations of the relative spacing differences. A low

	Mean			Variance		
	SC	S	VQ	SC	S	VQ
BalancedAvg	45.56	47.78	5.55	2465.30	4706.90	9.03
BalancedMax	40.00	45.56	11.11	2125.00	3346.50	73.61
PrimalAvg	28.89	28.89	14.44	236.11	236.11	165.28
PrimalMax	16.67	16.67	13.33	325.00	325.00	131.25
DualAvg	37.78	37.78	16.67	1563.20	1563.20	150.00
DualMax	36.67	36.67	15.56	1050.00	1050.00	84.03
InteriorAvg	45.56	47.78	5.56	2465.30	4706.90	9.03
InteriorMax	40.00	45.56	11.11	2125.00	3346.50	73.61
SC1	0.00	0.00	1.11	0.00	0.00	4.86
SC2	0.00	0.00	0.00	0.00	0.00	0.00
S	40.00	40.00	35.56	3481.20	3481.20	1909.00
Entropy	44.44	44.44	13.33	1259.00	1552.80	81.25
Mean	31.30	32.59	11.94	1424.60	2026.30	224.25

Table 6: The mean and standard deviation of the (minimum) difference between the 3 mm and 5 mm cases.

standard deviation indicates that the selected angles in one case are simply rotated versions of the other. For example, the VQ selector with the InteriorAvg angle values has a low standard deviation of 9.03, which means that we can nearly rotate the 3 mm angles of {30, 60, 90, 120, 155, 205, 255, 295, 340} to obtain the 5 mm angles of {40, 75, 105, 140, 185, 230, 270, 305, 345}. In fact, if we rotate the first set 15 degrees, the average discrepancy is the stated mean value of 5.56. A low mean value but a high standard deviation means that it is possible to rotate the 3 mm angles so that several of the angles nearly match but only at the expense of making the others significantly different. Methods with high mean and standard deviations selected substantially different angles for the 3 mm and 5 mm cases.

The last row of Table 6 lists the column averages. These values lead us to speculate that the VQ techniques are less susceptible to changes in the dose point resolution. We were surprised that the SC1 and SC2 angle values were unaffected by the dose point resolution, and that each

corresponding beam selector chose (nearly) the same angles independent of the resolution. In any event, it is clear that the dose point resolution generally affects each of the beam selectors.

Besides the numerical comparisons just described, a basic question is whether or not the beam selectors produce clinically adequate angles? Figure 10 and 11 depict the isodose contours and a DVH of a typical clinical 9-angle treatment. This is not a final design, but rather is what is typical of an initial estimate of angles. Treatment planners would adjust these angles in an attempt to improve the design. Using the BalancedAvg angle values, we used Nomos' commercial software to design 9-angle treatments with the angles produced by the three different techniques with 3 mm spacing. Figures 12 through 17 contain the isodose contours and DVHs from Nomos software.




Figure 10: The isodose contours for a clinically designed (initial) treatment.




Figure 11: The DVH for the balanced 72-angle treatment with 5 mm spacing.

The set cover and scoring treatments in Figures 12 through Figures 15 are inferior to the initial clinical design. The problem is that the 9 angles are selected too close to each other. The fact that these are similar treatments is not surprising since the angle sets only differed by one angle. The vector quantization treatment in Figures 16 and 17 is clinically relevant, indeed it is an improvement over the initial clinical design. Specifically, only about 50% of the brain stem in the VQ treatment received above 30 Gy, whereas about 80% of the brain stem in the clinical treatment received above 30 Gy. In fact, the DVHs in Figure 11 and 17 indicate that the 9 angles selected by the vector quantization technique fare well if compared to the 72 angle tomotherapy treatment.

5 Conclusions

We have implemented several heuristic beam selection techniques and tested them on a clinical case with two different dose point resolutions. We have also (for the first time) compared the results with those from a commercial planning system and studied (again for the first time) the influence of dose point resolution on beam selection techniques.

References

- Ahuja, R. and Hamacher, H. (2004). A network flow algorithm to minimize beam-on-time for unconstrained multileaf collimator problems in cancer radiation therapy. *Networks*, **45**(1), 36–41.

- Baatar, D., Hamacher, H., Ehrgott, M., and Woeginger, G. (2004). Decomposition of integer matrices and multileaf collimator sequencing. *Discrete Applied Mathematics*, **152**, 6–34.
- Bortfeld, T., Boyer, A., Kahler, D., and Waldron, T. (1994). X-ray field compansation with multileaf collimators. *International Journal of Radiation Oncology, Biology, Physics*, **28**(3), 723–730.
- Ehrgott, M. and Johnston, R. (2003). Optimisation of beam directions in intensity modulated radiation therapy planning. *OR Spectrum*, **25**(2), 251–264.
- Ehrgott, M., Holder, A., and Reese, J. (2005). Beam selection in radiotherapy design. Technical Report 95, Department of Mathematics, Trinity University , USA.
- Gersho, A. and Gray, R. (1991). *Vector Quantization and Signal Compression*. Kluwer Academic Publishers, Boston, MA.
- Hamacher, H. and Küfer, K.-H. (2002). Inverse radiation therapy planing – A multiple objective optimization approach. *Discrete Applied Mathematics*, **118**(1-2), 145–161.
- Holder, A. (2001). Partitioning multiple objective optimal solutions with applications in radiotherapy design. Technical report, Department of Mathematics, Trinity University, San Antonio, USA.
- Holder, A. (2003). Designing radiotherapy plans with elastic constraints and interior point methods. *Health Care Management Science*, **6**(1), 5–16.
- Holder, A. and Salter, B. (2004). A tutorial on radiation oncology and optimization. In H. Greenberg, editor, *Emerging Methodologies and Applications in Operations Research*. Kluwer Academic Press, Boston, MA.
- Kamath, S., Sahni, S., Li, J., Palta, J., and Ranka, S. (2003). Leaf sequencing algorithms for segmented multileaf collimator. *Physics in Medicine and Biology*, **48**(3), 307–324.
- Lee, E., Fox, T., and Crocker, I. (2003). Integer programming applied to intensity-modulated radiation therapy treatment planning. *Annals of Operations Research*, **119**, 165–181.
- Lim, J., Ferris, M., Shepard, D., Wright, S., and Earl, M. (2002). An optimization framework for conformal radiation treatment planning. Technical Report Optimization Technical Report 02-10, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin.
- Löf, J. (2000). *Development of a general framework for optimization of radiation therapy*. Ph.D. thesis, Department of Medical Radiation Physics, Karolinska Institute, Stockholm, Sweden.
- Nizin, P. (1998). On absorbed dose in narrow ^{60}Co gamma-ray beams and dosimetry of the gamma knife. *Medical Physics*, **25**(12), 2347–2351.
- Nizin, P. and Mooij, R. (1998). Tissue-air ratios for narrow ^{60}Co gamma-ray beams. *Medical Physics*, **25**(8), 1458–1463.
- Preciado-Walters, F., Rardin, R., Langer, M., and Thai, V. (2004). A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer. *Mathematical Programming, Series B*, **101**, 319–338.
- Pugachev, A. and Xing, L. (2001). Pseudo beam’s-eye-view as applied to beam orientation selection in intensity-modulated radiation therapy. *International Journal of Radiation Oncology, Biology, Physics*, **51**(5), 1361–1370.
- Romeijn, H., Ahuja, R., Dempsey, J., Kumar, A., and Li, J. (2003). A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning. *Physics in Medicine and Biology*, **48**, 3521–3542.

- Romeijn, H., Dempsey, J., and Li, J. (2004). A unifying framework for multi-criteria fluence map optimization models. *Physics in Medicine and Biology*, **49**, 1991–2013.
- Rosen, I., Lane, R., Morrill, S., and Belli, J. (1991). Treatment planning optimisation using linear programming. *Medical Physics*, **18**(2), 141–152.
- Schlegel, W. and Mahr, A. (2002). *3D-Conformal Radiation Therapy: A Multimedia Introduction to Methods and Techniques*. Springer-Verlag, Heidelberg. Springer Verlag, Berlin.
- Siochi, A. (1999). Minimizing static intensity modulation delivery time using an intensity solid paradigm. *International Journal of Radiation Oncology, Biology, Physics*, **43**, 671–689.
- Söderström, S. and Brahme, A. (1992). Selection of beam orientations in radiation therapy using entropy and fourier transform measures. *Physics in Medicine and Biology*, **37**(4), 911–924.
- Spirou, S. and Chui, C.-S. (1998). A gradient inverse planning algorithm with dose-volume constraints. *Medical Physics*, **25**(3), 321–333.
- Wang, C., Dai, J., and Hu, Y. (2003). Optimization of beam orientations and beam weights for conformal radiotherapy using mixed integer programming. *Physics in Medicine and Biology*, **48**, 4065–4076.
- Webb, S. (2001). *Intensity-modulated radiation therapy (Series in Medical Physics)*. Institute of Physics Publishing.
- Winz, I. (2004). *A Decision Support System for Radiation Therapy Treatment Planning*. Master’s thesis, The University of Auckland, Department of Engineering Science, Auckland, New Zealand.
- Xia, P. and Verhey, L. (1999). Multileaf collimator leaf sequencing algorithm for intensity modulated beams with multiple segments. *Medical Physics*, **25**, 1424–1434.




Figure 12: The isodose contours for a clinically designed treatment based on the 9 angles selected by the set cover method with BalancedAvg angle values and 3 mm spacing.




Figure 13: The DVH for a clinically designed treatment based on the 9 angles selected by the set cover method with BalancedAvg angle values and 3 mm spacing.




Figure 14: The isodose contours for a clinically designed treatment based on the 9 angles selected by the scoring method with BalancedAvg angle values and 3 mm spacing.




Figure 15: The DVH for a clinically designed treatment based on the 9 angles selected by the scoring method with BalancedAvg angle values and 3 mm spacing.




Figure 16: The isodose contours for a clinically designed treatment based on the 9 angles selected by the vector quantization method with BalancedAvg angle values and 3 mm spacing.




Figure 17: The DVH for a clinically designed treatment based on the 9 angles selected by the vector quantization method with BalancedAvg angle values and 3 mm spacing.