

3-2008

Beam Selection in Radiotherapy Design

M Ehrgott

Allen G. Holder

Trinity University, aholder@trinity.edu

Josh Reese

Trinity University, jreese@trinity.edu

Follow this and additional works at: https://digitalcommons.trinity.edu/math_faculty



Part of the [Mathematics Commons](#)

Repository Citation

Ehrgott, M., Holder, A., & Reese, J. (2008). Beam selection in radiotherapy design. *Linear Algebra and its Applications*, 428(5-6), 1272-1312. doi:10.1016/j.laa.2007.05.039

This Post-Print is brought to you for free and open access by the Mathematics Department at Digital Commons @ Trinity. It has been accepted for inclusion in Mathematics Faculty Research by an authorized administrator of Digital Commons @ Trinity. For more information, please contact jcostanz@trinity.edu.

Beam Selection in Radiotherapy Design

M. Ehrgott^a and A. Holder^b and J. Reese^c

August 11, 2005

Abstract

The optimal design of a radiotherapy treatment depends on the collection of directions from which radiation is focused on the patient. These directions are manually selected by a physician and are typically based on the physician's previous experiences. Once the angles are chosen, there are numerous optimization models that decide a fluency pattern (exposure times) that best treats a patient. So, while optimization techniques are often used to decide the length of time a patient is exposed to a high-energy particle beam, the directions themselves are not optimized. The problem with optimally selecting directions is that the underlying mixed integer models are well beyond our current solution capability. We present a rigorous mathematical development of the beam selection problem that provides a unified framework for the problem of selecting beam directions. This presentation provides insights into the techniques suggested in the literature and highlights the difficulty of the problem. We also compare several techniques head-to-head on clinically relevant, two-dimensional problems.

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

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

^b Trinity University Mathematics, San Antonio, Texas, USA, and the Department of Radiation Oncology, The University of Texas Health Science Center at San Antonio, San Antonio, TX, USA, aholder@trinity.edu.

^c Trinity University Mathematics, San Antonio, Texas, USA, jreese@trinity.edu. Supported by The Cancer Therapy and Research Center, San Antonio, TX, USA.

1 Introduction

Radiotherapy is the treatment of cancerous and displasiac tissues with ionizing radiation. These treatments are beneficial because non-cancerous cells have a *therapeutic advantage* over their cancerous counterparts, an advantage that allows healthy cells to correctly reproduce with slightly damaged DNA. Cancerous cells are in a heightened state of reproduction and do not have this correction mechanism, and hence, small amounts of DNA damage render the cell incapable of reproducing. There is a threshold at which the damage is so severe that both cancerous and non-cancerous tissues are destroyed, and the goal of treatment design is to focus the radiation so that 1) enough radiation is delivered to the targeted region to kill the cancerous cells and 2) surrounding anatomical structures are spared. It is worth pointing out that cancerous and non-cancerous cells are interspersed within the targeted region, and hence, physicians strive to deliver enough radiation to the target to destroy the cancerous cells but not enough to kill healthy cells.

Modern treatment technology allows patients to receive complicated treatments, and while it was possible in the past for a physician to manually design a treatment that took full advantage of the technology, the number of options that are available today places the optimal design of a treatment outside the realm of human awareness. As such, the use of optimization methods to systematically design appropriate treatments is now integral to the industry. There are numerous treatment paradigms that depend on the technology of the clinic, but the most modern procedures fall into the category called Intensity Modulated Radiotherapy (IMRT). These treatments use a *multileaf collimator* to shape the beam and control, or modulate, the dose that is delivered along a fixed direction of focus, see Figures 1 and 2. The plethora of clinical options, physical models, and optimization techniques are described in several survey articles, and we direct interested readers to [2, 19, 39, 43, 50].

Because of the significant complexity of the design process, treatment design is segmented into a three-phase process that 1) selects how to focus radiation on the patient, 2) decides a fluency pattern (exposure times) for the directions selected in phase one, and 3) chooses a delivery sequence that efficiently administers the treatment. The research and industrial communities have largely addressed the second phase because this is where optimization methods were directly and immediately useful. However, the directive of the science is to optimize the entire planning process, which is a monumental task that requires knowledge of sophisticated physical models that in turn lead to optimization problems that far exceed current limitations. Moreover, since there are a variety of clinical capabilities, each treatment must be designed with the knowledge of what is and is not possible in a specific clinical environment. So, while optimizing the entire process is the goal, researchers have initially directed themselves to the question of finding a fluency pattern because this was where they could help. With this said, researchers have recognized that optimization can aid the first and third phases of treatment design [25, 26, 34, 35], but instead of dealing with these questions within the global scope of calculating an optimal treatment, these problems have been considered sequentially.

This paper considers the first phase of IMRT design —i.e. we address the question of how to select the directions from which radiation is directed at the patient. This question is called the beam selection or geometry problem and is important for several reasons. First, re-directing the beam's path 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 often leads to inaccurate and potentially dangerous treatments. Most clinics treat patients steadily throughout the day, and to make sure that demand is satisfied, patients are usually treated in daily sessions of 15 - 30 minutes. Selecting



Figure 1: A Linear accelerator rotating through various angles. Note that the treatment couch is rotated.

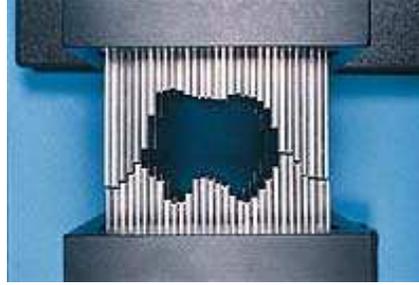


Figure 2: A multileaf collimator that is used to shape and modulate the intensity profile.

the beam directions is currently done manually, and hence, the process is time intensive and subject to the experience of the physician. The planning process proceeds as follows: the physician selects a collection of angles and waits 10 - 30 minutes while a fluency pattern is calculated. The resulting treatment is likely unacceptable, and the directions are adjusted and the process repeats. Finding a suitable collection of directions often takes several hours. The goal of using optimization methods to identify quality directions is to remove the dependency on a physician's experience and to alleviate the tedious repetitive process just described.

In this paper we present a rigorous mathematical development of the beam selection problem and compare different techniques on (non)clinical examples. Our mathematical development clearly describes the beam selection problem and provides notation that unifies the description of the techniques in the literature. The numerical comparisons are the first head-to-head comparisons in the field. A complete explanation of the underlying optimization models that are used to calculate a fluency pattern is beyond the scope of this paper, and we expect readers to be familiar with these models [2, 19, 39, 43, 50]. We also direct interested readers to the OR & Oncology web site at www.trinity.edu/aholder/HealthApp/oncology.

The mathematical presentation that follows depends heavily on collections of functions, and we use both functional and set notation. For example, we could refer to the collection of real-valued quadratics of the form ax^2 as $\{f : f(x) = ax^2, x \in \mathbb{R}\} : a \in \mathbb{R}$ or as $\{(x, ax^2) : x \in \mathbb{R}\} : a \in \mathbb{R}$. We use $\mathcal{W} \cong \mathcal{V}$ to mean that \mathcal{W} is isomorphic to \mathcal{V} , and we say that the two optimization problems

$$\min\{f(x) : x \in X\} \text{ and } \min\{g(y) : y \in Y\}$$

are *equivalent* if there is a bijection $h : X \rightarrow Y$ such that for all $x \in X$, $g(h(x)) = f(x)$. The notation $a := b$ means that a is defined to be b . A set subscript on a vector (matrix) represents the subvector (submatrix) whose elements (rows) correspond to the elements in the set. There are places where this notation is extended so that $D_{(T,B)}$ is the submatrix of D whose row indices are in T and whose column indices are in B . Other terms and notation are consistent with Greenberg's *Mathematical Programming Glossary* located at <http://carbon.cudenver.edu/~hgreenbe/glossary/>.

2 The Beam Selection Problem

The beam selection problem is to find N positions for the patient and gantry from which the treatment will be delivered. As an example, in Figure 1 we see that the gantry of the linear accelerator can rotate around the patient in a great circle and that 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 interfere with the patient. For geometrical and computational simplicity, we consider 2D treatment design instead of the entire 3D geometry. This means that we only consider a 2D slice of the anatomy, which is represented by a single MRI or CAT scan image. Although we only consider a 2D image, we use an accurate 3D dose model developed by Nizin [31, 32]. There are numerous dose models in the literature [48], with the gold standard being a Monte Carlo technique that simulates each particle’s path through the anatomy. Nizin’s work shows that his model is at least 97% accurate, making it clinically relevant. Figures 3 and 4 depict the situation. In Figure 3 the beam of radiation is directed toward the patient along angle a . To accommodate the multileaf collimator, the beam is divided into sub-beams, with the number and size of these sub-beams being dependent on the specific collimator. A sub-beam’s dimensions are stated by the manufacturer ‘at depth’. This is demonstrated in Figure 4. The 2D plane that contains the patient image is depicted by thicker lines, and the patient is positioned so that the point about which the gantry rotates (called the *isocenter*) is located in the center of the target. The angle defines an orthogonal plane that passes through the isocenter, and the dimensions of the sub-beams are defined on this plane.

Positions within the anatomy where dose is calculated are called *dose-points*. These are indicated in Figure 3 by the grid of dots. Placing dose-points in the anatomy is nontrivial, and many techniques have been suggested [1, 5, 26, 30]. Each dose-point represents a 3D hyper-rectangle whose length and width are decided by the spacing of the dose-points and whose height is decided by the spacing between the images (an average image represents a 3mm thickness). The 3D dose model is needed to accurately capture how the radiation along a sub-beam is deposited into a 3D hyper-rectangle that is represented by a 2D dose-point.

We mention that all of the results in this paper naturally extend to the 3D case, where thickness is accomplished by staking the images, and to the case where gantry rotation is not limited to rotations around a single image. However, these generalizations complicate notation and hinder numerical experimentation without adding benefit. For these reasons, we only consider single 2D images and gantry rotations in the image plane. Within this context, the terms beam and angle are used interchangeably.

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 $\mathcal{A} = \{i\pi/180 : i = 0, 1, 2, \dots, 359\}$. The beam selection problem depends on a judgment function that describes how well a patient can be treated with a set of N angles. Allowing $\mathcal{P}(\mathcal{A})$ to be the power set of \mathcal{A} and \mathbb{R}_+^* to be the nonnegative extended reals —i.e. $\mathbb{R}_+^* := \{x \in \mathbb{R} : x \geq 0\} \cup \{\infty\}$, we define a judgment function as follows.

Definition 1 *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 goal of a judgment function is to capture the nature of how a treatment planner decides between good and bad treatments, and f is the optimal value of an optimization problem that decides a fluency pattern. Thus, for any $\mathcal{A}' \in \mathcal{P}(\mathcal{A})$,

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

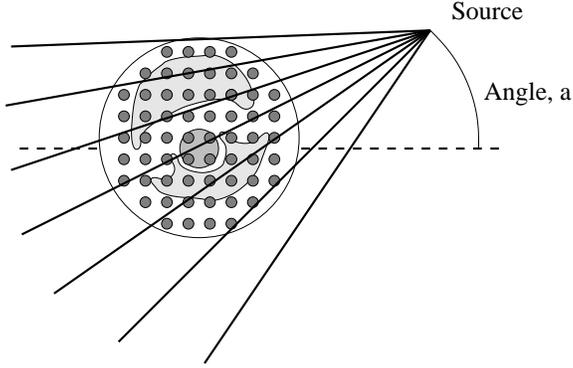


Figure 3: An angle’s sub-beams emanate from the gantry and pass through the anatomy that contains two sensitive regions and a circular target. The beam selection problem is to choose the best N angles.

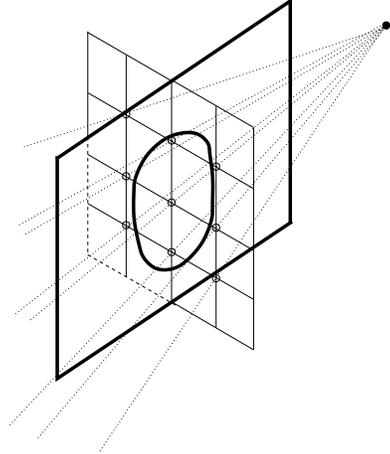


Figure 4: Although a 2D patient image is considered, an accurate 3D dose model is used.

where z maps a fluency pattern $x \in X(\mathcal{A}')$ into \mathbb{R}_+^* . If the demands of a physician are not possible with a collection of angles —i.e. $X(\mathcal{A}') = \emptyset$, then we assume that f assigns the value of ∞ to this collection. A fluency pattern is a vector x whose components $x_{(a,i)}$ represent the length of time that a patient is exposed to sub-beam i in angle a . The dependency of f on \mathcal{A}' is described by the constraints that define the feasible set, $X(\mathcal{A}')$. We ensure the requirement that $f(\mathcal{A}') \leq f(\mathcal{A}'')$ for $\mathcal{A}' \supseteq \mathcal{A}''$ by assuming that $X(\mathcal{A}') \supseteq X(\mathcal{A}'')$. We circumvent a dimensionality problem with the subset notation by assuming that if $\mathcal{A}' \subseteq \mathcal{A}$, then the sub-beams of angles not in \mathcal{A}' have a zero exposure time. Also, we only consider judgment functions with the property that $f(\mathcal{A}') < \infty$ for some $\mathcal{A}' \in \mathcal{P}(\mathcal{A})$, for if this were not the case, it would be impossible to design a treatment. We say that the fluency 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 that maps fluency patterns into anatomical dose, a prescription, and an objective function. Allowing D to be a dose operator, P to be a prescription, and z to be an objective function, we say that the triple (D, P, z) defines a judgment function. 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 positive for each (k, a, i) . It is worth mentioning that these rates are patient-specific constants, and hence, the operator that maps a fluency 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 fluency 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 fluency) 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 .

The majority of the literature is directed at modeling and calculating f , and while this is an

important issue, our goal is not to decide whether or not one method of finding a fluency pattern is better or worse than another. In fact, the authors' experience indicates that the manner in which a treatment is judged varies from patient to patient and from clinic to clinic. The problem here is that no single objective adequately measures the myriad of geometrical and numerical tools that a treatment planner uses to decide if a treatment is appropriate for a specific patient (some researchers have addressed the problem directly as a multiobjective problem [12, 16, 17, 22, 23, 45]). However, all commercial planning systems use an optimization routine to decide a fluency pattern, but the model and calculation method differ from system to system [50]. Our definition of a judgment function captures the single quality that is exhibited by all these methods, namely that if you allow a treatment to be designed with more angles, the quality of the treatment either improves or remains the same (low values of f are better than high values of f in the definition). This result is not immediately obvious from our definition, which says that the objective improves if we broaden any specific collection of angles. Our convention that angles outside of a candidate set, say \mathcal{A}' , have zero fluency allows us to conclude the stronger and more appropriate statement that a judgment function is non-increasing if we simply allow more angles than $|\mathcal{A}'|$ —i.e. the subset relationship is not necessarily required to (possibly) improve f .

Theorem 1 For all $0 \leq N < |\mathcal{A}|$ we have

$$\min\{f(\mathcal{A}') : \mathcal{A}' \in \mathcal{P}(\mathcal{A}), |\mathcal{A}'| = N + 1\} \leq \min\{f(\mathcal{A}'') : \mathcal{A}'' \in \mathcal{P}(\mathcal{A}), |\mathcal{A}''| = N\}.$$

Proof: From our assumption that angles outside of a candidate set of angles have zero fluency, we conclude

$$\begin{aligned} & \min\{f(\mathcal{A}'') : |\mathcal{A}''| = N, \mathcal{A}'' \in \mathcal{P}(\mathcal{A}')\} \\ &= \min\left\{f(\mathcal{A}') : |\mathcal{A}'| = N + 1, \mathcal{A}'' \in \mathcal{P}(\mathcal{A}'), \sum_i x_{(a,i)} = 0, \text{ for some } a \in \mathcal{A}'\right\} \\ &\geq \min\{f(\mathcal{A}') : |\mathcal{A}'| = N + 1, \mathcal{A}' \in \mathcal{P}(\mathcal{A})\}. \end{aligned}$$

■

There are a variety of forms that a prescription can have, each dependent on what the optimization problem is attempting to accomplish. As an example, the prescription P could be a vector whose elements p_k describe the ideal amounts of radiation for the dose-points. In this situation, we have that the judgment function is

$$f(\mathcal{A}') = \min\left\{z(x) : Dx = P, x \geq 0, \sum_i x_{(a,i)} = 0, a \in \mathcal{A} \setminus \mathcal{A}'\right\}.$$

We point out that the dependency on the collection of angles \mathcal{A}' is found in the requirement that $\sum_i x_{(a,i)} = 0$ for $a \in \mathcal{A} \setminus \mathcal{A}'$, which guarantees the non-use of angles outside of \mathcal{A}' . These constraints are needed since the dose matrix D includes a column for each sub-beam (a, i) , regardless of whether or not a is in \mathcal{A}' . This particular judgment function is not that useful because it requires the physician to describe an ideal prescription that is exactly attainable by a fluency pattern. However, since an ideal prescription would have the targeted dose-points uniformly receiving a high level of radiation

and all other tissues receiving no radiation, this optimization problem is often infeasible, and hence, $f(\mathcal{A}') = \infty$.

We partition the set of dose-points by letting T be the set of targeted dose points, C be the collection of dose-points in the critical regions, and N be the remaining dose-points. 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 fluency pattern x into the targeted region, the critical structures, and normal tissue, respectively. Although our focus is not on a specific judgment function, we use the following judgment functions to highlight how the techniques in the literature fit into our notation,

$$f(\mathcal{A}') = \min \left\{ w_T \|D_T x - TG\|_2 + w_C \|D_C x\|_2 + w_N \|D_N x\|_2 : x \geq 0, \sum_i x_{(a,i)} = 0, a \in \mathcal{A} \setminus \mathcal{A}' \right\} \quad (2)$$

and

$$\begin{aligned} f(\mathcal{A}') = \min \{ & l^T \alpha + u_C^T \beta + u_N^T \gamma : TLB - L\alpha \leq D_T x \leq TUB, \\ & D_C x \leq CUB + U_C \beta, D_N x \leq NUB + U_N \gamma, TLB \leq L\alpha, -CUB \leq U_C \beta, \\ & x \geq 0, \gamma \geq 0, \sum_i x_{(a,i)} = 0, a \in \mathcal{A} \setminus \mathcal{A}' \}. \end{aligned} \quad (3)$$

The prescription in the first problem is a vector of goal doses for the target, TG , and this judgment function uses the 2-norm in an attempt to attain the target dose and at the same time deliver no radiation to the remaining tissue. The weights w_T , w_C , and w_N place a preferential structure on these objectives. The second judgment function is linear and has a more complicated prescription. The vectors TLB and TUB are lower and upper bounds on the targeted dose points, CUB is a vector of upper bounds on the critical structures, and NUB is a vector of upper bounds on the normal tissue. The linear functions $L\alpha$, $U_C \beta$ and $U_N \gamma$ measure how the fluency pattern deviates from the goals of the prescription and l , u_C and u_N penalize these deviations. We typically consider the cases of L , U_C and U_N being either the identity matrix or the vector of ones, denoted by e , and l , u_C and u_N being positive.

For a fixed judgment function, the N-beam selection problem is

$$\min \{ f(\mathcal{A}') - f(\mathcal{A}) : \mathcal{A}' \in \mathcal{P}(\mathcal{A}), |\mathcal{A}'| = N \} = \min \{ f(\mathcal{A}') : \mathcal{A}' \in \mathcal{P}(\mathcal{A}), |\mathcal{A}'| = N \} - f(\mathcal{A}). \quad (4)$$

This problem is typically stated as a mixed integer (binary) extension of the optimization problem that defines f by adding binary variables y_a for all $a \in \mathcal{A}$ and constraints $\sum_i x_{(a,i)} \leq M y_a$ and $\sum_a y_a = N$ (where M is an arbitrarily large constant). Integer extensions are currently intractable because they are beyond modern solution capabilities. We point out that the difficulty arises because the number of binary variables is large and because f is evaluated by solving a large, continuous optimization problem. Attempts to solve these mixed integer problems are found in [13, 24, 33, 49], but $|\mathcal{A}|$ is severely restricted so that the number of binary variables is manageable. The feasible set for the beam selection problem has $\binom{|\mathcal{A}|}{N}$ elements, and since clinically relevant values of N range from 5 to 10 beams, there are between 4.9×10^{10} and 8.9×10^{18} subsets of $\{i\pi/180 : i = 0, 1, 2, \dots, 359\}$. This fact has lead researchers to investigate heuristics, and one of the contributions of this paper is that we define terminology that unifies the description of these heuristics. Our mathematical

presentation clearly and concisely relates techniques that are seemingly different from a clinical perspective. Our development provides a theoretical basis that explains the underlying difficulty of the different approaches, and it highlights how to merge heuristics to form new techniques.

We begin by defining a *beam selector*.

Definition 2 *The function $g : \mathcal{W} \rightarrow \mathcal{V}$ is a beam selector if*

1. \mathcal{W} and \mathcal{V} are subsets of $\mathcal{P}(\mathcal{A})$ and
2. $g(W) \subseteq W$ for all $W \in \mathcal{W}$.

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

There are several beam selection techniques in the literature, and we show how to represent them as beam selectors. These heuristics to the beam selection problem primarily fall into two categories, deterministic beam selectors that calculate a collection of beams in one step and iterative, stochastic beam selectors that randomly search for collections of angles and use the judgment function to decide a final collection.

3 Deterministic Beam Selectors

We first consider an approach that restricts the total exposure time of a beam. For each angle $a \in \mathcal{A}$, we select a nonnegative value M_a and require $\sum_i x_{(a,i)} \leq M_a$. We let M be the vector with components M_a and define

$$\mathcal{A}^M := \bigcup \left\{ \mathcal{A}' : \mathcal{A}' \in \mathcal{P}(\mathcal{A}), X(\mathcal{A}') \cap \left\{ x : \sum_i x_{(a,i)} \leq M_a, a \in \mathcal{A} \right\} \neq \emptyset \right\}.$$

So, \mathcal{A}_M excludes the angles that can not simultaneously satisfy the feasibility condition of being in $X(\mathcal{A}')$ and the limited total fluency constraint. As an example, if the lower bound TLB is positive in (3) and $L\alpha = TLB$ is inconsistent, then \mathcal{A}^M is empty for arbitrarily small M . But if f is defined by (2), then $\mathcal{A}_M = \mathcal{A}$ for any $M \geq 0$ because the only constraint in $f(\mathcal{A}')$ is the nonnegativity of x .

Definition 3 (Limited Fluency Beam Selector) *An LF-N beam selector is an N-beam selector*

$$g_f : \{\mathcal{A}^M\} \rightarrow \mathcal{P}(\mathcal{A}^M).$$

Letting $\mathcal{F}_{lf}^M(N)$ denote the collection of LF-N beam selectors, the limited fluency approach to the beam selection problem is

$$\min \{ f(g_f(\mathcal{A}^M)) : g_f \in \mathcal{F}_{lf}^M(N) \} - f(\mathcal{A}). \quad (5)$$

The insight into this approach is that small values of M_a restrict the feasible region of the underlying mixed integer models that directly solve the N-beam problem. Such restrictions are used in an

attempt to make the relaxed problems of a branch-and-bound method easier to solve [26]. The size of the search space in (5) depends on M , which controls the number of angles in \mathcal{A}^M . Since $\mathcal{A}^M \subseteq \mathcal{A}$, the search space for the limited fluency approach is no greater than the search space of the beam selection problem. The challenge is to find small enough values of M_a so that we can exactly solve (5) but large enough values so that the selected angles are meaningful in the unrestricted case —i.e. when $M_a = \infty$. As one would expect, the limited fluency problem approaches the beam selection problem as M increases, a result formally stated in Theorem 2.

Theorem 2 *There exists $\hat{M} \geq 0$ such that the LF- N beam selection problem in (5) and the beam selection problem in (4) have the same optimal value. Also, if M^j is a nondecreasing sequence in $\mathbb{R}^{|\mathcal{A}|}$ such that $\lim_{j \rightarrow \infty} M^j > \hat{M}$, then $f(\hat{g}_{lf}(\mathcal{A}^{M^j})) \downarrow f(\mathcal{A}^*)$, where \mathcal{A}^* is an optimal solution to the beam selection problem and \hat{g}_{lf} is any LF- N -beam selector that satisfies*

$$\hat{g}_{lf}(\mathcal{A}^{\hat{M}}) \in \operatorname{argmin}\{f(\mathcal{A}') : \mathcal{A}' \in \mathcal{P}(\mathcal{A}^{\hat{M}}), |\mathcal{A}'| = N\}.$$

Proof: Since M^j is nondecreasing, we have for any $\mathcal{A}' \in \mathcal{P}(\mathcal{A})$ that

$$\begin{aligned} X(\mathcal{A}') \cap \left\{ x : \sum_i x_{(a,i)} \leq M_a^j, a \in \mathcal{A} \right\} &\neq \emptyset \\ \Rightarrow X(\mathcal{A}') \cap \left\{ x : \sum_i x_{(a,i)} \leq M_a^{j+1}, a \in \mathcal{A} \right\} &\neq \emptyset. \end{aligned}$$

Hence, $\mathcal{A}^{M^j} \subseteq \mathcal{A}^{M^{j+1}}$, and we conclude that $X(\mathcal{A}^{M^j}) \subseteq X(\mathcal{A}^{M^{j+1}})$. We are subsequently guaranteed that $f(\mathcal{A}^{M^{j+1}}) \leq f(\mathcal{A}^{M^j})$. Let x^* be an optimal fluency pattern for \mathcal{A}^* , and for $a \in \mathcal{A}$, let $\hat{M}_a := \sum_i x_{(a,i)}^*$. Since $\lim_{j \rightarrow \infty} M^j > \hat{M}$, there is a j' such that $M^j \geq \hat{M}$ for all $j \geq j'$. We now have that $x^* \in X(\mathcal{A}^*)$ and that $\sum_i x_{(a,i)}^* \leq \hat{M}_a \leq M^j$, for $j \geq j'$. Hence, if $j \geq j'$, then $\mathcal{A}^* \subseteq \mathcal{A}^{\hat{M}} \subseteq \mathcal{A}^{M^j}$ and the beam selector $\{\mathcal{A}^{\hat{M}}\} \mapsto \mathcal{A}^*$ is contained in $\mathcal{F}_{lf}^{\hat{M}}$ and in $\mathcal{F}_{lf}^{M^j}$. From this we see that for $j \geq j'$,

$$\begin{aligned} f(\mathcal{A}^*) &\geq \min \left\{ f \left(g_{lf}(\mathcal{A}^{\hat{M}}) \right) : g_{lf} \in \mathcal{F}_{lf}^{\hat{M}}(N) \right\} \\ &\geq \min \left\{ f \left(g_{lf}(\mathcal{A}^{M^j}) \right) : g_{lf} \in \mathcal{F}_{lf}^{M^j}(N) \right\} \\ &\geq \min \{ f(\mathcal{A}') : \mathcal{A}' \in \mathcal{P}(\mathcal{A}), |\mathcal{A}'| = N \} \\ &= f(\mathcal{A}^*). \end{aligned}$$

■

Note, that it is not necessarily true that $\mathcal{A}^{\hat{M}}$ contains all the angles that can be used in an optimal treatment. However, each of the optimal N element collections of angles provides an \hat{M} as in the proof of Theorem 2. If we let M be greater than the largest of these, then \mathcal{A}^M is the entire set of angles that can be used in an optimal treatment. One may similarly think that the beam selection problem and the limited fluency approach are equal for large enough M , but this is not generally the case. The following theorem addresses the equivalence of the two problems.

Theorem 3 For any M , the search space of the limited fluency problem in (5) is no greater than the search space of the beam selection problem in (4). The two problems are equivalent for sufficiently large M if and only if for every $a \in \mathcal{A}$, there is an $\mathcal{A}' \in \mathcal{P}(\mathcal{A})$ such that $a \in \mathcal{A}'$, $|\mathcal{A}'| = N$, and $X(\mathcal{A}') \neq \emptyset$.

Proof: The fact that the search space for the limited fluency problem is smaller than the search space for the beam selection problem follows directly from the relationship that $\mathcal{A}^M \subseteq \mathcal{A}$, for any $M \in \mathbb{R}_+^{|\mathcal{A}|}$. Let $S = \{\mathcal{A}' : \mathcal{A}' \in \mathcal{P}(\mathcal{A}), |\mathcal{A}'| = N\}$ be the feasible region of the beam selection problem. The two problems are equivalent for sufficiently large M if and only if $\mathcal{A}^M = \mathcal{A}$, in which case the bijection from \mathcal{F}_{lf}^M onto S is defined by $g_{lf} \mapsto g_{lf}(\{\mathcal{A}^M\})$. First, notice that

$$X(\mathcal{A}') \cap \left\{ x : \sum_i x_{(a,i)} \leq M_a, a \in \mathcal{A} \right\} = \emptyset$$

for arbitrarily large M if and only if $X(\mathcal{A}') = \emptyset$. This follows since if $\hat{x} \in X(\mathcal{A}')$, then selecting $M_a = \sum_i x_{(a,i)}$ ensures that $\{x : \sum_i x_{(a,i)} \leq M_a, a \in \mathcal{A}\}$ contains \hat{x} . We now have from the definition of \mathcal{A}^M that $\mathcal{A}^M = \mathcal{A}$ for sufficiently large M if and only if for every $a \in \mathcal{A}$, there is an $\mathcal{A}' \in S$ such that $a \in \mathcal{A}'$ and $X(\mathcal{A}') \neq \emptyset$. ■

The significance of these results is that an optimal solution of the beam selection problem can be obtained by sequentially increasing M . An interesting question for further research is to find the smallest value of M so that the limited fluency approach identifies an optimal set of angles. Realistic upper bounds on M are obtained from clinical and physical considerations. A weakly informed heuristic for M is found in [26], where using the notation in (3), the authors set

$$M_a = \max \left\{ \frac{TUB_k}{\sum_i d_{(k,a,i)}} : k \in T \right\}.$$

This value of M_a makes sure that none of the tumor's upper bounds are violated for exposure times less than one unit. This is a sensible heuristic, but the authors of [26] still found it too large to solve (5) directly and resorted to an iterative stochastic beam selector that used this bound at each step, see Section 4.

We continue our investigation into beam selectors by considering the set covering approach found in [13]. An angle a covers the dose-point k if $\sum_i d_{(k,a,i)} \geq \varepsilon$, and for each $k \in T$, let $\mathcal{A}_k^\varepsilon = \{a \in \mathcal{A} : a \text{ covers dose-point } k\}$. It is important to notice that $\mathcal{A}_k^\varepsilon = \mathcal{A}$ for all $k \in T$ if and only if $0 \leq \varepsilon \leq \varepsilon^* := \min\{\sum_i d_{(k,a,i)} : k \in T, a \in \mathcal{A}\}$. Allowing $\mathcal{W}_{sc}^\varepsilon = \{\mathcal{A}_k^\varepsilon : k \in T\}$, we make the following definition.

Definition 4 (Set Cover Beam Selector) An SC- N -beam selector is an N -beam selector having the form

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

Notice that since g_{sc} can not map to \emptyset , the mapping has to select at least one angle to cover each targeted dose-point. The isocenter is commonly placed at the center-of-mass of the targeted region, making the most common scenario the one where each targeted dose-point is covered by every angle.

Of course, this depends on the threshold that decides whether or not an angle covers a dose-point, but this threshold is typically small. Allowing $\mathcal{F}_{sc}^\varepsilon(N)$ to be the collection of SC-N-beam selectors, we have that the set covering approach to the beam selection problem is

$$\min \left\{ f \left(\bigcup_{W \in \mathcal{W}_{sc}^\varepsilon} g_{sc}(W) \right) : g_{sc} \in \mathcal{F}_{sc}^\varepsilon(N) \right\} - f(\mathcal{A}). \quad (6)$$

As the following result shows, the set covering approach is typically equivalent to the beam selection problem.

Theorem 4 *If $0 \leq \varepsilon \leq \varepsilon^*$, the set covering problem in (6) is equivalent to the beam selection problem in (4).*

Proof: If each targeted dose point is covered by every angle, then $\mathcal{A}_k^\varepsilon = \mathcal{A}$ for all $k \in T$. Since the elements of $\mathcal{F}_{sc}^\varepsilon(N)$ are functions, the image of \mathcal{A} is unique. This means that

$$\begin{aligned} \mathcal{F}_{sc}^\varepsilon(N) &= \{ \{(\mathcal{A}_k^\varepsilon, \mathcal{V}_k) : k \in T\} : \mathcal{V}_k \in \mathcal{P}(\mathcal{A}_k^\varepsilon), |\cup_{k \in T} \mathcal{V}_k| = N \} \\ &= \{ \{(\mathcal{A}, \mathcal{V})\} : |\mathcal{V}| = N, \mathcal{V} \in \mathcal{P}(\mathcal{A}) \} \\ &\cong \{ \mathcal{V} : |\mathcal{V}| = N, \mathcal{V} \in \mathcal{P}(\mathcal{A}) \}. \end{aligned}$$

Notice that for any $\mathcal{A}' \in \mathcal{P}(\mathcal{A})$ such that $|\mathcal{A}'| = N$, there exists a unique $g_{sc}^\varepsilon \in \mathcal{F}_{sc}^\varepsilon(N)$ such that $g_{sc}(\mathcal{A}) \mapsto \mathcal{A}'$. Hence, $f(\mathcal{A}') = f(g_{sc}^\varepsilon(\mathcal{A}))$ and the proof is complete. \blacksquare

From Theorem 4 we see that under the common scenario of having each beam cover the entire target, the set cover problem solves the beam selection problem. However, since the feasible sets are isomorphic, there is no apparent numerical advantage to solving (6) instead of (4). This leads us to heuristic approaches, where the idea is to use information about the problem to selectively identify a subcollection of $\mathcal{F}_{sc}^\varepsilon(N)$, say $\hat{\mathcal{F}}_{sc}^\varepsilon(N)$, such that

$$\min \{ f(\mathcal{A}') : \mathcal{A}' \in \mathcal{P}(\mathcal{A}), |\mathcal{A}'| = N \} \approx \min \left\{ f \left(\bigcup_{W \in \mathcal{W}_{sc}} g_{sc}(W) \right) : g_{sc} \in \hat{\mathcal{F}}_{sc}^\varepsilon(N) \right\}. \quad (7)$$

Although not with this notation, this is the technique found in [13], where a traditional set covering problem is solved to identify a single g_{sc} . For each targeted dose-point k , let $q_{(k,a,i)}$ be 1 if $d_{(k,a,i)}$ is greater than some reasonably small threshold and 0 otherwise —i.e. $q_{(k,a,i)}$ is 1 if sub-beam i in angle a covers dose-point k . For each angle a , define

$$c_a = \begin{cases} \sum_{k \in C} \sum_i \frac{(u_C)_k \cdot q_{(k,a,i)}}{CUB_k}, & C \neq \emptyset, \\ 0, & C = \emptyset, \end{cases} \quad \text{and} \quad \hat{c}_a = \begin{cases} \sum_{k \in C} \sum_i \frac{(u_C)_k \cdot q_{(k,a,i)} \cdot d_{(k,a,i)}}{CUB_k}, & C \neq \emptyset, \\ 0, & C = \emptyset, \end{cases} \quad (8)$$

where u_C and CUB are part of the prescription in (3). Letting $q_{(k,a)}$ be 1 if angle a covers dose-point k and 0 otherwise, the authors of [13] use the following set covering problems to select a g_{sc} ,

$$\min \left\{ \sum_a c_a y_a : \sum_a q_{(k,a)} y_a \geq 1, k \in T, \sum_a y_a = N, y_a \in \{0, 1\} \right\} \quad \text{and} \quad (9)$$

$$\min \left\{ \sum_a \hat{c}_a y_a : \sum_a q_{(k,a)} y_a \geq 1, k \in T, \sum_a y_a = N, y_a \in \{0, 1\} \right\}. \quad (10)$$

The cost of selecting an angle is large if it contains sub-beams that intersect a critical structure with either a large penalty for violating its upper bound or a small upper bound. The \hat{c} coefficients are additionally scaled by the rate at which dose is deposited into dose-point k from sub-beam (a, i) . A solution to either problem defines an element of $\mathcal{F}_{sc}^\varepsilon(N)$ through the relationship that $a \in g_{sc}(\mathcal{A}_k^\varepsilon)$ if and only if $a \in \mathcal{A}_k^\varepsilon$ and $y_a = 1$. These N-beam selectors are weakly informed because they use information about the data that describes the judgment function, but they fail to be informed because they do not require an evaluation of f .

While the set covering problems in (9) and (10) are binary, they are generally easy to solve. In fact, in the ordinary situation of $\mathcal{A}_k^\varepsilon = \mathcal{A}$ for $k \in T$, we can completely classify the behavior of the set covering heuristic for selecting an N-beam selector.

Theorem 5 *Assume that $0 \leq \varepsilon \leq \varepsilon^*$, and order the angles so that $c_{a_1} \leq c_{a_2} \leq \dots \leq c_{a_{|\mathcal{A}|}}$. Let j' and j'' be the smallest and largest values, respectively, such that $c_{a_{j'}} = \dots = c_{a_N} = \dots = c_{a_{j''}}$. Then, y is an optimal solution to (9) if and only if*

- $y_{a_j} = 1$ for $1 \leq j < j'$,
- $\sum_{j' \leq j \leq j''} y_{a_j} = N - j'$, and
- $y_{a_j} = 0$ for $j > j''$.

Proof: From the assumption that $0 \leq \varepsilon \leq \varepsilon^*$, we see that y is feasible if and only if $\sum_a y_a = N$, and hence, the optimal solutions are those that satisfy $\sum_a c_a y_a = \sum_{j=1}^N c_{a_j}$. The result follows immediately since $c_{a_{j'}} = \dots = c_{a_N} = \dots = c_{a_{j''}}$. ■

The following corollary is immediate.

Corollary 1 *If $0 \leq \varepsilon \leq \varepsilon^*$, the problem in (9) has $\binom{j''-j'+1}{N-j'+1}$ optimal solutions.*

From Corollary 1 we see that for sufficiently small ε , the set covering heuristic for the beam selection problem reduces the search space of the beam selection problem from $\binom{|\mathcal{A}|}{N}$ to $\binom{j''-j'+1}{N-j'+1}$, which is dramatic if $j' \approx j''$. Notice that the success of the heuristic lies in the ability of the cost coefficients to accurately predict quality angles.

Although the result in Theorem 5 is simple, it is important because it allows us to interpret the set covering approach as a scoring technique in the common situation of $\mathcal{A}_k^\varepsilon = \mathcal{A}$, where each angle is scored with either c_a or \hat{c}_a . Other researchers have heuristically solved the beam selection problem by scoring (collections of) angles and selecting the best N angles [36, 37, 41, 47, 51]. Such techniques fit within the framework of the set covering approach. To make this precise, we define a scoring selector as follows.

Definition 5 (Scoring Beam Selector) *An S-N-beam selector is an N-beam selector such that*

$$g_s : \{\mathcal{A}\} \rightarrow \mathcal{P}(\mathcal{A}).$$

It may seem odd that the definition of a scoring selector does not take into account the score that an angle is assigned. However, this is done intentionally since the value of an angle's score is only

used to decide whether or not the angle is selected —i.e. whether or not a is in $g_s(\mathcal{A})$. Allowing $\mathcal{F}_s(N)$ to be the collection of S-N-beam selectors, we immediately have from Definition 5 that

$$\min\{f(\mathcal{A}') : \mathcal{A}' \subseteq \mathcal{A}, |\mathcal{A}'| = N\} - f(\mathcal{A}) = \min\{f(g_s(\mathcal{A})) : g_s \in \mathcal{F}_s(N)\} - f(\mathcal{A}). \quad (11)$$

So, the scoring approach is equivalent to the beam selection problem.

An example of a scoring approach is found in [37], where each angle is assigned the score

$$c_a = \frac{1}{|T|} \sum_{k \in T} \sum_i \left(\frac{d_{(k,a,i)} \cdot \hat{x}_{(a,i)}}{TG} \right)^2 \quad (12)$$

and

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

An angle's score increases as the sub-beams that comprise the angle become capable of delivering more radiation to the target without violating the restrictions placed on the non-targeted region(s). It is important to note that this weakly informed beam selector considers sub-beams individually and then aggregates this information to form a score for the entire angle. High scores are considered desirable since they indicate that it is possible to deliver large amounts of radiation to the target while maintaining the restrictions on the remaining tissues. So, this scoring technique uses the bounds on the non-targeted tissues to form constraints, and the score represents how well the target can be treated under these constraints. This is the reverse of the perspective in (9) and (10), where the constraints attempt to guarantee that the target is treated and the objective function strives to reduce the damage to the critical structures. These philosophies are different, but from Theorem 5 we see that selecting the N highest scoring angles is the same as solving the set covering problem in (9), provided that $\mathcal{A}_k^\varepsilon = \mathcal{A}$ for $k \in T$ and that the cost coefficients are replaced by the negatives of the scores in (12). The important observation is that every scoring technique is a set covering problem.

Theorem 6 *If $0 \leq \varepsilon \leq \varepsilon^*$, then $\mathcal{F}_{sc}^\varepsilon(N) = \mathcal{F}_s(N)$.*

Proof: As in the proof of Theorem 4, we have that if $0 \leq \varepsilon \leq \varepsilon^*$, then

$$\begin{aligned} \mathcal{F}_{sc}^\varepsilon(N) &= \{ \{(\mathcal{A}_k^\varepsilon, \mathcal{V}_k) : k \in T\} : \mathcal{V}_k \in \mathcal{P}(\mathcal{A}_k^\varepsilon), |\cup_{k \in T} \mathcal{V}_k| = N \} \\ &= \{ \{(\mathcal{A}, \mathcal{V})\} : |\mathcal{V}| = N \} \\ &= \mathcal{F}_s(N). \end{aligned}$$

■

The interpretation of Theorem 5 for a scoring technique is simply that if there is a tie for the N^{th} best score, then angles having this score can be interchanged. The score is only used to order the beams, but once the ordering is known, the scores could be converted into 1s and 0s to represent the beams that are and are not selected. The importance of Theorem 6 is that it shows how the techniques in the literature relate independent of how they address the clinically important question of how to treat the target without damaging the remaining anatomy.

Two other scoring methods are found in [44]. Letting x^* be an optimal fluency pattern for \mathcal{A} , the authors of [44] define the entropy of an angle by $e_a := -\sum_i x_{(a,i)}^* \ln x_{(a,i)}^*$ and the score of a is

$$c_a = 1 - \frac{e_a - \min\{e_a : a \in \mathcal{A}\}}{\max\{e_a : a \in \mathcal{A}\}}. \quad (13)$$

In this approach, an angle's score is high if the optimal fluency pattern of an angle's sub-beams is uniformly high. So, an angle with a single high-fluency sub-beam would likely have a lower score than an angle with a more uniform fluency pattern. Unlike the scoring procedure in [37], this technique is informed since it requires a solution to $f(\mathcal{A})$. An important observation is that there is not necessarily a uniquely optimal fluency pattern, which means that this scoring technique is solver dependent. We discuss this in detail at the conclusion of this section, and an example in Section 5 shows how radically different the scores can be from the same judgment function. The second score in [44] is the the absolute value of the Fourier transform of $x_{(a,\cdot)}^*$. Details for this approach are not provided, and we do not investigate this technique.

We now turn our direction to an approach that is based on the data compression technique called vector quantization [20] (we direct readers to the text of Gersho and Gray [14] for further information on vector quantization). We say that \mathcal{A}' is a contiguous subset of \mathcal{A} if \mathcal{A}' is an ordered subset of the form $\{a_j, a_{j+1}, \dots, a_{j+r}\}$. So, if $\mathcal{A} = \{i\pi/180 : i = 0, 1, 2, \dots, 359\}$, then $\{0, \pi/180, 2\pi/180\}$ is a contiguous subset but $\{0, \pi/180, 3\pi/180\}$ is not. 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} .

Definition 6 (Vector Quantization Beam Selector) *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 for notational simplicity, we consider $g_{vq}(W_j)$ to be a_j instead of $\{a_j\}$. Allowing $\mathcal{F}_{vq}(N)$ to be the collection of VQ- N beam selectors, we have that the vector quantization approach to the beam selection problem is

$$\min \left\{ f \left(\bigcup_{j=1}^N g_{vq}(W_j) \right) : g_{vq} \in \mathcal{F}_{vq}(N) \right\} - f(\mathcal{A}). \quad (14)$$

Unlike the previous set covering and scoring approaches, the vector quantization problem is not equivalent to the beam selection problem. In fact, the size of the search space in (14) is

$$\binom{|\mathcal{A}|}{N} < |\mathcal{F}_{vq}| = \sum_{\{W_j\} \in \mathcal{W}_{vq}(N)} \left(\prod_{j=1}^N |W_j| \right) < \binom{|\mathcal{A}|}{N}^2, \quad (15)$$

which is larger than the number of N element subsets of \mathcal{A} . For example, if $N = 5$, the vector quantization approach searches a space of between 4.9×10^{10} and 2.4×10^{21} functions. The reason

the search space is larger is that we have to select a contiguous partition in addition to selecting N beams.

The increased search space initially makes this approach look unattractive. However, the choice of a VQ-N-beam selector relies on the probability that an angle is used in an optimal treatment. So, although the vector quantization approach is deterministic, its heart is in modeling the probability that an angle is selected. Let $\alpha(a)$ be the probability that angle a is selected in an optimal collection of N beams. The *distortion* of a quantizer is

$$\sum_{j=1}^N \sum_{a \in W_j} \alpha(a) \cdot m(a, g_{vq}(W_j)),$$

where m measures the distance from $a \in W_j$ to $g_{vq}(W_j)$. It is typical in the data compression literature for $m(u, v) = \|u - v\|_2$, but this measure only makes limited sense in our setting. A treatment planner typically considers angles that are spaced at least 5° apart because angles that differ by a degree or two are not significantly different in the clinic unless such small differences place a critical structure at risk. So, $m(a, g_{vq}(W_j)) = \|a - g_{vq}(W_j)\|_2$ measures some but not all of the clinical relevance that such a metric should capture. A more meaningful metric is not known for the beam selection problem, and developing such a measure promises to be a fruitful avenue of research. In what follows, we assume that $m(a, g_{vq}(W_j)) = \|a - g_{vq}(W_j)\|_2$.

Once the probability distribution α is known, an element of $\mathcal{F}_{vq}(N)$ is calculated to minimize distortion. The following result is found in [14].

Theorem 7 *The VQ-N-beam selector g_{vq} minimizes distortion for the contiguous partition $\{W_j : j = 1, 2, \dots, N\}$ if and only if g_{vq} satisfies*

$$\sum_{j=1}^N \sum_{a \in W_j} \alpha(a) \cdot m(a, g_{vq}(W_j)) \leq \sum_{j=1}^N \min \left\{ \sum_{a \in W_j} \alpha(a) \cdot m(a, a') : a' \in W_j \right\}. \quad (16)$$

Theorem 7 states that once a contiguous partition is selected, the VQ-N-beam selector that minimizes distortion over the individual contiguous subsets actually minimizes the overall distortion. In the special case of a continuous \mathcal{A} , say the interval $[0, 2\pi)$, the authors of [14] show that a solution to the minimization problem in (16) is

$$g_{vq}(W_j) = \frac{\sum_{a \in W_j} a \cdot \alpha(a)}{\sum_{a \in W_j} \alpha(a)}. \quad (17)$$

So, once a contiguous partition is known, the image of a contiguous set that minimizes distortion is the center-of-mass of the contiguous set. 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. For example, if α is the uniform distribution over $\{i\pi/180 : i = 0, 1, 2, \dots, 359\}$ and $W_1 = \{0, \pi/180, 2\pi/180, 359\pi/180\}$, then the center-of-mass over this set is

$$\left(\frac{(1/360)(359 + 360 + 361 + 362)}{4/360} \right) \bmod 360 = \frac{1}{2}.$$

Instead of solving (16) exactly, we proceed by calculating the center-of-mass and mapping angles not in \mathcal{A} to their nearest neighbor. Angles with ties, such as $1/2$ in this example, are mapped to

the next largest element of \mathcal{A} . We denote the elements of $\mathcal{F}_{vq}(N)$ that satisfy (17) by g_{vq}^{opt} , and we let $\mathcal{F}_{vq}^{opt}(N)$ be the collection of these functions. Theorem 8 shows the relationships between the different search spaces.

Theorem 8 *If $0 \leq \varepsilon \leq \varepsilon^*$, then*

$$\left| \bigcup_{g_{vq}^{opt} \in \mathcal{F}_{vq}(N)} g_{vq}^{opt}(\mathcal{W}_{vq}^{opt}) \right| \leq |\mathcal{F}_{vq}^{opt}(N)| \leq |\mathcal{F}_{sc}^\varepsilon(N)| = |\mathcal{F}_s(N)| \leq |\mathcal{F}_{vq}(N)|, \quad (18)$$

where \mathcal{W}_{vq}^{opt} is the domain of g_{vq}^{opt} .

Proof: From (11), (15), and Theorem 6 we have that $|\mathcal{F}_{sc}^\varepsilon(N)| = |\mathcal{F}_s(N)| \leq |\mathcal{F}_{vq}(N)|$. Since there are $\binom{|\mathcal{A}|}{N}$ contiguous partitions of \mathcal{A} that contain N sets, we have from Theorem 7 that $|\mathcal{F}_{vq}^{opt}(N)| \leq \binom{|\mathcal{A}|}{N}$. The first inequality holds since $g_{vq}^{opt} \in \mathcal{F}_{vq}^{opt}(N)$, and because it is possible for different contiguous partitions to map to the same image set. ■

As previously stated, the fact that $|\mathcal{F}_s(N)| \leq |\mathcal{F}_{vq}(N)|$ makes it appear as though the vector quantization technique is fruitless. However, the inequalities in (18) show that if we only consider VQ-N-beam selectors that minimize distortion, then we have to consider at most $\binom{|\mathcal{A}|}{N}$ collections of angles. It is possible for $|\mathcal{F}_{vq}^{opt}(N)| < |\mathcal{F}_{sc}^\varepsilon(N)|$. For example, if W_j is a contiguous subset of \mathcal{A} with the property that $\alpha(a) = 0$ for all $a \in W_j$, then no VQ-N-beam selector satisfies (17). The first inequality in (18) further highlights the possibility that different partitions can map to the same image —i.e. it is possible to have different contiguous partitions, say $\{W_j : j = 1, 2, \dots, N\}$ and $\{W'_j : j = 1, 2, \dots, N\}$, such that $g_{vq}^{opt}(W_j) = g_{vq}^{opt}(W'_j)$, $j = 1, 2, \dots, N$. For example, assume that α is the uniform distribution over $\mathcal{A} = \{i\pi/180 : i = 0, 1, \dots, 359\}$ and let θ range from 1 to 89. Then, if we let

$$\begin{aligned} W_1 &= \{a : 0 \leq a \leq \theta\pi/180 \text{ or } 2\pi - \theta\pi/180 \leq a \leq 2\pi\}, \\ W_2 &= \{a : \theta\pi/180 < a < \pi - \theta\pi/180\}, \\ W_3 &= \{a : \pi - \theta\pi/180 \leq a \leq \pi + \theta\pi/180\} \text{ and} \\ W_4 &= \{a : \pi + \theta\pi/180 < a < 2\pi - \theta\pi/180\}, \end{aligned}$$

we have for any g_{vq}^{opt} that

$$g_{vq}^{opt}(\{W_1, W_2, W_3, W_4\}) = \{0, \pi/2, \pi, 3\pi/2\}.$$

So, there are 89 different contiguous partitions of \mathcal{A} , and hence 89 different mappings, with an image of $\{0, \pi/2, \pi, 3\pi/2\}$.

If we alter (14) so that it becomes

$$\min\{f(g_{vq}^{opt}(\mathcal{A})) : g_{vq}^{opt} \in \mathcal{F}_{vq}^{opt}\}, \quad (19)$$

then we have from Theorem 8 that the search space of the vector quantization approach is no larger than the size of the search space for the beam selection problem. Indeed, if we limit our search space to VQ-N-beam selectors with unique images, the size of the space may significantly reduce. The precise relationship between this reduction and the probability distribution α is an open question.

Theorem 7 provides a basis for vector quantization heuristics to be applied to the beam selection problem. The heuristics select a contiguous partition from which a single VQ-N-beam selector is created according to condition (17). The process in [20] selects the zero angle as the start of the first contiguous set. The end point of the first contiguous set is

$$\max \left\{ a' : |\mathcal{A}| \sum_{a \leq a'} \alpha(a) < 1/N \right\}, \quad (20)$$

and the first element of the second contiguous set is the next largest element of \mathcal{A} . The end of the next contiguous set is found by replacing $1/N$ in (20) by $2/N$, and the process continues until a contiguous partition is formed. This is the same as forming the cumulative density and evenly dividing its range into N intervals. This technique is not that intelligent because the contiguous sets are decided by zero being an end point of the first contiguous set. As an alternative, we could use the same rule and rotate the starting angle through the 360 candidates, which would produce a maximum of 360 different VQ-N-beam selectors. We could then evaluate f over the image of each of these beam selectors and take the smallest value. This heuristic reduces the search space from $\binom{\mathcal{A}}{N}$ to no more than 360.

We suggest a more thoughtful heuristic instead of simply calculating the possible 360 different VQ-N-beam selectors that minimize distortion. The key is to intelligently select the contiguous partition, and we use an algorithm first discovered by Lloyd [27]. This iterative process uses the m metric to sequentially select new contiguous partitions. The idea is to start with an initial contiguous partition, say $\{W_j^0 : j = 1, 2, \dots, N\}$, and then calculate the unique function $g_{vq}^{opt} \in \mathcal{F}_{vq}^{opt}$ that corresponds to this partition (we assume that $\alpha(a) > 0$ for some a in each contiguous set, which guarantees the existence of the VQ-N-beam selector). The m metric is used to form a new contiguous partition as follows

$$W_j^1 = \{a \in \mathcal{A} : m(a, g_{vq}^{opt}(W_j^0)) = \min \{m(a, g_{vq}^{opt}(W_t^0)) : t = 1, 2, \dots, N\}\}.$$

For example, if $m(a, g(W_j)) = \|a - g(W_j)\|_2$, then the algorithm first calculates the center of mass of each of the initial contiguous sets, and then forms a new contiguous partition by re-assigning the angles to their nearest center of mass, where nearest is interpreted naturally. The process repeats until the contiguous partition remains unchanged. The following result is found in [14].

Theorem 9 *If $\{W_j' : j = 1, 2, \dots, N\}$ and $\{W_j'' : j = 1, 2, \dots, N\}$ are consecutive contiguous partitions from the Lloyd algorithm, then*

$$\sum_{j=1}^N \sum_{a \in W_j''} \alpha(a) \cdot m(a, g_{vq}^{opt}(W_j'')) \leq \sum_{j=1}^N \sum_{a \in W_j'} \alpha(a) \cdot m(a, g_{vq}^{opt}(W_j')).$$

Theorem 9 guarantees that the distortion is non-increasing as the Lloyd algorithm selects VQ-N-beam selectors. The hope is that the Lloyd algorithm terminates with a VQ-N-beam selector that minimizes distortion, but this is not guaranteed. Also, there is no theoretical foundation on why a VQ-N-beam selector that minimizes distortion is a clinically relevant beam selector, and the relationships among the probability density, the m metric, and the ability of the Lloyd algorithm to accurately select quality angles is not well understood.

The success of vector quantization 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. This means the vector quantization technique is under the auspices of probability modeling, and we conclude this section by discussing some different probability models. The first idea is to make a weakly informed probability distribution by normalizing the scoring techniques in (8) and (12). Using these probabilities in a vector quantization heuristic is new, and we test them in Section 5.

Alternatively, an informed density is suggested in [20], where the authors assume that an optimal fluency pattern of $f(\mathcal{A})$ contains information about which angles should and should not be used. This idea is similar to the informed scoring technique in (13). We mention that the scores in (13) can be normalized to create an informed model of α . The authors of [20] similarly calculate an optimal fluency pattern, x^* for \mathcal{A} , and let

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

Evaluating $f(\mathcal{A})$ with this probability model allows the angles to compete for exposure time, with the goal being to optimally treat the patient. If there is a unique optimal fluency pattern corresponding to $f(\mathcal{A})$, then an angle's probability is proportional to the unique exposure time that is required of that angle to optimally treat the patient. However, if there are alternative optimal fluencies, the interpretation of the probability is not clear. In fact, an example in Section 5 shows that different solution techniques can produce significantly different fluencies from the same judgment function. In this situation, the probability model in [20] has the undesirable property of being dependent on the solver, just like the scoring method in (13). This flaw is recognized in [18] and [20], and a path-following interior point method is suggested to circumvent the problem. However, while the solution of a path-following technique is theoretically favorable, the implementations typically fall short of the promised theoretical advantages (the problem being that a path-following technique often terminates with a degenerate optimal solution).

To remove the dependency on a particular solver, we consider generating a balanced probability distribution —i.e. one that is as uniform as possible. Such a distribution results from solving

$$\text{lexmin} \left(z(x), \text{sort} \left(\sum_i x_{(a,i)} : a \in \mathcal{A} \right) \right), \quad (21)$$

where sort is a function mapping \mathbb{R}^A into \mathbb{R}^A that reorders the components of the vector $(\sum_i x_{(a,i)})_{a \in \mathcal{A}}$ in a nonincreasing order. This particular type of lexicographic optimization has been considered in multicriteria optimization [3, 28] and has interesting properties [11].

The optimization problem in (21) is solved by an iterative process. Like the techniques in [20] and [44], the procedure first calculates $f(\mathcal{A})$. Holding $f(\mathcal{A})$ constant, the procedure iteratively minimizes the infinity norm of the fluency pattern. So, after the first iteration we have found the value, say z_1^* , for which we are assured that any fluency pattern with each exposure time below z_1^* is non-optimal. The sub-beams that necessarily attain z_1^* are assigned this value, fixing the exposure times for these sub-beams. The process repeats with the remaining sub-beams and terminates once all beams are held constant. The algorithm is found in Figure 5.

The probability model formed by the algorithm in Figure 5 has the favorable quality that $\alpha(a)$ is large if and only if an optimal treatment requires a lengthy exposure time for angle a . Unfortunately,

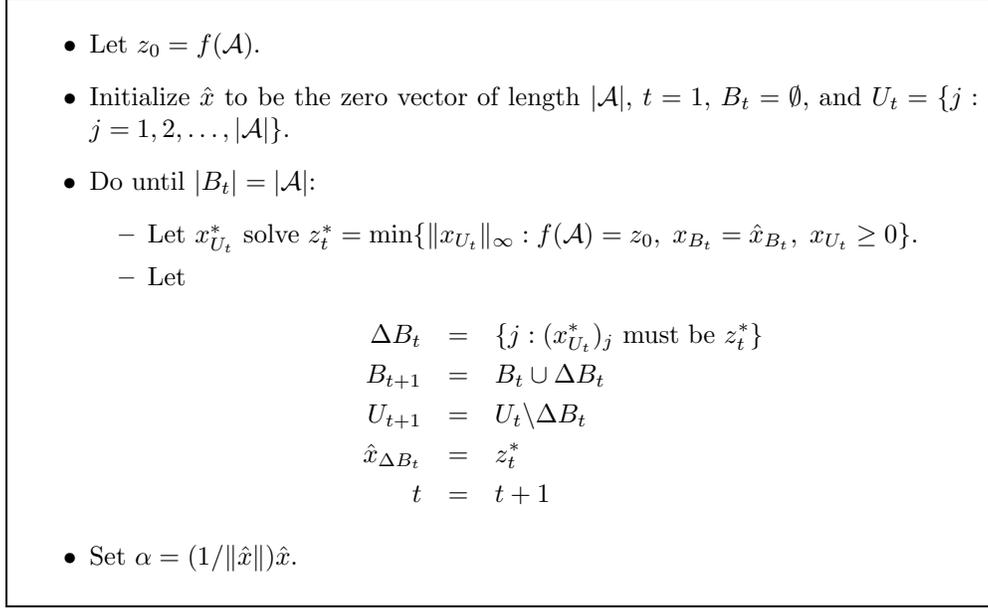


Figure 5: An iterative algorithm to calculate a solver-independent probability distribution that models an angles likelihood of being used in an optimal treatment.

we do not have a similar statement for the smaller components of α , and it is possible for $\alpha(a)$ to be zero even though the patient can be exposed to angle a in an optimal treatment. We make this precise by analyzing the algorithm with respect to the polytope

$$X^* = \{x : 0 < l \leq D_T x \leq u_1, D_{C \cup N} x \leq u_2, x \geq 0\}. \quad (22)$$

This polytope is general enough to represent the optimal fluency patterns of most representations of $f(\mathcal{A})$, including (2) and (3), and for this analysis we assume that

$$X^* = \{x : 0 < l \leq D_T x \leq u_1, D_{C \cup N} x \leq u_2, x \geq 0\} = \{x \in X(\mathcal{A}) : f(\mathcal{A}) = z(x)\}.$$

Theorem 10 explains how the algorithm terminates and how some sub-beams are assigned a zero probability.

Theorem 10 *With the notation of the algorithm in Figure 5, assume that f is a judgment function with the property that $x_{U_t}^*$ solves*

$$\begin{aligned} z_t^* &= \min\{\|x_{U_t}\|_\infty : f(\mathcal{A}) = z_0, x_{B_t} = \hat{x}_{B_t}, x_{U_t} \geq 0\} \\ &= \min\{\|x_{U_t}\|_\infty : x \in X^*, x_{B_t} = \hat{x}_{B_t}, x_{U_t} \geq 0\}. \end{aligned}$$

Then, the algorithm terminates with an α in the $(t+1)$ st iteration having the property that $\alpha(a) = 0$ for some $a \in \mathcal{A}$ if and only if either $D_{(T, B_t)} \hat{x}_{B_t} \geq l$ or

$$\begin{bmatrix} D_{(T, B_t)} \\ D_{(C \cup N, B_t)} \end{bmatrix} \hat{x}_{B_t} \leq \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}, \begin{bmatrix} D_{(T, B_t)} \\ D_{(C \cup N, B_t)} \end{bmatrix} \hat{x}_{B_t} \not\leq \begin{pmatrix} u_1 \\ u_2 \end{pmatrix}. \quad (23)$$

Moreover, $\alpha(a) = 0$ implies $a \in \Delta B_{t+1}$.

Proof: Let f be a judgment function with the stated property. Assume that the algorithm has entered the t^{th} iteration —i.e. that $B_t \neq \mathcal{A}$ and $U_t \neq \emptyset$. By design the optimization problem used to calculate z_t^* is feasible and $z_1^* \geq z_2^* \geq \dots \geq z_t^*$. If $D_{(T,B_t)}\hat{x}_{B_t} \geq l$, then

$$\begin{aligned} z_t^* &= \min\{\|x_{U_t}\|_\infty : x \in X^*, x_{B_t} = \hat{x}_{B_t}, x_{U_t} \geq 0\} \\ &= \min\{\|x_{U_t}\|_\infty : l - D_{(T,B_t)}\hat{x}_{B_t} \leq 0 \leq D_{(T,U_t)}x_{U_t} \leq u_1 - D_{(T,B_t)}\hat{x}_{B_t}, \\ &\quad D_{(C \cup N, U_t)}x_{U_t} \leq u_2 - D_{(C \cup N, B_t)}\hat{x}_{B_t}, x_{B_t} = \hat{x}_{B_t}, x_{U_t} \geq 0\}. \end{aligned}$$

Since $x_{U_t} = 0$ is feasible, it is optimal. So, $B_{t+1} = \mathcal{A}$ and $\alpha(a) = 0$ for $a \in U_t$.

Suppose that (23) holds. Then, for some $j \in T \cup C \cup N$ we have that $D_{(\{j\}, B_t)}\hat{x}_{B_t} = (u_1, u_2)_j^T$. This means that calculating z_t^* requires

$$x_{U_t} \geq 0 \quad \text{and} \quad D_{(\{j\}, U_t)}x_{U_t} \leq (u_1, u_2)_j^T - D_{(\{j\}, B_t)}\alpha_{B_t} = 0.$$

Since $D > 0$, the only solution is $x_{U_t} = 0$, and again we see that $B_{t+1} = \mathcal{A}$ and $\alpha(a) = 0$ for $a \in U_t$.

Alternatively, suppose that

$$l - D_{(T, B_t)}\hat{x}_{B_t} \geq 0, \quad l - D_{(T, B_t)}\hat{x}_{B_t} \neq 0, \quad \text{and} \quad \begin{pmatrix} u_1 \\ u_2 \end{pmatrix} - D_{B_t}\hat{x}_{B_t} > 0.$$

Then, $x_{U_t} = 0$ is not feasible to

$$\begin{aligned} z_t^* &= \min\{\|x_{U_{t+1}}\|_\infty : x \in X^*, x_{B_t} = \hat{x}_{B_t}, x_{U_t} \geq 0\} \\ &= \min\{\|x_{U_t}\|_\infty : 0 \leq l - D_{(T, B_t)}\hat{x}_{B_t} \leq D_{(T, U_t)}x_{U_t} \leq u_1 - D_{(T, B_t)}\hat{x}_{B_t}, \\ &\quad D_{(C \cup N, U_t)}x_{U_t} \leq u_2 - D_{(C \cup N, B_t)}\hat{x}_{B_t}, x_{B_t} = \hat{x}_{B_t}, x_{U_t} \geq 0\}. \end{aligned}$$

and we have that $z_t^* > 0$. There are two cases to consider. First, if $|B_{t+1}| = |\mathcal{A}|$, then the algorithm terminates with the smallest elements of \hat{x} being z_t^* . In this case no angle is assigned a zero probability. Second, if $|B_{t+1}| < |\mathcal{A}|$, then the algorithm continues to the $(t+1)$ st iteration. ■

The algorithm in Figure 5 iteratively reduces the maximum exposure time of the angles that are not fixed, which intuitively means that we are re-distributing fluency over the remaining angles. As the maximum exposure time decreases, the exposure times for some angles needs to increase to guarantee an optimal treatment. Theorem 10 shows that 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. Theorem 10 also helps us interpret the angles with a zero probability. 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 Iterative Beam Selectors

The previous beam selectors reduced the search space of the original beam selection problem and then selected a collection of angles from the reduced set, either by solving an optimization problem or through a rule. The beam selectors presented in this section are different in that they recognize the

immense size of the search space and iteratively select collections of angles, keeping track of which collection has the best judgment value. These iterative methods are either deterministic local-search techniques or stochastic global approaches.

The local search heuristics in [6, 13, 29] are naturally interpreted as iterative scoring methods. Starting with an initial collection of N angles, \mathcal{A}'_0 , a scoring selector is used to find a subsequent collection of angles, \mathcal{A}'_t , for $t \geq 1$. We let $c_{(a,t)}$ be the score of angle a in iteration t . Each of these methods sets $c_{(a,t)} = 0$ if a is not in a neighborhood of \mathcal{A}'_{t-1} , which makes these iterative approaches local search algorithms. The techniques differ in how they select \mathcal{A}'_0 and in how they update the scores.

The method in [13] assumes that the t^{th} angle set is $\mathcal{A}'_t = \{a_{(1,t)}, a_{(2,t)}, \dots, a_{(N,t)}\}$, where $a_{(i,t)} < a_{(j,t)}$ if $i < j$. The initial collection of angles is $\mathcal{A}'_0 = \{2\pi i/N : i = 1, 2, \dots, N\}$, and for $t = 1, 2, \dots, N$ the angles are scored by

$$c_{(a,t)} := \begin{cases} \gamma, & a \in \operatorname{argmin}\{f((\mathcal{A}'_{t-1} \setminus \{a_{(t,t-1)}\}) \cup \{a'\}) : a_{(t,t-1)} < a' < a_{(t+1,t-1)}, a' \in \mathcal{A}\} \\ 1, & a \in \mathcal{A}'_t \setminus \{a_{(t,t-1)}\} \\ 0, & \text{otherwise,} \end{cases}$$

where $0 < \gamma < 1$. With these scores, Theorem 5 implies that $a_{(t,t-1)}$ is replaced by an angle between $a_{(t,t-1)}$ and $a_{(t+1,t-1)}$ such that $c_{(a,t)} = \gamma$. The score of γ indicates that the replacement angle improves the judgment function as much as possible knowing that such a replacement is required. Notice that the score of the angle being replaced is $c_{(t,t)} = 0$. Obviously, this scoring technique is informed by evaluating f . If, after N iterations $f(\mathcal{A}'_N) < f(\mathcal{A}'_0)$, the heuristic repeats by setting $\mathcal{A}'_0 = \mathcal{A}'_N$. Otherwise it stops. The heuristics of [6] and [29] are similar but select \mathcal{A}'_0 differently and respectively score angles with

$$c_{(a,t)} := \begin{cases} \gamma, & a \in \operatorname{argmin}\{f((\mathcal{A}'_{t-1} \setminus \{a_{(t,t-1)}\}) \cup \{a'\}) : \\ & \quad a_{(t,t-1)} - u < a' < a_{(t,t-1)} + u, \\ & \quad a' \in \mathcal{A}, |a' - a''| \geq l, a'' \in \mathcal{A}'_{t-1} \setminus \{a_{(t,t-1)}\} \\ 1, & a \in \mathcal{A}'_{t-1} \setminus \{a_{(t,t-1)}\} \\ 0, & \text{otherwise} \end{cases} \quad (24)$$

and

$$c_{(a,t)} := \begin{cases} \gamma, & a \in \operatorname{argmin}\{f(\mathcal{A}'_{t-1} \cup \{a'\}) : a' \in \mathcal{A} \setminus \mathcal{A}'_{t-1}\} \\ 0, & a \in \operatorname{argmin}\{f(\mathcal{A}'_{t-1} \setminus \{a'\}) : a' \in \mathcal{A}'_{t-1}\} \\ 1, & a \in \mathcal{A}'_{t-1} \setminus \operatorname{argmin}\{f(\mathcal{A}'_{t-1} \setminus \{a'\}) : a' \in \mathcal{A}'_{t-1}\} \\ 0, & \text{otherwise.} \end{cases}$$

The parameters l and u in (24) guarantee that $a_{(t,t)}$ is replaced by an angle at least l degrees from any other angle in \mathcal{A}'_t but not more than u degrees from $a_{(t,t-1)}$ (the authors use $l = 5$ and $u = 50$). Similar to [13], the process in [6] repeats after N iterations as long as $\mathcal{A}'_N \neq \mathcal{A}'_0$. Otherwise it repeats once with $2u$ replacing u . Each of these iterative heuristics are informed because they require the evaluation of the judgment function (at each step).

Stochastic iterative beam selectors repetitively and randomly select candidate sets of angles and evaluate f , making them informed. This idea leads us to define a *stochastic- N -beam selector*.

Definition 7 *A stochastic- N -beam selector is an N -beam selector such that*

$$g_{st} : \{\mathcal{A}\} \rightarrow \mathcal{P}(\mathcal{A}),$$

where $g_{st}(\mathcal{A})$ is a random set according to some distribution.

As an example, the authors of [7] repeatedly select N angles uniformly from \mathcal{A} and evaluate f with each collection. The process stops once a fixed number of candidate sets have been generated, and the one with the best value of f is selected. This is an example of pure random search, and this iterative stochastic beam-selector is informed since it evaluates f .

The most popular stochastic heuristic is simulated annealing [4, 9, 35, 40, 46]. Allowing t to be an iteration counter, we choose $\mathcal{A}'_t = g_{st}(\mathcal{A})$ according to some stochastic process. The judgment function f is then evaluated and \mathcal{A}'_t is accepted as the incumbent if $f(\mathcal{A}'_t) < f(\mathcal{A}'_{t-1})$ or with some probability (that decreases as t increases) otherwise. The beam selectors g_{st} found in the literature often follow a uniform or Cauchy distribution.

We describe the simulated annealing heuristics in [9] and [46] (other papers report the use of simulated annealing but do not provide sufficient detail to reproduce their results). Let $\mathcal{A}'_{t-1} = \{a_j : j = 1, \dots, N\}$. In [9] and [46], a new set $\mathcal{A}'' = \{a_j + \phi : j = 1 \dots, N\}$ is chosen where ϕ is sampled from a Cauchy distribution with density function

$$p(\phi) = \frac{Y_t}{(\phi^2 + Y_t^2)^{\frac{N+1}{2}}}.$$

We note that the authors of [9] state that the numerator is instead 1, but this must be a typographical error since this function does not integrate to 1. In [46], $Y_t = Y_0/(1 + t/R)$, but values of Y_0 and R are not reported. In [9], $T_t = T_0/(1 + t)$ and $Y_t = \theta T_t^2$, where θ is chosen so that if $T_t = 1$, then each angle a has a 1% chance of being replaced by $a + \pi$. Neither a value of T_0 nor values of θ for which $T_t \geq 1$ are stated. The set of angles in the next iteration is defined probabilistically as

$$\mathcal{A}'_t = \begin{cases} \mathcal{A}'', & \text{with probability } \min\left(1, \exp\left(\frac{f(\mathcal{A}'_{t-1}) - f(\mathcal{A}'')}{T_t}\right)\right) \\ \mathcal{A}'_{t-1}, & \text{if } \mathcal{A}'' \text{ is not selected.} \end{cases} \quad (25)$$

Consequently, the acceptance probability follows a negative exponential distribution with parameter $1/T_t$. The algorithm in [46] stops as soon as $Y_t < 1$ and $(\bar{f} - f_{\min}/\bar{f}) < 0.01$, where $\bar{f} = \sum_{m=t-5N}^t f(\mathcal{A}'_m)$ is the average value of f over the last $5N$ iterations and $f_{\min} = \min\{f(\mathcal{A}'_m) : m = 1, \dots, t\}$ is the best value found so far. The stopping criterion in [9] is $T_t < 1$, which indicates that angles are only shifted by values up to π . The authors of [46] unfortunately set T_t to be 0 in (25), which means only new sets that reduce the judgment function are accepted, and hence, the simulated annealing approach reduces to pure random search with respect to the Cauchy distribution.

As previously stated, other papers reporting numerical results do not contain the information needed to reconstruct their technique. In [40], all that is known is that $T_t = 0$ (reducing to pure random search), that at least 100 iterations are performed, and that the stopping criterion is the same as in [46]. In [35], we only know that “ T_t decreases slowly” and that T_0 is bigger than the largest value of $f(\mathcal{A}')$ for several random beam orientations. In [38], we are simply told that $T_t = \theta T_{t-1}$, for $0 < \theta < 1$. The authors of this work combine their simulated annealing approach with a scoring technique to reduce the search space. Only collections of angles having a score as defined in (12) above a threshold are considered. The individual iterations of the stochastic beam selector in [38] are either informed or weakly informed since the evaluation of f is foregone if

$$\max \left\{ \left(\frac{c_a}{\max_{a \in \mathcal{A}} c_a} \right)^2 : a \in \mathcal{A}'_t \right\} < \varepsilon,$$

for some $\varepsilon > 0$. The value of ε is, however, not given. The idea here is that if the angles in the candidate set all score too low, then it is unlikely that this set is optimal, and hence, we should not waste time calculating f .

Evolutionary or genetic algorithms are used as a heuristic strategy in [15, 21, 42]. The main difference to the simulated annealing attack is that a set of beam configurations (a population) is maintained (and that a vector valued judgment function f is used), but the principle of iterating between stochastic beam selection and evaluation of f is the same. We explain the methods in [15] and [21]. As with the simulated annealing literature, the techniques in [42] do not contain enough information to reproduce their results (we are only told how the “package” in [8] is used).

Let \mathcal{A}^1 and \mathcal{A}^2 be N -element subsets of \mathcal{A} . The authors of [15] use the following genetic operators to get different sets of angles:

$$\begin{aligned}
(\mathcal{A}^1, \mathcal{A}^2) &\rightarrow \{a_i^1 + \gamma(a_i^2 - a_i^1) : i = 1, \dots, N\} \\
&\text{where } \gamma \in [-0.25, 1.25] \text{ is a random number} \\
\mathcal{A}^1 &\rightarrow (\mathcal{A}^1 \setminus \{a_i\}) \cup \{a_i + \delta\} \\
&\text{where } i \in \{1, \dots, N\} \text{ and } \delta \in [-\pi/12, \pi/12] \text{ are randomly chosen} \\
\mathcal{A}^1 &\rightarrow (\mathcal{A}^1 \setminus \{a_i\}) \cup \{a\} \\
&\text{where } i \in \{1, \dots, N\} \text{ and } a \in \mathcal{A} \text{ are randomly chosen} \\
\mathcal{A}^1 &\rightarrow (\mathcal{A}^1 \setminus \{a_i\}) \cup \{a_i + \pi\} \\
&\text{where } i \in \{1, \dots, N\} \text{ is randomly chosen} \\
\mathcal{A}^1 &\rightarrow \{a_i + (2\pi/N)k : k = 0, \dots, N-1\} \\
&\text{where } i \in \{1, \dots, N\} \text{ is randomly chosen.}
\end{aligned}$$

In Section 5 we report how a genetic algorithm based on these operators performs.

In [21], N_p sets of N angles are randomly chosen and this population is ranked in order of increasing values of f . Let $r(\mathcal{A}^j)$ be the rank of set \mathcal{A}^j in this order. N_r repetitions of the following procedure are performed.

1. Two sets \mathcal{A}^1 and \mathcal{A}^2 are chosen from the population according to the roulette-wheel selection mechanism, i.e. the probability of \mathcal{A}^j to be selected is

$$p = \frac{2(N_p - r(\mathcal{A}^j) + 1)}{N_p(1 + N_p)}.$$

With probability p_c a crossover operation is performed:

$$(\mathcal{A}^1, \mathcal{A}^2) \rightarrow (\{a_1^1, \dots, a_k^1, a_{k+1}^2, \dots, a_N^2\}, \{a_1^2, \dots, a_k^2, a_{k+1}^1, \dots, a_N^1\}),$$

where $k \in \{1, \dots, N\}$ is chosen randomly.

2. Both new sets undergo mutation with probability p_m . In that case Δ is randomly generated following a normal distribution with mean 0 and half-value-width $0.5|\mathcal{A}|$ and the mutation operator is

$$\mathcal{A} \rightarrow \{a_1, \dots, a_{k-1}, a_k + \Delta, a_{k+1}, \dots, a_N\},$$

where $k \in \{1, \dots, N\}$ is selected randomly.

The $2N_r$ newly created sets of beam angles replace the $2N_r$ lowest ranked ones in the previous population. We are told that this process is repeated for 500 iterations, that \mathcal{A} consists of 35 evenly spaced angles and that $N_p = 20$, but N_r and the probabilities for crossover and mutation are not given.

Before continuing with numerical comparisons, we mention the hybrid techniques in [10] and [26]. The process in [26] repetitively selects 10% of the candidate angles randomly and evaluates a judgment function. Based on which angles were used the most, the candidate set \mathcal{A} is pruned to a subset \mathcal{A}' . A limited fluency selector is then used to finalize the selected angles. So, this method uses a stochastic iterative method to seed a deterministic beam selector. In contrast, the method in [10] combines ideas of different deterministic selectors. The overriding goal is to solve an integer optimization problem that uses a neighborhood constraint to forbid selected angles from being too close. Another constraint disallows opposing angles. Each angle is scored by calculating the average damage to the critical structures observed by delivering a uniform dose of 2Gy to the target. This is somewhat the reverse of the scores in (12) since it treats the tumor and measures critical damage. These scores are used as the objective coefficients to solve an integer optimization problem that selects angles. The fact that it solves an integer problem makes it similar to a set cover technique, and if the neighborhoods are selected correctly, it may be interpreted as an SC-N-beam selector. The neighborhood constraint mimics the vector quantization approach since it attempts to spread the selected angles over the full rotation. If the neighborhood constraint were removed, the integer problem would reduce to a scoring method with the stipulation that opposing beams cannot both be selected.

5 Examples and Numerical Comparisons

In this section we numerically compare several of the beam selectors discussed in Sections 3 and 4. Everything was written in Matlab[®] and linked to CPLEX v. 6.6[®]. All examples were run on a dual 500 MHz workstation with 1 G of SRAM. Everything needed to reproduce the results in this section outside these commercial packages is available at <http://lagrange.math.trinity.edu/tumath/research/techreport.shtml>.

All examples represent a $10 \times 10 \times 0.3\text{cm}^3$ swath of the anatomy, with delivered dose being calculated every 2mm of the image plane. The dose calculations are 3D adaptations of the model in [31, 32] and are calibrated to a 6MV beam with a 25×25 field of $0.3 \times 0.3\text{cm}^2$ sub-beams (dimensions at isocenter - 100cm from source). A contour plot of a single sub-beam is seen in Figure 6. Notice that the beam delivers a maximum amount of radiation at a depth of 1.5cm and then attenuates as it continues through the anatomy. The model in [31, 32] does not account for the dose deposited in the first 1.5cm of the anatomy, and we linearly interpolate the from a depth of 0, which receives 60% of the maximum dose, to a depth of 1.5cm, which receives the maximum dose. While not exact, the interpolation is clinically sound. No adjustment for differing tissue densities was made —i.e. inhomogeneity correction was not used. The judgment function in (3) was used for all experiments, with the candidate set of angles being $\mathcal{A} = \{2\pi i/72 : i = 1, 2, \dots, 72\}$. This means that angles were selected from uniformly spaced angles every 5 degrees. We had hoped to select from 360 angles, but Matlab’s memory restriction did not allow this. Each beam selector was used to select angles for 5, 9 and 14 beam treatments.

The heart of many beam selectors is how values are assigned to angles. These values are used as scores in a scoring method, as objective coefficient in a set covering technique, and are normalized

to provide a probability for a vector quantizer. Value assignments were made in 3 basic ways:

Optimal Solutions the primal, dual, interior (barrier), and balanced solutions to (3) with the entire candidate set of angles,

Set Cover the unscaled and scaled values in (8) - labeled SC1 and SC2, and

Scoring the scoring assignment in (12) and the entropy assignment in (13).

Also, each of the value assignments from an optimal solution to (3) is constructed in two ways; by averaging the sub-beam exposures into an angle exposure and by using the maximum sub-beam exposure - denoted by avg and max, respectively. The entropy assignment is informed and could have been constructed with any of these optimal solutions. However, we only use the balanced solution because it is algorithm independent. Also, if $f(\mathcal{A})$ has a unique optimal fluency, then these value assignments are all the same. In total we explore 12 different value assignments.

We mention that not every situation warrants critical structures. As an example, radiosurgery differs from radiotherapy in that the entire therapeutic dose is delivered at once instead of being fractionated over several weeks. The intent of a radiosurgery is tissue ablation, and hence, is much like surgical removal. Conformality is the primary goal for radiosurgeries, and surrounding structures are destroyed if the dose is not precisely delivered to the target. For this reason, critical structures are often not delineated and the treatment is judged solely on how well the delivered dose conforms to the target. Two of our five examples do not have critical structures, and in these cases the SC1 and SC2 angle values do not make sense and are skipped.

The overall experimental design is highlighted in Table 1. For example, SC-N beam selectors were implemented for $N = 5, 9, 14$ —i.e. for 3 different beam sets. Each of these beam selectors was tested on 5 examples for the two thresholds of $0.50 \max_{p,q} \{(D_T)_{(p,q)}\}$ and $0.75 \max_{p,q} \{(D_T)_{(p,q)}\}$. This leads to the set cover method being tested 360 times. Surprisingly, the set cover techniques returned identical angle collections independent of the cover threshold, and we only report results for the 50% threshold (there were small fluctuations in run times, but otherwise these selectors were indistinguishable). The vector quantization and scoring ideas are similar but do not rely on a cover threshold, meaning that there is only one basic type of each. The selectors that were skipped because of no critical structures are included in these counts.

	SC	VQ	Scoring	GA	SA	LS
Angle Sets	3	3	3	3	3	3
Examples	5	5	5	5	5	5
Different Types	2	1	1	2	3	2
Angle Assignments	12	12	12	N/A	N/A	N/A
Total	360	180	180	30	45	30

Table 1: A taxonomy of the different beam selectors tested. The abbreviations are: SC - set cover, VQ - vector quantization, GA - genetic algorithm, SA - simulated annealing, and LS - local search.

The iterative beam selectors did not rely on assigning values to each angle, indicated by N/A in Table 1. The stochastic-N-beam selectors based on the genetic operations in [15] and [21] are initiated with random angle sets (populations) of size $2N$, for $N = 5, 9, 14$. At each iteration, $\lceil N/2 \rceil$ genetic operations are preformed on angle collections that are independently selected from

a uniform distribution. For example, if $N = 14$ and we are using the method in [15], then at each iteration we independently select 7 of the 5 operations, where the probability of selecting an operation is 0.20. These operations are performed on angle set(s) that are uniformly chosen from the current population, and because this selection is sequential and not simultaneous, it is possible for an angle collection to undergo several manipulations in a single iteration. The beam selector based on the operations in [21] is handled similarly, with all unspecified probabilities being uniform. Both selectors terminate when an update to the best known angle collection is not observed in 100 iterations. These selectors are denoted by GA1, based on [15], and GA2, based on [21].

We implemented 3 variations of a standard simulated annealing routine. In each iteration, an adjustment to the current set of angles is undertaken, and non-improving adjustments are accepted if a sample from the exponential distribution $-(1/T_t)e^{-t/T_t}$ is less than e^{-1/T_t} , where t is the iteration counter and T_t is the current temperature. The cooling structure is $T_{t+1} = T_t/(1 + c/10000)$, in which T_0 is 1000 and c is the number of iterations since the last successful update. This is a fairly generous cooling structure, and non-improving adjustments are often selected in the early iterations. All three variants begin with equally spaced angles but differ in how they make adjustments. All adjustments are selected from a Cauchy distribution with the scaling parameter being the current temperature and with median 0. The first method is the same as in [9], where the same adjustment is made to every angle in the current set. This technique maintains the initial spacing, making the initial guess extremely important. This dependence on the initial set did not seem natural, and the second implementation uniformly selects an angle to adjust. After N independent single angle updates, the new angle collection is accepted or rejected and the temperature is updated. Notice that a single angle may be updated several times in one iteration. The third algorithm adjusts each angle with its own independent sample from the Cauchy distribution. This is different than the second method because each angle is updated (unless the adjustment happens to be 0). All three simulated annealing selectors were initiated with N equally spaced angles, for $N = 5, 9, 14$, and are respectively referred to as SA1, SA2, and SA3.

Finally, the deterministic iterative beam selectors (local search methods) from [6] and [13] were implemented as described in Section 4. The value of l in (24) was 10, and two values of u were used to ensure that $\lceil a_{(t,t-1)} - (180/N) \rceil < a' < \lfloor a_{(t,t-1)} + (180/N) \rfloor$ (this was natural in our code because everything was handled with angle indices instead of angle values). The local search methods are referred to as LS1 and LS2, respectively. Both the genetic and simulated annealing selectors are capable of duplicating an angle in the selected sets, which occurred twice in our experiments. The only iterative selectors not capable of this nuisance are the local search methods.

Unfortunately, we were unable to solve the integer optimization problems corresponding to the beam selection problem and the limited fluency selectors. The problem again was Matlab's memory limitation. One of the authors' future goals is to port our implementations from Matlab so that the modeling is handled by a language such as AMPL or GAMS. This is a substantial undertaking, but we already have a Tcl script that manages problem instances and pipes directly to AMPL. The remaining significant hurdle is imaging, which is seamless in Matlab. The next phase of the numerical research is to use the developing code on realistic 3D clinical problems.

Before proceeding to the numerical results on clinical examples, we explain how the different heuristics behave in an 'ideal' situation of a spherical target that is centered within a cylindrical patient. While this situation is not clinical, it is an interesting case because medical physicists agree that the best treatment has a uniform fluency —i.e. one in which the exposure time for each angle is the same. Moreover, a best N-beam treatment has the property that the N angles are equidistant from each other. Since clinical consensus on a best treatment is rare, it is important to observe how

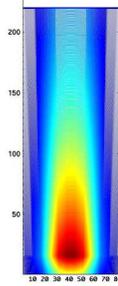


Figure 6: A contour plot of the dose delivered by a single sub-beam with a 1 second exposure.

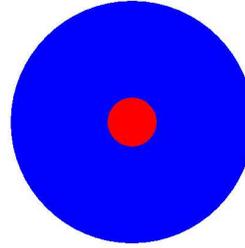


Figure 7: A fictitious example of a circular patient with a circular target in the center of the anatomy.

the beam selectors behave on an example where such consensus is known. The fictitious example is shown in Figure 7. The prescription for the target (center region) was between 80 and 88Gy.

The values assigned to the 72 angles for the 10 different value assignments are shown in Figures 8 through 17 (recall that 2 of the value assignments are not appropriate due to the absence of critical structures). The entropy assignments are on the order of 10^{-3} , causing our beam selector implementations to forgo this assignment because it was indistinguishable from zero. From the symmetry of this example, it is reasonable to expect that each angle should be assigned the same value. However, these graphs clearly show that this is only the case for the balanced solution. Moreover, the assignments from the primal and dual solvers have many of the angles receiving a value of 0, making it unlikely that these angles will be selected by a scoring, set cover, or vector quantization beam selector. Of course, this follows directly from the use of a simplex based algorithm on a degenerate problem. The interior solver nearly gets it correct since every angle has a positive value, a byproduct of CPLEX's barrier solver (with crossover turned off) returning a solution that is in the strict interior of the optimal set. The periodic behavior of the Interior-max assignment is interesting and unexpected.

In an attempt to compare the different beam selectors on this special example, we consider the spacing between the selected angles. This is a fair comparison because we have clinical consensus that an optimal treatment would have evenly spaced angles, and hence, the variance of the spacing should be zero. Tables 2 and 3 show how differently the beam selectors behaved and clearly shows that vector quantization consistently outperforms the scoring and set cover techniques. This is expected since the vector quantization methods naturally spread the angles over the candidate set, which is precisely what is desired. Notice that the first simulated annealing selector also did well, but this is a bit of a skulduggery. Recalling that this method maintains the initial uniform spacing, we see that this is pre-designed to do well on this example (the variance is not zero because we had to round to the nearest 5 degree angle in the candidate set). Contour maps of an optimal fluency for the SC-5-beam selector and the VQ-14-beam selector, both using the balanced-avg assignment, are shown in Figures 18 and 19.

Although the variance calculations indicate wide disparages in the beam selectors, we mention that the optimal treatment for many of the selectors appears satisfactory. For example, the angles

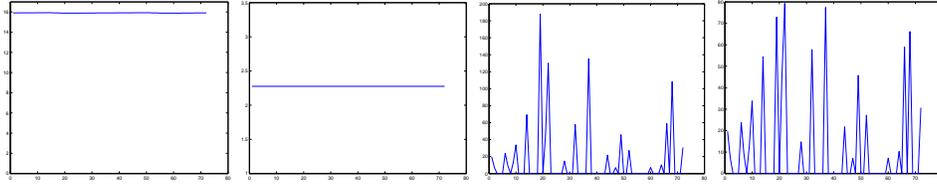


Figure 8: Balanced-avg angle values. Figure 9: Balanced-max angle values. Figure 10: Dual-avg angle values. Figure 11: Dual-max angle values.

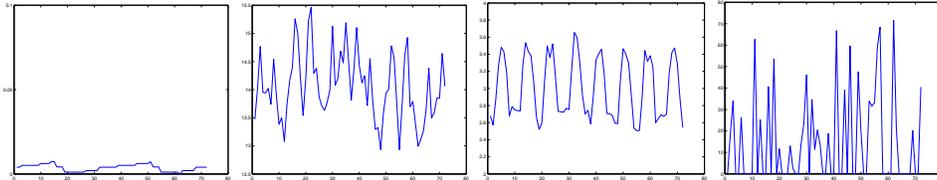


Figure 12: Entropy-balanced angle values. Figure 13: Interior-avg angle values. Figure 14: Interior-max angle values. Figure 15: Primal-avg angle values.

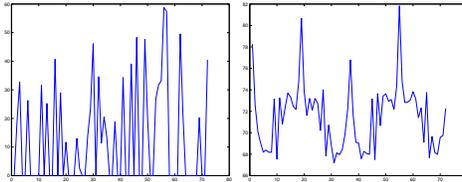


Figure 16: Primal-max angle values. Figure 17: Scoring angle values.

from the SC-5-beam selector have a huge variance, but in the end, the treatment in Figure 18 still largely treats the target. This leads to the possibility of low quality angle collections providing clinically acceptable treatments. Of course, the real problem is that the judgment function may not adequately capture the clinical desires, a topic of continued research.

Outside the fictitious example just discussed, the beam selectors were tested on four clinical examples: an acoustic neuroma, an arterial venous malformation (AVM), a pancreatic cancer, and a prostate. The examples are depicted in Figures 20 through 22. The AVM is a radiosurgical case and has no critical structures. This example is substantially simplified in our study since its full 3D representation shows a growth similar to a corkscrew, making conformality a significant challenge. The goal dose for the AVM target was between 80 and 88Gy.

The acoustic neuroma, pancreatic and prostate cases are typical radiotherapies. In the acoustic neuroma, the growth is to the left of the brain stem. Only the left eye socket is delineated, meaning that the treatment planner did not foresee a threat to the right eye socket. This is an interesting

	N	BA	BM	PA	PM	DA	DM	IA	IM	S
Scoring	5	8600	56	3040	1106	1356	3158	1490	1256	6
	9	4689	153	992	912	1070	1070	2313	257	1389
	14	2490	277	276	351	313	313	448	319	412
	avg	5260	162	1436	790	913	1514	1417	611	602
Vector Quantization	5	42	42	356	1173	1025	823	40	40	42
	9	7	10	186	143	513	298	7	10	14
	14	4	4	113	138	252	214	4	4	4
	avg	17	18	218	484	596	445	17	18	20
Set Cover	5	8092	56	3040	1106	313	3158	1490	1256	6
	9	5332	153	992	913	1070	1070	2313	257	1389
	14	2969	277	276	351	1356	313	1173	319	437
	avg	5464	162	1436	790	913	1514	1658	610	611

Table 2: Rounded variances for the scoring, vector quantization, and set cover beam selectors for the example in Figure 7. Acronyms: BA = Balanced-avg, BM = Balanced-max, PA = Primal-avg, PM = Primal-max, DA = Dual-avg, DM = Dual-max, IA = Interior-avg, IM = Interior-max, S = Scoring.

N	LS1	LS2	SA1	SA2	SA3	GA1	GA2
5	550	6	8	1425	5850	2050	3956
9	436	279	0	1542	1567	1158	614
14	104	4	4	174	140	305	276
avg	363	96	4	1047	2519	1171	1615

Table 3: Rounded variances for the iterative beam selectors.

observation, as it demonstrates that a treatment planner often guesses which organs are at risk. There is no doubt that the right eye socket would be distinguished at some clinics. For this problem, the goal dose for the tumor was between 76 and 85Gy and the eye socket and brain stem were restricted to 60 and 45Gy, respectively. The target in the pancreatic case is the large spherical growth in the middle of the image and is surrounded by the kidneys and spinal cord. The upper limit for the kidneys was 60Gy and for the spinal cord the limit was 45Gy. The target dose was between 80 and 88Gy. The prostate is sandwiched between the bladder (above) and the rectum (below). The prescription for the prostate example had the target receiving between 80 and 88Gy and the bladder and rectum being restricted to no more than 60 and 45Gy, respectively.

Unfortunately, consensus on a best treatment for these cases does not exist, making it difficult to compare the different beam selectors. The best rule-of-thumb to decide between good and bad treatments is the old adage, “You know one when you see one.” However, it is beyond the scope of this paper to display the numerous treatments provided by the various beam selectors. All treatments (including dose volume histograms) are available at www.trinity.edu/research/techreport.shtml. One thing we can measure is the computational time needed to select angles, and Table 4 contains the average time over all 5 examples for the different types of selectors. The time in minutes includes the time required to select angles and to solve for an optimal fluency. The longest running times were

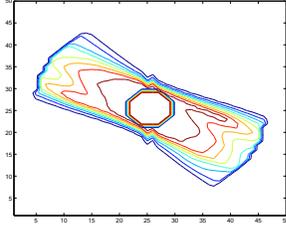


Figure 18: An optimal treatment for the angles selected by the SC-5-beam selector.

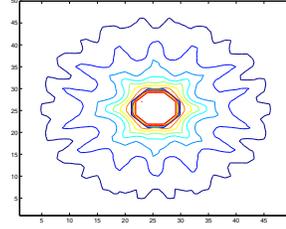


Figure 19: An optimal treatment for the angles selected by the VQ-14-beam selector.

observed with the iterative selectors, a consequence of them maneuvering through the entire search space while repetitively evaluating the judgment function. None of the scoring, vector quantization, or set cover techniques required more than a minute on average. The speed of the set cover selector was surprising, and these low values highlight the fact that the underlying integer problem is easy to solve.

S		VQ		SC		LS		SA		GA	
BA	0.43	BA	0.74	BA	0.55	LS1	10.40	SA1	36.73	GA1	14.46
BM	0.43	BM	0.74	BM	0.55	LS2	4.17	SA2	31.56	GA2	14.88
PA	0.46	PA	0.74	PA	0.58			SA3	32.06		
PM	0.46	PM	0.74	PM	0.58						
DA	0.46	DA	0.73	DA	0.58						
DM	0.47	DM	0.74	DM	0.58						
IA	0.47	IA	0.75	IA	0.58						
IM	0.47	IM	0.75	IM	0.61						
SC1	0.36	SC1	0.63	SC1	0.49						
SC2	0.37	SC2	0.63	SC2	0.49						
S	0.47	S	0.76	S	0.58						
ENT	0.36	ENT	0.58	ENT	0.48						

Table 4: The average run time for the different beam selectors over the 5 examples. Acronyms are as previously used.

Each beam selector did well at times, some more than others. The authors' general observations were:

- The vector quantizer with the balanced-avg angle assignments consistently produced quality treatments.
- The balanced and interior (both average and maximum) outperformed the primal and dual angle assignments, with the dual method often leading to poor quality angle collections.

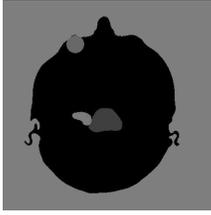


Figure 20: An acoustic neuroma: critical structures are the brain stem and the left eye socket (physician did not restrict the right eye socket).

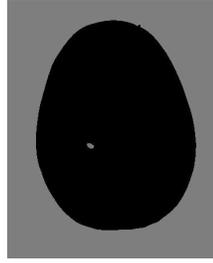


Figure 21: An arterial venous malformation: radiosurgical intent with no critical structures, conformality is extremely important.

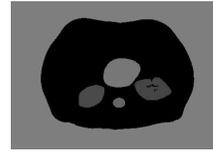


Figure 22: A pancreatic lesion: critical structures are the kidneys and the spinal cord.

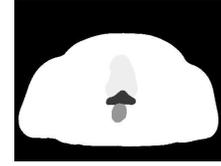


Figure 23: A prostate: critical structures are the bladder and rectum.

- Scoring methods were erratic.
- The local search algorithms worked well in general but were not as fast as vector quantization.
- The second and third simulated annealing selectors consistently produced quality treatments but were the most computationally heavy, often requiring hours to select beams.
- The genetic algorithm approaches did as well as simulated annealing in about half of the time.

The time needed by simulated annealing does not reward us with better angle collections. To demonstrate this, SA3 took 50.89 minutes to select 9 angles for the pancreatic case. In comparison, the scoring method with the SC1 scores and the vector quantization selector with the balanced-avg probability took only 0.45 and 0.77 minutes, respectively. The three treatments are shown in Figures 24 through 26. All three treatments are fair, but the simulated annealing treatment over irradiates the spinal cord. The scoring treatment irradiates the left kidney less than the vector quantization method, but treats the right kidney more than the vector quantization treatment.

The erratic behavior of the scoring selector is witnessed in the AVM and prostate examples. The scoring method used in conjunction with the primal or dual based scores only selected 2 angles in the AVM case. The problem was that the primal and dual fluencies of $f(\mathcal{A})$ used only 2 angles, and hence, the score for the rest of the angles was zero. In this case, our implementation only selected the two angles, the idea being that any other selection would have been random. This implies that these two angles are optimal for the beam selection problem for $N \geq 2$. In comparison, the other selectors returned a full set of N angles and the optimal fluencies used some subset of these. The optimal treatments for the S-5-beam selector with dual-max scores and the VQ-14-beam selector with the interior-max probability density are in Figures 27 and 28. The scoring selector additionally selected poor quality angles for the prostate case, with the optimal fluencies for the S-N-beam selectors all being zero for the SC1 and SC2 angle assignments - i.e. the angles selected by the scoring method led to no treatment being optimal. The prostate example was the most difficult to get good results

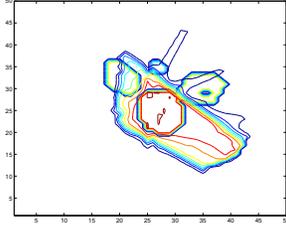


Figure 24: A 9 beam optimal treatment from SA3.

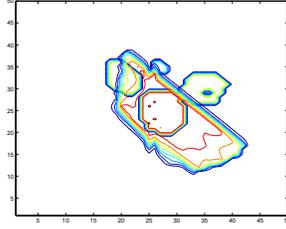


Figure 25: A 9 beam optimal treatment from scoring with the SC1 angle values.

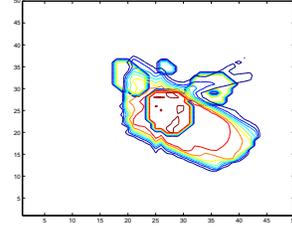


Figure 26: A 9 beam optimal treatment from vector quantization with the balanced-avg angle values.

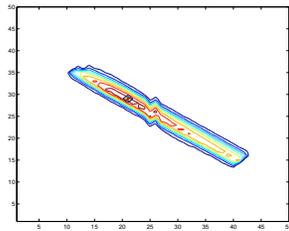


Figure 27: An optimal treatment for the AVM case using the scoring selector with dual-max scores.

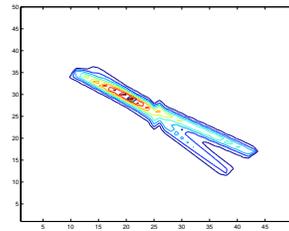


Figure 28: An optimal treatment for the AVM case using the vector quantization selector with the interior-max density.

for, with the majority of optimal fluencies over irradiating the critical structures. None of optimal fluencies were acceptable. We mention that the image was rescaled to achieve the $10 \times 10 \times 0.3\text{cm}^3$, which may have caused the poor results.

We conclude this section with the observation and the recognition that we could have used the judgment function to rank the beam selectors. First, we used the linear model in (3) because it was simple to link to CPLEX via Matlab and because the linear model allowed us to solve fairly large problems in a reasonable amount of time. As already stated, our goal was not to discuss what is a best judgment function. Indeed, a judgment function that is widely accepted in the clinic does not exist, and it seemed odd to rank selectors based on such speculative information. Besides, the objective values were close over the majority of examples, and deciding whether or not one treatment is better than another because this particular judgment function decreased by less than 1% did not make sense to the authors.

6 Conclusions

This research lays the mathematical foundation for a rigorous investigation into a long standing problem in medical physics. The previous attempts to solve the beam selection problem are unified in our development, showing that while their clinical perspectives may radically differ, these techniques are similar and intertwined. Moreover, we have implemented and combined several of the previous methods to demonstrate how they behave. These experiments are the first head-to-head comparisons in the area. Throughout this work we have pointed to interesting questions for future research. We conclude by mentioning a few that we feel are most interesting.

- The relationship between M and the angles in \mathcal{A}^M in the limited fluency approach begs for more analysis. This is potentially a critical step to numerically solving or estimating the beam selection problem.
- The convergence properties of the Lloyd algorithm in the vector quantization approach should be better understood. For example, it would be worthwhile to know the collection of initial contiguous partitions for which the Lloyd algorithm terminates with the same contiguous sets —i.e. what are the neighborhoods of stability?
- Although we have tested many of the previous heuristics, wide scale clinical testing was outside the scope of this work. Such testing is needed to compare how the different heuristics perform on different cancer types, different treatment modalities, and on different judgment functions. Moreover, combining the stochastic methods with the scoring, set covering, and vector quantization techniques is promising. Further numerical comparison is certainly warranted.
- Repeat the following n times: Choose a set \mathcal{A}' of N beams randomly and evaluate $f(\mathcal{A}')$. Then, let \bar{f}_n denote the mean of $f(\mathcal{A}')$ over the n repetitions. Is it true that $\lim_{n \rightarrow \infty} \bar{f} = \min\{f(\mathcal{A}') : \mathcal{A}' \in \mathcal{P}(\mathcal{A}), |\mathcal{A}'| = N\}$? If \bar{x}_n denotes the mean (balanced) optimal fluency pattern over the n repetitions, is it true that \bar{x}_n converges to an optimal fluency pattern for $f(\mathcal{A})$?
- The clinical consensus on the fictitious example makes one wonder if the desirable properties of an optimal treatment are the logical conclusion of model assumptions. Indeed, in an axiomatic development of fluency optimization one should be able to prove that such a treatment is optimal (unless this is one of the axioms). Such a development of fluency optimization does not exist, and research in this direction promises to be fruitful.

Acknowledgments

The authors are grateful for the assistance of Ryan Acosta and Daniel Nevin for their long hours making sure that the dose calculations were accurate. We thank Dr. Bill Salter for his clinical insights and advise. Lastly, we extend our gratitude to Denise Wilson for her help with L^AT_EX.

References

- [1] G. Bahr, J. Kereiakes, H. Horwitz, R. Finney, J. Galvin, and K. Goode. The method of linear programming applied to radiation treatment planning. *Radiology*, 91:686–693, 1968.

- [2] F. Bartolozzi et al. Operational research techniques in medical treatment and diagnosis. a review. *European Journal of Operations Research*, 121(3):435–466, 2000.
- [3] F. Behringer. A simplex based algorithm for the lexicographically extended linear maximin problem. *European Journal of Operational Research*, 7:274–283, 1981.
- [4] T. Bortfeld and W. Schlegel. Optimization of beam orientations in radiation therapy: Some theoretical considerations. *Physics in Medicine and Biology*, 38:291–304, 1993.
- [5] D. Cheek, A. Holder, M. Fuss, and B. Salter. The relationship between the number of shots and the quality of gamma knife radiosurgeries. Technical Report 84, Trinity University Mathematics, San Antonio, TX, 2004. to appear in *Optimization and Engineering*.
- [6] S. Das, T. Cullip, G. Tracton, S. Chang, L. Marks, M. Anscher, and 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):215–224, 2003.
- [7] S. Das and L. Marks. Selection of coplanar or noncoplanar beams using three-dimensional optimization based on maximum beam separation and minimized nontarget irradiation. *International Journal of Radiation Oncology, Biology, Physics*, 38(3):643–655, 1997.
- [8] K. Deb and T. Goel. Controlled elitist non-dominated sorting genetic algorithm for better convergence. In E. Zitzler, K. Deb, L. Thiele, C.A. Coello Coello, and D. Corne, editors, *Proceedings of the First International Conference on Evolutionary Multi-criterion Optimization EMO 2001, Zürich, Switzerland, 7-9 March 2001*, volume 1993 of *Lecture Notes in Computer Science*, pages 67–81. Springer Verlag, Berlin, 2001.
- [9] D. Djajaputra, Q. Wu, Y. Wu, and R. Mohan. Algorithm and performance of a clinical IMRT beam-angle optimization system. *Physics in Medicine and Biology*, 48:3191–3212, 2003.
- [10] W. D’Souza, R. Meyer, and L. Shi. Selection of beam orientations in intensity-modulated radiation therapy using single-beam indices and integer programming. *Physics in Medicine and Biology*, 49:3465–3481, 2004.
- [11] M. Ehrgott. Discrete decision problems, multiple criteria optimization classes and lexicographic max-ordering. In T. Stewart and R. van den Honert, editors, *Trends in Multicriteria Decision Making*, volume 465 of *Lecture Notes in Economics and Mathematical Systems*, pages 31–44. Springer Verlag, Berlin, 1998.
- [12] M. Ehrgott and M. Burjony. Radiation therapy planning by multicriteria optimization. In *Proceedings of the 36th Annual Conference of the Operational Research Society of New Zealand*, pages 244–253, 2001.
- [13] M. Ehrgott and R. Johnston. Optimisation of beam directions in intensity modulated radiation therapy planning. *OR Spectrum*, 25(2):251–264, 2003.
- [14] A. Gersho and R. Gray. *Vector Quantization and Signal Compression*. Kluwer Academic Publishers, Boston, MA, 1991.

- [15] O.C.L. Haas, K.J. Burnham, and J.A. Mills. Optimization of beam orientation in radiotherapy using planar geometry. *Physics in Medicine and Biology*, 43:2179–2193, 198.
- [16] H.W. Hamacher and K.-H. Küfer. Inverse radiation therapy planing – A multiple objective optimization approach. *Discrete Applied Mathematics*, 118(1-2):145–161, 2002.
- [17] A. Holder. Partitioning Multiple Objective Solutions with Applications in Radiotherapy Design. Technical Report 54, Trinity University Mathematics, 2001.
- [18] A. Holder. Designing radiotherapy plans with elastic constraints and interior point methods. *Health Care and Management Science*, 6(1):5–16, 2003.
- [19] A. Holder. Radiotherapy treatment design and linear programming. In M. Brandeau, F. Sainfort, and W. Pierskalls, editors, *Operations Research and Health Care: A Handbook of Methods and Applications*, chapter 29. Kluwer Academic Publishers, 2004.
- [20] A. Holder and B. Salter. A tutorial on radiation oncology and optimization. In H. Greenberg, editor, *Emerging Methodologies and Applications in Operations Research*, chapter 4. Kluwer Academic Press, Boston, MA, 2004.
- [21] Q. Hpu, J. Wang, Y. Chen, and J.M. Galvin. Beam orientation optimization for IMRT by a hybrid method of the genetic algorithm and the simulated dynamics. *Medical Physics*, 30:2360–2367, 2003.
- [22] K.-H. Küfer and H.W. Hamacher. A multicriteria optimization approach for inverse radiotherapy planning. In W. Schlegel and T. Bortfeld, editors, *XIIIth International Conference on the Use of Computers in Radiation Therapy*, pages 26–28, Heidelberg, Germany, 2000. Springer Verlag, Berlin.
- [23] K.-H. Küfer, A. Scherrer, M. Monz, F. Alonso, H. Trinkaus, T. Bortfeld, and C. Thieke. Intensity-modulated radiotherapy – A large scale multi-criteria programming problem. *OR Spectrum*, 25:223–249, 2003.
- [24] E.K. Lee, T. Fox, and I. Crocker. Integer programming applied to intensity-modulated radiation therapy treatment planning. *Annals of Operations Research*, 119:165–181, 2003.
- [25] J. Legras, B. Legras, and J. Lambert. Software for linear and non-linear optimization in external radiotherapy. *Computer Programs in Biomedicine*, 15:233–242, 1982.
- [26] J. Lim, M. Ferris, D. Shepard, S. Wright, and M. Earl. An optimization framework for conformal radiation treatment planning. Technical Report 02-10, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin, 2002. To appear in *INFORMS Journal on Computing*.
- [27] S. Lloyd. Least squares quantization in PCM. *IEEE Transaction on Information Theory*, 28:127–135, 1982.
- [28] E. Marchi and J.A. Oviedo. Lexicographic optimality in the multiple objective linear programming: The nucleolar solution. *European Journal of Operational Research*, 57:355–359, 1992.

- [29] G. Meedt, M. Alber, and F. Nüsslin. Non-coplanar beam direction optimization for intensity-modulated radiotherapy. *Physics in Medicine and Biology*, 48:2999–3019, 2003.
- [30] S. Morrill, I. Rosen, R. Lane, and J. Belli. The influence of dose constraint point placement on optimized radiation therapy treatment planning. *International Journal of Radiation Oncology, Biology, Physics*, 19:129–141, 1990.
- [31] P. Nizin, A. Kania, and K. Ayyangar. Basic concepts of corvus dose model. *Medical Dosimetry*, 26(1):65–69, 2001.
- [32] P. Nizin and R. Mooij. An approximation of central-axis absorbed dose in narrow photon beams. *Medical Physics*, 24(11), 1997.
- [33] F. Preciado-Walters, R. Rardin, M. Langer, and V. Thai. A coupled column generation, mixed integer approach to optimal planning of intensity modulated radiation therapy for cancer. *Mathematical Programming, Series B*, 101:319–338, 2004.
- [34] A. Pugachev. *Beam orientation optimization in intensity-modulated radiation therapy treatment planning*. PhD thesis, Stanford University, USA, 2002.
- [35] A. Pugachev, G.L. Li, A.L. Boyer, S.L. Hancock, Q.T. Le, S.S. Donaldson, and L. Xing. Role of beam orientation optimization in intensity-modulated radiation therapy. *International Journal of Radiation Oncology, Biology, Physics*, 50(2):551–560, 2001.
- [36] A. Pugachev and L. Xing. Computer-assisted selection of coplanar beam orientations in intensity-modulated radiation therapy. *Physics in Medicine and Biology*, 46:2467–2476, 2001.
- [37] A. Pugachev and L. Xing. Pseudo beam’s-eye-view as applied to beam orientation selection in intensity-modulated radiation therapy. *International Journal of Radiation Oncology, Biology, Physics*, 51(5):1361–1370, 2001.
- [38] A. Pugachev and L. Xing. Incorporating prior knowledge into beam orientation optimization in IMRT. *International Journal of Radiation Oncology, Biology, Physics*, 54(5):1565–1574, 2002.
- [39] I. Rosen, R. Lane, S. Morrill, and J. Belli. Treatment plan optimization using linear programming. *Medical Physics*, 18(2):141–152, 1991.
- [40] C.G. Rowbottom, C.M. Nutting, and S. Webb. Beam-orientation optimization of intensity-modulated radiotherapy: Clinical application to parotid gland tumours. *Radiotherapy and Oncology*, 59(2):169–177, 2001.
- [41] C.G. Rowbottom, S. Webb, and M. Oldham. Improvements in prostate radiotherapy from the customization of beam directions. *Medical Physics*, 25(7):1171–1179, 1998.
- [42] E. Schreibmann, M. Lahanas, L. Xing, and D. Baltas. Multi-objective evolutionary optimization of the number of beams, their orientations and weights for intensity-modulated radiation therapy. *Physics in Medicine and Biology*, 49:747–770, 2004.
- [43] D. Shepard, M. Ferris, G. Olivera, and T. Mackie. Optimizing the delivery of radiation therapy to cancer patients. *SIAM Review*, 41(4):721–744, 1999.

- [44] S. Söderström and A. Brahme. Selection of beam orientations in radiation therapy using entropy and Fourier transform measures. *Physics in Medicine and Biology*, 37(4):911–924, 1992.
- [45] D. Sonderman and P. Abrahamson. Radiotherapy treatment design using mathematical programming models. *Operations Research*, 33(4):705–725, 1985.
- [46] J. Stein, R. Mohan, X.H. Wang, T. Bortfeld, Q. Wu, K. Preiser, C.C. Ling, and W. Schlegel. Number and orientations of beams in intensity-modulated radiation treatments. *Medical Physics*, 24(2):149–160, 1997.
- [47] A. Sultan. *Inclusion-exclusion type evaluation criteria for beam arrangements in radiation therapy planning*. Diploma thesis, University of Kaiserslautern, Germany, 2002.
- [48] W. Ulmer and D. Harder. A triple Gaussian pencil beam model for photon beam treatment planning. *Zeitschrift für Medizinische Physik*, 5:25–30, 1995.
- [49] C. Wang, J. Dai, and Y. Hu. Optimization of beam orientations and beam weights for conformal radiotherapy using mixed integer programming. *Physics in Medicine and Biology*, 48:4065–4076, 2003.
- [50] I. Winz. A decision support system for radiation therapy treatment planning. Master’s thesis, The University of Auckland, Department of Engineering Science, Auckland, New Zealand, 2004.
- [51] E. Woudstra and P.R.M. Storchi. Constrained treatment planning using sequential beam selection. *Physics in Medicine and Biology*, 45:2133–2149, 2000.