

Fluence Map Optimization in IMRT Cancer Treatment Planning and A Geometric Approach

Yin Zhang* and Michael Merritt†

July, 2004 (Revised Oct., 2004 and May 2005)

Abstract

Intensity-modulated radiation therapy (IMRT) is a state-of-the-art technique for administering radiation to cancer patients. The goal of a treatment is to deliver a prescribed amount of radiation to the tumor, while limiting the amount absorbed by the surrounding healthy and critical organs. Planning an IMRT treatment requires determining fluence maps, each consisting of hundreds or more beamlet intensities. Since it is difficult or impossible to deliver a sufficient dose to a tumor without irradiating nearby critical organs, radiation oncologists have developed guidelines to allow tradeoffs by introducing so-called dose-volume constraints (DVCs), which specify a given percentage of volume for each critical organ that can be sacrificed if necessary. Such constraints, however, are of combinatorial nature and pose significant challenges to the fluence map optimization problem.

The purpose of this paper is two-fold. We try to introduce the IMRT fluence map optimization problem to a broad optimization audience, with the hope of attracting more interests in this promising application area. We also propose a geometric approach to the fluence map optimization problem. Contrary to the traditional view, we treat dose distributions as primary independent variables and beamlet intensities as secondary. We present theoretical and preliminary computational results for the proposed approach, while omitting excessive technical details to maintain an expository nature of the paper.

Key words: Cancer radiation therapy, Optimal treatment planning, Fluence map optimization, A geometric Approach.

*This author's work was supported in part by NSF Grants DMS-0405831 and DMS-0442065.

†This author's work was supported in part by DOE/LANL Contract 03891-99-23 and NSF Grant DMS-0240058.

Contents

1	Introduction	2
1.1	IMRT	2
1.2	Dose-Volume Constraints	3
2	Fluence Map Optimization	4
2.1	Discretizations	4
2.2	Dose Calculation	5
2.3	Prescriptions	6
2.4	Current Practice and Research	7
3	A Proposed Approach	9
3.1	Prescription and Physical Sets	10
3.2	Optimization Formulations	12
3.3	Alternating Projection Algorithm	13
3.4	Equivalence to a Greedy Algorithm	14
3.5	Convergence to Local Minimum	15
3.6	Preliminary Numerical Results	17
4	Final Remarks	18

1 Introduction

Using radiation to treat cancer requires careful planning. Bombarding malignant tumors with high-energy X-rays can kill cancerous cells (or hinder their growth), but it is usually impossible to deliver a terminal dose without damaging nearby healthy organs in the process. Serious patient complications can occur when the surrounding healthy tissues receive too much of this collateral radiation. On the other hand, sacrificing a modest number of healthy cells may be tolerable since many organs are resilient enough to sustain a certain degree of damage while still providing their anatomical function and can eventually recover. Therefore, research in radiation therapy seeks methods of delivering a sufficient dose to the tumor, while carefully controlling the dose received by neighboring critical organs and other healthy tissues.

1.1 IMRT

Intensity-modulated radiation therapy (IMRT) is a state-of-the-art method which delivers higher doses to tumors and allows more precise conformation than the conventional 3D conformal radiotherapy. The primary delivery tool for IMRT is a linear accelerator that rotates on a gantry around the patient, emitting “modulated” beams of X-rays. This modulation is accomplished by means of a device known as a multileaf collimator (MLC)

which is attached to the accelerator. Its adjustable heavy-metal leaves act as a filter, blocking or allowing radiation through in a precise manner controlled by a computer, in order to tailor the beam shape to the shape of the tumor volume while minimizing exposure of the neighboring structures.

Several mathematical problems arise in order to optimally administer IMRT. Treatment proceeds by rotating the accelerator around the patient and coordinating the leaf movements in the MLC so that the radiation delivered conforms to some desirable dose distribution at each gantry (beam) angle. We will assume in this paper that treatments are administered by fixing the accelerator at a finite number of given gantry angles, rather than emitting radiation while rotating through a continuous arc. We note that determining the number and the values of the gantry angles constitutes a higher-level optimization problem of a combinatorial nature, often called the *beam-angle optimization* problem. Typically, increasing the number of gantry angles would increase the quality and the cost of the treatments.

In addition to knowing the beam angles, one must also know how intense the beams should be at each point (x, y) on the MLC aperture for all gantry angles. These intensity profiles, or fluence maps, are represented by two-dimensional, nonnegative functions $I_a(x, y)$ for $a = 1, 2, \dots, k$, where k is the number of gantry angles in use. The process of determining the functions $I_a(x, y)$ is often called *fluence map optimization*.

Finally, once the fluence maps $I_a(x, y)$ are determined, one must convert these into MLC leaf sequences that attempt to realize them. The longer an MLC leaf is open at a certain position (x, y) , the more dose the tissue along a straight path from that position (plus some surrounding tissue) absorbs. The process of converting fluence maps into the opening and closing movements of leaves is called *leaf-sequencing*. There are many physical and mathematical issues that affect how successful MLC leaf sequences are at approximating the desired fluence maps.

In this paper, we will focus solely on the problem of computing the fluence maps $I_a(x, y)$, such that the tumor, or target, structures receive the prescribed doses and the healthy critical structures receive as little as possible. These conflicting goals are the primary cause of difficulty in fluence map optimization.

1.2 Dose-Volume Constraints

Besides some scattering, the radiation travels primarily in a straight line, so it must typically pass next to or even directly go through critical organs in order to reach and irradiate intended tumor targets. Since the doses that kill most cancers are much larger than those that kill most healthy tissue in the body, even though multiple angles are used in an attempt to focus radiation on the targets, more often than not one has no choice but to sacrifice some healthy tissues. The next sensible objective is to control the volume of the healthy tissues to be sacrificed. Toward this end, in standard practice oncologists prescribe dose-volume constraints (DVCs) that allow a certain percentage of volume in healthy tissues to be sacrificed in order to make sufficient progress in treating the cancer. A typical DVC

has the form, for example, that no more than 30% volume of the right lung can exceed a radiation dose of 20Gy, where “Gy” is the shorthand for “Gray” – the international unit for radiation dose absorption. In addition, oncologists may specify another level of constraint on the same organ, such as no more than 40% volume of the right lung can exceed 10Gy. These dose-volume constraints are natural for oncologists to specify and have become the *de facto* standard way to prescribe radiation therapy treatment in practice.

Clearly, dose-volume constraints provide the much needed flexibility necessary for the escalation of tumor doses. On the other hand, they also introduce a high degree of complexity to the underlying optimization problem. In the above example, for instance, which 30% of the right lung volume should be allowed to absorb more than 20Gy? This brings a combinatorial component to the optimization problem (once the problem is discretized). Mathematically, finding the globally optimal combination of critical organ cells to sacrifice in this way can be an extremely difficult problem.

2 Fluence Map Optimization

In this section, we provide details on the current practice of IMRT treatment planning, as they are relevant to the fluence map optimization.

2.1 Discretizations

To determine the fluence map functions $I_a(x, y)$, we first discretize the MLC aperture for each angle by putting a rectangular grid $\{(x_i, y_j)\}$ on it. As a result, the two-dimensional function $I_a(x, y)$ will be approximated by a set of discrete values $\{I_a(x_i, y_j)\}$. The actual number of these small rectangular elements, or “bixels,” will depend not only on the physical sizes of the MLC device (such as the width of the MLC leaves), but also on the gantry angles and the geometry of the region under treatment. For instance, if a beam emitting from a given grid point is determined not to have a significant intersection with or impact on the region of treatment, then this particular grid point will be omitted from consideration. With this discretization, each MLC aperture is broken into hundreds (or up to thousands) of discrete “bixels” and, correspondingly, each radiation beam is broken into as many discrete “beamlets.” The total number of beamlets in a given fluence map optimization problem is the sum of the beamlets for all the beam angles.

Let n be the total number of beamlets for all beam angles and let us index the bixels linearly. Instead of using the notation $I_a(x_i, y_j)$ for the unknown beamlet intensities, we denote the unknowns by a vector $x \in \mathbb{R}^n$. In addition, since the intensity values are nonnegative, we have $x \in \mathbb{R}_+^n$ where \mathbb{R}_+^n denotes the nonnegative orthant of \mathbb{R}^n .

Moreover, we also need to discretize the “region of interest” or “region of treatment.” This is the three-dimensional volume of the patient’s anatomy containing the target structures and any nearby critical structures that might be adversely affected by the radiation. Similarly, we will break this volume up into small three-dimensional rectangular elements

known as “voxels”, each of which is associated with a point $(x_i, y_j, z_k) \in \mathbb{R}^3$. Let m be the total number of voxels in the region of interest and let us index the voxels linearly. During the treatment, each voxel will absorb a dose of radiation. We denote the dose values absorbed by the voxels in the region of interest by a vector $d \in \mathbb{R}_+^m$. Furthermore, let m_t and m_h be the number of target and healthy voxels, respectively, so that $m_t + m_h = m$. Similarly, we will decompose the dose vector into two sub-vectors d_t and d_h , corresponding to dose values absorbed by the target and healthy voxels, respectively.

2.2 Dose Calculation

The standard IMRT model for dose absorbed at the i -th voxel in the region of interest is

$$d_i = \sum_{j=1}^n a_{ij} x_j, \quad (1)$$

where a_{ij} represents the amount of dose absorbed by the i -th voxel per unit intensity emission from the j -th beamlet. The values a_{ij} for all the voxels and beamlets form a matrix $A \in \mathbb{R}_+^{m \times n}$, known as the “influence matrix” (or kernel matrix). In the matrix notation, the dose calculation formula (1) becomes

$$d = Ax. \quad (2)$$

Figure 1 shows how each element a_{ij} relates to a beamlet emitted from the MLC and a voxel in the discretized region of interest.

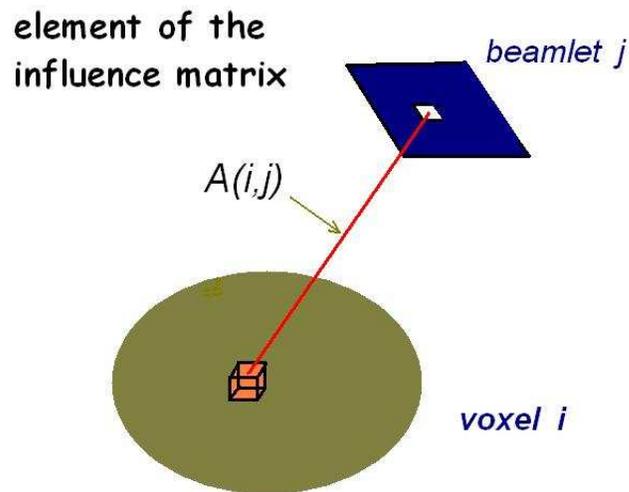


Figure 1: Influence Matrix Element

Assume A has no zero rows or columns. This means, respectively, that all voxels receive some nonzero amount of radiation and every beamlet influences at least one voxel's dose. These conditions can be easily met by pre-processing, if necessary. Typically, $m \gg n$ with m on the order of 10^5 or larger and n of 10^3 up to 10^4 . Note the entries a_{ij} are necessarily nonnegative. In fact, depending on how much scattering is included, the influence matrix A can be very sparse or fairly dense.

The dose calculation model using the influence matrix A can be considered as a first-order approximation. Radiation absorption as a function of the beamlet intensities can be modeled with linear Boltzmann transport equations [12]. Solving these equations can be complicated and computationally expensive, so many different approximation methods have been proposed for computing A . Monte Carlo sampling techniques are, for example, among the more popular methods because of their accuracy. However, Monte Carlo methods are also very slow and expensive. Some commercial planning systems include dose calculation engines with several levels of accuracy. In this way, dose calculations in early iterations, which are usually not required to be highly accurate, can be made less expensive [26], while more accurate (and more expensive) schemes can be used in later iterations.

Clearly, dose calculation is still an important research area on its own right. While acknowledging its importance, for now we will assume that a constant influence matrix A is provided to us *a priori*, and we will use it throughout our optimization process.

2.3 Prescriptions

A sample prescription for a lung cancer case is given below. As one can see, the prescription consists of a set of dose-volume constraints. Currently, this is the standard practice in prescribing radiation therapy treatments.

For each structure, tumorous or healthy, there is at least dose-volume constraint, specified by a percentage number on the left and a threshold value for dose on the right.

The first four lines of the prescription are for a tumor and a so-called extended area around the tumor that is introduced to account for uncertainties about the boundary of the tumor. The first two lines state that the target tumor dose should be higher than a threshold value of 63Gy, although 5% of the target volume may be lower than that. On the other hand, the target dose should be below 72Gy except for a portion of 1% of volume. We can similarly interpret the specifications on the other structures.

It should be pointed out that the dose-volume constraints for the target structures are very different in nature than those for the healthy structures. Obviously, they are not there for the purpose of sacrificing a part for the good of the whole. They are there because it is too difficult or impossible to achieve a uniform target dose, so some imperfections are allowed. For example, we may very well regard the first two lines of the prescription as a perturbation to a target dose specification of 65Gy. In the rest of the paper, that is precisely the approach we will take; namely, we will assume that a single target dose value will be given for each target structure, knowing that it is unlikely to be exactly satisfiable. This

```

-----
>= 95% of Tumor      receives >= 63 Gy
<= 1% of Tumor       receives >= 72 Gy
>= 95% of Ext_Tumor  receives >= 60 Gy
<= 1% of Ext_Tumor   receives >= 70 Gy
-----
<= 1% of Cord        receives >= 43 Gy
<= 15% of Heart       receives >= 30 Gy
<= 20% of Esophagus   receives >= 10 Gy
<= 2% of Lt_Lung      receives >= 20 Gy
<= 8% of Lt_Lung      receives >= 10 Gy
<= 30% of Rt_Lung     receives >= 19 Gy
<= 40% of Rt_Lung     receives >= 10 Gy
<= 50% of Norm_Tissue receives >= 54 Gy
-----

```

Figure 2: A sample prescription

assumption will significantly simplify our presentation, even though our formulation, to be introduced in the next section, can be extended to deal with dose-volume constraints for target structures.

2.4 Current Practice and Research

The IMRT fluence map optimization problem has been extensively studied for a number of years, mostly by medical physicists but more recently also by operations researchers and applied mathematicians. A survey on this subject from a mathematical viewpoint can be found in [20]. We now give a brief overview on the current practice and on-going research in IMRT fluence map optimization. Intended for a non-expert audience, this overview is by no means comprehensive. For a collection of recent survey papers on many aspects of IMRT treatment planning, including one on mathematical optimization by Censor [4], we refer the reader to the book [17] and the references thereof.

The fluence map optimization problem can be viewed as an inverse problem where one designates a desirable dose distribution and attempts to determine a beamlet intensity vector that best realizes the given distribution. There are different ways to formulate this problem into optimization problems using different objective functions, some biological and some physical. Biological models attempt to represent statistical knowledge of various biological responses, such as tumor control probability (see [3], for example). At present, however, the predominant formulation is the “weighted least squares” model as described below, which is being used in most commercial IMRT systems on the market.

If an objective function is associated with each anatomical structure, then this problem

can be naturally viewed as a multi-objective optimization problem (for example, see [7, 10]). However, due to the difficulties in directly solving multi-objective optimization problems, the prevalent approach in IMRT fluence map optimization is to use a weighted least squares fitting strategy. Although many variations exist, a typical form of the weighted least squares formulation is the following. For each voxel i , one tries to fit the calculated dose value d_i to some “desirable” value b_i . For a target voxel, this “desirable” value is just the prescribed dose value for the tumor structure to which the voxel belongs. For a healthy voxel (for which there is really no “desirable” dose other than zero), it is usually set to the threshold value of the dose-volume constraint, though sometimes adaptive values are used. If a calculated dose for a healthy voxel is less than its “desirable” value, then the corresponding error term is set to zero. This way, only those doses higher than their “desirable” values are penalized.

Then to each target and critical structure, one attaches a weight parameter that represents the relative priority of fitting its calculated doses to the desired one. In fact, different tradeoffs between structures can be made by adjusting these weights. To illustrate, suppose there are four structures S_j , $j = 0, 1, 2, 3$, each consisting of a set of voxels, where S_0 is a target structure and the other three are healthy ones. Then the objective function in the weighted least squares formulation takes the form

$$f(x) = \sum_{j=0}^3 w_j f_j(d(x)), \quad (3)$$

where $d(x) = Ax$ is the calculated dose vector corresponding to a beamlet intensity vector $x \geq 0$ (see (2)), w_j are the weights,

$$f_0(d) = \sum_{i \in S_0} (d_i - b_i)^2, \quad f_j(d) = \sum_{i \in S_j} \max(0, d_i - b_i)^2, \quad j = 1, 2, 3,$$

and b is the vector of “desirable values” for all voxels. In this case, the resulting error function $f(x)$ in (3) is a convex, piece-wise quadratic function of x . One seeks to minimize $f(x)$ subject to the nonnegativity of the beamlet intensities in x . Obviously, solutions to this weighted least squares problem vary with the *a priori* choice of the weights.

The weighted least squares model in (3) does not directly enforce the dose-volume constraints in a given prescription, which represent the most challenging aspect of the fluence map optimization problem. One approach to enforcing the dose-volume constraints is to try out different choices of the weights while using the dose-volume constraints as evaluation criteria for solutions. The fundamental difficulty in this approach is that there does not seem to exist any transparent relationships between weights and prescriptions. Hence, weight selection basically reduces to a trial-and-error process, which too often becomes overly time-consuming in terms of both human and computer times.

In addition to manipulating the weights, some formulations (e.g. [23]) add penalty terms to the weighted least squares objective to “encourage,” but not impose, dose-volume constraint feasibility. These penalty terms are inevitably non-convex, thus introducing the complexity of having to deal with local minima. To address this problem, some systems

include the option of using stochastic global optimization techniques such as simulated annealing and genetic algorithms to help escape from unsatisfactory local minima.

Since weighted least squares problems are usually quite large, gradient-type algorithms are often the methods of choice. Some implementations also employ conjugate gradient [21] or secant methods [10]. The nonnegativity of beamlet intensities are enforced either by projection [23] or some other means [10].

With all its shortcomings, the conceptual and algorithmic simplicity of the weighted least squares approach is still attractive to the practitioners. Indeed, this current state of IMRT treatment planning does represent a remarkable progress in cancer radiotherapy. On the other hand, many issues remain, ample room for improvement exists, and intensive research activities are still on-going. One of the research directions in this field is so-called *weight optimization* (see [28], for example), aimed at automating the weight selection process.

The IMRT fluence map optimization problem has attracted considerable attention from researchers in mathematical programming community who tend to formulate the problem into linear or mixed-integer linear programs (see [13, 8, 19, 18] for a sample of some recent works). Linear programming techniques for radiation therapy have been proposed and studied since the early days [1] and have also been considered for treating dose-volume constraints [11]. On the other hand, the introduction of mixed-integer programming techniques into radiotherapy planning for treating dose-volume constraints was a more recent event. Many contributions from the mathematical programming community seem encouraging and promising. Their impact on the clinical practice of radiotherapy, even though limited at this point, will hopefully be felt over time.

Finally, we reiterate that the above short overview is by no means comprehensive. Given the vast literature on this subject, many omissions have inevitably occurred, most notably works based on biological objective functions and works connecting the fluence map optimization to the beam angle optimization as well as to the multileaf sequencing.

3 A Proposed Approach

In our view, there are two levels of difficulties in IMRT fluence map optimization, as outlined below. Our focus in this paper will be on the first issue.

1. Given a prescription, one needs to find a beamlet intensity vector so that the calculated dose from it will satisfy the prescription as closely as possible. The difficulty for this problem lies in the fact that dose-volume constraints define a complicated non-convex feasibility set. This leads to a non-convex global optimization problem that is difficult to solve exactly. The traditional weighted least squares approach relies, to a large degree, on a trial-and-error weight-selection process to search for a good plan.
2. Due to variations from patient to patient even for the same kind of cancer, more often than not, oncologists themselves do not know *a priori* a “good and achievable” prescription for a particular patient. A highly desirable prescription could be too

good to be achievable, while an achievable one might not be close enough to the best possible. Development of procedures to assist oncologists in their decision-making is of paramount importance.

In this section we propose a geometric approach that is entirely prescription-driven and does not require any artificial weights. At the same time, we will retain the least squares framework. As such, one can consider our formulation as a “weightless least squares” approach. We consider two sets in the dose space: (i) the physical set consisting of physically realizable dose distributions, and (ii) the prescription set consisting of dose distributions meeting the prescribed tumor doses and satisfying the given dose-volume constraints. In the case where a prescription is given, we seek a suitable dose distribution by successively projecting between these two sets. A crucial observation is that the projection onto the prescription set, which is non-convex, can be properly defined and easily computed. The projection onto the physical set, on the other hand, requires solving a nonnegative least squares problem. We show that this alternating projection algorithm is actually equivalent to a greedy algorithm driven by local sensitivity information readily available in our formulation. Moreover, the availability of such local sensitivity information offers an opportunity to devise greedy algorithms to search for a desirable plan even when a “good and achievable” prescription is unknown.

To keep the expository flavor of the paper, we will not include long and overly technical proofs for some mathematical results stated. A more complete treatment, including extensive numerical results, will be presented in a subsequent paper in preparation.

3.1 Prescription and Physical Sets

We partition the rows of the influence matrix A into two groups: those for target voxels and those for healthy ones; that is,

$$A = \begin{bmatrix} A_t \\ A_h \end{bmatrix}, \quad (4)$$

where A_t is the submatrix consisting of the rows for target voxels and likewise A_h of those for healthy voxels. Recall that $A \in \mathbb{R}_+^{m \times n}$ where $m = m_t + m_h$ is the number of voxels, and n the number of bixels. Thus $A_t \in \mathbb{R}_+^{m_t \times n}$ and $A_h \in \mathbb{R}_+^{m_h \times n}$. With this notation, $A_t x$ gives the calculated doses for the target voxels and $A_h x$ those for the healthy ones. We start by defining two sets in the dose space.

Definition 1 (Prescription Set). *Let $b_t \in \mathbb{R}_+^{m_t}$ be the dose vector for target voxels in a given prescription, and $\mathcal{D}_v \subset \mathbb{R}_+^{m_h}$ be the set of dose vectors for healthy voxels that satisfy all the dose-volume constraints in the given prescription. We call the following set the prescription set*

$$\mathcal{H} = \left\{ \begin{bmatrix} b_t \\ u \end{bmatrix} : u \in \mathcal{D}_v \right\} \subset \mathbb{R}_+^m. \quad (5)$$

Clearly, any dose vector $d \in \mathcal{H}$ precisely meets the prescribed target doses and at the same time satisfies all the dose-volume constraints given in the prescription. If healthy tissue doses calculated from a beamlet intensity vector $x \in \mathbb{R}_+^n$ satisfy the dose-volume constraints, then we must have $A_h x \leq u$; or equivalently, $A_h x + s = u$ for some nonnegative slack variable $s \in \mathbb{R}_+^{m_h}$.

Definition 2 (Physical Set). *Let A be defined as in (4). We call the following set the physical set*

$$\mathcal{K} = \left\{ \begin{bmatrix} A_t x \\ A_h x + s \end{bmatrix} : (x, s) \geq 0 \right\} \subset \mathbb{R}_+^m. \quad (6)$$

Clearly, the physical set contains all the dose vectors that can be physically realized (disregarding the slack variable) under the standard dose calculation model.

Both \mathcal{H} and \mathcal{K} are closed sets in \mathbb{R}_+^m , and \mathcal{K} is a convex cone but \mathcal{H} is non-convex. In fact, \mathcal{D}_v is a non-convex union of convex “boxes.” For example, suppose we have only two healthy tissue voxels in the region of interest and one dose-volume constraint: at least 50% of voxel doses must be less than or equal to 1Gy. Then $\mathcal{D}_v = \{u \in \mathbb{R}_+^2 : u_1 \leq 1\} \cup \{u \in \mathbb{R}_+^2 : u_2 \leq 1\}$; i.e., either u_1 can be greater than one or u_2 , but not both. Clearly, this is the L-shaped (hence non-convex) region in the first quadrant along the two coordinate axes. Figure 3 shows the relation of \mathcal{D}_v to \mathcal{H} and \mathcal{K} for this case when there is one target voxel. Note the L-shaped region is elevated to a height corresponding to a given target dose value b_t . In this figure, the two sets \mathcal{H} and \mathcal{K} do not intersect.

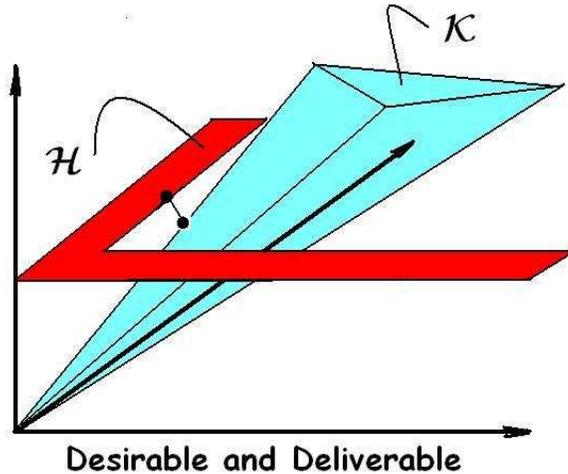


Figure 3: Prescription set \mathcal{H} and physical set \mathcal{K} in dose space \mathbb{R}_+^3

It is easy to imagine that with more voxels and more dose-volume constraints, the complexity of the geometry for \mathcal{D}_v grows quickly out of hand. However, \mathcal{D}_v always enjoys

a very nice geometric property. That is, despite its non-convexity, \mathcal{D}_v permits a trivial projection onto it once the issue of non-uniqueness is resolved (see below). For example, suppose that $\mathcal{D}_v \subset \mathbb{R}_+^{10}$ specifies only one dose-volume constraint: at least 70% of the voxels must have doses of no more than 5Gy. Then

$$\text{Proj}_{\mathcal{D}_v}((1, 2, 3, 4, 5, 6, 7, 8, 9, 10)^T) = (1, 2, 3, 4, 5, 5, 5, 8, 9, 10)^T.$$

where $\text{Proj}_{\mathcal{D}_v}$ is the projection onto \mathcal{D}_v . That is, we set the smallest two numbers greater than 5 equal to 5. Clearly, this is the closest point in \mathcal{D}_v as it affects the least change on the original point in \mathbb{R}_+^{10} . Since \mathcal{D}_v is non-convex, such a projection will not always be unique, but this issue can be resolved by setting some priority rules. It is not difficult to see that dose-volume constraints (DVCs) for multiple structures, and multi-level DVCs for the same structure, can be treated in a similar fashion.

Moreover, it is worth noting that projecting a point $d \in \mathbb{R}_+^m$ onto \mathcal{H} is tantamount to setting the first m_t components of d (the target voxel dose values) to b_t and projecting the last m_h components of d (the healthy voxel dose values) onto \mathcal{D}_v . On the other hand, projecting onto \mathcal{K} is substantially more difficult and will be discussed next.

3.2 Optimization Formulations

Given a prescription, ideally we would like to find $x \in \mathbb{R}_+^n, s \in \mathbb{R}_+^{m_h}$ and $u \in \mathcal{D}_v$ such that

$$A_t x = b_t, \quad A_h x + s = u,$$

but this system is generally over-determined and does not permit a solution. To see this, it suffices to examine the first equation $A_t x = b_t$ which has m_t equations with n unknowns. In practice, there are usually more target voxels than the total number of bixels, i.e., $m_t > n$. The reality of the IMRT fluence map problem is that there may be no physically achievable dose that both satisfies the DVCs and meets the target prescription. That is, $\mathcal{H} \cap \mathcal{K} = \emptyset$; or equivalently, $\text{dist}(\mathcal{H}, \mathcal{K}) > 0$ where $\text{dist}(\cdot, \cdot)$ is the Euclidean distance between two sets. Thus, we are motivated to find a prescription dose vector $d^T = [b_t^T \ u^T], u \in \mathcal{D}_v$, that is closest to the physical set \mathcal{K} (or vice versa). In this view, we have an optimization problem with a variable $u \in \mathcal{D}_v$ (because b_t is fixed):

$$\min_{u \in \mathcal{D}_v} \text{dist} \left(\begin{bmatrix} b_t \\ u \end{bmatrix}, \mathcal{K} \right). \quad (7)$$

The objective in the problem (7) describes the distance from a given prescription dose vector to the physical set \mathcal{K} which can be written as

$$\text{dist} \left(\begin{bmatrix} b_t \\ u \end{bmatrix}, \mathcal{K} \right) = \min_{x, s \geq 0} \left\| \begin{bmatrix} b_t \\ u \end{bmatrix} - \begin{bmatrix} A_t x \\ A_h x + s \end{bmatrix} \right\|,$$

where $\|\cdot\|$ is the Euclidean norm by default (though any other fixed, weighted norms can be used as well). Equivalently, we can replace the norm above by one half times the square of the norm and define the following objective function

$$f(u) = \min_{x, s \geq 0} \frac{1}{2} \|A_t x - b_t\|^2 + \frac{1}{2} \|A_h x + s - u\|^2. \quad (8)$$

Namely, $f(u)$ is itself the optimal value of a linear least squares problem with nonnegativity constraints. Using this notation, we can rewrite the problem (7) into the following equivalent form

$$\min_{u \in \mathcal{D}_v} f(u). \quad (9)$$

It is not difficult to show that $f(u)$ decreases monotonically as u increases. Let $(x(u), s(u))$ be the solution of the optimization problem defined in the right-hand side of (8) for a given $u \in \mathcal{D}_v$. Then under suitable conditions, it can be proved that $f(u)$ is differentiable and

$$\nabla f(u) = -\max(0, A_h x(u) - u) \leq 0, \quad (10)$$

where the maximum is taken component-wise. The derivation of this formula is rather long and we omit it here for the sake of space.

Formulation (9) readily provides the sensitivity of a planning quality measure, $f(u)$ in (8), with respect to the dose upper bounds u . This information can potentially be utilized to search the “prescription space” in the more realistic case where a definitive “optimal” prescription is not available, but a set of guidelines are.

3.3 Alternating Projection Algorithm

The method of alternating (or successive) projections is well known in convex optimization. It is simple and effective, though it can be quite slow. Given two closed convex sets E and F with a nonempty intersection, one can obtain a point in $E \cap F$ by successively projecting points in E and F onto each other. The convergence of this procedure when $E \cap F \neq \emptyset$ was first proven by von Neumann for closed convex sets in Hilbert space satisfying certain properties [22]. In finite dimensional cases, convergence is guaranteed (see [2, 5]). One can further show that if the intersection is empty, the algorithm still converges to a pair of points $(\hat{x}, \hat{y}) \in E \times F$ such that $\|\hat{x} - \hat{y}\| = \text{dist}(E, F) > 0$. For our problem, it is usually the case that $\text{dist}(\mathcal{H}, \mathcal{K}) > 0$.

Successive or simultaneous projection algorithms have been applied to different formulations of the IMRT fluence map optimization problem, see [6, 24, 27] for example, where the sets involved in projections are all convex sets (in some cases convex approximations to non-convex sets). To our best knowledge, projections have not been directly applied to the non-convex DVC feasibility set \mathcal{D}_v defined in (9). The set \mathcal{D}_v in (9) consists of a large number of “branches,” where one or more of the voxels has a dose that exceeds its threshold dose. Obtaining or verifying a global minimum on such a set can be excessively difficult. We will instead seek a local minimum in one of the branches.

We propose to apply an alternating projection algorithm to find a local minimum of (7) by successively projecting iterates in dose space back and forth between the prescription set \mathcal{H} and the physical set \mathcal{K} . Specifically, given $d_0 \in \mathcal{H}$ and for $k = 0, 1, 2, \dots$, do

$$d_{k+1/2} = \text{Proj}_{\mathcal{K}}(d_k), \quad d_{k+1} = \text{Proj}_{\mathcal{H}}(d_{k+1/2}). \quad (11)$$

In this algorithm, the iterates are in the prescription set, while the intermediate iterates are in the physical set corresponding to a sequence of beamlet intensity vectors $\{x_{k+1/2} : k = 0, 1, 2, \dots\}$. As mentioned earlier, the projection onto \mathcal{H} is easy, and the projection onto \mathcal{K} requires solving a nonnegative linear least squares problem as defined in the right-hand of (8).

We emphasize that the starting point $d_0 \in \mathcal{H}$ should be chosen to satisfy threshold values of all the dose-volume constraints; i.e., d_0 should be in the intersection of all the “branches” (or “boxes”). For example, if a dose-volume constraint for a given structure is “no more than 30% of voxels can have dose values greater than or equal to 20Gy,” then we should require that every component of d_0 corresponding to a voxel of that structure to be set to the threshold value 20 (or possibly lower). As the iterations progress, the algorithm will then automatically select voxels where the threshold value of 20 will be exceeded. This way we avoid arbitrarily selecting which “branch” to enter at the outset.

3.4 Equivalence to a Greedy Algorithm

We now consider a gradient projection algorithm directly applied to the problem (9): given $u_0 \in \mathcal{D}_v$,

$$u_{k+1} = \text{Proj}_{\mathcal{D}_v}(u_k - \alpha_k \nabla f(u_k)), \quad k = 0, 1, 2, \dots \quad (12)$$

This is a steepest-descent type algorithm, or a greedy algorithm. At each step, the movement is based on the local sensitivity information – the gradient of $f(u)$. Likewise, we select the initial iterate u_0 to be at or below the threshold values of all the dose-volume constraints, ensuring $u \in \mathcal{D}_v$. Then the algorithm will automatically increase u (recall that $\nabla f(u) \leq 0$) in proportion to the sensitivity of the objective function at the current iterate. Moreover, the projection $\text{Proj}_{\mathcal{D}_v}$ that follows each move will keep the iterate within the feasibility set \mathcal{D}_v . Thus, this algorithm can be considered as a sensitivity-driven greedy algorithm for solving (9).

We note that $f(u)$ is monotone in the direction $-\nabla f(u) \geq 0$. Hence the step length selection in (12) seems not as critical as in general situations. We now show that the constant step length $\alpha_k \equiv 1$ will lead to an algorithm that is equivalent to the alternating projection algorithm (11).

Theorem 1. *Let $\{d_k\}$ and $\{u_k\}$ be generated by algorithms (11) and (12), respectively, where $d_0^T = [b_t^T \ u_0^T]$ and $\alpha_k \equiv 1$ in (12). Then*

$$d_k = \begin{bmatrix} b_t \\ u_k \end{bmatrix}, \quad k = 1, 2, 3, \dots \quad (13)$$

Proof. Let us drop the iteration subscript k . Define $x(u)$ and $s(u)$ as the solutions associated with the subproblem in the right-hand side of (8). By the definitions of the relevant projections,

$$\text{Proj}_{\mathcal{H}} \left(\text{Proj}_{\mathcal{K}} \begin{pmatrix} b_t \\ u \end{pmatrix} \right) = \text{Proj}_{\mathcal{H}} \begin{pmatrix} A_t x(u) \\ A_h x(u) + s(u) \end{pmatrix} = \begin{pmatrix} b_t \\ \text{Proj}_{\mathcal{D}_v}(A_h x(u) + s(u)) \end{pmatrix}.$$

Therefore, it suffices to show that $u - \nabla f(u) = A_h x(u) + s(u)$. By the gradient formula (10) of $\nabla f(u)$,

$$u - \nabla f(u) = u + \max(0, A_h x(u) - u) = \max(u, A_h x(u)).$$

So, it remains to show that $A_h x(u) + s(u) = \max(u, A_h x(u))$ for all $u \in \mathcal{D}_v$.

In the following, we use subscripts to denote components of vectors. If $[A_h x(u)]_i \leq u_i$, then necessarily the slack variable $s(u)_i \geq 0$ must satisfy

$$[A_h x(u) + s(u)]_i = u_i = \max(u_i, [A_h x(u)]_i).$$

On the other hand, if $[A_h x(u)]_i > u_i$, then necessarily the slack variable $s(u)_i = 0$ and

$$[A_h x(u) + s(u)]_i = [A_h x(u)]_i = \max(u_i, [A_h x(u)]_i).$$

This completes the proof. □

The equivalence between these two algorithms allows us to view the problem geometrically and apply the alternating projection algorithm (11) with the confidence that locally, reasonable choices are being made as to which dose bounds to relax to take advantage of the flexibility in the dose-volume constraints.

3.5 Convergence to Local Minimum

Since the prescription set \mathcal{H} is non-convex, the classic convergence theory for alternating projection algorithm is not directly applicable. In our limited computational experience, the algorithm has never failed to converge so far. We observe that despite \mathcal{H} being non-convex, it is the union of finitely many (simple) convex sets. This “local convexity” of \mathcal{H} seems to likely allow a modified convergence proof for the alternating projection algorithm in our case, which remains to be a further research topic.

In the meantime, if we introduce a minor modification to the alternating projection algorithm, then the convergence of the algorithm to a local minimum will be guaranteed. For simplicity, let us assume that there is only one healthy structure with a single dose-volume constraint that defines the feasibility set \mathcal{D}_v :

“No more than P -percent of the healthy voxels can receive doses exceeding a given threshold value $\gamma > 0$.”

Let us work with the simpler algorithmic form (12). Suppose that at iteration k , the dose vector $u_k \in \mathcal{D}_v$ is such that $(\tau_k P)$ -percent of the voxels have already exceeded the threshold value γ for some $\tau_k \in [0, 1]$. Then we define \mathcal{D}_v^k to be the set of dose vectors for the healthy structure that satisfy the following dose-volume constraint:

“No more than $(1 - \tau_k)P$ -percent of the healthy voxels corresponding to $[u_k]_i \leq \gamma$ can receive doses exceeding the threshold value γ .”

In this setting, once a voxel has taken a dose value $[u_k]_i > \gamma$ at some iteration k , it will be allowed to take dose values greater than γ for all the subsequent iterations. Moreover, once $\tau_k = 1$ at some iteration k , then in all subsequent iterations no more dose upper-bounds will be allowed to exceed γ except those already having been allowed.

The modified algorithm will take a projection at iteration k onto the set \mathcal{D}_v^k instead of onto \mathcal{D}_v ; that is,

$$u_{k+1} = \text{Proj}_{\mathcal{D}_v^k}(u_k - \nabla f(u_k)), \quad k = 0, 1, 2, \dots \quad (14)$$

In essence, this algorithm provides a greedy scheme to select a set of healthy voxels that are allowed to receive higher doses. We now state the following convergence result for this algorithm without a proof.

Theorem 2. *Let the iteration sequence $\{u_k\}$ be generated by the algorithm (14) with $u_0 \leq \gamma$. Then $\{u_k\} \subset \mathcal{D}_v$ and satisfies that (i) $u_{k+1} \geq u_k$ component-wise, (ii) $f(u_{k+1}) \leq f(u_k)$, and (iii) $\{u_k\}$ converges to a local minimum u_* of $f(u)$ in \mathcal{D}_v .*

We emphasize that the proposed algorithms in this paper are designed for quickly finding a good local optimum instead of locating a global optimum. A number of studies [25, 14, 9] indicate that the existence of multiple local minima due to dose-volume constraints does not appear to notably affect the quality of treatment plans obtained by the weighted least squares approach. A plausible interpretation of this phenomenon is that there exist many easily reachable local minima with function values very close to the global minimum value. Given the presence of various errors in mathematical models (such as dose calculation models) and in data measurements, finding a global optimum for the underlying non-convex optimization problem does not seem necessary nor practically meaningful, as long as a good local minimum is found. Of course, it is still important to carefully assess the quality of obtained solutions from a clinical viewpoint.

Let us examine the dose values calculated at the solution u_* . Clearly, $u_* \in \mathcal{D}_v$ and $Ax(u_*) \in \mathcal{K}$ (corresponding to $s = 0$). However, in general one should not expect that $A_h x(u_*) \in \mathcal{D}_v$. That is, the locally optimal physical dose calculated by the algorithm generally does not satisfy the dose-volume constraints, because such constraints are not explicitly imposed in our “weightless least-squares” formulation, just as in weighted least-squares formulations. While this lack of a direct control over the dose-volume constraint satisfaction could be viewed as a potential disadvantage on one hand, it does usually allow fast solution times on the other hand.

3.6 Preliminary Numerical Results

In this section, we demonstrate the potential of our algorithm on some two-dimensional phantom cases. The region of interest is a 101×101 voxel cross section of a simulated treatment area. Each of the three test cases have different geometries for the “tumor” and “critical organ”, or organ at risk (OAR). The simplest involves a C-shaped tumor that has grown around a small OAR. More challenging is a small OAR completely surrounded by an “O”-shaped tumor. In the third case, we add further complication to the “O” configuration by having the OAR also include a rectangular region just outside the tumor. The geometries of these cases, as outlined in the left-side pictures of Figures 4-6, are nontrivial and, in our view, sufficient for preliminary proof-of-principle studies.

In all the test cases, we specify the prescribed target dose to be 80Gy for all the tumors, and consider the dose-volume constraint: “at most 30% of the critical organ voxels can have doses greater than 25Gy.” We label as “normal” all the tissue which is neither the tumor nor the organ at risk. Although not as concerning as injury to the critical organ, we would always like to prevent this normal tissue from receiving too high a dose. Therefore, we also specify an upper bound of 75Gy for the normal tissue, equivalent to a dose-volume constraint: “0% of the normal tissue can have doses greater than 75Gy.” Additionally, each plan uses 9 coplanar beams with dose absorption governed by an influence matrix (i.e., A in (2)) that we have obtained from The University of Wisconsin-Madison Tomotherapy Research Group.

We implemented the algorithm (14) in *Matlab*. To perform the minimization in (8) (projection onto \mathcal{K}) at each iteration, we used an interior-point scaled gradient algorithm [16]. In anticipation that least squares solutions will allow calculated doses to vary both above and below their desired values, to be on the safer side we adjust b_t to be 5% higher than the desired tumor dose 80Gy, and similarly the OAR threshold value to be 15% lower than the desired 25Gy. We stop the algorithm once the relative change from u_k to u_{k+1} becomes less than 1%. In our experiments, we have observed that the algorithm took very few (usually two) iterations to terminate in all the tested cases.

Our computational results are presented in Figures 4-6 corresponding to the three test cases. In each figure, we have included a dose distribution on the left, and a dose-volume histogram (DVH) on the right. The dose distribution indicates the level of calculated radiation intensity (in gray scale) deposited in the region of interest. As can be seen, the calculated doses are well focused on the tumors while more or less sparing the critical organs. The dose-volume histograms show the relationship between a given dose value (in x -axis) and the volume percentage (in y -axis) of an anatomical structure receiving that level of radiation or higher. For instance, in Figure 4 or 5 the point (40, 0.1) on the “Normal” curve means that 10% of the normal tissue has received a radiation 40Gy or higher.

We have previously performed computational experiments on the same set of phantom cases with a weighted least squares (WLS) approach and a successive linear programming (SLP) approach [15]. Given these experiences, drawing some comparison would be useful.

The weighted least squares approach requires trying out multiple sets of weights and solving a non-negative least squares problem for each set of weights. As such, it generally requires considerably more computation than the new approach. In [15], we used an exhaustive search, designed only for cases with one critical organ and one tumor, to find a set of optimal weights. With such optimal weights, the WLS approach produced solutions of a quality similar to that of the new approach. The SLP approach enforces the exact satisfaction of the dose-volume constraints and solves a sequence of linear programs. It obtained slightly better quality solutions than the new approach, but required far more computation than the new approach. In addition, the beamlet intensity distributions generated by the SLP approach are generally less smooth, creating difficulties for the later leaf-sequencing stage. For more details on the WLS and SLP approaches, we refer interested readers to [15]. These preliminary numerical results, as encouraging as they may appear, constitute only a first step towards validating the viability of the proposed approach.

4 Final Remarks

The IMRT fluence map optimization problem arises, along with a few other optimization problems, from the state-of-the-art technologies of radiation therapy for cancer treatment. The problem has been extensively studied by medical physicists, and has also attracted considerable on-going research from the operations research and optimization communities. Currently, the predominant methodology in practice is the “classic” weighted least squares (WLS) approach, which focuses on determining an optimal beamlet intensity vector.

In this paper, we take a different view to treat dose distributions as the primary variables, resulting in a formulation based on the geometry in “dose space.” It is our purpose to retain the popular “least squares” framework, while doing away with the burden of having to select weights in the classic WLS approach. The proposed formulation is free of weights, prescription-driven, sensitivity guided, and still shares basic characteristics of a least-squares approach such as not having a precise control over the dose-volume constraint satisfaction and, at the same time, being much less computationally demanding. It is designed for quickly finding a good locally optimal plan associated with a given prescription. Preliminary computational results indicate that the approach is potentially capable of producing solutions of a quality at least comparable to that obtainable by the classic WLS approach. Encouraged by these proof-of-principle results, we are currently working towards more realistic testings on three-dimensional clinical cases.

The approach presented in this paper is only one of many on-going research efforts in helping optimize IMRT cancer treatment planning. It is hopeful that an active participation of the operations research and optimization communities in this important application field will bring about an advancement to cancer treatment planning.

Acknowledgment

The first author would like to thank his colleagues in the Optimization Collaborative Working Group, sponsored by the National Cancer Institute and the National Science Foundation, and Dr. Helen Liu of M. D. Anderson Cancer Center, from whom he has learned the basics about cancer radiation therapy.

References

- [1] G. K. Bahr, J. G. 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] H. Bauschke and J. Borwein. “On the Convergence of von Neumann’s Alternating Projection Algorithm for Two Sets,” *Set-Valued Analysis*, 1: pp. 185–212 (1993).
- [3] A. Brahme. “Optimized radiation therapy based on radiobiological objectives,” *Sem. in Rad. Oncol.*, Vol. 9, No. 1: pp. 35–47, (1999).
- [4] Y. Censor. “Mathematical optimization for the inverse problem of intensity modulated radiation therapy.” In: J.R. Palta and T.R. Mackie (Editors), *Intensity-Modulated Radiation Therapy: The State of The Art*. American Association of Physicists in Medicine, Medical Physics Monograph No. 29, Medical Physics Publishing, Madison, Wisconsin, USA, 2003, pp. 25-49.
- [5] W. Cheney and A. Goldstein. “Proximity maps for convex sets.” *Proceedings of the AMS*, Vol. 10: pp. 448–450 (1959).
- [6] Cho PS, Lee S, Marks RJ II, Oh S, Sutlief SG and Phillips MH. “Optimization of intensity modulated beams with volume constraints using two methods: cost function minimization and projections onto convex sets”. *Medical Physics* 25:435-443 (1998).
- [7] C. Cotrutz, M. Lahanas, C. Kappas, and D. Baltas. “A multiobjective gradient-based dose optimization algorithm for external beam conformal radiotherapy,” *Phys. Med. Biol.* 46: pp. 2161–2175 (2001).
- [8] A. Holder. “Designing Radiotherapy Plans with Elastic Constraints and Interior Point Methods,” *Health Care and Management Science*, Vol. 6, Num. 1, pp. 5-16 (2003).
- [9] R. Jeraj, C. Wu and T. R. Mackie. “Optimizer convergence and local minima errors and their clinical importance.” *Phys. Med. Biol.* 48: 2809–2827 (2003).
- [10] M. Lahanas, M. Schreibmann, and D. Baltas. “Multiobjective inverse planning for intensity modulated radiotherapy with constraint-free gradient-based optimization algorithms,” *Phys. Med. Biol.*, 48(17): 2843–2871 (2003).

- [11] R. G. Lane, S. M. Morrill, I. I. Rosen, and J. A. Wong. “Dose volume considerations with linear programming optimization.” *Med. Phys.*, 18(6):1201-1210 (1991).
- [12] E. W. Larsen. “Tutorial: The Nature of Transport Calculations Used in Radiation Oncology,” *Transport Theory Statist. Phys.*, 26: pp. 739 (1997).
- [13] E. Lee, T. Fox, and I. Crocker. “Integer Programming Applied to Intensity-Modulated Radiation Treatment Planning Optimization,” *Annals of Operations Research, Optimization in Medicine*, 119: pp. 165–181 (2003).
- [14] J. Llacer, J. Deasy, T. Bortfeld, T. Solberg and C. Promberger. “Absence of multiple local minima effects in intensity modulated optimization with dose-volume constraints.” *Phys. Med. Biol.* 48 (2003) 183210.
- [15] M. Merritt, Y. Zhang, Helen Liu, Xiaodong Zhang, Xiaochun Wang, Lei Dong, Radhe Mohan. “A successive linear programming approach to the IMRT fluence map optimization problem.” Manuscript, 2004.
- [16] M. Merritt and Y. Zhang. “Interior-Point Gradient Method for Large-Scale Totally Nonnegative Least Squares Problems,” To appear in *JOTA*, Vol. 126, No. 1 (2005) pp. 191-202.
- [17] J. R. Palta and T. R. Mackie, eds. “Intensity-Modulated Radiation Therapy: The State of the Art,” Medical Physics Publishing, 2003.
- [18] R. Rardin, M. Langer, F. Preciado-Walters, V. Thai. “A coupled column generation, mixed integer approaches to optimal planning of intensity-modulated Radiation therapy for cancer.” To appear in *Mathematical Programming*, 2004.
- [19] H. Romeijn, R. Ahuja, J. Dempsey, A. Kumar and J. Li. “A novel linear programming approach to fluence map optimization for intensity modulated radiation therapy treatment planning,” *Phys. Med. Biol.*, Vol. 48: pp. 3521–3542 (2003)
- [20] D. Shepard, M. Ferris, G. Olivera, and T. Mackie. “Optimizing the delivery of radiation therapy to cancer patients,” *SIAM Review*, 41: pp. 721–744 (1999).
- [21] S. V. Spirou and C. Chui. “A gradient inverse planning algorithm with dose-volume constraints,” *Med. Phys.*, 25(3): pp. 321–333 (1998).
- [22] J. von Neumann. “The geometry of orthogonal spaces, Functional operators - vol. II.” *Annals of Math. Studies*, no. 22, Princeton University Press, 1950. (This is a reprint of mimeographed lecture notes, first distributed in 1933.)
- [23] Q. Wu and R. Mohan. “Algorithms and functionality of an intensity modulated radiotherapy optimization system,” *Med. Phys.*, 27(4): pp. 701–711 (2000).

- [24] C. Wu, R. Jeraj, W. Lu and T. Mackie. “Fast treatment plan modification with an over-relaxed Cimmino algorithm.” *Med. Phys.* 31:191-200 (2004).
- [25] Q. Wu and R. Mohan. “Multiple local minima in IMRT optimization based on dose-volume criteria.” *Med. Phys.* 29 (2002) 151427.
- [26] Q. Wu, D. Djajaputra, M. Lauterbach, Y. Wu, and R. Mohan. “A fast dose calculation method based on table lookup for IMRT optimization,” *Phys. Med. Biol.*, 48(12): pp. N159–N166 (2003).
- [27] Y. Xiao, D. Michalski, Y. Censor and J. Galvin. “Inherent smoothness of intensity patterns for intensity modulated radiation therapy generated by simultaneous projection algorithms.” *Phys. Med. Biol.* 49 (2004) 3227-3245.
- [28] L. Xing, J. G. Li, S. Donaldson, Q. T. Le, and A. L. Boyer. “Optimization of importance factors in inverse planning,” *Phys. Med. Biol.*, 44(10): pp. 2525–2536 (1999).

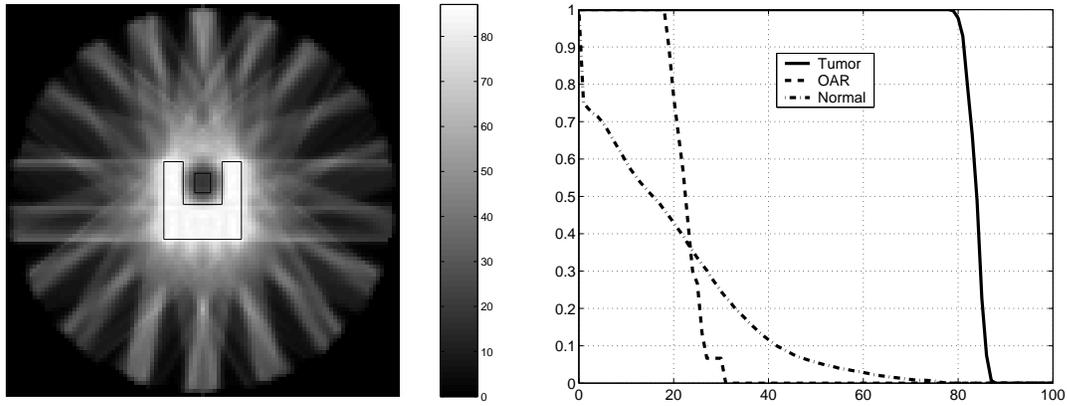


Figure 4: C-Shape Dose Distribution and DVH

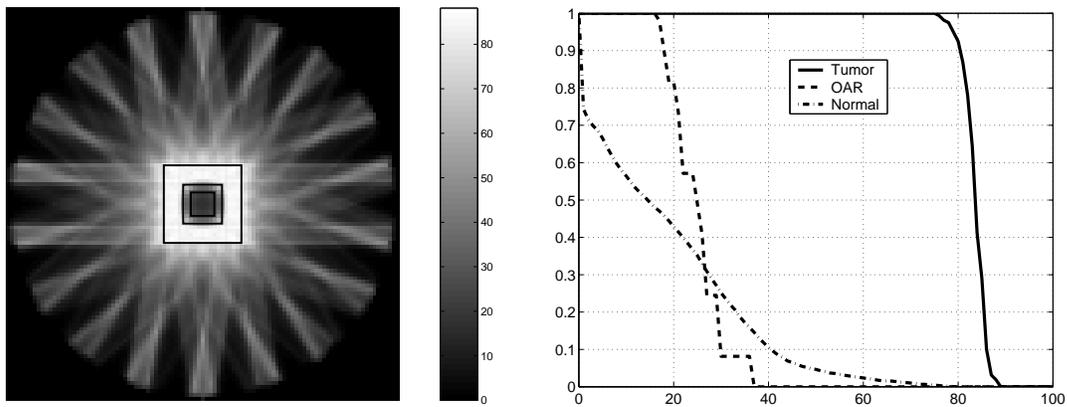


Figure 5: O-Shape Dose Distribution and DVH

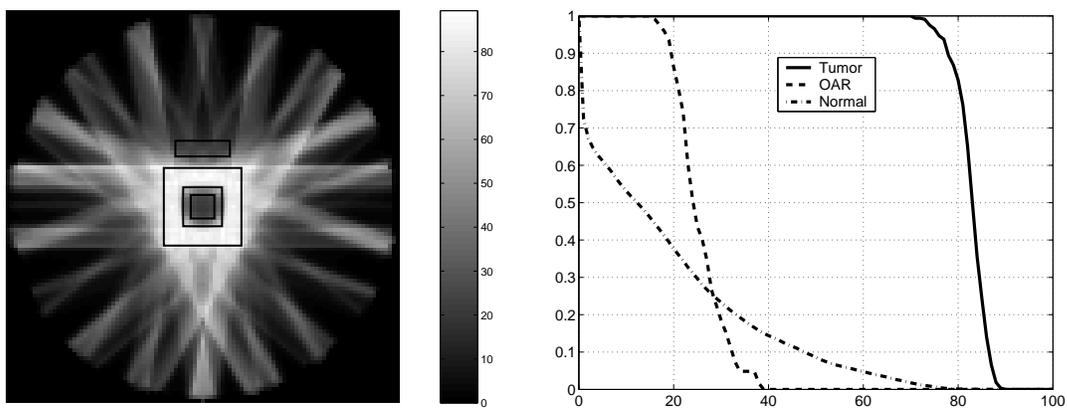


Figure 6: OA-Shape Dose Distribution and DVH