OPTIMIZING THE DELIVERY OF RADIATION THERAPY TO CANCER PATIENTS *

DAVID M. SHEPARD†, MICHAEL C. FERRIS‡, GUSTAVO H. OLIVERA† , AND T. ROCKWELL MACKIE†

Abstract.
In the field of radiation therapy, much of the research is aimed at developing new and innovative techniques for treating cancer patients with radiation. In recent years, new treatment machines have been developed that provide a much greater degree of computer control than was available with the machines of previous generations. One innovation has been the development of an approach called "tomotherapy". Tomotherapy can be defined as computer-controlled rotational radiotherapy delivered using an intensity-modulated fan beam of radiation.

The successful implementation of the new delivery techniques requires the development of a suitable approach for optimizing each patient's treatment plan. One of the challenges is to quantify optimality in radiation therapy. We have tested a variety of objective functions and constraints in pursuit of a formulation that performs well for a wide variety of disease sites. An additional challenge stems from sizable amount of data and the large number of variables that are involved in each optimization. This paper presents several approaches to optimizing treatment plans in radiation therapy, and the advantages and disadvantages of a number of formulations are explored.

1. Introduction. This year in the United States more than 1,200,000 people will be diagnosed with cancer [55]. More than half of these cancer patients will be treated with radiation at some point during the course of their disease [66]. Radiotherapy refers to the use of radiation as a means for treating disease.

A beam of radiation is attenuated as it passes through a patient, and energy is deposited at each location along the path of the beam. The patient's dose is the amount of energy deposited per unit mass. High doses of radiation can kill cells or prevent them from growing and dividing. This is true of both cancer cells and normal cells. Radiotherapy is a useful tool for treating cancer, because the repair mechanism of cancer cells is less efficient than that of normal cells. Thus, normal cells are more likely to fully recover from the effects of radiation. The treatments are carefully planned so that the radiation damage is concentrated within the cancerous region [62]. If all of the tumor cells are killed, the patient is cured.

In recent years, advances in radiation therapy have led to the development of delivery techniques with a high degree of computer control. These techniques offer many new opportunities in the delivery of radiation therapy. Due to the complexity of the treatments, however, an optimization approach is needed in order to develop the best plan for treating each patient.

When designing a patient's treatment plan, the goal is to provide the highest probability of tumor control while maintaining an acceptable probability of radiation-induced complications. Unfortunately, it is very difficult to quantify optimality in radiation therapy. Over years of practice, physicians develop insights into what is "optimal" for a particular type of cancer in a particular location. The physician must

---

*This material is based on research supported by Air Force Office of Scientific Research Grant F49620-98-1-0417, National Science Foundation Grant CCR-9619765 and National Institute of Health Grant CA 48902

† Department of Medical Physics, University of Wisconsin – Madison, 1300 University Ave, Rm 1530, Madison, Wisconsin 53706 (shepard@madrad.radiology.wisc.edu) T.R. Mackie has financial interest in TomoTherapy Incorporated which may participate in the commercialization of tomotherapy.

‡ Computer Sciences Department, University of Wisconsin – Madison, 1210 West Dayton Street, Madison, Wisconsin 53706 (ferris@cs.wisc.edu)

also consider the patient's medical history and the institution's capabilities. Thus, it is important that the optimization approach provides the physician with sufficient flexibility so that he or she can always produce an acceptable treatment plan. In pursuit of a suitable approach for optimizing each patient's treatment plan, we have tested a variety of mathematical programming approaches to this problem. This paper outlines the advantages, the disadvantages, and the quality of the solutions that are produced using each of these approaches. It is hoped that this research will provide insights to both developers of optimization technology and to physicians about the potential capabilities of a carefully devised optimization approach.

This project made use of two mathematical software packages, MATLAB [48] and GAMS [9]. MATLAB was chosen because of its capabilities in both numeric computation and advanced graphic visualization. GAMS, the General Algebraic Modeling System, was chosen because of its strength as a tool for efficiently performing large scale optimizations. This environment provided a means for efficiently testing models that varied widely in term of both the complexity of formulation and the time required to obtain an optimal solution.

A simulated "patient" was represented by an all-water cylinder, and the optimizations were performed for a single transverse cross-section of the patient. This simple setup was used in order to speed up the dose computations [70]. Although actual patient data can be used with our algorithms, our initial investigation into a variety of optimization techniques did not require that level of complexity. Throughout this work, we have used pre-computed photon beamlets as a means of testing a large number of objective functions and constraints [60, 7].

Section 2 of this paper provides an overview of radiation therapy and conventional delivery techniques. Section 3 introduces the more advanced delivery techniques that have been developed in recent years. These delivery techniques provide the user with a large number of "beamlets" of radiation, each of which is assigned a separate beam intensity. One of these techniques, tomotherapy, is described in detail.

Section 4 of this paper presents the results of an investigation into a variety of optimization approaches for determining each patient's treatment plan. Linear programming formulations are discussed in Section 4.1. One linear programming approach is to minimize the total dose delivered to the patient subject to a lower bound on the dose to the tumor and an upper bound on the mean dose to the region at risk. Linear programming formulations such as this can be solved very efficiently, particularly when modeling tools and optimization packages are used. Nonlinear optimization approaches are described in Section 4.2. These formulations are primarily based on a weighted least squares approach. Section 4.3 presents two implementations of partial volume constraints. One implementation is based on the use of a nonlinear ramp function, and the second is based on mixed integer programming. Extensions of our models to conformal avoidance problems are provided in Section 4.4. Section 5 draws some conclusions and outlines research challenges in this area.

2. Conventional Radiotherapy. An appreciation of how a beam of radiation is attenuated in a human body helps one to better understand radiation therapy. Figure 2.1 presents the distribution of dose from a single beam of radiation. More specifically, it shows the dose to the central slice of a 20 cm diameter all-water cylinder irradiated with 2 MeV x-ray photons. This figure can be thought of as a simplistic representation of a transverse cross-section through a patient. The inner circle denotes the outline of the tumor. If we were to treat a patient with a single beam of radiation, it might be possible to kill all of the tumor cells. Unfortunately, however, we would

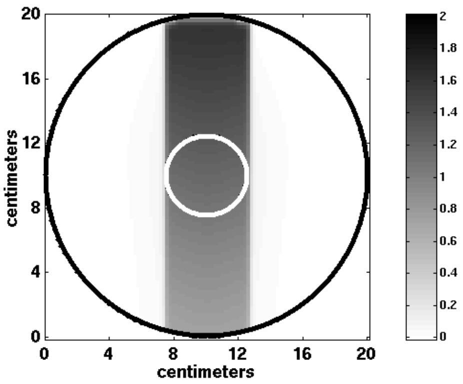

FIG. 2.1. *The dose distribution from a single beam of radiation. The highest dose level appears at the top of the picture, because the beam enters from the top. The outer circle represents the outline of the patient, and the inner circle denotes the outline of the tumor.*

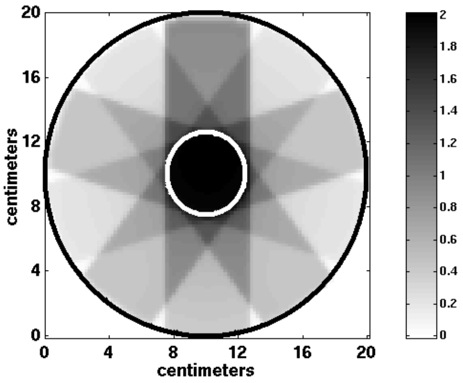

FIG. 2.2. *A dose distribution created with 5 equispaced beams. Note that the dose is concentrated in the tumor where the beams intersect.*

also risk severe damage to any normal tissues located within the path of the beam of radiation. To avoid this, beams are delivered from a number of different angles spaced around the patient.

In Figure 2.2, a dose distribution is shown from the delivery of five equispaced beams of radiation. Note that the dose to the tumor is significantly higher than the dose to the surrounding normal tissues. This is a result of the fact that the beams intersect within the tumor.

The vast majority of patients who receive radiation are treated on a machine called a linear accelerator (see Figure 2.3). The patient lies on the couch of the treatment unit. The x-rays produced by the linear accelerator are filtered so that the beam of x-rays is uniform in intensity as it exits the machine. The x-rays are then shaped into a beam of radiation that is directed toward the patient. The head of the linear accelerator can be rotated around the patient. This provides the ability to treat the patient from a wide range of beams directions.

Radiotherapy patients are typically treated with between 3 and 7 beams of radiation. Through an iterative process of trial and error, a radiation oncologist and

FIG. 2.3. *A clinical linear accelerator used to treat cancer patients*

a physicist determine an acceptable set of beam angles and beam weights that produce an acceptable treatment plan. The beam angles, the beam intensities, and the predicted dose distribution together make up the treatment plan. In conventional treatments, a uniform or linearly-varying beam intensity is delivered from each beam angle.

Figure 2.4 shows three treatment plans developed using conventional uniform beam delivery. In all three cases, seven equispaced beam angles were used. Figure 2.4a shows the optimized dose distribution for a simple square shaped target with a smaller square shaped region at risk located adjacent to the target. Figure 2.4b shows the dose distribution for a target of intermediate complexity. The target is L-shaped with a region at risk located between the arms of the L. Figure 2.4c, shows the dose distribution for a complex target shape. The target is U-shaped with a region at risk placed in the concavity of the U. This type of geometry would be seen clinically if a tumor had grown around a patient's spinal cord or brain stem.

In Figure 2.4, the dashed line represents the 90% isodose curve. This line indicates a dose level that is equal to 90% of the dose prescribed to the target. Ideally, this line should closely match the boundary of the target. The 90% isodose curve is shown as a dashed line a number of figures throughout this paper. It can be seen in Figure 2.4a that uniform beam delivery is able to closely match the 90% isodose line with the border of the square shaped target. Figure 2.4b and 2.4c reveal that uniform beam delivery does not permit the sparing of a region at risk located in the concavity of a more complex target shape. Because of the high dose to the region at risk, this treatment could lead to a high probability of complications.

A simple method for decreasing the probability of normal tissue complications is to block any portion of a beam that passes through a region at risk. This is referred to as segmented beam delivery. Figure 2.5 displays three treatment plans developed using segmented beam delivery. Segmented beam delivery is of particular importance in cases where it is essential to avoid complications. For example, physicians often choose to block out any portion of a beam that passes through a patient's spinal cord. This is done in order to eliminate the possibility of radiation-induced paralysis.

In all three cases shown in Figure 2.5, segmented beam delivery was successful at sparing the region at risk. It should be noted, however, that with the intermediate and complex target shapes, portions of the tumor were significantly underdosed. The areas of underdosage could permit the survival of some of the cancer cells, and this could eventually lead to a recurrence of the disease. More complex treatment plans to alleviate these problems will be described in the sequel.

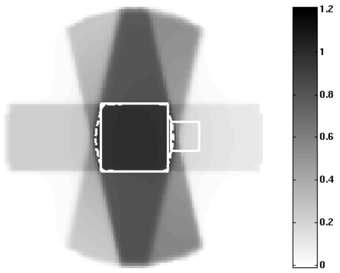

(a) Simple: Square Target

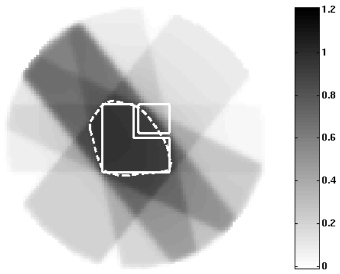

(b) Intermediate: L-Shaped Target

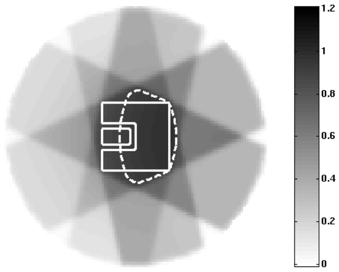

(c) Complex: U-Shaped Target

FIG. 2.4. *Three treatment plans developed using uniform beam delivery. Seven equispaced beam angles were used. Unmodulated fields were used with the field size set to match to target dimensions. The dashed line represents the 90% isodose curve. This line indicates a dose level equal to 90% of the prescribed dose to the target.*

3. IMRT and Tomotherapy. In recent years, the ability to pinpoint the size, shape, and location of a patient's tumor has been dramatically improved by the development of computed tomography (CT) and magnetic resonance imaging (MRI). In radiotherapy, much of the current research is aimed at developing better hardware that will make it possible to take full advantage of these improved imaging capabilities.

Independent investigation performed by Brahme [5] and Cormack [16, 17] demonstrated that the use of nonuniform beam intensities permits the delivery of dose distributions which conform to the target shape while simultaneously sparing any neighboring regions at risks. Recent technological advances have made possible the delivery of intensity modulated radiation therapy (IMRT) [71, 26, 4, 40, 3, 8, 22, 49, 11, 75]. In IMRT, each treatment beam has an individually defined intensity distribution. In other words, one can think of each beam as being divided into a large number of smaller beamlets of radiation. The weights of these beamlets can be optimized so as to produce the most favorable dose distribution [70, 67, 6, 2, 29, 30, 31, 50, 61, 60, 43, 15, 12, 20].

Figure 3.1 shows three dose distributions produced using IMRT with 7 equispaced

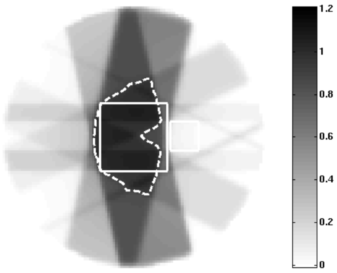

(a) Simple: Square Target

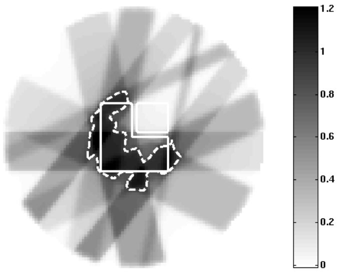

(b) Intermediate: L-Shaped Target

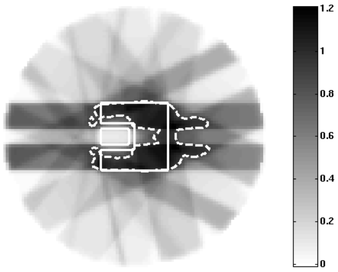

(c) Complex: U-Shaped Target

FIG. 2.5. *Three treatment plans developed using segmented beam delivery. Seven equispaced beam angles were used. Open and unmodulated fields were used, and the portion of the field passing through the region at risk was blocked out.*

beam angles. Approximately 80 beamlets were used in each of these cases. It can be seen that for all three setups, IMRT was able to successfully conform the 90% isodose line to the outline of the target shape. Thus, the region at risk was spared without reducing the probability of tumor control.

The development of IMRT represents an important step in the realization of conformal therapy. The goal of conformal therapy is the delivery of a high dose of radiation to the tumor while simultaneously sparing any neighboring region at risks. Successful delivery of conformal therapy should allow for an escalation of the tumor dose and may enhance local tumor control [41]. The dosimetric advantage of conformal therapy can also be used as a means of providing a reduced probability of normal tissue complications [28].

One approach to the delivery of IMRT is called tomotherapy (see Figure 3.2) [45, 46, 77, 1, 23]. A number of key features distinguish tomotherapy from conventional treatments. One characteristic of tomotherapy is the use of a rotational treatment delivery. The linear accelerator is mounted on a ring-shaped gantry, and it progresses around the patient at a constant speed.

Tomotherapy is also distinguished by the size of the beam that is used to treat

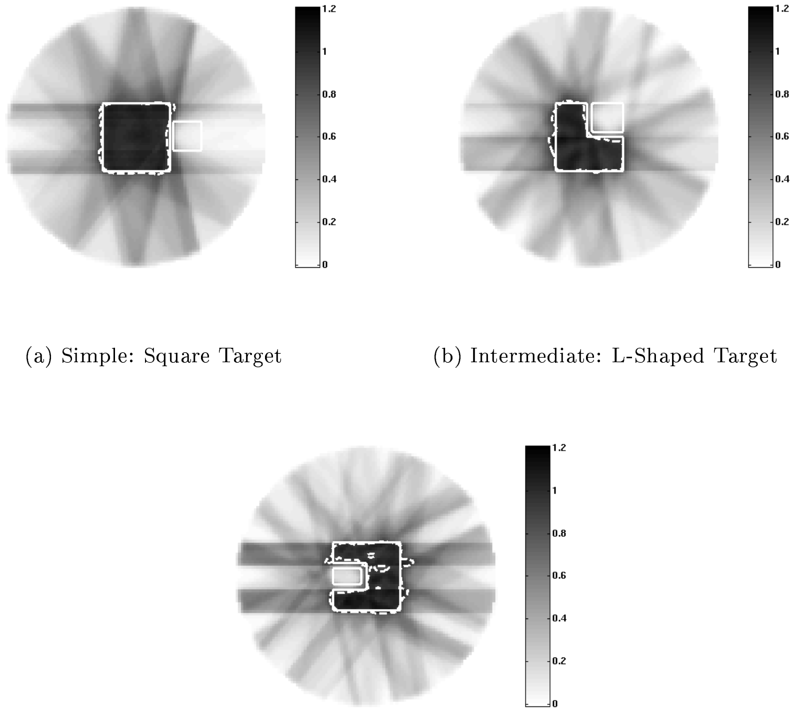

(a) Simple: Square Target

(b) Intermediate: L-Shaped Target

(c) Complex: U-Shaped Target

FIG. 3.1. *Three treatment plans developed using IMRT. Seven equispaced beam angles were used with a beamlet size of 6 mm.*

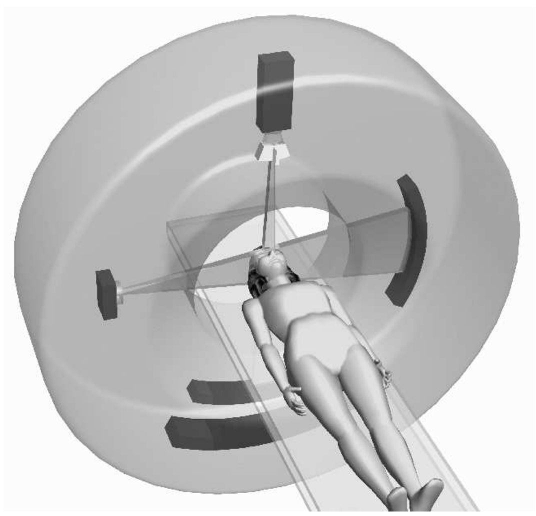

FIG. 3.2. *An artist's rendering of a tomotherapy machine.*

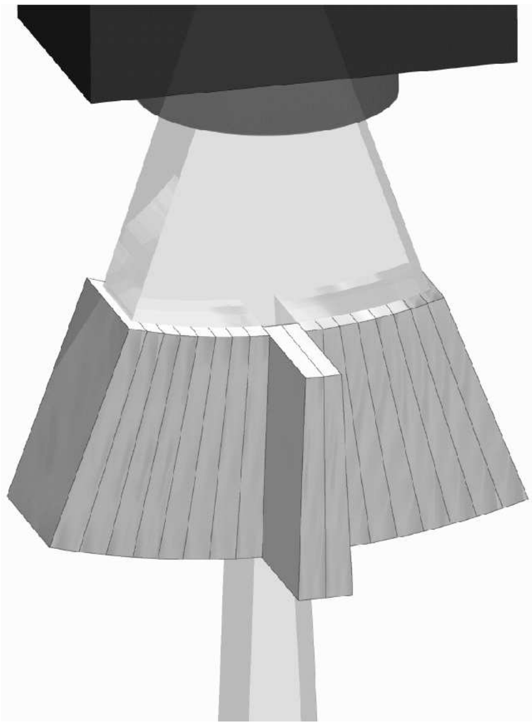

FIG. 3.3. *In the housing of the tomotherapy machine, the fan beam of radiation passes through a multi-leaf collimator. Each leaf of the collimator can be used to block a portion of the beam. The collimator is binary in nature, meaning each leaf is either in the field or out of the field.*

the patient. Tomotherapy does not make use of a broad beam of radiation. Instead, the beam is collimated into a fan beam that delivers radiation to a narrow slice of the patient. Over the course of the treatment the fan beam is able to treat the entire tumor volume. This is accomplished through the use of a moving couch. The patient lies on the couch of the tomotherapy unit, and the patient is slowly advanced through the rotating fan beam of radiation.

Within the housing of the machine, the fan beam passes through a device called a multi-leaf collimator (see Figure 3.3). The collimator has a series of high density leaves, and each leaf can be used to block out a portion of the fan beam. As the linear accelerator progresses around the patient, the leaves move into and out of the beam.

Each full rotation of the linear accelerator can be approximated as a series of discrete steps. The amount of time that an individual leaf is open during a step is a variable that is referred to as a beam weight. Through optimizing the beam weights, it is possible deliver a dose distribution that closely conforms to the target shape.

Tomotherapy promises to provide a number of enhancements in addition to its improved beam delivery capabilities. For example, a tomotherapy machine will be able to produce CT images of the patient [69]. Another feature of tomotherapy is that it will provide precise registration capabilities designed to ensure that the location of the patient on the treatment couch is consistent from day to day [25]. The hardware in the tomotherapy machine will also provide a means of obtaining an accurate reconstruction of the dose delivered the patient. Thus, any variation between the delivered and the predicted dose distributions can be compensated for in

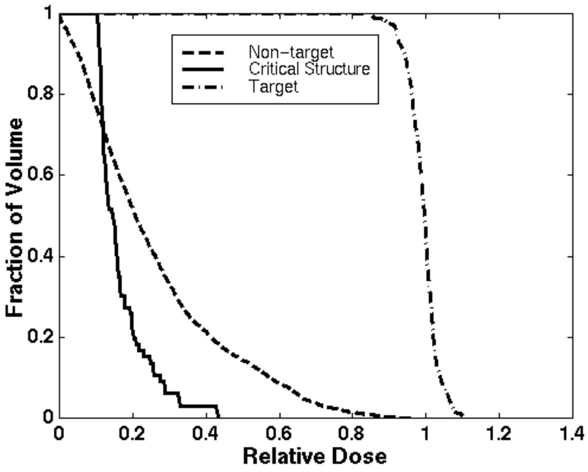

FIG. 4.1. *A dose volume histogram from an optimized treatment plan.*

subsequent treatments [63].

4. Optimization Approaches. IMRT treats each patient with thousands or even tens of thousands of beamlets of radiation. Due to the large number of beams and the range of intensities involved, a computer based optimization routine is required in order to develop each patient's optimal treatment plan. The remainder of this paper is devoted to investigating this issue.

As outlined in Section 1, it is difficult to quantify the goals of radiation therapy. Radiation oncologists often use a cumulative dose volume histogram as a means of determining the quality of a treatment plan (see Figure 4.1). A cumulative dose volume histogram displays the fraction of the patient that receives at least a specified dose level.

The first step in the development of a treatment plan is to define the location of the tumor and the regions at risk (critical structures). Next, the dose to the patient from each individual beamlet is computed. The matrix D represents the total dose delivered to a transverse cross-section of the patient, and it equals the sum over all of the beamlets p of the beam weight multiplied by its corresponding dose matrix D^p:

$$D_{ij} = \sum_{p=1}^{n} w_p D_{ij}^p \tag{4.1}$$

n is the number of beamlets, i and j are the pixel coordinates, and w_p are the beamlet weights. The dose matrices were computed using the convolution/superposition technique. This technique was chosen because of its ability to accurately model complex dose distributions [47, 65].

Originally, the dose optimizations were performed using the routines provided in MATLAB's Optimization Toolbox. MATLAB's algorithms performed reasonably well for small scale optimization problems. As the number of beamlets increased, however, the problems became too large for MATLAB to solve in a reasonable amount of time. Another concern was the fact that physicians may need to make a number of changes in the parameters of an optimization before they are satisfied with the treatment plan.

Therefore, the speed of each technique is an important concern. We thus turned to GAMS in hopes of increasing our optimization capabilities.

The development of an optimal treatment plan is made particularly challenging as the result of the large number of variables and the sizable amount of data that is involved in each optimization. For our simulated treatments, the dose distribution for each beamlet was stored in a relatively dense 100 by 100 matrix. Up to 7000 beamlets were used in an optimization. Currently, each simulated treatment is set up within MATLAB and a simple interface [24] is used to pass the data to GAMS. Selecting or changing the problem formulations is very simple in a modeling language environment. The user can choose between linear, non-linear, and mixed integer models. GAMS also provides the ability to switch between a number of commercial solvers including OSL[32], CPLEX [33], CONOPT [21], and MINOS [56]. All of these solvers are very sophisticated large scale implementations of state-of-the-art algorithms for their respective problem classes. The advantages of easy problem formulation, modification, and automatic differentiation far outweigh the slight overhead imposed by a modeling language. We believe that only after a particular optimization approach becomes prevalent would it make sense to use one of these solvers exclusively.

At the end of each simulated treatment, the optimized beam weights were returned from GAMS to MATLAB using the aforementioned interface. MATLAB was then used to visualize and analyze the results. Many of the figures shown in this paper were produced using standard MATLAB graphics utilities along with the optimization results generated using GAMS. The analysis of the results included a determination of the mean, minimum, and maximum dose to each region. The standard deviation in dose over the target was also determined and the dose values were displayed using a histogram.

The advantages and disadvantages of a number of approaches were tested and the results are described in the following subsections.

4.1. Linear Programming Approaches. The most widely used mathematical programming formulation is undoubtedly the linear program [13, 18, 57]. The linear programming (LP) problem is of the form:

$$\begin{array}{ll} \min & c^T x \\ \text{subject to} & Ax \geq b \\ & \ell \leq x \leq u \end{array}$$

where x is the vector of continuous decision variables, $c^T x$ describes the linear objective function, $Ax \geq b$ represents the set of constraints, and ℓ and u are vectors of lower and upper bounds on the variables. Linear programming codes have successfully solved many large scale optimization problems, using either sophisticated implementations of the simplex method [13, 18, 59, 64] or adaptations of primal-dual interior point methods [76].

There are number of potential linear programming approaches to radiotherapy dose optimization [68]. One LP formulation is to minimize the total integral dose subject to a lower bound on the dose to the tumor. The integral dose is the total dose summed over all of the pixels. Using the notation from (Equation 4.1), this can be reformulated as follows:

$$\begin{array}{lll} \min_w & \sum_{(i,j)} D_{ij} & \\ \text{subject to} & D_{ij} = \sum_{p=1}^{n} w_p D_{ij}^p & \forall (i,j) \\ & \gamma \leq D_{kl} & \forall (k,l) \in \mathcal{T} \\ & w_p \geq 0 & \end{array}$$

$\mathcal{T}$ is the subset of the pixels located in the tumor, n is the number of beamlets, and γ is the lower bound on the dose to the tumor. A non-negativity constraint is placed on the beam weights. The above formulas are designed so that the optimizer will drive the tumor dose down to the specified lower bound. The results from a simulation using this approach are shown in Figure 4.2a.

The objective function can be modified so as to minimize a weighted integral dose:

$$\theta_{\mathcal{T}} \sum_{(i,j)\in\mathcal{T}} D_{ij} + \theta_{\mathcal{R}} \sum_{(i,j)\in\mathcal{R}} D_{ij} + \theta_{\mathcal{N}} \sum_{(i,j)\in\mathcal{N}} D_{ij}$$

$\mathcal{R}$ is the subset of pixels in the region at risk, and $\mathcal{N}$ is the subset of normal tissue pixels (those located outside of both the target and the sensitive structures). In this formulation, a weighting factor is assigned to each region of the patient. $\theta_{\mathcal{T}}$ is the target weight, $\theta_{\mathcal{R}}$ is the region at risk weight, $\theta_{\mathcal{N}}$ is the normal tissue weight. The objective function equals the sum over the entire volume of each pixels' dose multiplied by its weighting factor. By increasing the relative weight assigned to a region at risk, the user can place a greater emphasis upon reducing the dose to that region. By increasing the relative weight assigned to the tumor, the user can place a greater emphasis upon achieving a uniform target dose.

Another LP approach is to place both an upper and a lower bound on the dose to the tumor. In this case, the objective might be to minimize a weighted integral dose summed over all of the non-tumor pixels:

$$\begin{array}{lll} \min_w & \theta_{\mathcal{N}} \sum_{(k,l)\in\mathcal{N}} D_{kl} + \theta_{\mathcal{R}} \sum_{(k,l)\in\mathcal{R}} D_{kl} & \\ \text{subject to} & D_{ij} = \sum_{p=1}^{n} w_p D_{ij}^p & \forall (i,j) \\ & \gamma_L \leq D_{kl} \leq \gamma_U & \forall (k,l) \in \mathcal{T} \\ & w_p \geq 0 & \end{array}$$

γ_L and γ_U indicate the lower and upper bounds on the dose to the target. Figure 4.2b displays a dose distribution produced using this approach with $\theta_{\mathcal{N}}$ equal to one and $\theta_{\mathcal{R}}$ equal to two.

When minimizing the integral dose, the highest weights are assigned to the beamlets that deposit the greatest fraction of their integral dose within the tumor. The results can become unsatisfactory if any of the beam weights are made exceedingly large. Heavily weighted beams produce streaks of high dose through the patient, and this could lead to patient complications. A solution to this problem is to place an upper bound on the ratio between the maximum beam weight and the average beam weight.

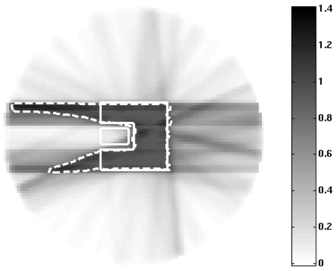

(a) the optimizer sought to minimize the integral dose subject to a lower bound on the tumor dose of 0.9.

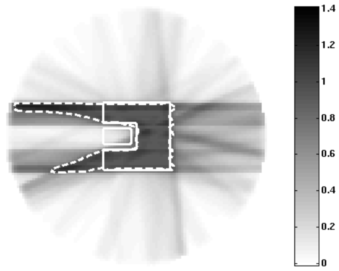

(b) an upper bound of 1.1 and a lower bound of 0.9 were placed on the target dose. The objective function was a weighted integral dose.

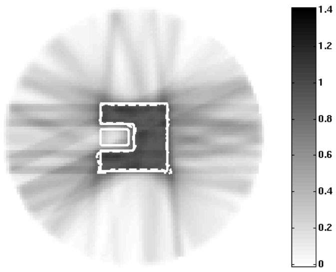

(c) the dose is more evenly distributed outside of the tumor as a result of a constraint that the maximum beam weight could not be more than a factor of 5 greater than the mean beam weight.

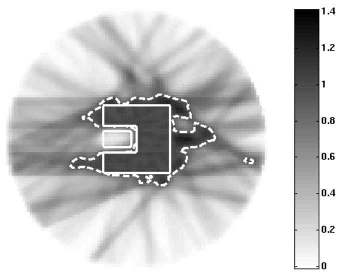

(d) the maximum deviation in dose over the tumor was minimized subject to the constraint that the mean dose to the region at risk had to be kept below 0.5.

FIG. 4.2. *A series of four dose distributions produced using linear programming techniques. In all cases, 15 beam angles were used along with a collimator size of 6 mm; the dashed line denotes the 90% isodose curve. The U-shaped target was cut out of a 5 cm by 5 cm square.*

$$
\begin{array}{lcr}
\min_w & \sum_{(k,l)\in\mathcal{N}} D_{kl} + \theta \sum_{(k,l)\in\mathcal{R}} D_{kl} & \\
\text{subject to} & D_{ij} = \sum_{p=1}^{n} w_p D_{ij}^p & \forall (i,j) \\
 & \gamma_L \le D_{kl} \le \gamma_U & \forall (k,l) \in \mathcal{T} \\
 & w_m \le \frac{\alpha}{n} \sum_{p=1}^{n} w_p & m = 1, 2, \ldots, n \\
 & w_p \ge 0 &
\end{array}
$$

In this case, the maximum beam weight could not exceed α times the mean beam weight. In Figure 4.2c, it can be seen that this approach resulted in a more evenly distributed integral dose as compared to the previous LP optimizations.

A final LP approach is to minimize the maximum deviation from the prescribed target dose, δ, subject to one or more constraints. This can be accomplished using the following model:

$$\begin{array}{lll}
\min_w & \max_{(i,j)\in\mathcal{T}} |D_{ij} - \delta_{ij}| & \\
\text{subject to} & D_{ij} = \sum_{p=1}^{n} w_p D_{ij}^p & \forall (i,j) \\
& \gamma_L \le D_{kl} \le \gamma_U & \forall (k,l) \in \mathcal{T} \\
& \sum_{(r,s)\in\mathcal{R}} D_{rs} \le n_{\mathcal{R}}\beta & \\
& w_m \le \frac{\alpha}{n}\sum_{p=1}^{n} w_p & m = 1, 2, \ldots, n \\
& w_p \ge 0 &
\end{array}$$

In the above equation, $n_{\mathcal{R}}$ is the number of pixels in $\mathcal{R}$. An upper limit of β was placed on the mean dose to the region at risk. Once again, the maximum beam weight was constrained to be less than α times the mean beam weight. The technique used to reformulate a minimax problems as a linear program has been well documented, see [13] for example. Figure 4.2d shows an optimized dose distribution produced with this model. Overall, the dose distribution Figure 4.2c appears to be best LP result for this particular target shape.

The two primary advantages to the use of an LP approach to treatment planning optimization are its speed and the ease of formulation. A disadvantage of the LP formulations is a lack of flexibility. One can only devise a relatively limited number of objective functions and constraints that fall within the linear realm. With any particular LP formulation, it is unlikely that a physician could always achieve an acceptable result.

4.2. Nonlinear Programming Approaches. A concise mathematical description of the nonlinear programming problem (NLP) is as follows:

$$\begin{array}{ll}
\min & f(x) \\
\text{subject to} & g(x) \le 0 \\
& \ell \le x \le u
\end{array}$$

Here, x is a vector of variables that are continuous real numbers, $f(x)$ is the objective function, and $g(x)$ represents the set of constraints. ℓ and u are vectors of lower and upper bounds placed on the variables. With a nonlinear formulation [27], there is a expanded range of possible objective functions and constraints as compared with linear programming.

For many of our simulations, we have used a weighted least squares objective function. In these cases, the optimizer sought to minimize the weighted squared differences between the prescribed and the actual doses summed over all of the pixels. The objective function is shown below:

$$\min_{w\ge 0} \theta_{\mathcal{T}} \sum_{(i,j)\in\mathcal{T}} (D_{ij}(w) - \delta_{ij})^2 + \\ \theta_{\mathcal{R}} \sum_{(i,j)\in\mathcal{R}} (D_{ij}(w) - \delta_{ij})^2 + \theta_{\mathcal{N}} \sum_{(i,j)\in\mathcal{N}} (D_{ij}(w) - \delta_{ij})^2$$

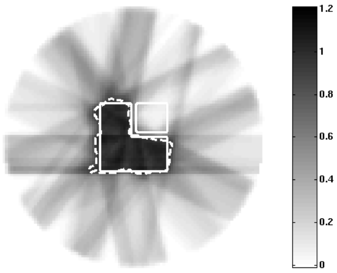

(a) the target, region at risk, and normal tissue were assigned weights of 25,1, and 1 respectively.

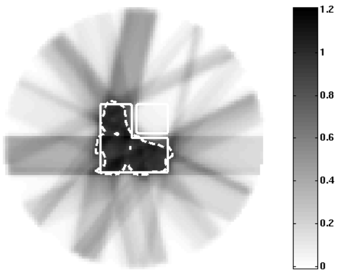

(b) the target, region at risk, and normal tissues were assigned weights of 10,5,2.

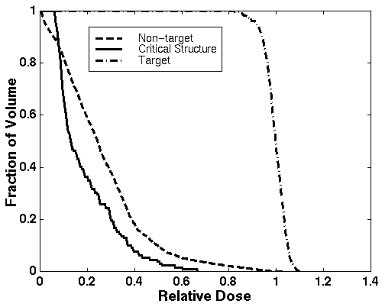

(c) cumulative dose volume histogram for (a).

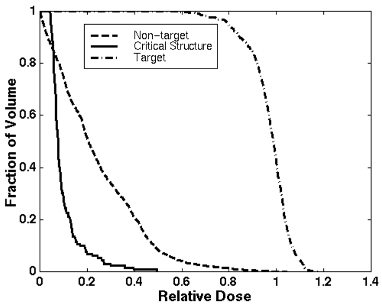

(d) cumulative dose volume histogram for (b).

FIG. 4.3. *The results from two simulations are shown. In both cases, a weighted quadratic objective function was used. Note that the tumor dose is more uniform when the tumor was assigned a higher relative weight. The cumulative dose volume histogram plots the fraction of each structure that receives at least the dose specified by the abscissa.*

As before, $\theta_{\mathcal{T}}$ is the target weight, $\theta_{\mathcal{R}}$ is the region at risk weight, $\theta_{\mathcal{N}}$ is the normal tissue weight. The values of $D_{ij}(w)$ are determined using (4.1). For our simulations, each weighting factor was normalized by the number of pixels in the region to which the weight was assigned. This helps to eliminate difficulties with a large structure dominating the optimization. The matrix δ describes the prescribed dose. Outside of the target, δ is typically set equal to zero. This problem is a bound constrained weighted least squares problem and can therefore be solved by various specialized large scale optimization algorithms [14, 42, 51, 52, 58].

Figure 4.3 displays two optimized dose distributions produced using different weighting schemes. In both cases, the target was prescribed a dose of one. For a particular patient, the best choice of weighting factors is not always intuitive. Thus, in order to obtain an acceptable result, one may need to run a series of optimizations.

Ratio constraints on average and maximum beam intensities and other bounding constraints outlined in Section 4.1 give rise to general constrained nonlinear programs.

It should also be mentioned that nonlinear formulations have been used by a number of researchers in an attempt to use mathematical models as a means of predicting biological response. For example, researchers have devised mathematical formulas that are designed to predict the probability of tumor control [72, 36]. Formulas have also been designed in hopes of predicting the probability of normal tissue complications [38, 10, 37, 44, 34].

Ideally, this type of biological modeling could serve as a very useful tool in radiotherapy optimization. For example, an oncologist could choose to maximize the probability of tumor control subject to a cap placed upon the probability of complications for each sensitive structure. Unfortunately, the current biological models require input parameters that are not known with great certainty. The ability to predict the probability of normal tissue complications is particularly suspect. Because of the degree of uncertainty involved, we have chosen not to include these models in our current study.

From our investigations into nonlinear programming approaches, we have concluded that the primary advantage of using this approach is that it is possible to devise numerous complicated objective functions and constraints. A drawback to this approach is that increasing the complexity of the formulation will often lead to more time consuming optimizations.

Another weakness of the NLP approach is that for general functions, the optimizer can only guarantee that the solution is locally optimal. The importance of multiple local minima in radiotherapy optimization has been investigated elsewhere [19, 35, 60]. Because of the potential to become stuck in a local minima, a number of researchers in radiotherapy optimization have made use of simulated annealing approaches [74, 73, 54]. It is our hope that a deterministic method will provide a more rapid approach for developing a patient's treatment plan. We believe that it is possible to develop convex objective functions that adequately describe the problem. Note that the least squares formulation presented above was convex. Thus, in this case, a locally optimal solution is also globally optimal.

4.3. Partial Volume Constraints. Recall that the dose volume histogram (see Figure 4.1) displays the fraction of each region of the patient that receives at least a specified dose level. In some cases, the radiation oncologist is willing to sacrifice a portion of a region at risk in order to improve the probability of curing the disease. Oncologists often specify constraints of the form: "no more than x% of this region at risk can exceed a dose of y." Thus, for a particular region at risk, the oncologist determines both a dose limit and a fraction of the structure that can exceed the dose limit. This type of requirement is called a partial volume constraint [53, 39].

For our formulas, the dose limit will be denoted by Λ and the fraction of the volume allowed to exceed this limit will be denoted as Ω. We have tested both a nonlinear and a mixed integer approach to the implementation of partial volume constraints.

The nonlinear formulation represents a new approach to partial volume constraints:

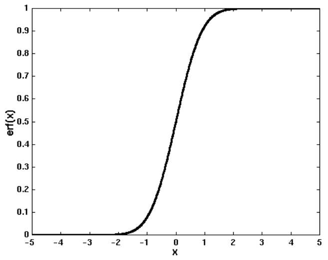

(a) The error function used in our implementation, $f(x) = erf(x)$.

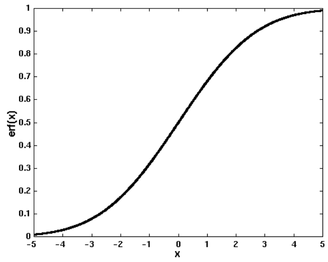

(b) A modified error function (erf(x/3)); the factor of 3 in the denominator creates a more sloped ramp $f(x) = erf(x/3)$.

FIG. 4.4. *Error functions used for partial volume constraints*

$$
\begin{array}{lcr}
\min_w & \sum_{(i,j)\in\mathcal{T}} (D_{ij} - \delta_{ij})^2 & \\
\text{subject to} & D_{ij} = \sum_{p=1}^{n} w_p D_{ij}^p & \forall (i,j) \\
 & \sum_{(k,l)\in\mathcal{R}} \mathrm{erf}(D_{kl} - \Lambda_{kl}) \leq n_{\mathcal{R}} \Omega_{\mathcal{R}} & \\
 & \sum_{(k,l)\in\mathcal{N}} \mathrm{erf}(D_{kl} - \Lambda_{kl}) \leq n_{\mathcal{N}} \Omega_{\mathcal{N}} & \\
 & w_p \geq 0 &
\end{array}
$$

In these formulas, $n_{\mathcal{R}}$ is the number of pixels in the region at risk and $n_{\mathcal{N}}$ is the number of pixels in the normal tissue. The partial volume constraints were realized through the use of a ramp function. In this case the error function (erf) was used (see Figure 4.4a). For each partial volume constraint, the error function was shifted so that the center of the ramp matched the dose limit. With other shifts and scalings it may be possible to improve the solution quality and numerical performance. The goal of the optimizer was to minimize the sum of the squared difference between the prescribed and the actual doses over all of the pixels in the tumor subject to two partial volume constraints.

Figure 4.5 show the results of an optimization performed using this approach. The geometry is the inverse of the "L" shaped target used previously. In this case, the arms of an "L" embrace a smaller square shaped target. This geometry was chosen in order to effectively test this implementation of partial volume constraints. For this simulated treatment, the prescription dose level in the tumor was set at 70 Gy. The first partial volume constraint specified that 80% of the region at risk must be kept below 30 Gy. The second partial volume constraint specified that 95% of the normal tissue region must be kept below 50 Gy.

A scalar value can be placed in the denominator of the error function in order to alter the slope of the curve (see Figure 4.4b). As the ramp function is made steeper, the constraint is becomes more absolute. It may be advisable to first solve a problem using a ramp function with a gradual slope. That solution can then be used as an initial guess in an optimization using a more steep ramp function. One must be

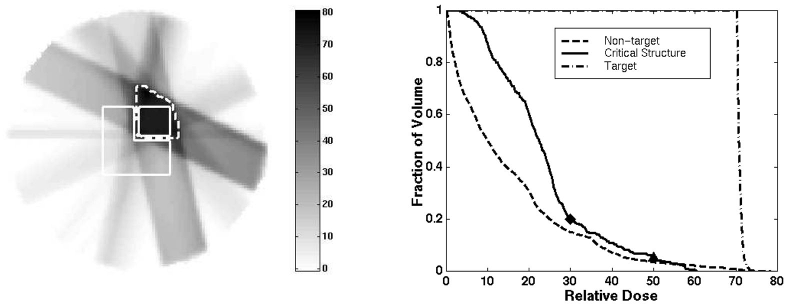

(a) optimized dose distribution (b) cumulative dose-volume histogram

FIG. 4.5. *Optimization using partial volume constraints implemented via a ramp function. In this case, 20% of the region at risk was allowed to exceed 30 Gy and 5% of the region at risk was allowed to exceed 50 Gy. The diamond on the DVH indicate the DVH constraint applied to the critical structure, and the triangle indicates the DVH constraint applied to the normal tissue.*

careful, however, when choosing the scaling factor. If the slope of the ramp is too gradual, the solution may differ considerably from the desired results. On the other hand, a ramp that is too steep may prove to be too difficult for the optimization algorithm. Both of these possibilities occurred in our experimentation. The results presented in Figure 4.5 were produced using the error function depicted in Figure 4.4a.

Mixed Integer Programming (MIP) makes possible a second approach to the implementation of partial volume constraints. Mathematically, the mixed integer programming problem appears as follows:

$$\begin{array}{ll} \min & c^T x + d^T y \\ \text{subject to} & Ax + By \geq b \\ & \ell \leq x \leq u \\ & y_i \in \{0, 1, 2, \ldots\} \\ & w_p \geq 0 \end{array}$$

In this formulation, x is a vector of variables that are continuous real numbers and y is a vector of variables that can only take integer values, $c^T x + d^T y$ is the linear objective function, and $Ax + By \geq b$ represents the set of constraints. ℓ and u are vectors of lower and upper bounds placed on the continuous variables, and $y \in \{0, 1, 2, \ldots\}$ is the integrality requirement. Some related work using MIP formulations is described in [39].

An optimized dose distribution produced using a mixed integer model is shown in Figure 4.6. The goal of this optimization was to minimize the maximum deviation in dose over the target subject to a partial volume constraint that specified that 85% ($\Omega_{\mathcal{R}}$) of the region at risk must be maintained below the dose limit (Λ) of 30 Gy. This constraint was modeled through the use of a binary variable O_{ij} that equaled 1 if the pixel in the region at risk exceeded the prescribed limit and 0 otherwise.

The following formulation was used to generate the simulated treatment:

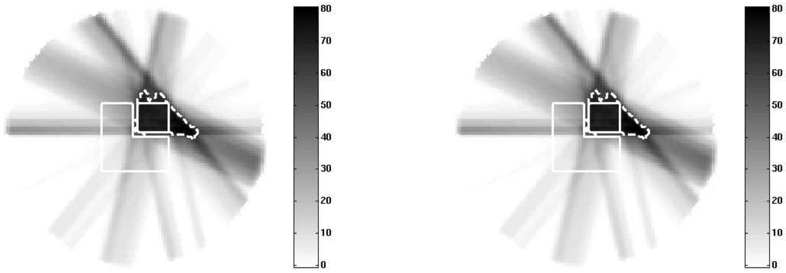

(a) optimized dose distribution (b) dose-volume histogram

FIG. 4.6. *A MIP implementation of partial volume constraints; here 15% of the critical structure was allowed to exceed 30 Gy.*

$$
\begin{array}{lll}
\min_w & \max_{(i,j)\in\mathcal{T}} |D_{ij}\text{-}\delta_{ij}| & \\
\text{subject to} & D_{ij} = \sum_{p=1}^{n} w_p D_{ij}^p & \forall (i,j) \\
& D_{kl} \leq \Lambda_{kl} + M * O_{kl},\ O_{k,l} \in \{0,1\} & \forall (k,l) \in \mathcal{R} \\
& \sum_{(r,s)\in\mathcal{R}} O_{rs} \leq n_{\mathcal{R}} \Omega_{\mathcal{R}} & \\
& w_p \geq 0 &
\end{array}
$$

The constraint

$$D_{kl} \leq \Lambda_{kl} + M * O_{kl}, O_{k,l} \in \{0,1\}$$

is a standard MIP technique for modeling an "if-then" constraint. The value of M is chosen to be large enough so that the corresponding constraint is trivially satisfied when $O_{kl} = 1$. In our case, the dose at any pixel is always less than 100 Gy, so we typically set M equal to 100. Note that if D_{kl} exceeds Λ_{kl}, then O_{kl} must be set equal to 1. The variables O_{kl} can then be used to enforce the partial volume constraint. In the above case, the optimizer summed O_{rs} over all of the pixels in the region at risk, and it required that this value be less than $\Omega_{\mathcal{R}}$ times the total number of pixels in the region at risk.

One appealing quality of partial volume constraints is that physicians are experienced with this type of requirement. Partial volume constraints are much more intuitive than assigning relative weights to each region of the patient. A limitation of the MIP implementation of partial volume constraints is that only linear objective functions and linear constraints can be included in the optimization. With the ramp function technique, however, any linear or nonlinear objective function or constraint can be used. This, however, can lead to more complex and difficult nonlinear programs that increase the probability of finding local solutions that are not globally optimal. The primary disadvantage of partial volume constraints is that due to the complex formulation, each constraint can add considerable time to the optimization.

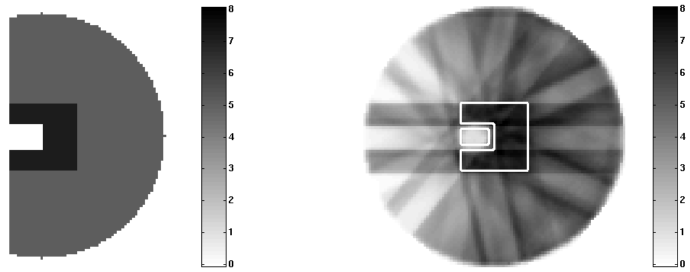

(a) prescription (b) 7 beam angles

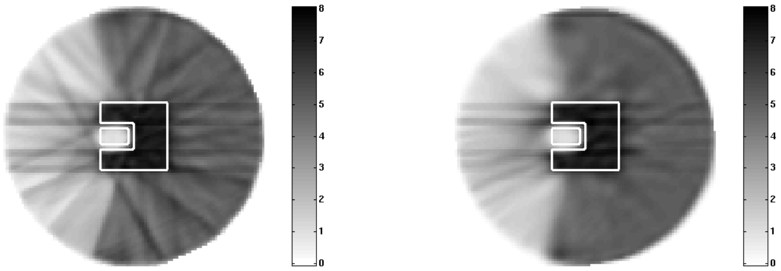

(c) 15 beam angles (d) 71 beam angles

FIG. 4.7. *A series of conformal avoidance dose distributions. Note that a large number of beam angles are needed to approach this difficult prescription.*

4.4. Conformal Avoidance. For a large number of radiotherapy patients, the tumor volume is not clearly defined. Thus, there is little rationale for attempting to create a dose distribution that conforms tightly to the tumor volume. Instead, it may be beneficial to design a treatment that treats a regional field while "conformally avoiding" any critical structures in that field. We have studied the possibility of using intensity-modulated beams in order to deliver conformal avoidance radiotherapy.

The results from a series of simulations are shown in Figures 4.7 and 4.8. This case utilized a U-shaped target with a prescribed relative dose of 7. A large nodal region was prescribed a relative dose of 5. The goal of the optimization was to produce a treatment plan that provided a uniform dose over both the target and the nodal regions. Simultaneously, the optimizer attempted to both conformally avoid a sensitive structure in the concavity of the U and minimize the dose to the non-nodal region.

It should be noted that finite source size and the phenomenon of electron transport both cause a blurring of each beamlet. This blurring makes it impossible to exactly match the prescription. This is particularly true when the prescription calls for adjacent pixels to receive dramatically different dose levels.

From the perspective of optimization, conformal avoidance simulations represent

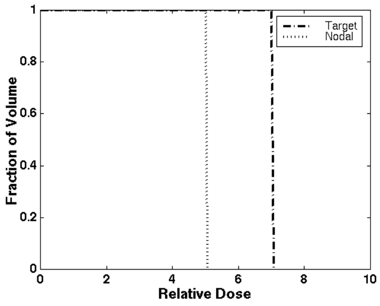

(a) prescription

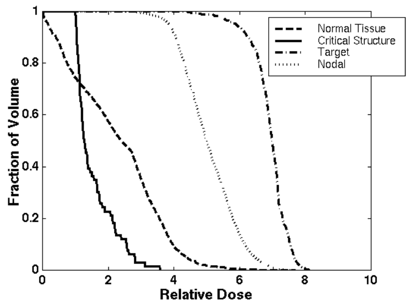

(b) 7 beam angles

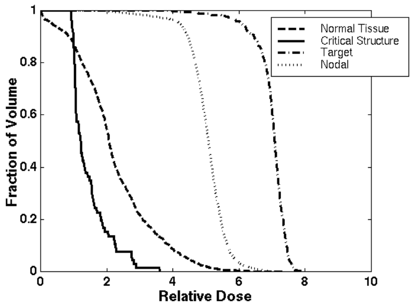

(c) 15 beam angles

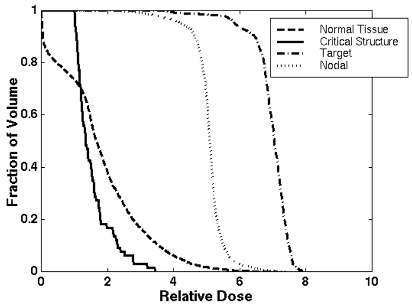

(d) 71 beam angles

FIG. 4.8. *A series of cumulative dose-volume histograms corresponding to the conformal avoidance results shown in Figure 4.7.*

a particular challenge. In conventional treatments, the majority of potential beamlets are removed from the optimization because they do not pass through the tumor. In conformal avoidance, a large regional field is treated. Thus, virtually all of the potential beamlets are used in the optimization. While the general form of the optimization problems is the same as those described above, typically the weighting parameters are changed and the amount of data involved is greatly increased.

One method to speed up conformal avoidance optimizations is to optimize over a subset of the pixels. The optimization using 72 beam angles was performed over a grid of points. A computation point was placed at every other pixel in both the x and y direction. Thus, 75% of the pixels were excluded from the optimization.

From Figures 4.7 and 4.8, it can be seen that a large number of beam angles are needed to successfully deliver conformal avoidance treatment plans.

5. Conclusions and Future Work. The field of radiotherapy is currently in a period of dramatic change. New treatment machines are providing the opportunity to treat cancer patients with thousands of independently weighted beamlets of radiation. The goal of this work has been to examine a number of mathematical models for determining the optimal beam weights. Our simulated treatments have helped to further our understanding of the advantages and disadvantages of linear, nonlinear, and mixed integer approaches to optimizing each patient's treatment plan.

The flexibility provided by nonlinear programming is particularly appealing. The use of nonlinear constraints in conjunction with a quadratic objective function should provide an oncologist with the flexibility to produce an acceptable treatment plan for almost any disease site.

A number of critical challenges remain. Clinical implementation of an optimization approach requires a technique that is suitable for solving very large problems. Actual patient cases will involve as many as 100,000 beamlets of radiation in a three dimensional setting. We believe both model formulation and algorithmic improvements are necessary if problems of this size are to be solved in an acceptable amount of time.

The ideal optimization approach should provide that ability to include partial volume constraints in an optimization model. We have tested both a nonlinear and a mixed integer approach to partial volume constraints. Although these two approaches have proven successful for small scale problems, either technique can become excessively time consuming in optimizing a real patient case. Similarly, the enormous amounts of data involved in the optimization of conformal avoidance treatments necessitates substantial computational advancements.

It is our hope that the community of optimization experts will be able to offer further insights that will improve our ability to solve these difficult problems. For readers interested in experimenting with these problems, the beamlet data sets along with the code used to compute them are available at the web site:

`http:\\www.madrad.radiology.wisc.edu`

REFERENCES

[1] J. Balog. Dose computation for a one-dimensional multileaf collimator. In *Proceedings of the XII International Conference on the Use of Computers in Radiation Therapy*, Salt Lake City, Utah, USA, 1997. Medical Physics Publishing.

[2] T. Bortfeld and W. Schlegel. Optimization of beam orientations in radiation therapy: some theoretical considerations. *Physics in Medicine and Biology*, 38(2):291–304, 1993.

[3] T.R. Bortfeld, D.L. Kahler, T.J. Waldron, and A.L. Boyer. X-ray field compensation with multileaf collimators. *International Journal of Radiation Oncology, Biology, Physics*, 28(3):723–730, 1994.

[4] A.L. Boyer. Present and future developments in radiotherapy treatment units. *Seminars in Radiation Oncology*, 5(2):146–155, 1995.

[5] A. Brahme. Optimization of stationary and moving beam radiation therapy techniques. *Radiother. Oncol.*, 12:129–140, 1988.

[6] A. Brahme. Optimization of radiation therapy and the development of multileaf collimation. *International Journal of Radiation Oncology Biology and Physics*, 25:373–375, 1993.

[7] A. Brahme. Treatment optimization: Using physical and radiobiological objective functions. In A.R. Smith, editor, *Radiation Therapy Physics*, pages 209–246. Springer-Verlag, Berlin, 1995.

[8] L. Brewster, R. Mohan, G. Mageras, C. Burman, S. Leibel, and Z. Fuks. Three dimensional conformal treatment planning with multileaf collimators. *International Journal of Radiation Oncology, Biology, Physics*, 33(5):1081–1089, 1995.

[9] A. Brooke, D. Kendrick, and A. Meeraus. *GAMS: A User's Guide*. The Scientific Press, South San Francisco, CA, 1988.

[10] C. Burman, G.J. Kutcher, B. Emami, and M. Goitein. Fitting of normal tissue tolerance data to an analytical function. *Int. J. Radiation Oncoloby Biol. Phys.*, 1991.

[11] M. Carol, H. Targovnik, D. Smith, and D. Cahill. 3-d planning and delivery system for optimzed conformal therapy. *International Journal of Radiation Oncology, Biology, Physics*, 24, 1992.

[12] P.S. Cho, S. Lee, R.J. Marks, S. Oh, S. Sutlief, and H. Phillips. Optimization of intensity

modulated beams with volume constraints using two methods: cost function minimization and projection onto convex sets. *Medical Physics*, 25(4):435–443, 1998.

[13] V. Chvátal. *Linear Programming.* W. H. Freeman and Company, New York, 1983.

[14] A. R. Conn, N. I. M. Gould, and Ph. L. Toint. *LANCELOT: A Fortran package for Large–Scale Nonlinear Optimization (Release A).* Number 17 in Springer Series in Computational Mathematics. Springer Verlag, Heidelberg, Berlin, 1992.

[15] R.E.M. Cooper. A gradient method of optimizing external-beam radiotherapy treatment plans. *Radiology*, 128:235–243, 1978.

[16] A.M. Cormack. A problem in roatation therapy with x-rays. *International Journal of Radiation Oncology, Biology, Physics*, 13:623–630, 1987.

[17] A.M. Cormack. A problem in roatation therapy with x-rays2: dose distributions with an axis of symmetry. *International Journal of Radiation Oncology, Biology, Physics*, 13:1921–1925, 1987.

[18] G. B. Dantzig. *Linear Programming and Extensions.* Princeton University Press, Princeton, New Jersey, 1963.

[19] J.O. Deasy. Multiple local minima in radiotherapy optimization problems with dose volume constraints. *Med. Phys.*, 1997.

[20] B.G. Fallone D.H. Hristov. A continuous penalty function method for inverse treatment planning. *Medical Physics*, 25(2):208–223, 1998.

[21] A. Drud. CONOPT: A GRG code for large sparse dynamic nonlinear optimization problems. *Mathematical Programming*, 31:153–191, 1985.

[22] M. N. Du, C. X. Yu, M. Symons, D. Yan, R. Taylor, R. C. Matter, G. Gustafson, A. Martinez, and J. W. Wong. A multileaf collimator field prescription preparation system for conventional radiotherapy. *International Journal of Radiation Oncology, Biology, Physics*, 30(3):707–714, 1994.

[23] G. Fang, B. Geiser, and T.R. Mackie. Software system for uw/ge tomotherapy prototype. In *Proceedings of the XII International Conference on the Use of Computers in Radiation Therapy*, Salt Lake City, Utah, USA, 1997. Medical Physics Publishing.

[24] M. C. Ferris. MATLAB and GAMS: Interfacing optimization and visualization software. Mathematical Programming Technical Report 98-19, Computer Sciences Department, University of Wisconsin, Madison, Wisconsin, 1998.

[25] E.E. Fitchard, J.S. Aldridge, P.J. Reckwerdt, and T.R. Mackie. Registration of synthetic tomographic projection data sets using cross-correlation. *Phys. Med. Biol.*, 43:1645–1657, 1998.

[26] A. Fraass. The development of conformal radiation therapy. *Medical Physics*, 22(11):1911–1921, 1995.

[27] P. E. Gill, W. Murray, and M. H. Wright. *Practical Optimization.* Academic Press, London, 1981.

[28] W.I. Grant and S.Y. Woo. Clinical and financial issues for intensity modulated radiation therapy delivery. *Seminars in Radiation Oncology*, 1999.

[29] A. Gustafsson, B. K. Lind, R. Svensson, and A. Brahme. Simultaneous optimization of dynamic multileaf collimation and scanning patterns or compensation filters using a generalized pencil beam algorithm. *Medical Physics*, 22(7):1141–1156, 1995.

[30] T. Holmes and T.R. Mackie. A comparison of three inverse treatment planning algorithms. *Physics in Medicine & Biology*, 39(1):91–106, 1994.

[31] T. Holmes and T.R. Mackie. A filtered backprojection dose calculation method for inverse treatment planning. *Medical Physics*, 21(2):303–313, 1994.

[32] M. S. Hung, W. O. Rom, and A. D. Warren. *OSL Handbook.* Boyd and Fraser, 1995.

[33] ILOG CPLEX Division, Incline Village, Nevada. *CPLEX Version 6.0.* http://www.cplex.com/.

[34] A. Jackson, G.J. Kutcher, and E.D. Yorke. A probability of radiation-induced complications for normal tissues with parallel architecture subject to non-uniform irradiation. *Med. Phys.*, 1993.

[35] A. Jackson, X.H. Wang, and R. Mohan. Optimization of conformal treatment planning and quadratic dose objectives (abstract). *Med. Phys.*, 1994.

[36] P. Kallman, A. Agren, and A. Brahme. Tumor and normal tissue responses to fractionated non-uniform dose delivery. *Int. J. Radiation Oncoloby Biol. Phys.*, 1992.

[37] G.J. Kutcher and C. Burman. Calculation of complication probability factors for non-uniform normal tissue irradiation: The effective volume method. *Int. J. Radiation Oncoloby Biol. Phys.*, 1989.

[38] G.J. Kutcher, C. Burman, L. Brewster, M. Goitein, and R. Mohan. Histogram reduction method for calculating complication probabilities for three-dimensional treatment planning evaluation. *Int. J. Radiation Oncology Biol. Phys.*, 1991.

[39] M. Langer and S. Morrill. A comparison of mixed integer programming and fast simulated annealing for optimized beam weights in radiation therapy. *Medical Physics*, 23(6):957–964, 1996.

[40] D.D. Leavitt, M. Martin, J.H. Moeller, and W.L. Lee. Dynamic wedge field techniques through computer-controlled collimator motion and dose delivery. *Medical Physics*, 17(1):87–92, 1990.

[41] S.A. Leibel, G.J. Kutcher, L.B. Harrison, D.E. Fass, C. Burman, M. Hunt, R. Mohan, L.J. Brewster, C.C. Ling, and Z.Y. Fuks. Improved dose distributions for 3d conformal boost treatments in carcinoma of the nasopharynx. *Int. J. Radiation Oncology Biol. Phys.*, 20:823–833, 1991.

[42] D. C. Liu and J. Nocedal. On the limited memory method for large scale optimization. *Mathematical Programming*, 45:503–528, 1989.

[43] J. Llacer. Inverse radiation treatment planning using the dynamically penalized likelihood method. *Medical Physics*, 24(11):1751–1764, 1997.

[44] J.T. Lyman and A.B. Wolbarst. Optimization of radiation therapy iv: A dose volume histogram reduction algorithm. *Int. J. Radiation Oncoloby Biol. Phys.*, 1989.

[45] T.R. Mackie, T. Holmes, S. Swerdloff, P. Reckwerdt, J.O. Deasy, J. Yang, B. Paliwal, and T. Kinsella. Tomotherapy: a new concept for the delivery of dynamic conformal radiotherapy. *Medical Physics*, 20(6):1709–1719, 1993.

[46] T.R. Mackie, T.W. Holmes, P.J. Reckwerdt, and J. Yang. Tomotherapy: Optimized planning and delivery of radiation therapy. *Int. J. of Imaging Systems and Technology*, 6:43–55, 1995.

[47] T.R. Mackie, J.W. Scrimger, and J.J. Batista. A convolution method of calculating dose from 15 mev x-rays. *Medical Physics*, 12:188–196, 1985.

[48] MATLAB. *User's Guide*. The MathWorks, Inc., 1992.

[49] R. Mohan. Field shaping for three-dimensional conformal radiation therapy and multileaf collimation. *Seminars in Radiation Oncology*, 5(2):86–99, 1995.

[50] R. Mohan, G.S. Mageras, B. Baldwin, L.J. Brewster, G.J. Kutcher, S. Leibel, C.M. Burman, C.C. Ling, and Z. Fuks. Clinically relevant optimization of 3-d conformal treatments. *Med. Phys.*, 19(4):933–944, 1992.

[51] J. J. Moré and G. Toraldo. On the solution of large quadratic programming problems with bound constraints. *SIAM Journal on Optimization*, 1:93–113, 1991.

[52] J. J. Moré and S. J. Wright. *Optimization Software Guide*. Number 14 in Frontiers in Applied Mathematics. SIAM, Philadelphia, Pennsylvania, 1993.

[53] S. Morrill, R. Lane, J. Wong, and I. Rosen. Dose-volume considerations with linear programming. *Medical Physics*, 18(6):1201–1210, 1991.

[54] S.M. Morrill, K.S. Lam, R.G. Lane, M. Langer, and I.I. Rosen. Very fast simulated annealing in radiation therapy treatment plan optimization. *Int. J. Radiation Oncoloby Biol. Phys.*, 1995.

[55] G.P. Murphy, W.L. Lawrence, and R.E. Lenhard. *American Cancer Society Textbook of Clinical Oncology*. The American Cancer Society, Inc., 1995.

[56] B. A. Murtagh and M. A. Saunders. MINOS 5.0 user's guide. Technical Report SOL 83.20, Stanford University, Stanford, California, 1983.

[57] K. G. Murty. *Linear Programming*. John Wiley & Sons, New York, 1983.

[58] S. G. Nash and J. Nocedal. A numerical study of the limited memory BFGS method and the truncated-Newton method for large scale optimization. *SIAM Journal on Optimization*, 3:358–372, 1991.

[59] J. L. Nazareth. *Computer Solution of Linear Programs*. Oxford University Press, Oxford, 1987.

[60] A. Niemierko. Optimization of 3d radiation therapy with both physical and biological end points and constraints. *Int. J. Radiation Oncology Biol. Phys.*, 23:99–108, 1992.

[61] A. Niemierko. Treatment plan optimization. In *3-D Radiation Treatment Planning and Conformal Therapy*. Medical Physics Publishing, St. Louis, MO, 1993.

[62] N.C.I. National Institutes of Health. *Radiation Therapy and You - A Guide to Self-Help During Treatment*. NIH Publication, Bethesda, Maryland, 97-2227 edition, 1997.

[63] G.H. Olivera, D.M. Shepard, P.J. Reckwerdt, J. Zachman, E.E. Fitchard, and T.R. Mackie. Maximum likelihood as a common computational framework in tomotherapy. *Phys. Med. Biol.*, 1998 (accepted for publication).

[64] W. Orchard-Hays. *Advanced Linear Programming Computing Techniques*. McGraw–Hill, New York, 1968.

[65] N. Papanikolaou, T.R. Mackie, C. Meger-Wells, M. Gehring, and P. Reckwerdt. Investigation of the convolution method for polyenergetic spectra. *Medical Physics*, 20(5):1327–1336, 1993.

[66] C.A. Perez and L.W. Brady. *Principles and practice of radiotherapy.* Lippincott-Raven, Philadelphia, 3rd edition, 1998.

[67] P.J. Reckwerdt, T.R. Mackie, J.P. Balog, and T.R. McNutt. Three-dimensional inverse treatment optimization for tomotherapy. In *XII International Conference on the Use of Computers in Radiation Therapy*, Salt Lake City, Utah, USA, 1997. Medical Physics Publishing.

[68] I. Rosen, R. Lane, S. Morrill, and J. Belli. Treatment plan optimization using linear programming. *Medical Physics*, 18(2):141–152, 1990.

[69] K.J. Ruchala, G.H. Olivera, and T.R. Mackie. A comparison of maximum-likelihood and iterative filtered backprojection reconstruction algorithms for megavoltage ct on a tomotherapy system. *Phys. Med. Biol.*, 1999 (submitted).

[70] D.M. Shepard, L. Angelos, O.A. Sauer, and T.R. Mackie. A simple model for examining radiotherapy optimization. In *Proceedings of the XII International Conference on the Use of Computers in Radiation Therapy*, Salt Lake City, Utah, USA, 1997. Medical Physics Publishing.

[71] L.J. Verhey. Immobilizing and positioning patients for radiotherapy. *Seminars in Radiation Oncology*, 5(2):100–113, 1995.

[72] S. Webb. Optimum parameters in a model of tumor control probability including interpatient heterogeneity. *Phys. Med. Biol.*, 1994.

[73] S. Webb. *The Physics of Conformal Radiotherapy - Advances in Technology.* Institute of Physics Publishing, 1997.

[74] S. Webb. *The Theory and Practice of Intensity Modulated Radiation Therapy*, chapter Inverse Planning for IMRT: the Role of Simulated Annealing. Advanced Medical Publishing, 1997.

[75] S. Y. Woo, M. Sanders, W. Grant, and E. B. Butler. Does the "peacock" have anything to do with radiotherapy? *International Journal of Radiation Oncology, Biology, Physics*, 29(1):213–214, 1994.

[76] S. J. Wright. *Primal–Dual Interior–Point Methods.* SIAM, Philadelphia, Pennsylvania, 1997.

[77] J.N. Yang, T.R. Mackie, P. Reckwerdt, J.O. Deasy, and B.R. Thomadsen. An investigation of tomotherapy beam delivery. *Medical Physics*, 24(3):425–436, 1997.