

Creating a Transparent and Convex Objective Function for Intensity-Modulated Radiotherapy

Alexander Kendall*

Honors Thesis Advisor: Doron Levy
Mathematics Department, Stanford University

June 3, 2004

Abstract

This paper considers the objective functions currently being used in Intensity-Modulated Radiotherapy. In particular, we argue that there are three key features that any objective function should have. It should be convex, biologically relevant, and easily understandable. None of the current methods possess these features. To address this problem, we develop a new, two-tiered approach, based on intra- and inter-structural scoring functions. The Equivalent Uniform Dose (EUD) is one such intra-structural scoring function, so we develop the Linear Expenditure System as an inter-structural scoring function and show that this function, when combined with the EUD satisfies the three key properties. We then consider an alternative to the EUD, the n-part EUD. This function allows physicians to have greater control over patients' dose distributions.

*There are many people to whom I owe a deep debt of gratitude for their assistance in this project. Doron Levy has been an excellent thesis advisor, first pointing me in the direction of this research topic, then supporting me in any way he could. The expertise of Lei Xing and Edi Schreibmann has been invaluable. Without them I would know nothing about radiation oncology. For her clinical wisdom that has motivated much of my work, I am indebted to Cathy Su. I am also grateful to Jay Bhattacharya for putting me in touch with Dr. Su and discussing my project with me at great length. Finally, I would never accomplish anything were it not for the support of my friends and family. In particular, I wish to honor Elizabeth Madjlessi, whose volunteer work has done more for the care of cancer patients than any research I have pursued. All errors are, of course, my own. Please address correspondence to akendall@stanfordalumni.org.

1 Introduction

Radiation oncology is a complicated venture. In order to destroy tumors in patients' bodies, doctors irradiate them, depriving them of their ability to divide and multiply. Of course, tumors are often deep inside patients' bodies, so it is impossible to hit only malignant cells with radiation—doctors must irradiate healthy cells as well. Thus one of the medical practitioner's main jobs is to determine how best to balance the trade-offs between destroying tumors and saving critical organs from collateral damage.

To that end, doctors employ complicated computer programs to help them find the best possible way to treat patients. The most cutting edge method is Intensity Modulated Radiotherapy (IMRT), in which doctors input prescribed dosages into their systems, which in turn produce treatment plans based upon the doctors' intentions.

In this paper, we examine the primary methods employed in IMRT and conclude that they suffer from several flaws. First and foremost, conventional IMRT methods do not maximize patient health, but rather minimize the deviation from the prescribed dose. While these two approaches are closely related, the distinction leaves open the possibility that suboptimal plans are being selected. Second, the primary parameters that doctors have at their disposal to guide the planning process are not easily interpretable. The importance of this seemingly academic criticism is evidenced by the fact that, in practice, doctors must go through several iterations with the computer systems in order to reach outcomes they consider acceptable. Finally, many of the more biologically promising techniques are hindered by the fact that they require non-convex optimization in order to attain a solution.

To address these flaws, we argue that IMRT should be redefined in a two-tiered framework. In the first tier, comparisons should be made within organs and tumors to determine what the biologically relevant level of exposure they receive is. In the second tier, comparisons should be made between organs and tumors in order to determine how much exposure the cancer should receive and at what cost to healthy tissue.

The primary contribution of this paper is the introduction of a useful objective func-

tion for the second tier of this process, the Linear Expenditure System (LES). Conditional on the necessary assumptions' holding true, maximizing this function maximizes patient health. Unlike current functions, the LES's parameters are easily interpretable. Perhaps most importantly, the function permits convex optimization.

We also consider the first, intra-organ tier. To that end, we acknowledge the potential of the Equivalent Uniform Dose (EUD) function that has already been proposed. However, we propose a refinement, the n -part Equivalent Uniform Dose (nEUD), which holds the potential to make the measure more easily manipulated by practitioners.

We proceed in seven steps. In Section 2, we give a general background on IMRT, detailing the methods by which doctors use radiation to treat cancer. In Section 3, we detail the primary current methods available to doctors. In Section 4, we argue that a two-tiered mechanism is necessary. In Section 5, we develop the LES function to allow comparisons between organs. In Section 6, we consider what benefits a new intra-organ radiobiological measure, the nEUD might have. In Section 7, we offer suggestions for further research. In Section 8, we conclude.

2 Background

We begin by describing the basic mechanisms of radiation oncology, focusing specifically on the role of IMRT. Radiation destroys tumors in two ways. First, high enough dosages can kill cancerous cells. This effect, however, is much less important than the second, in which radiation prevents cancerous cells from dividing and multiplying, thus eliminating the tumor's ability to regenerate itself from natural attrition.

If radiation could be directed so as to only affect malignant cells, the story would be over right there. However, there is no way to irradiate tumorous cells without damaging healthy cells as well. Fortunately, it is an empirical fact that healthy cells are more resilient to radiation than cancerous cells are. Healthy cells are better able to repair their DNA—the

component of cells damaged by radiation—than their malignant counterparts. Furthermore, radiation impacts cells that rapidly divide much more strongly than cells that do not. Thus cancer’s own weapon is turned against itself.

However, healthy organs still suffer at the hands of radiation, so it is important to minimize the exposure to healthy tissue. To that end, the fundamental strategy is to have several radiation beams, each approaching from a different direction, converge at the tumor. That way the tumor is hit by much more radiation compared to the healthy tissue. Such an orientation of radiation beams is illustrated in Figure 1.




Figure 1: A Typical Beam Orientation

Determining where exactly to source the beams and what strength and shape they should be is a complicated, multi-step procedure. The first step is to take a computerized tomographic scan (CT) of the area of the patient affected by the tumor. The output of such a scan is a three dimensional model of the patients body made up of voxels, the three-dimension equivalent of pixels. Each voxel is assigned to a particular tissue type. Some voxels will be marked as part of the tumor, some as part of the liver, some as part of the stomach, etc.

Based on the three-dimensional map of the patient's body, the radiation oncologist then determines which parameters he will enter into the inverse treatment planning platform. The computer then optimizes the treatment plan subject to the physician's commands and the physical constraints of the problem. The output of this step is a map of the radiation exposure to each voxel and directions for how to obtain the given dose distribution.

After this first optimization, doctors often decide that their initial parameter entries were insufficient, so they iterate the steps of entering parameters and viewing output until they are satisfied. Once satisfied, the doctors follow the instructions of the computer and irradiate the patient.

In this paper, we focus exclusively on the steps of the physician entering parameters and the computer finding the conditionally optimal plan. Thus, it is worthwhile to consider this phase in greater detail. In particular, we now examine the nature of the solution space faced by the optimization algorithm.

Radiation beams can be fine tuned. Indeed, each radiation beam is a function of thousands of parameters. In particular, every beam is made up of beamlets, which are portions of the beam that can have different intensities. So if one were to look at a cross-section of a beam, he could map different parts of the beam to different strengths of radiation. The fundamental control variable of the planning process is therefore beamlet intensity. This fact is the derivation of the term for this procedure, Intensity-Modulated Radiotherapy.

The fact that beamlets are the control variables gives rise to the constraints in this problem. In particular, each beamlet must have a nonnegative intensity.

Once emitted from the source, beamlets decay at a predictably exponential rate. Thus, if one knows all of the beamlet intensities, one can determine the total radiation exposure for each voxel in the patient's body. It is a mathematical fact that directly follows exponential decay that total exposure for a voxel is a linear function of each beamlet intensity. Knowing the radiation exposure at every voxel, the computer then maps this extremely high dimensional vector into a real-valued function to be maximized or minimized, depending on the particulars of the strategy being implemented. Given such a function, the computer finds the optimal plan using a strategy suited to the particular function.

Thus, Intensity-Modulated Radiotherapy is essentially a constrained optimization over an extremely large (tens of thousands) space of control variables which map into one real-valued function. This paper will consider potential objective functions to be used for this optimization.

3 Current Methods

We now turn to considering the methods currently available for doctors to employ in generating IMRT plans. First, we develop a system of notation we will use throughout. Then we move on to describe the traditional method: the quadratic objective function. Third, we examine its most promising successor, optimizations based upon the equivalent uniform dose. Finally, we detail the drawbacks of the current methods.

3.1 Notation

We begin by defining notation that we will use throughout this paper.

As we have already stated, the control variable is the beamlet intensity vector. We denote this vector as \vec{b} and define the intensity of the i^{th} beamlet as b_i . Thus the length of the vector \vec{b} is the dimension of the space which are optimizing over. Furthermore, each beamlet defines

a constraint:

$$b_i \geq 0. \quad (1)$$

Because these are the only constraints, the space over which the objective function is to be optimized is convex.

Beamlet intensities, of course, are not biologically relevant. The medically important vector is \vec{d} , the dose vector. We denote the total radiation exposure of the j^{th} voxel of the i^{th} organ as $d_{i,j}$. For ease of notation, we will denote the exposure to the j^{th} voxel of the tumor as $d_{T,j}$. We denote the total number of voxels in the tumor as N_T and the total number of voxels in organ i as N_i .

As radiation intensity decays exponentially, it follows that $d_{i,j}$ and $d_{T,j}$ are linear functions of \vec{b} (Choi and Deasy 2002). As a result, we apply the simple notation:

$$d_{i,j} = d_{i,j}(\vec{b}) \quad (2)$$

$$d_{T,j} = d_{T,j}(\vec{b}), \quad (3)$$

where $d_{i,j}(\cdot)$ and $d_{T,j}(\cdot)$ are linear functions from the beamlet-dimensional intensity space to a one dimensional exposure space.

3.2 Quadratic Objective Function

We are now ready to consider the tradition objective function used in IMRT: the quadratic objective function. As we will do throughout this paper, we begin by mathematically describing the function. We then consider the strengths of the approach. We conclude by pointing out the weaknesses.

3.2.1 Mathematics

The standard practice in IMRT is to minimize a quadratic distance function from an ideal dose distribution (Webb, 1989; Bortfeld *et al*, 1990; Mageras and Mohan, 1993; Xing and

Chen, 1996; and Cotrutz and Xing, 2002). Thus, the optimization problem is given by the following:

$$\min_{\vec{b}} \frac{\alpha_T}{N_T} \sum_j \beta_{T,j} \left(d_{T,j}(\vec{b}) - p_{T,j} \right)^2 + \sum_i \frac{\alpha_i}{N_i} \sum_j \beta_{i,j} \left(d_{i,j}(\vec{b}) - p_{i,j} \right)^2 \mid b_i \geq 0 \forall i. \quad (4)$$

The equation above requires some unpacking. In the literature, α_T and α_i are known as structural importance factors, $\beta_{T,j}$ and $\beta_{i,j}$ are known as voxel-dependent importance factors, and $p_{T,j}$ and $p_{i,j}$ are known as dose prescriptions. Denote the collection of these parameters in vectors as $\vec{\alpha}$, $\vec{\beta}$, and \vec{p} . As we have already stated, N_T and N_i are the number of voxels in given structures, $d_{T,j}(\vec{b})$ and $d_{i,j}(\vec{b})$ are dose functions, and the b_i are beamlet intensities.

The fundamental mechanics of this objective function are simple. The function calculates the deviation from the prescribed dosage for each voxel, weights it appropriately, taking care to make sure large organs are not overweighted for having more voxels, and sums up every weighted, squared deviation.

At this point we should note that, for the sake of simplicity, $\beta_{T,j}$ and $p_{T,j}$ are usually constant over all voxels within the tumor. Similarly, $\beta_{i,j}$ and $p_{i,j}$ are usually constant over all voxels in organ i .

3.2.2 Strengths

The standard quadratic objective function has three important strengths. First, it possesses a rather logical derivation, upon which the very name of the field, inverse treatment planning, rests. Under this derivation, we assume that the oncologist knows an ideal dose distribution and consequently inputs this distribution as the prescription vector \vec{p} . Thus, the goal of the computer is to invert this dose description into a treatment plan. Problematically, such an inverse function is, at best, difficult to generate. At worst, and most likely, such a function does not even exist. In other words, subject to the physical constraints it is most likely impossible that $\vec{p} = \vec{d}$ for any \vec{b} .

To deal with the difficulty of finding an inverse function, we redefine the nature of the problem. Instead of trying to find a distribution exactly equal to the one prescribed, we find the one closest to the prescription. Doing so allows well-developed optimization techniques to be used. Of course, in order to find the closest distribution, it is important to define a notion of closeness. To that end, any distance on \mathbb{R}^n would suffice. In the standard IMRT methodology, the distance function used is a weighted version of the standard euclidean distance function, with the weights given by the product $\frac{\alpha_T}{N_T}\beta_{T,j}$ for voxels in the tumor and $\frac{\alpha_i}{N_i}\beta_{i,j}$ for voxels in healthy tissue. The number of voxel's being in the denominator of the weights means that if all weights are equal, all structures will be treated equally, regardless of how many voxels are included. Regardless, minimizing the described distance is equivalent to minimizing the given objective function.

The second advantage of the standard method is tautological. In short, it is the standard method, so doctors are familiar with its workings and have an intuition on what the inputs and outputs of platforms using this method mean and how they can be used.

Finally, because the objective function is a distance, and thus a norm on the dose distribution space, it is a convex function on this space. Because compositions of convex functions with linear functions are convex, this objective function is, therefore, convex (Choi and Deasy, 2002). Since the constraints restrict the space to a convex subset of \mathbb{R}^n , the problem thus becomes a convex optimization. As a result, techniques that are well-suited to such problems may be applied to speed computation and ensure accuracy.

3.2.3 Weaknesses

Unfortunately, the quadratic objective function has two related weaknesses which give great cause for concern. First, it is entirely possible that minimizing this objective function will reject attainable and undeniably superior plans. For example, suppose two plans to be compared are one in which $d_{T,j}(\vec{b}) = p_{T,j}$ and $d_{i,j}(\vec{b}) = p_{i,j}$ for all i and j and another in which $d_{T,j}(\vec{b}) = p_{T,j} + 1$ and $d_{i,j}(\vec{b}) = p_{i,j} - 1$ for all i and j . In other words, one plan is

exactly what the physician input into the computer. The other plan is uniformly superior to the plan the physician entered, as, for each voxel in the tumor, the exposure is greater than prescribed and for each voxel in healthy tissue, the exposure is less than prescribed. The problem here is clear, as the computer would choose the prescribed plan, even though there is an undeniably better one available.

Though the problem is most pronounced in the case described above, whenever there exist some i and j such that $d_{T,j}(\vec{b}) > p_{T,j}$ or $d_{i,j}(\vec{b}) < p_{i,j}$, the quadratic optimization makes little sense, as plans will seem better as $d_{T,j}(\vec{b})$ decreases or $d_{i,j}(\vec{b})$ increases. Such can simply not be the case as the very goal of IMRT is to maximize exposure to the tumor, while minimizing it to healthy tissue.

This problem does have a simple—and often implemented—remedy. Many routines redefine the objective function as:

$$\frac{\alpha_T}{N_T} \sum_j \beta_{T,j} H(p_{T,j} - d_{T,j}(\vec{b})) (d_{T,j}(\vec{b}) - p_{T,j})^2 + \sum_i \frac{\alpha_i}{N_i} \sum_j \beta_{i,j} H(d_{i,j}(\vec{b}) - p_{i,j}) (d_{i,j}(\vec{b}) - p_{i,j})^2, \quad (5)$$

where $H(\cdot)$ is the Heaviside function given by:

$$H(x) = \begin{cases} 1 & \text{if } x \geq 0 \\ 0 & \text{if } x < 0 \end{cases}. \quad (6)$$

Thus, only voxels which are on the wrong side of the prescription contribute to the objective function.

Such a redefinition does not solve the problem. The reason for this fact is our second primary cause for concern: The objective function is not built upon reasonable medical assumptions. In order to explore this fact, it is helpful to define the level curves of the objective function. In the standard case, the level curves are ellipsoids centered at \vec{p} . Two-dimension projections of these level curves are given in Figure 2. To find which level curve a point is on in the redefined case, the easiest way is to apply the following algorithm: For all




Figure 2: A Two-Dimensional Projection of Level Curves for the Standard Quadratic Objective Function (f)

voxels in the tumor such that $d_{T,j}(\vec{b}) > p_{T,j}$, reset $d_{T,j}(\vec{b})$ to $p_{T,j}$ and for all voxels in health tissues such that $d_{i,j}(\vec{b}) < p_{i,j}$, reset them to $p_{i,j}$, then find the level curve the point is on. Two-dimensional projections of these level curves are given in Figures 3, 4, and 5.

This algorithm makes crystal clear one of the problems with this strategy: All points in the region where $d_{T,j}(\vec{b}) > p_{T,j}$ and $d_{i,j}(\vec{b}) < p_{i,j}$ for i and j lie on the same level curve. Thus, the optimization cannot distinguish between plans that might be undeniably better than other plans.

The projections illustrate well the problem of introducing the Heaviside function into the objective function: The biological assumptions required for the objective function to be correct are unreasonable. For example, suppose one normal tissue voxel is below prescription, and one is above. Then even the redefined quadratic objective function would reduce a




Figure 3: A Two-Dimensional Projection of Level Curves for the Redefined Quadratic Objective Function (f): $D1$ and $D2$ are both exposures for voxels in healthy tissue

change in beamlet intensities that would bring exposure in the voxel that is already below prescription to zero while only increasing the exposure in the other voxel by a small fraction of a percent.

Even if we were in the region where no voxel satisfies its prescription, the objective function flounders. There is simply no biological reason the metric of health is the same as the metric of distance from the prescription. Regardless, it is highly unlikely that not a single voxel will satisfy its prescription, so the function's merits—or lack thereof—in that region are probably irrelevant.

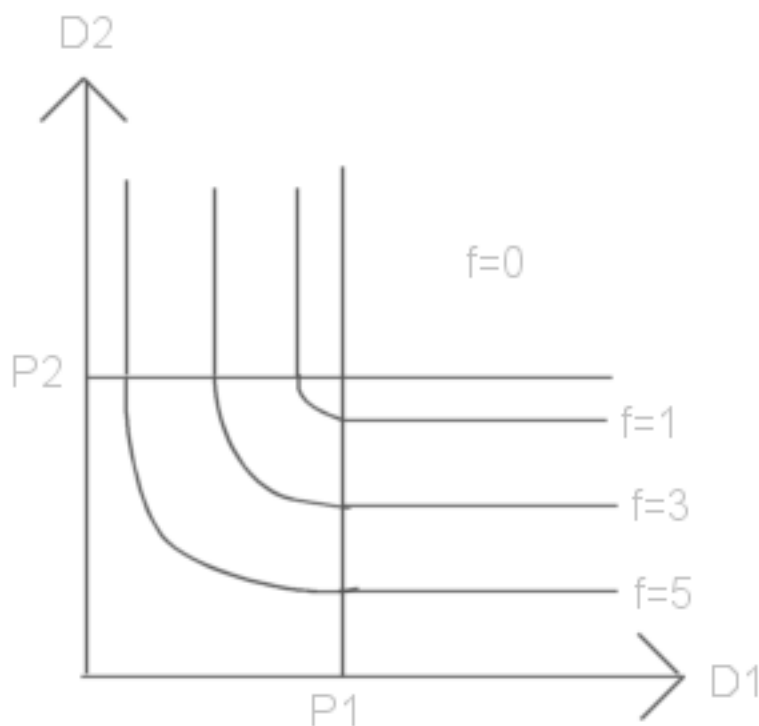


Figure 4: A Two-Dimensional Projection of Level Curves for the Redefined Quadratic Objective Function (f): $D1$ and $D2$ are both exposures for voxels in the tumor

3.3 Equivalent Uniform Dose

A much more reasonable approach than the one described above is that of the equivalent uniform dose. This approach makes an attempt to address the problem of biological relevancy. We first detail the mathematical structure of this approach. Then we consider the strengths of the EUD. Finally, we consider the weaknesses.

3.3.1 Mathematics

The EUD approach is a crude form of the two-tiered approach I will advocate in Section 4. Essentially, a biologically relevant score is determined for each structure—tumor or organ—and then the objective function combines these scores into a real-valued function.

The fundamental building block of this approach is the EUD function developed by




Figure 5: A Two-Dimensional Projection of Level Curves for the Redefined Quadratic Objective Function (f): D1 is an exposure in healthy tissue and D2 is an exposure in the tumor

Niemierko (1997, 1999). In mathematics this function is known as the generalized mean or power mean (Abramowitz, 1964). Regardless of the nomenclature, the EUD function for the tumor is given by:

$$EUD_T(a_T, \vec{b}) = \left(\frac{1}{N_T} \sum_j (d_{T,j}(\vec{b}))^{a_T} \right)^{\frac{1}{a_T}}, \quad (7)$$

and the EUD for the i^{th} healthy organ is given by:

$$EUD_i(a_i, \vec{b}) = \left(\frac{1}{N_i} \sum_j (d_{i,j}(\vec{b}))^{a_i} \right)^{\frac{1}{a_i}}. \quad (8)$$

The parameters a_T and a_i are the fundamental guiding parameters of the EUD functions. As these parameters adjust from $-\infty$ to ∞ , they give a large number of useful and familiar means. For example, when $a = 1$, the EUD function gives the arithmetic mean dose to

an organ. When $a = -\infty$, the EUD function maps to the minimum dose in a structure. Symmetrically, when $a = \infty$, the EUD is equal to the maximum dose in the organ. When $a = 2$, the EUD equals the root-mean-square. Finally, the limit as a approaches 0 of the EUD function is the geometric mean dose. Examining these special cases gives rise to the correct intuition that as a increases, the EUD is more a function of the maximum dose, while as a decreases the EUD is more a function of the minimum dose.

As we will discuss in the next section, different types of structures demand different types of EUDs. In particular, the tumor is usually given a_T large and negative. Serially functioning organs, such as the spinal cord or brain are given large and positive a_i s. Finally, parallel functioning organs, such as the lungs and liver, are given small values for a_i , which are larger than 1. We examine the rationale later.

Of course, the EUD functions do not constitute an objective function. Rather, a new function must be fashioned from them in order to give something to optimize. The main methodology of this optimization is given by Wu *et al* (2002). They formulate the optimization problem as a constrained maximization given by:

$$\max_{\vec{b}} \frac{1}{1 + \left(\frac{P_T}{EUD_T(a_T, \vec{b})}\right)^{\alpha_T}} \times \prod_i \frac{1}{1 + \left(\frac{EUD_i(a_i, \vec{b})}{P_i}\right)^{\alpha_i}} | b_i \geq 0 \forall i. \quad (9)$$

This objective function introduces two new types of parameters. First, P_T and P_i correspond to prescribed doses. Second, keeping in line with prior notation, α_T and α_i are importance factors, though they behave differently than in the quadratic case.

3.3.2 Strengths

The EUD formulation has four major strengths. We consider them in the context of the weaknesses of the quadratic objective function.

The primary motivation for the EUD formulation is the fundamental fact that the important building blocks of the patient's health are not voxels, but rather tangible biological

structures. Thus, radiation exposure to a given voxel means little, except insofar as it will help kill a tumor or cause damage to healthy tissue. The EUD seeks to develop a score for each organ or tumor based upon the dosages received by each voxel.

One simple score would be to simply take the arithmetic mean dose for a given structure. The problem is that doing so would disregard several important biological features of human physiology. First, though it is rather easy to kill parts of a tumor, because cancerous cells can grow very quickly, it is important that therapy destroy an entire tumor. To that end, then, the biologically relevant measure of radiation exposure is not the mean dose. If it were, then it would be seen as better to irradiate half the tumor at a level of 100 Gray and half at a level of 0 Gray, rather than all of the tumor at a level of 40 Gray. In this case, of course, the part of the tumor receiving 0 Gray would not be affected at all, so the treatment would be ineffective. As a result of this empirical fact, then, the most important determinant of how effective radiation will be in destroying a tumor is the minimum dose a part of that tumor receives. This fact is the reason a_T is usually large and negative.

There are two types of healthy organs: Those that function serially and those that function parallel. Serially functioning organs are ones for whom a failure in just one small part could be catastrophic. Such organs include the brain and spinal cord. Because the goal in these organs is to minimize the probability of any failure at all, the goal is then to minimize the maximum dose to these organs. Alternatively, the goal is to minimize average dose, as well as maximize dose homogeneity. As a result of this requirement, the parameter a_i is large and positive in such organs.

Parallel functioning organs, such as the liver and lungs, behave in a manner different than both tumors and serially functioning organs. In parallel functioning organs, losing part of the organ is clinically unimportant, as long as enough remains to allow proper functioning. As a result, the clinically relevant measure of exposure is close to the arithmetic mean. Thus, for such organs a_i are usually small and greater than 1.

Regardless of the type of structure, the EUD has a simple interpretation. It is simply the

dose level that, if applied equally to each voxel in a structure, would give the same medical result as the inputted dose distribution. The only problem is that it remains to be shown that this medical equivalence exists.

The second important strength of the EUD formulation stems from this fact. For tumors, Niemierko has derived the biological equivalence between the EUD and the actual dose (1997). This equivalence is based upon a probabilistic model of tumor elimination. Thus, at least for the tumor, the EUD function is biologically relevant.

The third strength of the EUD formulation is that the objective function is well defined on the whole space of dose distribution. No matter what, increasing the dose to a cancerous voxel or decreasing the dose to a healthy voxel will increase the value of the objective function. This fact allows the computer to achieve plans better than the oncologist might even ask for. These last three strengths effectively address the fundamental problem of the quadratic objective function's rejection of superior solutions.

Finally, the EUD approach follows the two-tiered algorithm I will advocate in a later section. Essentially, this approach demands that the primary objective function be built from a secondary objective function representing the health of each biological structure.

3.3.3 Weaknesses

Unfortunately, the EUD function is not without weaknesses. First of all, Amols and Ling claim that the EUD functions for healthy organs do not have a medical basis (2002). In particular, they point to the fact that all of the different biological structures are algebraically equivalent and claim that simply generalizing the EUD function may not be sufficient. This criticism most likely misses the point, as it is probably a good thing that the EUD function is algebraically equivalent for different structures.

Secondly, on the level of an individual biological structure, it is hard to know what the a parameter truly means. While the meaning is clear for certain special values, and it is true that the EUD will count the highest values the most as a increases and the lowest values

the most as a decreases, the degree and mechanism by which this occurs is unclear. For example, it is hard to know what the difference between setting a_T equal to -10 and setting it equal to -20 are, other than the fact that the latter places more of an emphasis on dose homogeneity. This criticism is also rather minimal.

The third criticism is much more severe. Wu *et al* claim that the P_T and P_i parameters correspond to prescriptions (2002). To that effect, their formulation makes sense, as their objective function is convex when the prescription is not met, and concave when it is met. However, their claim that α_T and α_i represent importance factors is completely false. To see this, we consider a very simply IMRT problem.

Suppose there are two voxels. One is the tumor and one is healthy tissue. Furthermore, let us assume that there is just one beamlet. Then, because all doses are linear functions of beamlet intensities, we have the following relationship:

$$Cd_T = d_1, \quad (10)$$

where d_T is the dose to the tumor, d_1 is the dose to the healthy tissue, and C is a positive constant. In this case, the EUD functions are simply the doses themselves, because each structure has only one voxel. Thus the maximization problem is given by:

$$\max_{d_T} \frac{1}{1 + \left(\frac{P_T}{d_T}\right)^{\alpha_T}} \times \frac{1}{1 + \left(\frac{d_1}{P_1}\right)^{\alpha_1}}. \quad (11)$$

Using simple algebra, the objective function may be reduced to:

$$\frac{d_T^{\alpha_T} P_1^{\alpha_1}}{d_T^{\alpha_T} P_1^{\alpha_1} + P_T^{\alpha_T} P_1^{\alpha_1} + d_T^{\alpha_T} d_1^{\alpha_1} + P_T^{\alpha_T} d_1^{\alpha_1}}, \quad (12)$$

which becomes:

$$\frac{d_T^{\alpha_T} P_1^{\alpha_1}}{d_T^{\alpha_T} P_1^{\alpha_1} + P_T^{\alpha_T} P_1^{\alpha_1} + C^{\alpha_1} d_T^{\alpha_T} d_T^{\alpha_1} + C^{\alpha_1} P_T^{\alpha_T} d_T^{\alpha_1}}, \quad (13)$$

when the linear relationship between d_T and d_1 is noted. The first-order condition thus gives:

$$0 = \frac{\begin{pmatrix} (d_T^{\alpha_T} P_1^{\alpha_1} + P_T^{\alpha_T} P_1^{\alpha_1} + C^{\alpha_1} d_T^{\alpha_T + \alpha_1} + C^{\alpha_1} P_T^{\alpha_T} d_T^{\alpha_1}) (d_T^{\alpha_T} P_1^{\alpha_1})' \\ -d_T^{\alpha_T} P_1^{\alpha_1} (d_T^{\alpha_T} P_1^{\alpha_1} + P_T^{\alpha_T} P_1^{\alpha_1} + C^{\alpha_1} d_T^{\alpha_T + \alpha_1} + C^{\alpha_1} P_T^{\alpha_T} d_T^{\alpha_1})' \end{pmatrix}}{(d_T^{\alpha_T} P_1^{\alpha_1} + P_T^{\alpha_T} P_1^{\alpha_1} + C^{\alpha_1} d_T^{\alpha_T + \alpha_1} + C^{\alpha_1} P_T^{\alpha_T} d_T^{\alpha_1})^2}, \quad (14)$$

so

$$\begin{aligned} & (d_T^{\alpha_T} P_1^{\alpha_1} + P_T^{\alpha_T} P_1^{\alpha_1} + C^{\alpha_1} d_T^{\alpha_T + \alpha_1} + C^{\alpha_1} P_T^{\alpha_T} d_T^{\alpha_1}) (d_T^{\alpha_T} P_1^{\alpha_1})' \\ &= d_T^{\alpha_T} P_1^{\alpha_1} (d_T^{\alpha_T} P_1^{\alpha_1} + P_T^{\alpha_T} P_1^{\alpha_1} + C^{\alpha_1} d_T^{\alpha_T + \alpha_1} + C^{\alpha_1} P_T^{\alpha_T} d_T^{\alpha_1})'. \end{aligned} \quad (15)$$

Expanding the derivatives gives:

$$\begin{aligned} & \alpha_T P_1^{\alpha_1} (d_T^{2\alpha_T - 1} P_1^{\alpha_1} + d_T^{\alpha_T - 1} P_T^{\alpha_T} P_1^{\alpha_1} + C^{\alpha_1} d_T^{2\alpha_T + \alpha_1 - 1} + C^{\alpha_1} P_T^{\alpha_T} d_T^{\alpha_1 + \alpha_T - 1}) \\ &= P_1^{\alpha_1} (\alpha_T d_T^{2\alpha_T - 1} P_1^{\alpha_1} + (\alpha_T + \alpha_1) C^{\alpha_1} d_T^{2\alpha_T + \alpha_1 - 1} + \alpha_1 C^{\alpha_1} P_T^{\alpha_T} d_T^{\alpha_1 + \alpha_T - 1}), \end{aligned} \quad (16)$$

which simplifies to:

$$\alpha_T \left(\frac{P_1}{C} \right)^{\alpha_1} = \frac{\alpha_1 d_T^{\alpha_T + \alpha_1}}{P_T^{\alpha_T}} + (\alpha_1 - \alpha_T) d_T^{\alpha_1}. \quad (17)$$

Any choice of d_T that maximizes the objective function will satisfy Equation 17. As a result, we can use implicit differentiation on Equation 17 to obtain the derivative of d_T^* with respect to α_T , where d_T^* is the optimized value of d_T^* . Implicit differentiation gives:

$$\left(\frac{P_1}{C} \right)^{\alpha_1} = \frac{\alpha_1 (d_T^*)^{\alpha_T + \alpha_1} \left(\frac{\partial d_T^*}{\partial \alpha_T} \frac{\alpha_T + \alpha_1}{d_T^*} + \ln d_T^* \right) - \alpha_1 (d_T^*)^{\alpha_T + \alpha_1} \ln P_T}{P_T^{\alpha_T}} - (d_T^*)^{\alpha_1} + \alpha_1 (\alpha_1 - \alpha_T) (d_T^*)^{\alpha_1 - 1} \frac{\partial d_T^*}{\partial \alpha_T}. \quad (18)$$

Solving for $\frac{\partial d_T^*}{\partial \alpha_T}$, we obtain:

$$\frac{\partial d_T^*}{\partial \alpha_T} = \frac{P_T^{\alpha_T} \left(\frac{P_1}{C} \right)^{\alpha_1} + P_T^{\alpha_T} (d_T^*)^{\alpha_1} + \alpha_1 (d_T^*)^{\alpha_T + \alpha_1} \ln P_T - \ln d_T^* \alpha_1 (d_T^*)^{\alpha_T + \alpha_1}}{P_T^{\alpha_T} \alpha_1 (\alpha_1 - \alpha_T) (d_T^*)^{\alpha_1 - 1} + \alpha_1 (\alpha_T + \alpha_1) (d_T^*)^{\alpha_T + \alpha_1 - 1}}. \quad (19)$$

This function can have support over both positive and negative values. For example, suppose $P_T = 1$ and C , α_1 , and α_T are such that $d_T^* = 1$. Then the derivative is given as:

$$\frac{\partial d_T^*}{\partial \alpha_T} = \frac{\left(\frac{P_1}{C}\right)^{\alpha_1} + 1}{\alpha_1^2}, \quad (20)$$

which is positive. This is as expected, as the the dose to the tumor should increase as its importance is increased.

Now suppose that the following hold: $P_1 = P_T = 10$, $\alpha_T = \alpha_1 = 10$, and $C = \frac{1}{4}$. Then Equation 17 gives $d_T^* = 20$. Then the effect on d_T^* of increasing α_T is given by:

$$\frac{\partial d_T^*}{\partial \alpha_T} = \frac{10^{10}40^{10} + 10^{10}20^{10} + 10 \times 20^{20} \ln \frac{1}{2}}{200 (20)^{19}} \quad (21)$$

$$= \frac{10^{20}}{200 (20)^{19}} \left(2^{20} + 2^{10} + 10 \times 2^{10} \ln \frac{1}{2} \right) \quad (22)$$

$$< \frac{10^{20}}{200 (20)^{19}} (2^{20} + 2^{10} - 5 \times 2^{20}) \quad (23)$$

$$< 0. \quad (24)$$

Thus, increasing α_T actually decreases the radiation exposure to the tumor. This effect is replicable for more complicated voxel arrangements and multiple organs. The reason for this anomaly is the logistic nature of the objective function, which places a high premium on changes right around the prescription, but does not heavily weight other changes. As a result, when an EUD is well within its prescription, it is possible that increasing its importance factor will heighten this effect, thus effectively lowering the importance.

The fact that increasing importance factors can actually decrease importance is nontrivial. Indeed, this fact completely destroys and semblance of transparency in this objective function.

As if that flaw were not deadly enough, the fourth and final flaw puts the nail in the coffin. According to Choi and Deasy (2002), the EUD function is convex when $a \geq 1$ and concave when $a \leq 1$. These are good properties, as the goal is to minimize the EUDs of

healthy tissues, which always have $a \geq 1$ and maximize the EUD of cancerous tissue which always has $a \leq 1$. Unfortunately, the objective function we have been working with is not convex in the beamlet space, so techniques for convex optimization are unavailable in this case. As a result, the accuracy and speed of IMRT is significantly reduced.

To address this problem Choi and Deasy reformulate the IMRT problem as a series of convex optimizations. However, their method is non-deterministic and consequently uncontrollable by practitioners, so we reject their approach.

3.4 The Need for Improvement

Having considered the primary methods used in conventional IMRT, we can now outline the ways in which they need to be improved. As we have already stated, in the real world IMRT planning is an iterative process. In that regard oncologists repeatedly input refined parameters into their planning computers in an attempt to converge on the clinically optimal plan. For this process to work well, several conditions must hold.

First, it must be easy to pick a good set of starting parameters. Doing so ensures that one can find the right locally optimal solution. Furthermore, choosing good initial values decreases the time required for the process. This time saved is indeed an important variable, as doctors are significantly constrained in that regard.

Second, clinical practitioners have stated that they often desire to make minor adjustments to plans. To ensure that such adjustments are possible, the speed of the iterative steps of the process must be quick. In the same vein, doctors must be able to easily understand the parameters they are changing.

Third, in the same vein, solutions must be stable. In other words, small changes in parameters should lead to small, predictable changes in dose distributions. Such is not the case with current methods.

Fourth, as Amols and Ling (2002) argue eloquently, it is not enough to simply choose an objective function, maximize it, and claim it gives the optimal plan. For this statement

to hold true, the objective must truly be equal to the patient's health. Thus, biological relevancy is important.

Fifth, it is important that parameters be so transparent that patients, as well as doctors, can understand them. Amols *et al* have showed that doctors and patients have significantly different preferences over outcomes (1997). In particular, doctors tend to be more risk-averse than patients. In order to allow patients to have more control over their treatment, it must be possible for patients to easily communicate their preferences to their physicians. One way to ensure such is possible is to make the parameters so transparent that anyone can understand them.

The above requirements boil down to three essential features of an objective function for IMRT. First and foremost, in order to gain desirable computational properties, the objective function must be convex. Second, the function must be founded upon biological principles. Third, the inner workings of the objective function must be easily understandable. These first two requirements eliminate the methods currently in use. In the rest of this paper, we develop a new method that satisfies all three requirements.

Before proceeding, we should note that, by some analyses, the objective function is irrelevant. As long as it is reasonably well-conditioned, by changing the parameters of an objective function, a physician could get to any feasible plan he desires. However, if the objective function is not well-formulated, this process would be incredibly time-consuming. Indeed, even when well-behaved objective functions are used, physicians admit that they are usually treating patients with suboptimal plans. Thus, it is our belief that the objective function does indeed matter.

4 A Two-Tiered Approach

In this section we argue for a general, two-tiered approach to the IMRT problem. Essentially, we argue that the objective function should be composed of three steps. First, beamlet

intensities should be mapped to a dose distribution. The dose distribution should then be mapped to individual scores for each structure. These two steps make up the first, intra-structural tier. Finally, these structural scores should be mapped into a real-valued objective function. This step is the second, inter-structural tier. We proceed by first describing some simple properties we want our objective function to have. From these properties, we show that a two-tiered approach is necessary. We then consider the significant advantages of this approach. We finish section by assessing current methods in this light,

Any reasonable objective function for IMRT should satisfy several basic conditions. Regardless of whether or not the standard formulation of the EUD function is correct, any dose distribution for any biological structure can be mapped to another dose distribution, which satisfies two key properties. First, the dose is uniform. Second, from a medical standpoint, the two dose distributions are medically equivalent. Mathematically, this means we can define the following functions:

$$EUD_T^*(\vec{d}_T), \quad (25)$$

which is the biologically correct equivalent uniform dose function for the tumor and

$$EUD_i^*(\vec{d}_i), \quad (26)$$

which is the biologically correct equivalent uniform dose function for organ i . In these functions, \vec{d}_T is a vector containing the exposure to every voxel in the tumor and \vec{d}_i is a vector containing the exposure to every voxel in organ i .

Suppose the objective function for IMRT with n healthy organs is given by:

$$f\left(\vec{d}_T, \vec{d}_1 \dots \vec{d}_i \dots \vec{d}_n, \vec{r}\right), \quad (27)$$

where \vec{r} is the parameter vector. By the existence of the EUD functions, we can create the

following vectors:

$$\vec{d}_T = \left\langle EUD_T^*(\vec{d}_T) \dots EUD_T^*(\vec{d}_T) \right\rangle \quad (28)$$

$$\vec{d}_i = \left\langle EUD_i^*(\vec{d}_i) \dots EUD_i^*(\vec{d}_i) \right\rangle, \quad (29)$$

which are uniform dose vectors that have the same health effects as the original dose vectors.

Since the health effects of exposing a patient to the two different sets of dose vectors are equivalent, the following equality must hold:

$$f\left(\vec{d}_T, \vec{d}_1 \dots \vec{d}_n, \vec{r}\right) = f\left(\vec{d}_T, \vec{d}_1 \dots \vec{d}_n, \vec{r}\right). \quad (30)$$

This equality implies the following fact about the structure of f :

$$f\left(\vec{d}_T, \vec{d}_1 \dots \vec{d}_n, \vec{r}\right) = g\left(EUD_T^*(\vec{d}_T), EUD_1^*(\vec{d}_1) \dots EUD_n^*(\vec{d}_n)\right), \quad (31)$$

where g is a real-valued function of $n + 1$ parameters.

In other words, to know the proper objective function, one must know $n + 2$ functions. The first $n + 1$ functions, the EUD^* s are the intra-structural scoring functions that reduce a dose distribution in a single organ to one real variable representing the effective dose. They populate the first tier. The final function, g , combines these $n + 1$ functions into one, overall health score. Thus g is known as the inter-structural comparison function. This function is the second tier.

This two-tiered approach has two important advantages. First and foremost, this approach does not preclude any of the three vital characteristics described in the last section. Indeed, it is very easy to imagine that f can be convex in this approach. Furthermore, the simple composite nature of this function will make it easy to check for convexity. This approach is derived from biological assertions, so it should be easy to define a biologically sound objective function using this method. Finally, this approach lends itself very well to

comprehension, as the user need only understand how the EUD^* s and g make comparisons.

Second, under certain assumptions, this approach can be made particularly fast. In particular, under the derivation given above, we assumed that the EUD^* s are biologically defined functions. As a result, a medical practitioner should not want to change these functions iteratively in the IMRT process. The only function the oncologist might want to change is g , as that function also represents preferences over risks to various organs. Given this structure, the IMRT planning process should be a two-stage process.

In the first stage, the doctor defines the EUD^* s and gives them to the computer. The computer then calculates all of the efficient dose distributions, where an efficient dose distribution is defined as one such that the EUD^* of one structure cannot be made better without making the EUD^* of another worse. Finally, the computer can interpolate to find an n -dimensional curve over these efficient dose distributions. This process would be very computationally intensive, and would have to be done over night.

In the second stage, the doctor inputs the parameters for the g function. At this point we can make use of the fact that we know all of the efficient allocations. In particular, we know that g can only be optimized if it is on this efficiency frontier. Thus, the optimization is reduced to a much simpler problem, as it is only a maximization or minimization of g , a function of $n + 1$ variables subject to one binding constraint (that the allocation is efficient). Since n is very small, if f is convex this computation can be done more or less instantly. The speed consequently allowed would give physicians ample time to find the best possible plan.

The current methods provide an adequate proxy for EUD^* , the standard EUD function. However, they do not provide an adequate inter-structural comparison function, as the one described by Wu *et al* suffers from many drawback previously described. As a result of these two facts, we believe that it is critical that a better inter-structural comparison function be developed.

5 The Linear Expenditure System

We now turn to developing a new inter-structural scoring function, the linear expenditure system. We first present the mathematical formulation of the LES function. Then we consider the strengths of the function. In particular, we give an intuitive way to understand the LES before we enumerate conditions under which the LES function is the biologically correct function. We also prove the LES function is concave. We then shed light on the non-trivial weakness of the function. Finally we give an example of an optimization based upon the LES function.

5.1 Mathematics

We now turn to defining the LES function. To that end it is necessary to have EUD functions to work with. We denote them EUD_T^* for the tumor and EUD_i^* for organ i . The LES function is given by:

$$g\left(EUD_T^*(\vec{d}_T), EUD_1^*(\vec{d}_1) \dots EUD_n^*(\vec{d}_n)\right) = \left(EUD_T^*(\vec{d}_T) - \overline{EUD}_T\right)^{1-\sum_i \alpha_i} \times \prod_i \left(\overline{EUD}_i - EUD_i^*(\vec{d}_i)\right)^{\alpha_i}. \quad (32)$$

The goal is to maximize this function of $2n + 1$ parameters. The α_i are analogous to importance factors. They must satisfy the following conditions:

$$\alpha_i \geq 0 \forall i \quad (33)$$

$$\sum_i \alpha_i \leq 1. \quad (34)$$

The second condition ensures that the implied importance of the tumor, $1 - \sum_i \alpha_i$ is also nonnegative. As we shall see in the next section, the α_i are better referred to as effort shares than importance factors. Regardless, if the doctor wishes to decrease the dose to organ i , one way to do so is by increasing α_i .

The other parameters are \overline{EUD}_T and \overline{EUD}_i . These parameters are analogous to pre-

scribed doses, but with an important difference. The dose to organ i will never be allowed above \overline{EUD}_i . Similarly, the dose to the tumor will never be below \overline{EUD}_T . Thus, rather than prescriptions, these parameters are more akin to the maximum acceptable values for exposure to organs and a minimal acceptable value for exposure to the tumor. These definitions have a nontrivial implication for the optimization process. In particular, the objective function is only meaningfully defined when the following conditions hold:

$$\overline{EUD}_i \geq EUD_i^*(\vec{d}_i) \forall i \quad (35)$$

$$\overline{EUD}_T \leq EUD_T^*(\vec{d}_T). \quad (36)$$

These conditions imply two important facts about the optimization. First, the doctor must input physically realizable values for \overline{EUD}_T and \overline{EUD}_i . Second, the first guess of the optimization program must lie in the appropriate region. If the parameters are physically realizable, ensuring that this condition holds can be done using a simple distance minimization.

5.2 Strengths

We now enumerate the important strengths of the LES inter-structural comparison function. In particular, we highlight three undeniable strengths of the LES function. First, we give an intuitive motivation for the function. Then we detail the conditions under which this function is biologically correct. Finally, we prove that the LES function admits convex optimization if used properly.

Suppose we know the entire set of efficient dose distributions. As we have stated before, this set gives all dose distributions which possess the following property: no change can be made that improves the EUD for one organ without harming the EUD for another. We know that this set is an n -dimensional manifold with boundary in the $n + 1$ -dimensional space of

possible EUD distributions. Thus this set is defined by the following:

$$P(EUD_T^*, EUD_1^* \dots EUD_n^*) = 0. \quad (37)$$

Because Equation 37 defines the set of efficient dose distributions, we know that several conditions hold:

$$\text{sign} \left(\frac{\partial P}{\partial EUD_i^*} \right) = \text{sign} \left(\frac{\partial P}{\partial EUD_j^*} \right) \forall i, j \quad (38)$$

$$\text{sign} \left(\frac{\partial P}{\partial EUD_i^*} \right) = \text{sign} \left(\frac{\partial P}{\partial EUD_T^*} \right) \forall i. \quad (39)$$

These conditions hold because it is always better to decrease EUD_i^* for any organ and always better to increase EUD_T^* . If these conditions did not hold, it would be possible to decrease exposure to two organs while keeping $P = 0$ or to do the same increase EUD_T^* while decreasing EUD_i^* . If doing either of these things were possible, the efficiency conditions that define the function P would be violated.

We know that any optimal dose distribution will satisfy Equation 37. Thus, it is possible to redefine the optimization of the LES function as follows:

$$\begin{aligned} \max_{EUD_T^*, EUD_i^*} (EUD_T^* - \overline{EUD_T})^{1 - \sum_i \alpha_i} \times \prod_i (\overline{EUD_i} - EUD_i^*)^{\alpha_i} \\ \text{s.t. } P(EUD_T^*, EUD_1^* \dots EUD_n^*) = 0. \end{aligned} \quad (40)$$

This function achieves its maximum when its logarithm achieves its maximum. Using the method of Lagrange multipliers, we can then easily derive the first-order conditions for a maximum:

$$\frac{1 - \sum_i \alpha_i}{EUD_T^* - \overline{EUD_T}} = \lambda \frac{\partial P}{\partial EUD_T^*} \quad (41)$$

$$\frac{-\alpha_i}{\overline{EUD_i} - EUD_i^*} = \lambda \frac{\partial P}{\partial EUD_i^*} \forall i \quad (42)$$

$$P(EUD_T^*, EUD_1^* \dots EUD_n^*) = 0, \quad (43)$$

where λ is the traditional Lagrange multiplier times the following quantity. These first-order conditions can be transformed to the following equations that must hold at an optimum:

$$\overline{EUD}_i - EUD_i^* = \frac{\alpha_i \frac{\partial P}{\partial EUD_j^*}}{\alpha_j \frac{\partial P}{\partial EUD_i^*}} (\overline{EUD}_j - EUD_j^*) \forall i, j \quad (44)$$

$$\overline{EUD}_i - EUD_i^* = \frac{-\alpha_i \frac{\partial P}{\partial EUD_T^*}}{(1 - \sum_i \alpha_i) \frac{\partial P}{\partial EUD_i^*}} (EUD_T^* - \overline{EUD}_T) \forall i. \quad (45)$$

Now note the fact that $\frac{\frac{\partial P}{\partial EUD_j^*}}{\frac{\partial P}{\partial EUD_i^*}}$ is the rate at which reductions to EUD_j^* can be converted to reductions in EUD_i^* . By that we mean that if one decides EUD_j^* is too low relative to EUD_i^* , this is the rate at which EUD_i^* can be decreased by increasing EUD_j^* and holding everything else constant. Thus we can refer to $\frac{\frac{\partial P}{\partial EUD_j^*}}{\frac{\partial P}{\partial EUD_i^*}}$ as the inverse of the cost of decreasing EUD_i^* relative to EUD_j^* .

We can also define a notion of total reduction of EUD to an organ. This is the reduction of exposure from the maximum allowable, in other words, $\overline{EUD}_i - EUD_i^*$. Thus, we can define a very simple interpretation to what the LES function means: The total reduction of exposure to organ i is equal to the total reduction to organ j divided by the cost of reducing EUD_i^* relative to reducing EUD_j^* times the ration of importance factors. Similar logic may be applied to trade-offs between the tumor and healthy tissue.

We can better see the full implications of this logic by considering a simple example. In particular, let us assume that P is a linear function. By Equations 38 and 39, we know we can rewrite Equation 37 as:

$$EUD_T^* = \sum_i C_i EUD_i^*, \quad (46)$$

where $C_i \geq 0$. The C_i s have a simple interpretation: If we increase EUD_i^* by one Gray, that means we get to increase EUD_T^* by C_i Grays.

Regardless, the maximization problem is now given by:

$$\begin{aligned} \max_{EUD_T^*, EUD_i^*} (EUD_T^* - \overline{EUD_T})^{1-\sum_i \alpha_i} \times \prod_i (\overline{EUD_i} - EUD_i^*)^{\alpha_i} \\ \text{s.t. } EUD_T^* = \sum_i C_i EUD_i^*. \end{aligned} \quad (47)$$

The first-order conditions are given by:

$$\frac{1 - \sum_i \alpha_i}{EUD_T^* - \overline{EUD_T}} = -\lambda \quad (48)$$

$$\frac{\alpha_i}{\overline{EUD_i} - EUD_i^*} = \lambda C_i \forall i \quad (49)$$

$$\sum_i C_i EUD_i^* = EUD_T^*. \quad (50)$$

After some algebra, it can be shown that the solution is given by:

$$EUD_T^* = \overline{EUD_T} + \left(1 - \sum_i \alpha_i\right) \left(\sum_i C_i \overline{EUD_i} - \overline{EUD_T}\right) \quad (51)$$

$$EUD_i^* = \overline{EUD_i} - \frac{\alpha_i}{C_i} \left(\sum_i C_i \overline{EUD_i} - \overline{EUD_T}\right) \quad (52)$$

To give intuition as to the meaning of this solution, we now turn to a highly stylized example.

Let us imagine a world in which cancer patients are treated in the following way: Their doctors take them to the Health Swapmeet, rather than the gantry table. There they serve as brokers for their patients, helping them buy clonogen reduction from the market. This clonogen reduction comes in the form of radiation to the tumor—just the tumor, so no other part of the patient's body is irradiated. Of course, that would be too good to be true, so the patients have to finance their purchase this irradiation of the tumor by selling something they have: normal tissue health. Specifically, let's imagine that there are shopkeepers who will pay patients for the right to irradiate their organs—and, as before, nothing else. So to summarize, doctors lead patients around the market, helping them find buyers who are willing to pay high prices to irradiate their healthy organs and then help them find sellers

who will irradiate their tumorous cells for the lowest prices.

Now suppose we normalize the price in this market of increasing EUD_T^* is normalized to 1, and the normalized price of radiation to healthy organs is given by C_i . Since the patient can only buy as much radiation to the tumor as he finances by increasing radiation to his healthy organs, he again faces the constraint:

$$EUD_T^* = \sum_i C_i EUD_i^*. \quad (53)$$

The solution described in Equations 51 and 52 define the appropriate strategy for the patient. In particular, he should follow a two step process. First, he should sell as much radiation to his healthy tissue as he can and buy as little radiation to his tumor as required. In other words, he should set $EUD_T^* = \overline{EUD_T}$ and $EUD_i^* = \overline{EUD_i}$. At this point, the patient should have a surplus of funds given by:

$$B = \sum_i C_i \overline{EUD_i} - \overline{EUD_T}. \quad (54)$$

In the second, step, the consumer should allocate his surplus B in the following manner: He should spend $\alpha_i \times B$ to reduce the radiation to each healthy organ i and $(1 - \sum_i \alpha_i) B$ to increase the radiation to the tumor. Pursuing this strategy gives the solution described by Equations 51 and 52.

The above example gives the motivation for referring to α_i s as effort weights. Essentially the LES function performs a two-stage optimization. First, it ensures that the minimum thresholds for radiation exposure for each structure are met. Second, it expends exactly α_i of its effort to reduce radiation to organ i , for each organ, and uses whatever is left over to increase radiation to the tumor.

In sum, the \overline{EUD} s represent minimal allowable exposures, and the α_i s represent effort shares, so we conclude that the LES function is easily understandable.

We can now turn to deriving the conditions under which the LES function is biolog-

ically correct. The first major assumption required is that the EUD^* s are the proper intra-structural scoring functions. Second, we assume that the \overline{EUD} s are correctly specified. In other words, no treatment plan would ever be acceptable if $EUD_T^* \leq \overline{EUD}_T$ or $EUD_i^* \geq \overline{EUD}_i$. Subject to these two assumptions, we can derive the final condition under which the LES function gives the biologically correct objective function.

A sufficient condition for the LES function to be correct is for the first derivatives to be correct. To that end, we consider the following quantities:

$$\frac{\partial LES}{\partial EUD_i^*} \frac{EUD_i^*}{LES} = -\alpha_i \quad (55)$$

$$\frac{\partial LES}{\partial EUD_T^*} \frac{EUD_T^*}{LES} = 1 - \sum_i \alpha_i, \quad (56)$$

where the function LES denotes the objective function. The quantities above are described as elasticities. The first represents the percentage change in LES for a small percentage change in EUD_i^* . The second represents the percentage change in LES for a small percentage change in EUD_T^* . Interpreting the values given, a sufficient condition for the correctness of the LES function is the following: The LES function is correct if a patient is made α_i percent worse off for every one percent increase to EUD_i^* and $1 - \sum_i \alpha_i$ percent better off for every one percent decrease in EUD_T^* . While we do not have an empirical or theoretical justification for this condition, it seems reasonable. Thus we conclude that it is at least possible that the LES function is biologically correct.

We now demonstrate the final important feature of the LES function: its concavity. This feature will allow convex optimization techniques to be used. We begin by proving a simple lemma on the composition of concave functions.

Lemma 1 *If the functions $f(\vec{x})$ and $g(\vec{x})$ are positive and concave, then the following function is concave:*

$$f(\vec{x})^\alpha g(\vec{x})^{1-\alpha}, \quad (57)$$

where $0 \leq \alpha \leq 1$.

Proof. It is sufficient to show that the second derivative of $f(\vec{x})^\alpha g(\vec{x})^{1-\alpha}$ is negative in an arbitrary direction. We define the first derivatives of $f(\vec{x})$ and $g(\vec{x})$ in this direction by $f'(\vec{x})$ and $g'(\vec{x})$ and the second derivatives by $f''(\vec{x})$ and $g''(\vec{x})$. The second derivative of the function in question in this direction is given by the following:

$$\begin{aligned} & - (1 - \alpha) \alpha f(\vec{x})^{\alpha-2} g(\vec{x})^{1-\alpha} f'(\vec{x})^2 + 2(1 - \alpha) \alpha f(\vec{x})^{\alpha-1} g(\vec{x})^{-\alpha} f'(\vec{x}) g'(\vec{x}) - (1 - \alpha) \alpha f(\vec{x})^\alpha g(\vec{x})^{-1-\alpha} g'(\vec{x})^2 \\ & \quad + \alpha f(\vec{x})^{\alpha-1} g(\vec{x})^{1-\alpha} f''(\vec{x}) + (1 - \alpha) f(\vec{x})^\alpha g(\vec{x})^{-\alpha} g''(\vec{x}) \end{aligned} \quad (58)$$

By simple algebra, we can rearrange the above as follows:

$$- (1 - \alpha) \alpha f(\vec{x})^{\alpha-2} g(\vec{x})^{-1-\alpha} (g(\vec{x}) f'(\vec{x}) - f(\vec{x}) g'(\vec{x}))^2 + \alpha f(\vec{x})^{\alpha-1} g(\vec{x})^{1-\alpha} f''(\vec{x}) + (1 - \alpha) f(\vec{x})^\alpha g(\vec{x})^{-\alpha} g''(\vec{x}). \quad (59)$$

Noting the fact that $f''(\vec{x})$ and $g''(\vec{x})$ are both negative, it is then clear that the above function—and thus the second derivative—are negative. Thus the function $f(\vec{x})^\alpha g(\vec{x})^{1-\alpha}$ is concave. ■

We now prove that the LES function is concave if the EUD_i^* s are convex and EUD_T^* is concave. Because the goal is to maximize this function, this is the desirable property.

Theorem 2 *If EUD_i^* is convex for all i and EUD_T^* is concave, then the following function is concave:*

$$\left(EUD_T^*(\vec{b}) - \overline{EUD_T} \right)^{1-\sum_i \alpha_i} \times \prod_i \left(\overline{EUD_i} - EUD_i^*(\vec{b}) \right)^{\alpha_i}, \quad (60)$$

provided $0 \leq \alpha_i \leq 1$ for all i and $\sum_i \alpha_i \leq 1$.

Proof. We begin by noting that the conditions of the theorem mean that the following functions are concave and positive:

$$EUD_T^*(\vec{b}) - \overline{EUD_T} \quad (61)$$

$$\overline{EUD_i} - EUD_i^*(\vec{b}). \quad (62)$$

Thus, it is sufficient to show that for any set of positive and concave f_i the following must be concave:

$$(f_n)^{1-\sum_i \alpha_i} \times \prod_{i=1}^{n-1} (f_i)^{\alpha_i}. \quad (63)$$

We prove this by induction on n . Suppose $n = 2$. Then this condition holds by the lemma. Now suppose we know it holds for $n - 1$. Then we know the following function is concave:

$$g = \left(\prod_{i=1}^{n-1} (f_i)^{\alpha_i} \right)^{\frac{1}{\sum_i \alpha_i}}. \quad (64)$$

Now consider the function given by:

$$(f_n)^{1-\sum_i \alpha_i} g^{\sum_i \alpha_i}. \quad (65)$$

By the lemma, this function is concave. But

$$(f_n)^{1-\sum_i \alpha_i} g^{\sum_i \alpha_i} = (f_n)^{1-\sum_i \alpha_i} \times \prod_{i=1}^{n-1} (f_i)^{\alpha_i}, \quad (66)$$

so we are finished. ■

Thus the concavity of the LES function holds if the EUD_i^* s are convex and EUD_T^* is concave. In particular, Choi and Deasy have shown that the standard EUD_i functions are indeed convex in both the dosage and beamlet intensity spaces (2002). Likewise, they have shown that the EUD_T function is concave. Thus, if the standard EUD functions are used, the LES function is suitable for convex optimization techniques.

5.3 Weaknesses

The one important weakness of the LES function is the fact that one of its sets of parameters— $\overline{EUD_T}$ and the $\overline{EUD_i}$ s—require great care in their entry. In particular, the LES function fails if these dose levels are unattainable. However, given the fact that the efficient frontier

of possible dose distributions should have been calculated prior to the entry of the LES parameters, we believe that this problem should not pose many difficulties, as it would be easy to check if the inputted parameters were attainable.

5.4 Testing the LES objective function

We now consider an empirical application of the LES objective function. In particular, we test the function on a brain cancer case. The healthy tissue to be protected is the brainstem, the spinal cord, and the parotids. We compare our results to the standard quadratic method.

As we have stated before, any objective function, if manipulated correctly can be made to produce an attainable dose distribution. The relevant question, therefore, is how easily the distribution may be manipulated. To that end, we present two LES optimizations to show how easily the function is to manipulate. The dose value histograms are given in Figures 6 and 7.

For both of the optimizations given, the EUD^* s are based upon the standard EUD functions. The parameters for these functions are chosen from clinically supported variables. The effort shares are equal for all structures. The only difference between the two optimizations is the value of \overline{EUD}_i for the spinal cord, as this value has not been biologically determined yet. The \overline{EUD} s for the other structures are clinically determined.

In Figure 6, \overline{EUD}_i has been set to 40 for the spinal cord. In figure 7, it has been set to 30. The effect of this change is simple. The portion of the spinal cord that was being irradiated above 30 is reduced to below thirty. More or less nothing else changes for the spinal cord. The other structures bear the rest of the burden of this change more or less equally. This is exactly as expected. First, the total reduction amount to the spinal cord from 30 is very similar to that from 40. Much more encouraging than that is the fact that the burden of these shifts is borne equally by the other structures, with the protection of each of these structures reduced by roughly half.

Thus, we conclude that the LES function behaves as expected.

6 n -Part Equivalent Uniform Dose

Though the conventional EUD function is most likely good enough for all practical purposes, we now consider one new alternative, which is essentially a composite of multiple EUD functions. We thus call this new function the n -Part Equivalent Uniform Dose and denote by $nEUD$. We proceed in three steps as usual. First we give the mathematics of the function. Then we consider the strengths of the method. Then we consider the weaknesses.

6.1 Mathematics

The $nEUD$ function is quite simple. It is given by the following:

$$nEUD(\vec{d}) = \prod_{k=1}^n EUD(a_k, \vec{d})^{\beta_k}, \quad (67)$$

where $EUD(a_k, \vec{d})$ is the standard EUD function with parameter a_k . The β_k parameters control how important each different EUD function is. To that end, we require the following:

$$0 \leq \beta_k \leq 1 \forall k \quad (68)$$

$$\sum_k \beta_k = 1. \quad (69)$$

Finally, we require that if the $nEUD$ function is for a tumor the following holds:

$$a_k \leq 1 \forall k. \quad (70)$$

If the function is for healthy tissue we require that the following holds:

$$a_k \geq 1 \forall k. \quad (71)$$

Though the $nEUD$ function can have many permutations, we restrict our analysis to

what we believe the most useful ones are. In particular, we will henceforth define the $nEUD$ function for a tumor as:

$$nEUD = EUD(-\infty, \vec{d})^{\beta_{\min}} \times EUD(1, \vec{d})^{\beta_{mean}} \times EUD(a, \vec{d})^{\beta_a}, \quad (72)$$

where a is the clinically correct EUD parameter. Thus, for a tumor, the $nEUD$ function is a weighted geometric mean of the minimum dose, the average dose, and the biologically relevant equivalent uniform dose.

We define the $nEUD$ function for an organ as:

$$nEUD = EUD(\infty, \vec{d})^{\beta_{\max}} \times EUD(1, \vec{d})^{\beta_{mean}} \times EUD(a, \vec{d})^{\beta_a}, \quad (73)$$

where where a is the clinically correct EUD parameter. Thus, for a tumor, the $nEUD$ function is a weighted geometric mean of the maximum dose, the average dose, and the biologically relevant equivalent uniform dose. Note the symmetry for organs and tumors.

6.2 Strengths

Before detailing the specific strengths of the $nEUD$ formulation, it is important to show that the $nEUD$ function is concave for tumors and convex for normal tissue. We now prove this fact.

Theorem 3 *The $nEUD$ function is concave for tumors and convex for normal tissue if it is given by:*

$$nEUD = EUD(-\infty, \vec{d})^{\beta_{\min}} \times EUD(1, \vec{d})^{\beta_{mean}} \times EUD(a, \vec{d})^{\beta_a}, \quad (74)$$

for the tumor and

$$nEUD = EUD(\infty, \vec{d})^{\beta_{\max}} \times EUD(1, \vec{d})^{\beta_{mean}} \times EUD(a, \vec{d})^{\beta_a}, \quad (75)$$

for normal tissue.

Proof. The proof is trivial for the tumor. Choi and Deasy have shown that all *EUD* functions with parameter a less than 1 are concave (2002). Thus we can apply Lemma 1 once to show that

$$\left(EUD(-\infty, \vec{d})^{\beta_{\min}} \times EUD(1, \vec{d})^{\beta_{\text{mean}}} \right)^{\frac{1}{\beta_{\min} + \beta_{\text{mean}}}} \quad (76)$$

is concave and twice to show that

$$\left(EUD(-\infty, \vec{d})^{\beta_{\min}} \times EUD(1, \vec{d})^{\beta_{\text{mean}}} \right) \times EUD(a, \vec{d})^{\beta_a} \quad (77)$$

is concave.

The proof for normal tissue is more complex. We proceed in three steps. First we show that if $f(\vec{x})$ is convex, then $f(\vec{x})^\alpha g(\vec{x})^{1-\alpha}$ is convex if $g''(\vec{x}) = 0$ where defined and g is convex. As in Lemma 1, the second derivative of $f(\vec{x})^\alpha g(\vec{x})^{1-\alpha}$ is given by:

$$-(1-\alpha)\alpha f(\vec{x})^{\alpha-2} g(\vec{x})^{-1-\alpha} (g(\vec{x})f'(\vec{x}) - f(\vec{x})g'(\vec{x}))^2 + \alpha f(\vec{x})^{\alpha-1} g(\vec{x})^{1-\alpha} f''(\vec{x}) + (1-\alpha) f(\vec{x})^\alpha g(\vec{x})^{-\alpha} g''(\vec{x}). \quad (78)$$

Simple algebra shows that the sign of this function is the sign of the function:

$$-(1-\alpha) f(\vec{x})^{-1} g(\vec{x})^{-2} (g(\vec{x})f'(\vec{x}) - f(\vec{x})g'(\vec{x}))^2 + f''(\vec{x}). \quad (79)$$

This function is at its most negative when $\alpha = 0$. But when $\alpha = 0$, $f(\vec{x})^\alpha g(\vec{x})^{1-\alpha}$ is simply $g(\vec{x})$, which is, by assumption, convex.

Thus, plugging in the standard *EUD* function for $f(\vec{x})$ and the maximum function for $g(\vec{x})$, we can see that

$$\left(EUD(a, \vec{d})^{\beta_a} \times EUD(\infty, \vec{d})^{\beta_{\max}} \right)^{\frac{1}{\beta_{\max} + \beta_a}} \quad (80)$$

is convex. redefining $f(\vec{x})$ equal to $\left(EUD(a, \vec{d})^{\beta_a} \times EUD(\infty, \vec{d})^{\beta_{\max}}\right)^{\frac{1}{\beta_{\max} + \beta_a}}$ and $g(\vec{x})$ equal to the arithmetic mean, we can see that

$$EUD(\infty, \vec{d})^{\beta_{\max}} \times EUD(1, \vec{d})^{\beta_{mean}} \times EUD(a, \vec{d})^{\beta_a} \quad (81)$$

is convex. ■

In addition to having the desired convexity properties, the $nEUD$ has three important and related strengths. First, it allows doctors greater control over treatment plans. Often-times, doctors decide that plans have too many hot or cold spots (extremely high or low exposures). The $nEUD$ function allows the doctor to easily control such hot or cold spots. This aim cannot necessarily be accomplished by simply changing the a parameter, as the conventional EUD function heavily weights all hot or cold spots for large values of a , depending on the sign. Of course, hot or cold voxels are likely to be close to each other, so the additional information provided by considering additional hot or cold voxels other than the extremes is probably low. Likewise, doctors often wish to play with average exposure. The $nEUD$ function allows this as well.

This nature of the greater control provided for doctors is important because it is, in some ways, limited. In the standard formulation, the only way for a doctor to change intra-structural scoring is to change the a parameter. If the optimization is based upon calculating an efficiency frontier, then this must be completely redone. The $nEUD$ function does not suffer from this flaw. In this case, the efficiency frontier can be defined by demanding that none of the component EUD functions of the $nEUD$ can be made better without making another component EUD function (either for the same or for a different structure) worse.

Finally, the $nEUD$ function is easy to understand. The logic from the LES discussion applies here. In particular, the β s are effort shares, so of the total effort expended on a given structure, β_{\min} or β_{\max} corresponds to the percentage of the total effort on that structure to make the minimum or maximum dose better. Similar logic applies to β_{mean} and β_a .

6.3 Weaknesses

Unfortunately, the $nEUD$ formulation suffers from two flaws. First, it is not based on biological fact. Rather, it is essentially based upon making life easier for doctors. As this can consequently lead to better plans, this is a noble aim, but it is still troubling that the function is not biologically based.

Secondly, the inclusion of the minimum and maximum functions introduces discontinuous first derivatives of the objective function. This can lead to corner solutions. However, as it is easy to predict where such corner solutions will occur, we believe that this flaw is not catastrophic.

Because of these two flaws, we remain agnostic on the relative merits of the $nEUD$ function versus the standard EUD .

7 Suggestions for Further Research

We now offer several suggestions for further research. First, we detail some questions that remain theoretically unanswered. Then we consider some empirical problems that remain unsolved.

7.1 Theoretical

In the framework of the two-tiered approach we have advocated, it is important to know the biologically correct EUD^* functions. To that end, it is important to develop theoretically justifiable EUD^* functions. Currently, such a function exists only for tumors. It is important to develop a reasonable model for healthy organs as well. This is the most theoretical topic for research, as currently not even a minimally acceptable justification of the best EUD functions exists.

In a similar vein, it could be possible to develop inter-structural scoring functions other than the LES function. Potential functions should be evaluated theoretically to ensure that

the critical features enumerated for such functions are met.

On a more complicated—and interesting—level, research should be conducted into specific optimization techniques based upon the two-tiered approach. To that end, it is entirely plausible that nuanced numerical methods that take into account specific features of the *EUD* and *LES* functions could improve solutions. The method based upon determining an efficiency frontier is one such method described in this paper that merits further analysis.

Finally, and most generally, theoretical research should be conducted into what to do when doctors are unsure what is best for their patients. To that end, objective functions that include uncertainty over outcomes or even preferences should be researched.

7.2 Empirical

There are also several important empirical avenues for research. First of all, *EUDs* can be empirically determined. To that end, researchers should gather statistics on outcomes for various treatments plans and find what dose distributions are statistically equivalent. Conducting such research will allow the fundamental building block of the two-tiered approach, the *EUD** functions to be more accurate. This, in turn, will allow superior objective functions to be used.

Relatedly, the link between effective dosages structures receive and final patient outcomes should be examined. If statistics are effectively gathered on this subject, then researchers will be able to more effectively formulate inter-structural scoring functions. If the true mapping from the *EUD** space is found, and the *EUD** functions are correct, then maximizing the objective function based upon this mapping will truly maximize patient health.

In a different vein, it would be useful to collect data on the final treatment plans chosen by doctors. Assuming that the medical profession is, on average, correct in determining health outcomes from dose distributions, examining such data could be a different way to determine the proper intra- and inter-structural scoring functions. Armed with this data, it could also be possible for computers to learn what features doctors like in plans. If this is

the case, it could be possible to speed the planning process.

Finally, it would also be useful to conduct surveys on patient preferences, as they might not be to strictly maximize health. If such is the case and we assume that the role of the doctor is to maximize the patient's actual objective function, knowing patient preferences will be necessary in the development of an appropriate inter-structural scoring function.

8 Conclusion

In this paper, we have asked and answered the following question: What objective function should be used for IMRT planning? To address this question we proceeded in several steps. We began by giving reasons why current methods are fatally flawed. Then we outlined three key features that any objective function for IMRT should have. First, the function should be convex. Second, it should be biologically relevant. Third, it should be easily understandable. To give an objective function that satisfies these features, we argued that a two-tiered approach should be used. First, dose distributions should be scored for each biological structure. Then the intra-structural scores should be combined into an inter-structural score. Because no reasonable inter-structural scoring mechanism currently exists in the literature, we then turned to developing the LES function. When the intra-structural scores satisfy certain conditions that all such scores considered in this paper do, this function allows convex optimization. When certain assumptions hold, this function is medically correct. Finally, this function is very easy to understand. We finally turned to considering an alternative to the *EUD* function which is currently the best available. To that end, we developed the *nEUD* function, which has several desirable properties. However, as we could not give a biological justification for this function we remain agnostic as to whether the *nEUD* or *EUD* should be used. We finished by offering suggestions for further research.

It is our hope that the techniques developed in this paper will be implemented in IMRT to allow more optimal plans for be chosen to treat patients, at less cost to the physicians.

References

- [1] Abramowitz, M. "Elementary analytical methods." *Handbook of Mathematical Functions with Formulas, Graphs, and Mathematical Tables*. ed. A. Stegun (Washington, DC: US GPO: 1964) p. 10.
- [2] Amols, Howard and C Clifton Ling. "EUD but not QED." *Int. J. Radiation Oncology Biol. Phys.* 52 (2002) 1-2.
- [3] Amols, Howard, Marco Zaider, M.K. Hayes, and P.B. Schiff. "Physician/patient-driven risk assignment in radiation oncology: reality or fancy?" *Int. J. Radiation Oncology Biol. Phys.* 38 (1997) 455-461.
- [4] Bortfeld T., J Burkelbach, R Boesecke, and W Schlegel. "Methods of image reconstruction from projections applied to conformal radiotherapy." *Phys. Med. Biol.* 35 (1990) 1423-34.
- [5] Choi, Beong and Joseph O. Deasy. "The generalized equivalent uniform dose function as a basis for intensity-modulated treatment planning." *Phys. Med. Biol.* 47 (2002) 3579-3589.
- [6] Cotrutz, Cristian and Lei Xing. "Using voxel-dependent importance factors for interactive DVH-based dose optimization." *Phys. Med. Biol.* 47 (2002) 1659-1669.
- [7] Mageras, G.S. and R Mohan. "Application of fast simulated annealing to optimization of conformal radiation treatments." *Med. Phys.* 20 (1993) 639-647.
- [8] Niemierko, Andrzej. "Reporting and analyzing dose distributions: A concept of equivalent uniform dose." *Med. Phys.* 24 (1997) 103-110.
- [9] Niemierko, Andrzej. "A generalized concept of equivalent uniform dose (EUD)." *Med. Phys.* 26 (1999) 1100 (abstract).

- [10] Webb, S. "Optimization of conformal radiotherapy dose distributions by simulated annealing." *Phys. Med. Biol.* 34 (1989) 1349-1370.
- [11] Wu, Qiuwen, Radhe Mohan, Andrzej Niemierko, and Rupert Schmidt-Ullrich. "Optimization of intensity-modulated radiotherapy plans based on the equivalent uniform dose." *Int. J. Radiation Oncology Biol. Phys.* 52 (2002) 224-235.
- [12] Xing L. and G.T.Y. Chen. "Iterative algorithms for inverse treatment planning." *Phys. Med. Biol.* 41 (1996) 2107-2103.




Figure 6: Comparing an LES Optimization with a Quadratic Optimization: Dotted DVH lines are for the quadratic case, solid lines are for the LES case. The structures are, from left to right, the brainstem, the parotids, the spinal cord, and the tumor.

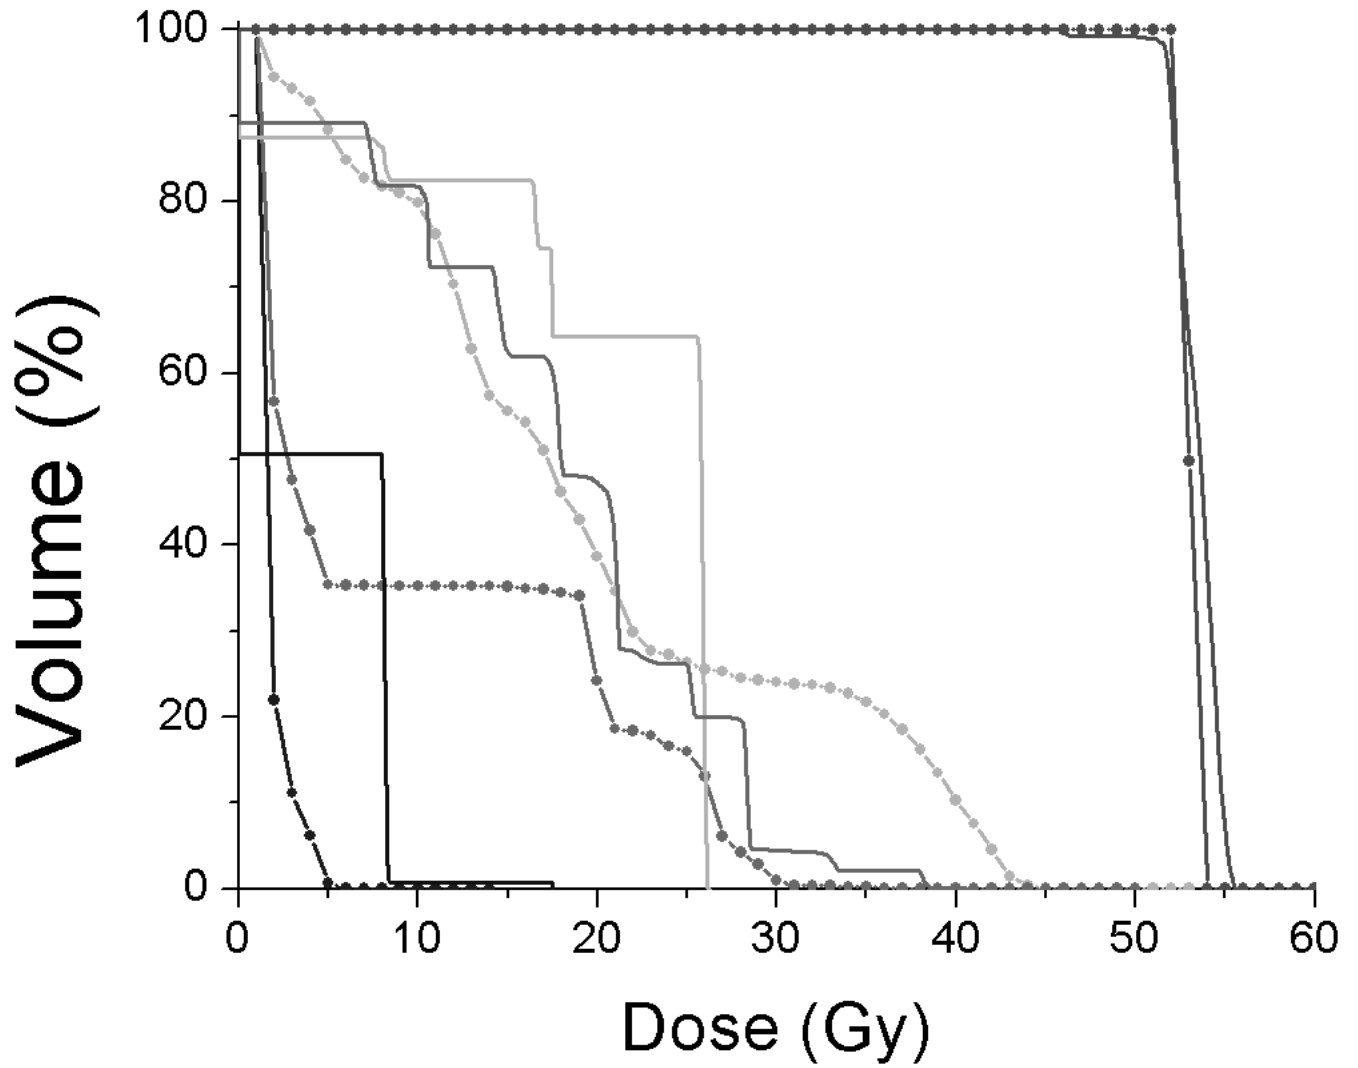


Figure 7: Comparing an LES Optimization with a Quadratic Optimization: Dotted DVH lines are for the quadratic case, solid lines are for the LES case. The structures are, from left to right, the brainstem, the parotids, the spinal cord, and the tumor.